UnitNorm: Rethinking Normalization for Transformers in Time Series

Anonymous authors
Paper under double-blind review

Abstract

Normalization techniques are crucial for enhancing Transformer models' performance and stability in time series analysis tasks, yet we originally identify that traditional methods like batch and layer normalization often lead to issues such as token shift, attention shift, and sparse attention. We propose UnitNorm, a novel normalization approach that scales input vectors by their norms and modulates attention patterns, effectively circumventing these challenges. Grounded in existing normalization frameworks, UnitNorm's effectiveness is demonstrated across diverse time series analysis tasks, including forecasting, classification, and anomaly detection, via a rigorous evaluation on 6 state-of-the-art models and 10 datasets. UnitNorm demonstrates superior performance, particularly where robust attention and contextual understanding are vital, achieving up to a 1.46 MSE decrease in long-term forecasting (ETTh2, PatchTST) and a 4.89% accuracy increase in classification (UWaveGestureLibrary). This work not only calls for a re-evaluation of normalization strategies in time series Transformers but also sets a new direction for enhancing model performance and stability. The source code is available at `https://anonymous.4open.science/r/UnitNorm-5B84`.

1 Introduction

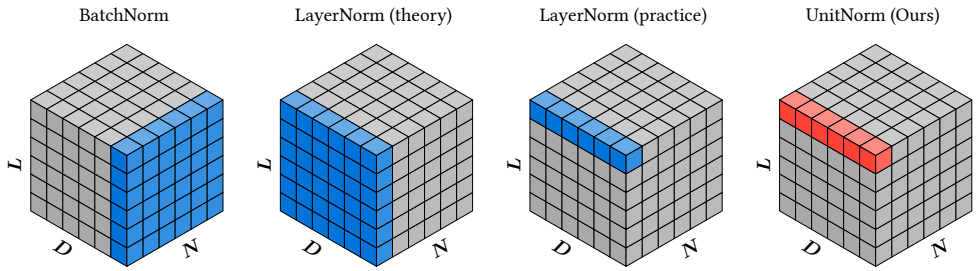

Figure 1: Scheme of different normalization methods applied to batched sequences of **time series tokens** $\mathbf{X} \in \mathbb{R}^{N \times L \times D}$, where N is the batch size, L is the sequence length (or historical window length) and D is the feature dimension (number of variates) of each token vector. The blue sections demonstrate a single slice of the input tensor for computing the mean μ and variance σ^2, while the red section shows a single slice of data for computing the vector norm $\|\mathbf{x}\|$ (see Section C.1).

Transformers have emerged as powerful tools for time series analysis (TSA), offering new capabilities for modeling complex temporal dependencies (Wen et al., 2023). However, adapting them from domains like NLP (Wolf et al., 2020) and CV (Han et al., 2023) presents challenges, especially regarding how normalization interacts with time series characteristics like **seasonality, trends, and autocorrelation**. Central to these models is the representation of data as sequences of tokens, denoted by $\mathbf{X} \in \mathbb{R}^{N \times L \times D}$, where N stands for batch size, L is the sequence length and D represents the dimensionality of each token.

Time series data presents unique challenges that distinguish it from other domains. These include capturing **complex/multiscale temporal dependencies**, **periodic patterns**, and **variable sampling frequencies**.

Unlike text or images, time series often exhibit **strong autocorrelation**, **seasonal/trend components**, and **non-stationarity**. These properties mean data distributions can shift significantly, making normalization easily distorting. While Transformers have shown promise in time series forecasting, classification, and anomaly detection, they were not originally designed with these specific characteristics of time series data in mind. This is particularly evident in how normalization techniques, developed for more stationary data, interact with and potentially disrupt the learning of temporal patterns and periodic structures.

The core mechanism facilitating the Transformers' ability to model complex dependencies is the attention mechanism. It computes a weighted sum of value vectors $\mathbf{V}$, capturing the sequential relationships between tokens through a scalable dot-product operation of queries $\mathbf{Q}$ and keys $\mathbf{K}$: $\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}}\right)\mathbf{V}$. Vaswani et al. (2017)

To mitigate issues during the training process of Transformers related to vanishing or exploding gradients Lubana et al. (2021); Yang & Schoenholz (2017), Layer Normalization (LayerNorm, LN, Ba et al. 2016) plays a significant role and is therefore incorporated at each sub-layer of the architecture (Figure 3)[1]. The LayerNorm operation follows the center-and-scale standardization paradigm, by first centering the means to 0 and then rescaling the variances of the input vectors to 1 such that $\text{LN}(\mathbf{X}) = \frac{\mathbf{X}-\boldsymbol{\mu}}{\sqrt{\boldsymbol{\sigma}^2+\varepsilon}}$, where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the mean and standard deviation of the input vector $\mathbf{X}$, respectively. Ba et al. (2016)

While LayerNorm, compared to other normalization strategies such as batch normalization Ioffe & Szegedy (2015); Shen et al. (2020); Wang et al. (2022), has established itself as the dominant normalization strategy in Transformers, RMSNorm (Zhang & Sennrich, 2019) has emerged as the standard in modern large language models. In contrast, most time series Transformer architectures continue to rely on LayerNorm as their default (Wu et al., 2023; Wen et al., 2023). Dedicated normalization-specific research has mostly focussed on its impact on model convergence Wang et al. (2019), its inner dynamics Wang et al. (2022); Shen et al. (2020) or its location Xiong et al. (2020) within the architecture. On the other hand, few works address normalization's interaction with the attention mechanism Kobayashi et al. (2021) (Section 5), a key challenge in TSA (Section 2) due to attention's dot product.

In this work, we provide a new viewpoint on these challenges by first identifying and formalizing Transformer-specific challenges of normalization techniques, highlighting three key issues. Building on these insights, we introduce UnitNorm, a normalization method designed to address these challenges through a unified architectural parameter that jointly controls attention sparsity and signal magnitude.

Our contributions lie in: 1) We **originally identify** two challenges, namely *token shift* and *attention shift*, and reassess the challenge of *sparse attention* Zhai et al. (2023) in Transformers for time series analysis; 2) We propose UnitNorm, a normalization method that mitigates these issues through a single hyperparameter k jointly governing attention sparsity and signal magnitude, thereby **better preserving crucial temporal information**; 3) We empirically validate the effectiveness of UnitNorm on ten datasets spanning three downstream TSA tasks.

2 Challenges in Normalization

Transformers rely on attention mechanisms to achieve remarkable performance in time series analysis tasks. However, the interplay between attention and normalization methods introduces critical, unaddressed challenges. This paper aims to reveal the complexities of token shift, attention shift, and sparse attention, which arise from such interaction between normalization and the attention mechanism. Our theoretical and empirical analysis demonstrates these challenges are intrinsic to conventional normalization, impacting self-attention in time series Transformers.

Time series data presents distinct challenges for Transformers due to its inherent temporal properties. Unlike text or image data, time series often contain critical periodic patterns, trends, and seasonal components that require a balanced attention distribution to capture effectively. Conventional normalization, effective elsewhere,

[1]The LayerNorm used in Transformers, referred to as LayerNorm (practice), computes the statistics within each token rather than over the whole batch as LayerNorm (theory) does (Figure 1). In this paper, we will refer to the LayerNorm (practice) as LayerNorm if no distinction is made.

can disrupt these temporal relationships, *e.g.*, by altering token vector orientations and obscuring long-range dependencies. This disruption is particularly problematic in applications like forecasting periodic signals, detecting anomalies in regular patterns, or classifying time series based on their temporal characteristics, all tasks relying on accurately interpreting temporal token relationships.

2.1 Pilot Study: Normalization Impact on Capturing Periodicity

To illustrate normalization's impact on capturing periodicity, a key aspect of TSA, we conducted a pilot study on a synthetic two-channel sine wave dataset with varying periods, amplitudes, and Gaussian noise (details and full results in Table A2). We employ PatchTST (Nie et al., 2022) as the base model for this study. Patch-based tokenization has become the dominant paradigm in modern time series Transformers, adopted by various architectures including PatchTST and Crossformer (Wu et al., 2023). Among these, PatchTST represents the most canonical patch-based design, making it the natural choice for a controlled pilot study. Furthermore, its patch-based tokenization makes the effect of normalization on individual token representations particularly transparent.

In this study (Table A2), traditional normalization techniques such as BatchNorm and LayerNorm significantly impair the model's ability to capture periodic patterns, with BatchNorm yielding an MSE of 2.721 on this simple synthetic task. This strongly illustrates how **traditional techniques can distort attention, via mechanisms like token shift and by inducing sparse attention** (detailed in Sections 2.2 to 2.4), thus impairing the capture of vital periodic patterns.

The subsequent sections will now dissect these challenges, token shift (Section 2.2), attention shift (Section 2.3), and sparse attention (Section 2.4), both theoretically and empirically, before detailing UnitNorm's methodology (Section 3).

We will further explore the relationship between normalization and attention by examining a simplified equivalent attention process, with normalization preceding attention (Zhang et al. 2022, Figure 3). This perspective allows for a detailed exploration of how normalization influences the attention scores derived from the query and key vectors. For simplicity, our discussion will center on a singular instance of self-attention within the encoder layer, assuming identical query and key vectors to streamline our analysis. Since both normalization and the linear projections for queries and keys are affine transformations, their order does not affect the theoretical analysis; we assume normalization is applied after projection throughout this section.[2]

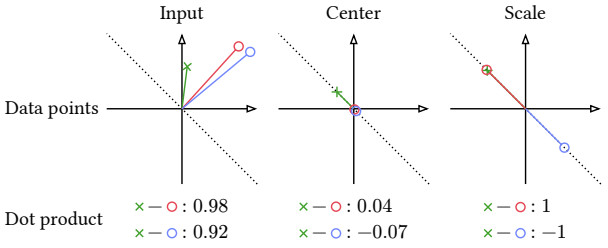

Figure 2: Case of token shift in LayerNorm. The green cross denotes a query vector, the red and blue circles denote two key vectors. Token shift at the centering step causes dot product sign flips; scaling does not.

2.2 Token shift

Previous study Brody et al. (2023) has attributed LayerNorm's efficacy to its center-and-scale operations: centering projects the input vectors to a hyperplane orthogonal to $\mathbb{1}$ vector, and scaling normalizes the vectors to a unit sphere to prevent any token vector being contained in the convex hull of the others. However, this can significantly alter the orientation of input vectors, especially for those that are near parallel to the hyperplane's norm vector $\mathbb{1}$. This impacts dot products, potentially causing sign flips (Figure 2), severely disrupting softmax outputs (Table 1), and **catastrophically altering token importance** (Table 2). By contrast, scale-only normalization such as RMSNorm omits the centering step, preserving vector directions and thereby avoiding token shift. This motivates the design of UnitNorm, which builds upon scale-only normalization and further introduces controlled attention modulation (Section 3). The issue of significant deviations in attention weight distributions caused by token shift will be further explored in Section 2.3. Here, we will focus on the probability of such sign flips.

[2]A detailed justification is provided in Section C.2.

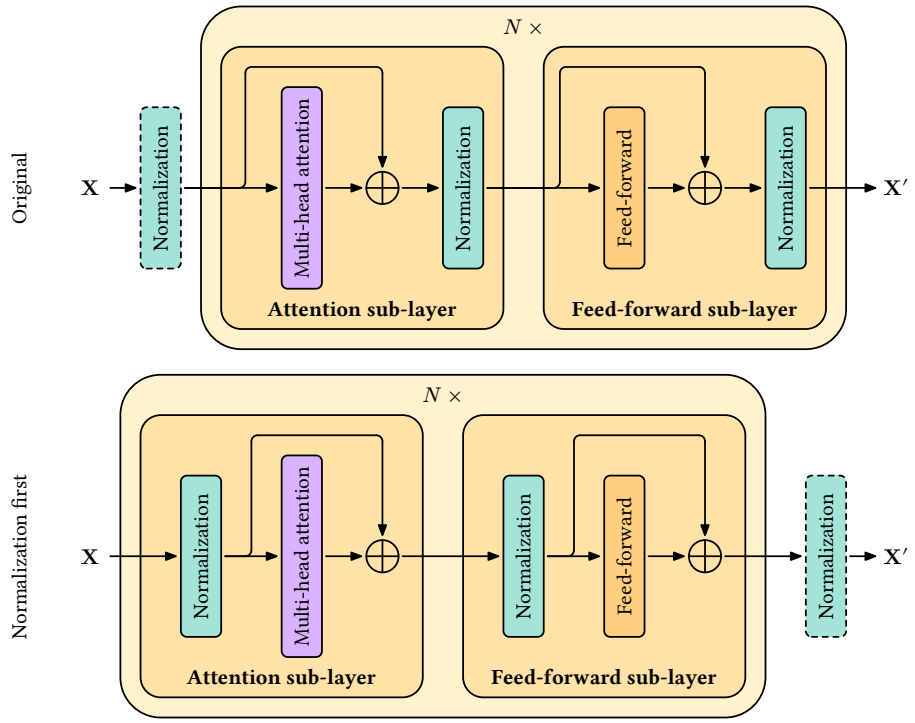

Figure 3: Transformer layer architecture. The original Post-LN architecture is equivalent to a normalization-first sub-layer design when normalizing before attention/FFN projections, enabling a simpler analysis of normalization's impact on attention scores.

Table 1: Effect of input transformations on the softmax function output. Importance order invariant refers to whether the relative importance of the tokens is preserved. Of all possible input transformations, only the reflection transformation will definitely change the importance order of the tokens.

Type	Function	Input			Output			Order invariant?
None	$f : x \mapsto x$	−2	1	3	0.01	0.12	0.88	
Stretch	$f : x \mapsto k \cdot x, k \in \mathbb{R}^+$	−4	2	6	0	0.02	0.98	✓
Translate	$f : x \mapsto x + a, a \in \mathbb{R}$	−1	2	4	0.01	0.12	0.88	✓
Jitter	$f : x \mapsto x + \varepsilon, \varepsilon \sim \mathcal{N}\left(0, \sigma^2\right)$	−2.1	1.1	3	0.01	0.13	0.87	✓/✗
Reflection	$f : x \mapsto -x$	2	−1	−3	0.95	0.05	0.01	✗

The high probability of "center-and-scale" normalization inducing such sign flips is not merely theoretical, as Theorem 2.1 elucidates. Proof in Section B.

Theorem 2.1 (High probability of sign flip due to center operation)**.** *Assume that* $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_x, \operatorname{diag}\left(\boldsymbol{\sigma}_x^2\right))$, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_y, \operatorname{diag}\left(\boldsymbol{\sigma}_y^2\right))$ *are two independent token vectors, with* $\boldsymbol{\mu}_x, \boldsymbol{\mu}_y, \boldsymbol{\sigma}_x, \boldsymbol{\sigma}_y \in \mathbb{R}^D$*. Let* $\tilde{\mathbf{x}} = \frac{\mathbf{x}-\boldsymbol{\mu}_x}{\boldsymbol{\sigma}_x}$ *and* $\tilde{\mathbf{y}} = \frac{\mathbf{y}-\boldsymbol{\mu}_y}{\boldsymbol{\sigma}_y}$ *be the normalized vectors. If*

$$
\begin{aligned}
&|\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y| \geq 12\left(\sqrt{\boldsymbol{\sigma}_x^{2\top}\boldsymbol{\sigma}_y^2} + \|\boldsymbol{\sigma}_x \circ \boldsymbol{\sigma}_y\|_\infty\right) + \\
&5\left(\sqrt{\boldsymbol{\sigma}_y^{2\top}\boldsymbol{\mu}_x^2} + \sqrt{\boldsymbol{\sigma}_x^{2\top}\boldsymbol{\mu}_y^2} + \|\boldsymbol{\sigma}_y \circ |\boldsymbol{\mu}_x|\|_\infty + \|\boldsymbol{\sigma}_x \circ |\boldsymbol{\mu}_y|\|_\infty\right)
\end{aligned}
\tag{1}
$$

Table 2: Effect of normalization on the attention weight distribution based on empirical results (Figures A5 and A6). Normalization methods are categorized by their operation: center-and-scale (LayerNorm, BatchNorm) and scale-only (RMSNorm). Center-and-scale normalization significantly alters the attention weights to an extreme extent, while scale-only normalization largely preserves the original distribution.

Normalization	Chebyshev Distance $\downarrow$	Cosine Similarity $\uparrow$	KL Divergence $\downarrow$
None (original)	/	/	/
Center-and-scale	High	Low	High
Scale-only (RMSNorm)	Low	High	Low

then the probability that the signs of $\mathbf{x}^\top \mathbf{y}$ *and* $\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}$ *do not coincide is at least* 40%*, i.e.,*

$$\Pr(\text{sign}\left(\mathbf{x}^\top \mathbf{y}\right) \neq \text{sign}\left(\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}\right)) \geq 0.40. \tag{2}$$

Remark 2.2. Derived from the computational methodologies for the statistics of vectors $\mathbf{x}$ and $\mathbf{y}$ (Section C.1), BatchNorm posits that the mean vectors are the same so that $\boldsymbol{\mu}_x = \boldsymbol{\mu}_y = \boldsymbol{\mu}$, and similarly $\boldsymbol{\sigma}_x^2 = \boldsymbol{\sigma}_y^2 = \boldsymbol{\sigma}^2$, while LayerNorm assumes that the mean and standard deviation are shared across feature dimension: $\boldsymbol{\mu}_x = \mu_x \mathbb{1}, \boldsymbol{\mu}_y = \mu_y \mathbb{1}$ and $\boldsymbol{\sigma}_x^2 = \sigma_x^2 \mathbb{1}, \boldsymbol{\sigma}_y^2 = \sigma_y^2 \mathbb{1}$. Given these assumptions, the condition equation 1 outlined in Theorem 2.1 is satisfied for many token vector distributions. In fact, we show that in the setup of LayerNorm, the condition equation 1 allows for the quotients of token means and standard deviations, i.e., for μ_x/σ_x and μ_y/σ_y, to decay as $\Omega(D^{-1/4})$ while still implying a high sign flip probability, cf. Section A.

Theorem 2.1 highlights how "center-and-scale" normalization can inadvertently alter attention. The potential for such sign flips, demonstrated with significant likelihood, poses a serious risk to the integrity of the attention scores, as it can lead to a complete reordering of the tokens' importance.

2.3 Attention shift

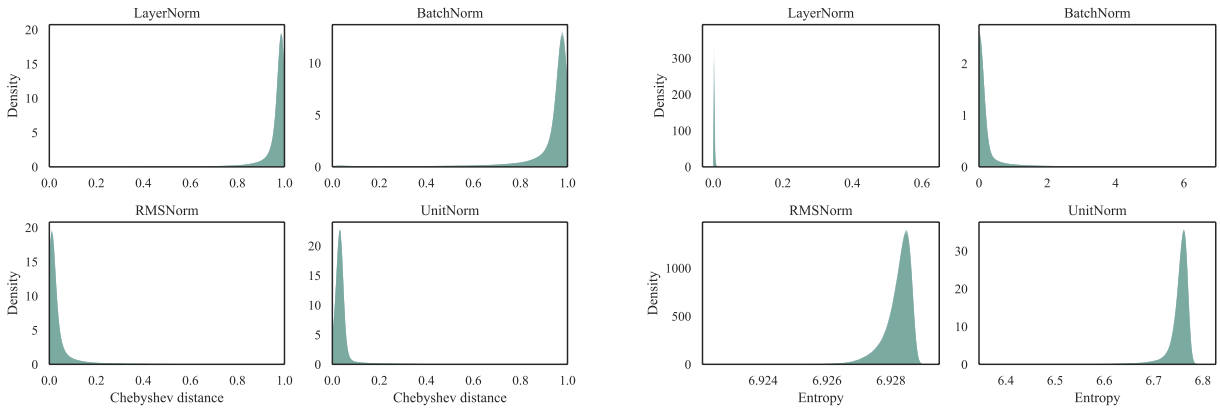

(a) Distribution of Chebyshev distance. RMSNorm preserves the distribution of attention scores; center-and-scale methods significantly alter it.

(b) Distribution of entropy. RMSNorm preserves the high entropy of attention scores; center-and-scale methods result in a collapsed distribution.

Figure 4: Empirical statistics for attention scores after each normalization method. Results from 10 independent experiments are overlaid. UnitNorm is included for reference and introduced in Section 3.

Attention shift represents a critical challenge in Transformer models, directly stemming from the token shift issue. This shift perturbs the relative significance of tokens, leading to discrepancies between attention weights from normalized versus original inputs. To validate the prevalence of attention shift across normalization techniques, we conduct a study utilizing pre-trained Word2Vec embeddings Fares et al. (2017). Our analysis

includes a comparison of batch normalization (BatchNorm, BN, Ioffe & Szegedy 2015), layer normalization (LayerNorm, LN, Ba et al. 2016; Vaswani et al. 2017), root-mean-square layer normalization (RMSNorm, RMSN, Zhang & Sennrich 2019), and our proposed unit normalization (UnitNorm, UN; see Section 3).

Our investigation utilizes sequences of token vectors, $\mathbf{X} \in \mathbb{R}^{N \times L \times D}$, as inputs to the normalization layer, where N is the batch size, L is the sequence length, and D is the dimensionality of each token. The attention scores $\mathbf{A} \in \mathbb{R}^{N \times L \times L}$, given as Equation (3), are computed for 10 independent sets of 32 batches, each containing 1,024 randomly sampled embeddings from a total of 2 million. The primary goal is to assess the impact of normalization on the fidelity of attention scores $\mathbf{A}$ and $\tilde{\mathbf{A}}$, pre- and post-normalization, using the Chebyshev distance as a metric (Table A10).

$$\mathbf{A}_{n,i} = \text{softmax}\left(\frac{\mathbf{X}_{n,i}\mathbf{X}_n^T}{\sqrt{D}}\right) \tag{3}$$

where $\mathbf{A}_{n,i} \in \mathbb{R}^L$ is the attention scores for the i-th anchor token $\mathbf{X}_{n,i}$ to the context sequence $\mathbf{X}_n$ from the n-th batch; $\tilde{\mathbf{A}}$ is computed similarly from normalization output $\tilde{\mathbf{X}}$.

Chebyshev distance distributions (Figure 4(a)) reveal that center-and-scale methods struggle to maintain faithful attention. For them, distances cluster near 1, suggesting profound alteration of attention weights. Conversely, scale-only normalization (RMSNorm) demonstrates a distribution concentrated around zero, indicating minimal disruption to the original attention scores. This aligns with the token-shift analysis: by omitting centering, scale-only methods preserve vector directions and the resulting dot-product relationships.

The empirical evidence underscores a fundamental issue with current normalization practices in Transformers: they compromise attention score fidelity, distorting relational dynamics. This harms interpretability and learning of complex dependencies. As demonstrated in our pilot study with periodic time series data (Table A2), this issue directly affects the model's ability to capture important periodic patterns that are common in time series analysis tasks.

2.4 Sparse attention

The challenge of sparse attention further complicates the normalization landscape in Transformer models. Traditional "center-and-scale" methods often cause undesirable attention concentration (skewing towards single point distributions). This is due to fact that centering removes a degree of freedom from the vectors, and only query that are tightly around the $\mathbb{1}$ vector can produce uniform attention scores Brody et al. (2023). This can be depicted by the entropy of the attention scores $\mathbf{A}_i$: $H(\mathbf{A}_i) = -\sum_{j=1}^{L} \mathbf{A}_{i,j} \log \mathbf{A}_{i,j}$. A higher entropy value suggests a more uniform attention distribution, enabling models capturing periodicity in time series. Conversely, lower entropy, or a trend towards single point distributions, limits its attention to narrow ranges of tokens. While some studies Hyeon-Woo et al. (2022); Zhai et al. (2023) in other fields have shown that Transformer models may benefit from capturing longer-range, denser connections, we will show later that such sparse attention is particularly problematic in TSA tasks.

Analysis of normalization methods through the lens of attention score entropy (Figure 4(b)) reveals a stark contrast between normalization categories. BatchNorm and LayerNorm, as center-and-scale methods, significantly skew attention distributions towards minimal entropy. This condition not only narrows the model's focus but may also precipitate training instability Zhai et al. (2023). In contrast, scale-only normalization (RMSNorm) maintains higher entropy levels, suggesting a more balanced and contextually aware attention mechanism. However, RMSNorm lacks a mechanism to control the degree of attention concentration, producing entropy consistently near the theoretical upper bound $\log L$ regardless of the task. This motivates UnitNorm's k parameter, which enables controlled modulation of attention sparsity within the scale-only framework (Section 3).

2.5 From Static Analysis to Training Dynamics

The analyses in Sections 2.2 to 2.4 examine token-shift, attention-shift, and sparse attention under static assumptions: treating token vectors as fixed samples rather than learnable features that co-evolve with model

parameters. We note that such static analysis carries implications beyond initialization; a body of work demonstrates that initialization-time properties of normalization layers shape training dynamics in lasting ways.

The underlying mechanism is the gradient scaling property of LayerNorm: $\|\partial \mathrm{LN}(\mathbf{x})/\partial \mathbf{x}\| = O(\sqrt{D}/\|\mathbf{x}\|)$ (Xiong et al., 2020). This creates a feedback loop in which large parameter updates inflate downstream input norms, diminishing gradients to earlier layers and destabilizing optimization. The onset of this loop is determined by the initial token geometry and attention sparsity, the very quantities characterized by our static analysis. Huang et al. showed that a static Taylor expansion of the model output at initialization yields an initialization scheme (T-Fixup) enabling 200-layer Transformers to train without explicit normalization layers or warm-up. Liu et al. identified that the amplification effect underlying Post-LN instability is captured by statically computed dependency coefficients between residual branches. Wang et al. derived static bounds on model update magnitudes to stabilize training at 1,000 layers.

In the context of UnitNorm, the analysis of token-shift (Section 2.2) quantifies how center-and-scale normalization can reorder token importance at initialization, while the attention-shift and sparse attention analyses (Sections 2.3 and 2.4) characterize the resulting concentration of attention mass. These effects are not transient: sparse attention at initialization yields sparse gradients that propagate through the LN feedback loop, potentially starving earlier layers of learning signal, particularly damaging for time series tasks that depend on long-range periodic dependencies. By preserving token direction and providing an input-independent entropy lower bound (Theorem 3.2), UnitNorm ensures that optimization begins with gradients flowing evenly across tokens and layers. We note, consistent with He et al., that initialization analysis alone does not fully capture training behavior as normalization continues to shape dynamics throughout optimization, but it provides a necessary foundation for understanding the pathologies we identify.

3 Methodology

To mitigate the challenges identified with traditional normalization methods, we introduce unit normalization (UnitNorm, UN), formulated such that

$$\mathrm{UN}(\mathbf{X}) = D^{\frac{k}{2}} \frac{\mathbf{X}}{\|\mathbf{X}\|_2}. \tag{4}$$

UnitNorm omits centering, diverging from center-and-scale. Like RMSNorm, it scales inputs by their ℓ^2 norm, but further scales by $D^{\frac{k}{2}}$, where k controls attention sparsity.

3.1 Theoretical foundation

UnitNorm is theoretically grounded as a variant of LayerNorm and RMSNorm. Specifically, when taking $k = 1$, UnitNorm is effectively acting as LayerNorm with asserted zero mean, and the RMSNorm can be seen as a special case of UnitNorm with $k = 1$.

This suggests UnitNorm inherits benefits from LayerNorm and RMSNorm (e.g., mitigating vanishing/exploding gradients, stabilizing training). It ensures consistent forward/gradient propagation irrespective of learnable parameter scaling, and scales down gradients to large parameters (Proof: Section B), ensuring stability:

Theorem 3.1 (UnitNorm preserves the gradient to the input and stabilizes the gradient to the learnable parameters)**.** *Given the output of an affine transformation* $\mathbf{x} = \mathbf{W}\mathbf{v} + \mathbf{b}$*, where* $\mathbf{W}$ *and* $\mathbf{b}$ *are learnable parameters. If* $\mathbf{x}' = (\alpha\mathbf{W})\mathbf{v} + (\alpha\mathbf{b})$*, then the output of UnitNorm is unchanged, i.e.,* $\tilde{\mathbf{x}}' = \tilde{\mathbf{x}}$*, while the gradients to loss* $\mathcal{L}$ *are given as follows:*

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial (\alpha \mathbf{W})} &= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{W}} &&= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J}\mathbf{v}^\top \\
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial (\alpha \mathbf{b})} &= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{b}} &&= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J} \\
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial \mathbf{v}} &= \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{v}} &&= \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J}\mathbf{W}^\top
\end{aligned} \tag{5}$$

where $\mathbf{J}$ *is the Jacobian matrix of* $\tilde{\mathbf{x}}$ *w.r.t.*$\mathbf{x}$*.*

3.2 Selection of k values

While learnable k offers flexibility, studies show fixed k values often yield optimal performance Section G. Specifically, values in the range of $0.5 \sim 0.7$ have been found to be particularly effective for time series data with periodic patterns (detailed results are provided in Section G).

This optimal range can be explained by examining the entropy lower bound (ELB) characteristics. With $k \approx 0.5 \sim 0.7$, UnitNorm maintains sufficient attention diversity to capture complex patterns while still allowing for the focus on relevant tokens needed for effective periodicity recognition. This balance is critical for time series tasks where models must simultaneously recognize periodic patterns and focus on specific point of interest. Values of k that are too low may lead to overly dense attention, diluting the model's focus, while values that are too high may result in excessively sparse attention, missing important contextual information.

In practice, we recommend starting with $k = 0.7$ for datasets with strong periodicity, as it typically provides an excellent trade-off between sparse and dense attention distributions. For applications where the optimal k value is uncertain, both fixed values ($k = 0.5, 0.7$) and learnable k implementations can be evaluated to determine the best configuration.

3.3 Overcoming defects

By omitting centering, UnitNorm preserves input vector directions, addressing token/attention shift by maintaining dot product signs (Figure A2). It's a drop-in replacement for LayerNorm/RMSNorm in time series Transformers, needing no structural changes.

Additionally, UnitNorm confronts the sparse attention issue by introducing an entropy lower bound (ELB) for attention scores, modulated by the hyperparameter k (proved in Section B). This feature enables the control of attention patterns, from dense as uniform to sparse as single point, offering versatile attention modeling:

Theorem 3.2 (UnitNorm guarantees an entropy lower bound independent of the input)**.** *For a given set of L, D and a given k, there exists an entropy lower bound (ELB) of the attention scores, i.e.*

$$\mathrm{ELB}(k; L, D) = \log\left(L - 1 + e^d\right) - \frac{de^d}{L - 1 + e^d}, \quad (6)$$

where $d = 2D^{k-\frac{1}{2}}$.

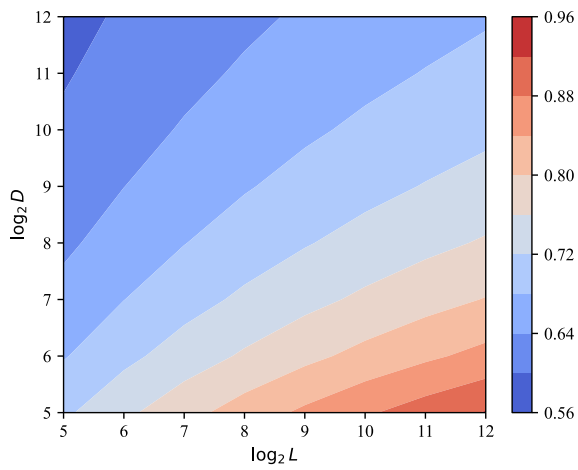

Figure 5: Landscape of k_{50} for different L, D. The k_{50} is the value of k that achieves an ELB of half of the theoretical maximum $\log L$ for a given L, D pair. The landscape of k_{50} is rather smooth and insensitive to the sequence length L, indicating UnitNorm with fixed k can be applied to sequences with variable length without significant change in the attention pattern.

Corollary 3.3 (The ELB of UnitNorm can be any possible value by modulating k)**.** *The ELB is a monotonically decreasing function of k for a given L, D. Furthermore, it is bounded that $\forall k$:*

$$0 < \mathrm{ELB}(k; L, D) < \log L \quad (7)$$

The adaptability of UnitNorm is further exemplified by its applicability across variable sequence lengths, with the entropy lower bound's sensitivity to k remaining relatively consistent irrespective of sequence length (Figure A3), along with the smooth landscape of k_{50}, the value of k that achieves an ELB of $\frac{1}{2}\log L$ for a given L, D pair (Figure 5), particularly with larger D. This property, combined with the option of setting k as a learnable parameter, empowers the model to dynamically adjust its attention pattern, optimizing performance across different scenarios.

3.4 Relationship to Temperature Scaling

UnitNorm exhibits a connection to temperature scaling in the attention mechanism. In the static analysis framework of Section 2, where attention logits retain the standard $1/\sqrt{D}$ scaling and omit learnable projection weights, the attention logits after UnitNorm are $\mathbf{A}_{ij}^{\mathrm{UN}} = D^{k-\frac{1}{2}} \cdot (\mathbf{e}_i^\top \mathbf{e}_j)$, where $\mathbf{e}_i = \mathbf{X}_i / \|\mathbf{X}_i\|$ are unit-norm token vectors. Applying RMSNorm (Zhang & Sennrich, 2019) without learnable affine parameters, followed by a temperature parameter τ in the softmax, yields $\mathbf{A}_{ij}^{\mathrm{RMSN},\tau} = \frac{D^{0.5}}{\tau} \cdot (\mathbf{e}_i^\top \mathbf{e}_j)$. These coincide when $\tau = D^{1-k}$, which interprets k as modulating attention concentration analogously to temperature scaling: $k > 1$ ($\tau < 1$) sharpens attention, while $k < 1$ ($\tau > 1$) yields more uniform distributions.

This equivalence, however, is confined to the pre-softmax attention logits under static assumptions. In the full Transformer, the same normalized representations feed the query, key, *and value* projections, as well as the feed-forward network, where no temperature compensation exists. UnitNorm scales the value projection inputs by $D^{k/2}$ rather than $\sqrt{D}$, and similarly scales the FFN inputs, affecting activation thresholds and residual stream magnitudes. Consequently, k plays a *dual role*: it controls attention sparsity via the ELB (Theorem 3.2) and regulates signal magnitude through the residual stream via the $D^{k/2}$ factor. Rather than requiring separate tuning of temperature and normalization scale, UnitNorm provides a single parameter that jointly governs both attention behavior and feature propagation. A head-to-head ablation (Section J) confirms this empirically: UnitNorm($k = 0.7$) outperforms RMSNorm + $\tau = D^{0.3}$ by +5.67% combined across PatchTST and Crossformer on ETTh2 and Weather (24 runs, 3 seeds per condition). The gap is largest on the attention-sensitive dataset ETTh2 (up to +9.67%), confirming that k's effect extends beyond what temperature scaling alone can capture.

4 Experiments

In our experimental evaluation, UnitNorm is rigorously tested across a spectrum of TSA tasks to illustrate its theoretical advantages in practical applications, including long term forecasting (ETTh1, ETTh2, ECL, Exchange), classification (FaceDetection, Heartbeat, PEMS-SF, UWaveGestureLibrary) and anomaly detection (MSL). We integrate UnitNorm into various Transformer models, namely Crossformer Zhang & Yan (2022), FEDformer Zhou et al. (2022), Informer Zhou et al. (2021), PatchTST Nie et al. (2022), and the vanilla Transformer Vaswani et al. (2017), all with same set of hyperparameter as described in Wu et al. (2023). For comparison, we also include BatchNorm, LayerNorm, RMSNorm, and various settings of UnitNorm (see figure legends). By doing so, we aim to demonstrate its superior ability to address normalization-related challenges, enhancing model performance in these tasks. Detailed experimental settings and full results are provided in Tables A3 to A5 and A7 to A9. Below, we outline the significance of these tasks and the specific benefits UnitNorm brings.

Long-term forecasting: Long-term forecasting represents a significant challenge for Transformer models, primarily due to the difficulty in maintaining periodic pattern recognition over extended sequences Li et al. (2023). The conventional normalization methods often exacerbate the sparse attention problem, hindering the model's capability to capture periodicity. In contrast, UnitNorm excels here, with superior rank and slower error increase over longer horizons across datasets (Figure 6). With a maximum increase of 1.46/0.45 in MSE/MAE on ETTh2, and 1.27/0.36 in MSE/MAE on Exchange at the longest prediction horizon, it substantiates UnitNorm's ability to preserve the attention mechanism's effectiveness, even with increasing prediction horizons, by maintaining balanced attention and avoiding token/attention shift. Our analysis of dataset periodicity (see Appendix H.1) reveals that datasets with stronger periodic patterns, such as ETTh2, particularly benefit from UnitNorm's attention-preserving properties.

Classification: In classification tasks, the key challenge lies in effectively capturing long-range dependencies within sequences Vyas et al. (2022), a task at which Transformers excel. However, the efficacy of this capability can be significantly impacted by the choice of normalization method. UnitNorm, with its unique approach to normalization, has been shown to enhance model performance across multiple datasets, outperforming traditional methods in 3 out of 4 datasets on average (Figure A7), with a significant increase in accuracy of up to 4.90% on UWaveGestureLibrary, 1.95% on Heartbeat and 0.48% on FaceDetection. This underscores

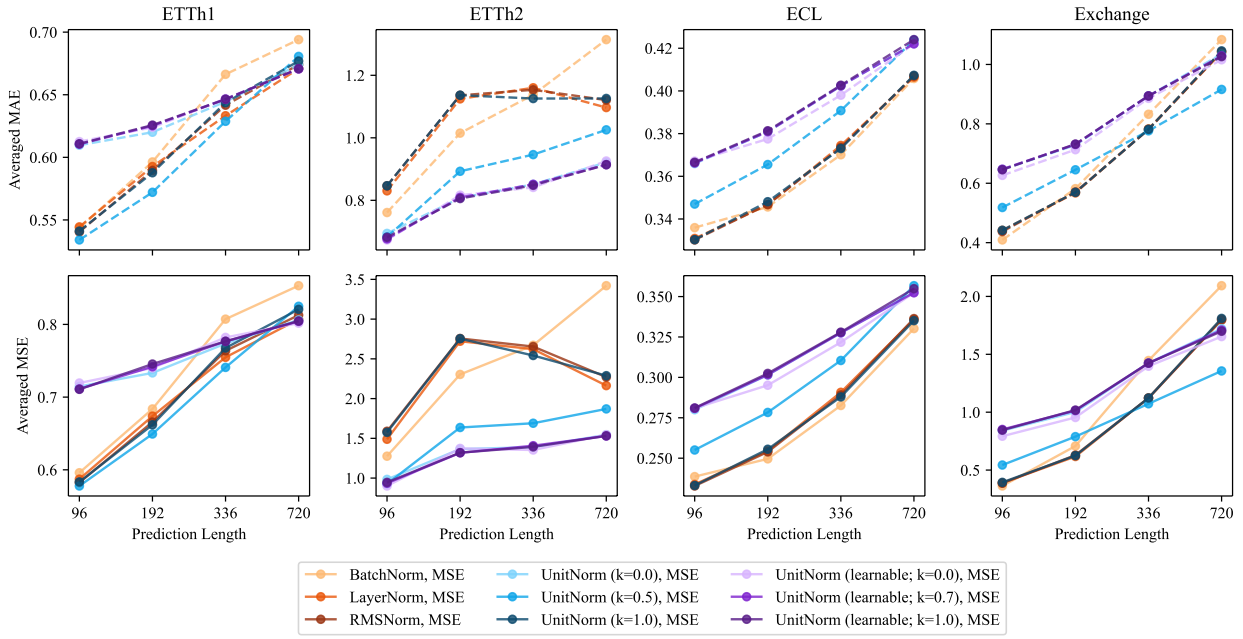

Figure 6: Average rank of normalization methods on the long-term forecasting tasks. X-axis: number of tokens to forecast, Y-axis: average rank over models. Ranks are computed based on the MAE or MSE of each model on each task with different normalization methods (lower is better). UnitNorm and UnitNorm (learnable) achieve better results with the increase of prediction horizon, and have a slower increase in prediction error compared to other normalization methods.

the versatility of UnitNorm in adapting to varied datasets, offering improved accuracy by enabling a more robust, contextually aware attention mechanism.

Anomaly detection: Anomaly detection in time series data demands robust model sensitivity to subtle deviations Haq & Lee (2023); Yang et al. (2023), a requirement often compromised by normalization-induced shifts in attention. The token and attention shift problems, in particular, pose significant challenges in learning stable representations. UnitNorm addresses these challenges head-on, providing a more stable foundation for anomaly detection models to operate on, therefore gaining a maximum of 7.32% in recall, 5.58% in F-score, and 2.81% in precision. Its effectiveness is dominant in all accuracy, recall, precision, and F-score metrics (Figure A8), showcasing its capacity to facilitate more accurate and reliable modeling for anomaly detection.

4.1 Extension to Large-Scale Dataset and Modern Architecture

To validate UnitNorm's generalizability beyond standard benchmarks, we also evaluated its performance on larger, more complex datasets with modern Transformer architectures. Experiments with the Pathformer model on the Solar dataset (137 channels, 52K samples) showed that UnitNorm maintained its effectiveness even at scale, achieving the best MSE and competitive MAE scores compared to other normalization methods. This suggests that UnitNorm's benefits extend to real-world, large-scale applications and remain compatible with newer Transformer architectures. Refer to Appendix H.2 for detailed results and analysis.

5 Discussion

Related work. The development of effective normalization techniques is crucial in the optimization of Transformer models training Wang et al. (2019). Previous research has primarily explored two avenues: the optimal placement of normalization layers, highlighted by the Post-Layer Normalization (Post-LN) and Pre-Layer Normalization (Pre-LN) debate, which impacts training stability Xiong et al. (2020); and the normalization of model parameters, exemplified by Weight Normalization Salimans & Kingma (2016). These methods aim to improve training dynamics by adjusting either the model architecture or the weights.

In contrast, our proposed UnitNorm shifts the focus from placement or parameters to the fundamental role of normalization within the attention mechanism. UnitNorm is distinguished by its emphasis on preserving information of token vectors, a core principle applicable to both Post-LN and Pre-LN configurations. This focus on normalizing layer inputs to maintain vector integrity presents a novel perspective that diverges from prior work centered on architectural adjustments or parameter optimization.

Another line of work has explored adaptive post-normalization strategies. Adaptive LayerNorm (Xu et al.) learns layer-specific scaling and shifting parameters applied after the normalization operation, adapting the degree of normalization per layer. This is complementary to UnitNorm: Adaptive LayerNorm adjusts what happens after normalization, while UnitNorm changes the normalization operation itself. In principle, the two could be combined.

Adopting UnitNorm in Foundation Models. The mainstream of existing Foundation Models, like Large Language Models and Vision-Language Models, already employ RMSNorm in their Transformer architectures. Raschka (2025); Jiang et al. (2023) UnitNorm invites reconsideration of normalization practices in these Transformers, suggesting an alternative that enhance model performance and stability with minimal changes in calculation and number of parameters. Its simplicity and versatility suggest it could be readily adopted across various Transformer applications. The broader impact of UnitNorm lies in its potential to improve the applicability and efficiency of Transformers in fields where precision and model stability are paramount. By addressing specific normalization-related challenges, UnitNorm can make Transformers more suitable for tasks with complex sequential relationships.

Limitations. While UnitNorm represents a significant advancement in normalization techniques for Transformers, several areas still warrant further investigation:

- **Broader Application Scope**: Extending the application of UnitNorm beyond Transformers to other neural network architectures could provide valuable insights into the fundamental principles of normalization across deep learning models.
- **Cross Domain Validation**: Applying UnitNorm across diverse domains and challenging datasets beyond TSA, e.g., NLP Brown et al. (2020); Devlin et al. (2019) and CV Dosovitskiy et al. (2020), will further elucidate its effectiveness and generalizability, providing insights into its broad utility in deep learning.
- **Problem characterization**: Understanding how and what certain dataset characteristics influence the efficacy of normalization methods, including quantitatively assess the presence of token shift, attention shift, and sparse attention in the dynamic interplay of attention mechanisms and normalization during training, can guide the community in selecting appropriate techniques for varied deep learning challenges.

Much as UnitNorm marks a promising advancement in normalization for Transformers, its exploration is far from complete. The limitations identified herein not only highlight the need for further empirical validation across domains but also the potential for refining and extending the methodology to accommodate a wider array of architectures and applications.

Conclusion. UnitNorm challenges prevailing normalization norms in Transformers for TSA, underscoring the need for tailored approaches. By avoiding centering, it directly addresses **token shift, attention shift, and sparse attention**, issues overlooked by traditional methods. Our contribution extends beyond the theoretical introduction of UnitNorm; it includes empirical evidence showcasing its efficacy across various tasks, setting a new precedent for normalization techniques within the Transformer architecture. UnitNorm enables more **stable and faithful representation learning**, paving the way for enhanced Transformer performance and applicability in complex sequential data analysis.

References

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization, July 2016.

Anthony Bagnall, Hoang Anh Dau, Jason Lines, Michael Flynn, James Large, Aaron Bostrom, Paul Southam, and Eamonn Keogh. The UEA multivariate time series classification archive, 2018, October 2018.

Shaked Brody, Uri Alon, and Eran Yahav. On the Expressivity Role of LayerNorm in Transformers' Attention. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association*

for Computational Linguistics: ACL 2023, pp. 14211–14221, Toronto, Canada, July 2023. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *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, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

Sjoerd Dirksen. Tail bounds via generic chaining. *Electronic Journal of Probability*, 20(none):1–29, January 2015.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. October 2020.

Murhaf Fares, Andrey Kutuzov, Stephan Oepen, and Erik Velldal. Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In Jörg Tiedemann and Nina Tahmasebi (eds.), *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pp. 271–276, Gothenburg, Sweden, May 2017. Association for Computational Linguistics.

Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, Zhaohui Yang, Yiman Zhang, and Dacheng Tao. A Survey on Vision Transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):87–110, January 2023. ISSN 1939-3539.

Ijaz Ul Haq and Byung Suk Lee. TransNAS-TSAD: Harnessing Transformers for Multi-Objective Neural Architecture Search in Time Series Anomaly Detection, December 2023.

Bobby He, Lorenzo Noci, Daniele Paliotta, Imanol Schlag, and Thomas Hofmann. Understanding and Minimising Outlier Features in Transformer Training.

Xiao Shi Huang, Felipe Perez, Jimmy Ba, and Maksims Volkovs. Improving Transformer Optimization Through Better Initialization. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 4475–4483. PMLR.

Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom. Detecting Spacecraft Anomalies Using LSTMs and Nonparametric Dynamic Thresholding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, pp. 387–395, New York, NY, USA, July 2018. Association for Computing Machinery. ISBN 978-1-4503-5552-0.

Nam Hyeon-Woo, Kim Yu-Ji, Byeongho Heo, Dongyoon Han, Seong Joon Oh, and Tae-Hyun Oh. Scratching Visual Transformer's Back with Uniform Attention, October 2022.

Sergey Ioffe and Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, pp. 448–456, Lille, France, July 2015. JMLR.org.

Zixuan Jiang, Jiaqi Gu, Hanqing Zhu, and David Pan. Pre-rmsnorm and pre-crmsnorm transformers: Equivalent and efficient pre-LN transformers. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 45777–45793. Curran Associates, Inc., 2023.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. Incorporating Residual and Normalization Layers into Analysis of Masked Language Models. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 4547–4568, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.

Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling Long- and Short-Term Temporal Patterns with Deep Neural Networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18, pp. 95–104, New York, NY, USA, June 2018. Association for Computing Machinery. ISBN 978-1-4503-5657-2.

Yan Li, Xinjiang Lu, Haoyi Xiong, Jian Tang, Jiantao Su, Bo Jin, and Dejing Dou. Towards Long-Term Time-Series Forecasting: Feature, Pattern, and Distribution. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pp. 1611–1624, April 2023.

Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the Difficulty of Training Transformers. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5747–5763. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.463.

Ekdeep S Lubana, Robert Dick, and Hidenori Tanaka. Beyond batchnorm: towards a unified understanding of normalization in deep learning. *Advances in Neural Information Processing Systems*, 34:4778–4791, 2021.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. September 2022.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, number 721, pp. 8026–8037. Curran Associates Inc.

Sebastian Raschka. The Big LLM Architecture Comparison, July 2025.

Tim Salimans and Diederik P. Kingma. Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks, June 2016.

Sheng Shen, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. PowerNorm: rethinking batch normalization in transformers. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML'20*, pp. 8741–8751. JMLR.org, July 2020.

Artur Trindade. ElectricityLoadDiagrams20112014, 2015.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2018. ISBN 978-1-108-41519-4.

Jayant Vyas, Nishit Bhardwaj, Bhumika, and Debasis Das. TransDBC: Transformer for Multivariate Time-Series based Driver Behavior Classification. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, July 2022.

Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, and Furu Wei. DeepNet: Scaling Transformers to 1,000 Layers. 46(10):6761–6774. ISSN 0162-8828. doi: 10.1109/TPAMI.2024.3386927.

Jiaxi Wang, Ji Wu, and Lei Huang. Understanding the Failure of Batch Normalization for Transformers in NLP, October 2022.

Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. Learning Deep Transformer Models for Machine Translation. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1810–1822, Florence, Italy, July 2019. Association for Computational Linguistics.

Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in Time Series: A Survey. volume 6, pp. 6778–6786, August 2023.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-Art Natural Language Processing. In Qun Liu and David Schlangen (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *The Eleventh International Conference on Learning Representations*, 2023.

Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. On layer normalization in the transformer architecture. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML'20*, pp. 10524–10533. JMLR.org, July 2020.

Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. Understanding and Improving Layer Normalization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Chaocheng Yang, Tingyin Wang, and Xuanhui Yan. DDMT: Denoising Diffusion Mask Transformer Models for Multivariate Time Series Anomaly Detection, October 2023.

Greg Yang and Samuel S. Schoenholz. Mean Field Residual Networks: On the Edge of Chaos, December 2017.

Shuangfei Zhai, Tatiana Likhomanenko, Etai Littwin, Dan Busbridge, Jason Ramapuram, Yizhe Zhang, Jiatao Gu, and Joshua M Susskind. Stabilizing Transformer Training by Preventing Attention Entropy Collapse. In *International Conference on Machine Learning*, pp. 40770–40803. PMLR, 2023.

Biao Zhang and Rico Sennrich. Root Mean Square Layer Normalization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Lily Zhang, Veronica Tozzo, John Higgins, and Rajesh Ranganath. Set Norm and Equivariant Skip Connections: Putting the Deep in Deep Sets. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 26559–26574. PMLR, June 2022.

Yunhao Zhang and Junchi Yan. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. September 2022.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):11106–11115, May 2021. ISSN 2374-3468.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. FEDformer: Frequency Enhanced Decomposed Transformer for Long-term Series Forecasting. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 27268–27286. PMLR, July 2022.

A Dimension Dependence of Sign-Flip Probability

We recall that Theorem 2.1 provided a condition for token vector means and variances, condition equation 1, to imply that the sign of the token dot product $\mathbf{x}^\top \mathbf{y}$ is flipped by center-and-scale standardization as in LayerNorm Ba et al. (2016).

In this section, we elucidate the dimension dependence of the required relationship between token means and standard deviations implied by this condition in the case of shared means and standard deviations across feature dimensions, such as implicitly assumed by LayerNorm.

Corollary A.1. *Assume that the mean and variance vectors of independent token vectors* $\mathbf{x}$ *and* $\mathbf{y}$ *satisfy* $\boldsymbol{\mu}_x = \mu_x \mathbb{1}, \boldsymbol{\mu}_y = \mu_y \mathbb{1}$ *and* $\boldsymbol{\sigma}_x^2 = \sigma_x^2 \mathbb{1}, \boldsymbol{\sigma}_y^2 = \sigma_y^2 \mathbb{1}$. *Then the mean-variance condition equation 1 of Theorem 2.1 is satisfied for all* $L \geq 77$ *if*

$$\frac{\mu_x}{\sigma_x} \geq \frac{6}{D^{1/4}} \quad \text{and} \quad \frac{\mu_y}{\sigma_y} \geq \frac{6}{D^{1/4}}, \tag{8}$$

Furthermore, if additionally the independent token vectors are distributed as $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_x, \operatorname{diag}\left(\boldsymbol{\sigma}_x^2\right))$, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_y, \operatorname{diag}\left(\boldsymbol{\sigma}_y^2\right))$, *then the dot product* $\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}$ *of normalized vectors* $\tilde{\mathbf{x}} = \frac{\mathbf{x}-\boldsymbol{\mu}_x}{\boldsymbol{\sigma}_x}$ *and* $\tilde{\mathbf{y}} = \frac{\mathbf{y}-\boldsymbol{\mu}_y}{\boldsymbol{\sigma}_y}$ *attains a sign flip with respect to the original inner products* $\mathbf{x}^\top \mathbf{y}$ *with probability of at least* 40%.

Theorem A.1 implies that for high-dimensional token vectors with $D \gg 1$, it might become easier to satisfy equation 8 given an empirical token distribution, which means that sign flips of dot products after LayerNorm-style normalization might become even more prevalent in that case.

Proof of Theorem A.1. For the case of $\boldsymbol{\mu}_x = \mu_x \mathbb{1}, \boldsymbol{\mu}_y = \mu_y \mathbb{1}$ and $\boldsymbol{\sigma}_x^2 = \sigma_x^2 \mathbb{1}, \boldsymbol{\sigma}_y^2 = \sigma_y^2 \mathbb{1}$, it follows that

$$\begin{aligned}
&12\left(\sqrt{\boldsymbol{\sigma}_x^{2\top}\boldsymbol{\sigma}_y^2} + \|\boldsymbol{\sigma}_x \circ \boldsymbol{\sigma}_y\|_\infty\right) + 5\left(\sqrt{\boldsymbol{\sigma}_y^{2\top}\boldsymbol{\mu}_x^2} + \sqrt{\boldsymbol{\sigma}_x^{2\top}\boldsymbol{\mu}_y^2} + \|\boldsymbol{\sigma}_y \circ |\boldsymbol{\mu}_x|\|_\infty + \|\boldsymbol{\sigma}_x \circ |\boldsymbol{\mu}_y|\|_\infty\right) \\
=&12\left(\sqrt{D\sigma_x^2\sigma_y^2} + \sigma_x\sigma_y\right) + 5\left(\sqrt{D\sigma_y^2\mu_x^2} + \sqrt{D\sigma_x^2\mu_y^2} + \sigma_y|\mu_x| + \sigma_x|\mu_y|\right) \\
\leq&12\left(\sqrt{D\frac{D\mu_x^2\mu_y^2}{36^2}} + \frac{D^{1/4}\mu_x}{6}\frac{D^{1/4}\mu_y}{6}\right) \\
&+ 5\left(\sqrt{D\frac{D^{1/2}\mu_y^2}{36}\mu_x^2} + \sqrt{D\frac{D^{1/2}\mu_x^2}{36}\mu_y^2} + \frac{D^{1/4}\mu_y}{6}|\mu_x| + \frac{D^{1/4}\mu_x}{6}|\mu_y|\right) \\
=&12\left(D\frac{\mu_x\mu_y}{36} + D^{1/2}\frac{\mu_x\mu_y}{36}\right) + \frac{5}{6}\left(\sqrt{DD^{1/2}\mu_y^2\mu_x^2} + \sqrt{DD^{1/2}\mu_x^2\mu_y^2} + D^{1/4}\mu_y|\mu_x| + D^{1/4}\mu_x|\mu_y|\right) \\
=&\mu_x\mu_y\left(\frac{1}{3}D + \frac{1}{3}D^{1/2} + \frac{5}{3}(D^{3/4} + D^{1/4})\right) \leq D\mu_x\mu_y = |\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y|.
\end{aligned}$$

Here, we used in the first inequality the assumption Equation (8) and the fact that $\frac{1}{3}D^{1/2} + \frac{5}{3}(D^{3/4} + D^{1/4}) \leq \frac{2}{3}D$ for $D \geq 77$ in the last inequality. The last assertion of the theorem then follows by application of Theorem 2.1. □

B Proofs

In this section, we detail the proofs of the theoretical results of this paper. In particular, we present the proofs of Theorem 2.1, Theorem 3.1, Theorem 3.2, Theorem 3.3, as well as of auxiliary lemmas.

B.1 Proof of Theorem 2.1

Proof of Theorem 2.1. Let $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_x, \operatorname{diag}\left(\boldsymbol{\sigma}_x^2\right))$ and $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_y, \operatorname{diag}\left(\boldsymbol{\sigma}_y^2\right))$ be independent, and write $\mathbf{x} = (X_1, \quad \ldots, X_D)$ and $\mathbf{y} = (Y_1, \quad \ldots, Y_D)$, respectively.

then we can compute the expectation $\mathbb{E}\left[\mathbf{x}^\top \mathbf{y}\right]$ of the dot product of $\mathbf{x}$ and $\mathbf{y}$ as

$$\begin{aligned}\mathbb{E}\left[\mathbf{x}^\top \mathbf{y}\right] =& \mathbb{E}\left[\sum_{i=1}^{D} X_i Y_i\right] \\ =& \sum_{i=1}^{D} \mathbb{E}\left[X_i Y_i\right] \\ =& \sum_{i=1}^{D} \mathbb{E}\left[X_i\right] \mathbb{E}\left[Y_i\right] \\ =& \sum_{i=1}^{D} (\boldsymbol{\mu}_x)_i (\boldsymbol{\mu}_y)_i \\ =& \boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y.\end{aligned}$$

$$\begin{aligned}\mathrm{Var}\left(\mathbf{x}^\top \mathbf{y}\right) =& \mathbb{E}\left[\left(\mathbf{x}^\top \mathbf{y}\right)^2\right] - \left(\mathbb{E}\left[\mathbf{x}^\top \mathbf{y}\right]\right)^2 \\ =& \mathbb{E}\left[\left(\sum_{i=1}^{D} X_i Y_i\right)^2\right] - \left(\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)^2 \\ =& \sum_{i,j=1}^{D} \mathbb{E}\left[X_i Y_i X_j Y_j\right] - \left(\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)^2 \\ =& \sum_{i,j=1}^{D} \mathbb{E}\left[X_i X_j\right] \mathbb{E}\left[Y_i Y_j\right] - \left(\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)^2\end{aligned} \tag{9}$$

By definition of covariance, we have $\boldsymbol{\sigma}_x = \mathbb{E}\left[\mathbf{x}\mathbf{x}^\top\right] - \boldsymbol{\mu}_x \boldsymbol{\mu}_x^\top$, and here $\boldsymbol{\sigma}_x = \mathrm{diag}(\boldsymbol{\sigma}_x^2)$, then Equation (9) can be simplified as follows:

$$\begin{aligned}\mathrm{Var}\left(\mathbf{x}^\top \mathbf{y}\right) =& \sum_{i=1}^{D} \left(\boldsymbol{\sigma}_x + \boldsymbol{\mu}_x \boldsymbol{\mu}_x^\top\right)_{ij} \left(\boldsymbol{\sigma}_y + \boldsymbol{\mu}_y \boldsymbol{\mu}_y^\top\right)_{ij} - \left(\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)^2 \\ =& \left\langle \boldsymbol{\sigma}_x, \boldsymbol{\sigma}_y \right\rangle_F + \left\langle \boldsymbol{\mu}_x \boldsymbol{\mu}_x^\top, \boldsymbol{\sigma}_y \right\rangle_F + \left\langle \boldsymbol{\sigma}_x, \boldsymbol{\mu}_y \boldsymbol{\mu}_y^\top \right\rangle_F + \left\langle \boldsymbol{\mu}_x \boldsymbol{\mu}_x^\top, \boldsymbol{\mu}_x \boldsymbol{\mu}_x^\top \right\rangle_F - \left(\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)^2 \\ =& \left(\boldsymbol{\sigma}_x^2\right)^\top \left(\boldsymbol{\sigma}_y^2\right) + \left(\boldsymbol{\sigma}_y^2\right)^\top \left(\boldsymbol{\mu}_x^2\right) + \left(\boldsymbol{\sigma}_x^2\right)^\top \left(\boldsymbol{\mu}_y^2\right)\end{aligned} \tag{10}$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product.

Consider now the normalized random vectors $\tilde{\mathbf{x}} = \frac{\mathbf{x} - \boldsymbol{\mu}_x}{\boldsymbol{\sigma}_x}$ and $\tilde{\mathbf{y}} = \frac{\mathbf{y} - \boldsymbol{\mu}_y}{\boldsymbol{\sigma}_y}$. Due to the Gaussianity assumption on $\mathbf{x}$ and $\mathbf{y}$, it follows that the normalized vectors are also Gaussian, and in particular, are distributed as $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$. Plugging the respective mean and variance values into the formulas for the expectation and variance for dot products above, we obtain that

$$\mathbb{E}\left[\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}\right] = 0 \quad \text{and} \quad \mathrm{Var}\left(\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}\right) = D \tag{11}$$

As $\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}}$ is a symmetric random variable, it follows that

$$\Pr\left(\tilde{\mathbf{x}}^\top \tilde{\mathbf{y}} \leq 0\right) = 0.5. \tag{12}$$

Next, due to the definition of the random vectors $\mathbf{x}$ and $\mathbf{y}$, it holds that $\mathbf{x}^\top \mathbf{y} = \sum_{i=1}^{D} X_i Y_i$, where $X_i \sim \mathcal{N}((\boldsymbol{\mu}_x)_i, (\boldsymbol{\sigma}_x)_i^2)$ and $Y_i \sim \mathcal{N}((\boldsymbol{\mu}_y)_i, (\boldsymbol{\sigma}_y)_i^2)$ are independent normal random variables. Going forward, we will use the ψ_1-Orlicz norm

$$\|X\|_{\psi_1} := \inf\{t > 0 : \mathbb{E}[\exp(|X|/t)] \leq 2\}, \tag{13}$$

cf. Definition 2.7.5 of Vershynin (2018). We call a random variable for which $\|\cdot\|_{\psi_1}$ is finite sub-exponential, following, e.g., Vershynin (2018).

Define now $Z_i := X_iY_i - (\boldsymbol{\mu}_x)_i(\boldsymbol{\mu}_y)_i$. We observe that

$$Z_i = X_iY_i - (\boldsymbol{\mu}_x)_i(\boldsymbol{\mu}_y)_i = X_i(Y_i - (\boldsymbol{\mu}_y)_i) + (X_i - (\boldsymbol{\mu}_x)_i)(\boldsymbol{\mu}_y)_i = Z_i^{(1)} + Z_i^{(2)}$$

with $Z_i^{(1)} := X_i(Y_i - (\boldsymbol{\mu}_y)_i)$ and $Z_i^{(2)} := (X_i - (\boldsymbol{\mu}_x)_i)(\boldsymbol{\mu}_y)_i$. To bound the ψ_1-norm of Z_i, we bound this norm for $Z_i^{(1)}$ and $Z_i^{(2)}$ separately.

Indeed, due to Lemma 2.7.7 of Vershynin (2018), it holds that

$$\|Z_i^{(1)}\|_{\psi_1} \leq \|X_i\|_{\psi_2}\|Y_i - (\boldsymbol{\mu}_y)_i\|_{\psi_2},$$

where

$$\|X\|_{\psi_2} := \inf\{t > 0 : \mathbb{E}[\exp(X^2/t^2)] \leq 2\}, \tag{14}$$

is the ψ_2-Orlicz norm Vershynin (2018) characterizing sub-Gaussian random variables X. From Lemma B.2, it follows therefore that

$$\|Z_i^{(1)}\|_{\psi_1} \leq \max\left(2(\boldsymbol{\sigma}_x)_i, \sqrt{\frac{(\boldsymbol{\mu}_x)_i^2}{\log 2} + (\boldsymbol{\sigma}_x)_i^2}\right)\sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_y)_i.$$

For the second part, since $\|\cdot\|_{\psi_2}$ is a norm, we estimate that

$$\|Z_i^{(2)}\|_{\psi_1} \leq \|X_i - (\boldsymbol{\mu}_x)_i\|_{\psi_2}\|(\boldsymbol{\mu}_y)_i\|_{\psi_2} \leq \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_x)_i\|(\boldsymbol{\mu}_y)_i\|_{\psi_2} \leq \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_x)_i\frac{|(\boldsymbol{\mu}_y)_i|}{\sqrt{\log 2}},$$

where we used again Lemma 2.7.7 and (2.17) of Vershynin (2018) in the first and last inequality, respectively, and Lemma B.2 in the second inequality.

From this, it follows that

$$\begin{aligned}
\|Z_i\|_{\psi_1} \leq \|Z_i^{(1)}\|_{\psi_1} + \|Z_i^{(2)}\|_{\psi_1} &\leq \max\left(2(\boldsymbol{\sigma}_x)_i, \sqrt{\frac{(\boldsymbol{\mu}_x)_i^2}{\log 2} + (\boldsymbol{\sigma}_x)_i^2}\right)\sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_y)_i + \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_x)_i\frac{|(\boldsymbol{\mu}_y)_i|}{\sqrt{\log 2}} \\
&\leq \left(2(\boldsymbol{\sigma}_x)_i + \sqrt{\frac{(\boldsymbol{\mu}_x)_i^2}{\log 2} + (\boldsymbol{\sigma}_x)_i^2}\right)\sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_y)_i + \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_x)_i\frac{|(\boldsymbol{\mu}_y)_i|}{\sqrt{\log 2}} \\
&\leq 2\sqrt{6}(\boldsymbol{\sigma}_x)_i(\boldsymbol{\sigma}_y)_i + \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_y)_i\frac{|(\boldsymbol{\mu}_x)_i|}{\sqrt{\log 2}} + \sqrt{\frac{8}{3}}(\boldsymbol{\sigma}_x)_i\frac{|(\boldsymbol{\mu}_y)_i|}{\sqrt{\log 2}},
\end{aligned} \tag{15}$$

using that $\sqrt{a^2+b^2} \leq a + b$ for any non-negative $a, b \geq 0$ in the last inequality. We next establish a lower bound on the probability of a sign flip through normalization, i.e., for $\Pr(\text{sign}\left(\mathbf{x}^\top\mathbf{y}\right) \neq \text{sign}\left(\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}}\right))$. Assuming without loss of generality that $|\boldsymbol{\mu}_x^\top\boldsymbol{\mu}_y| = \boldsymbol{\mu}_x^\top\boldsymbol{\mu}_y$, we observe that

$$\begin{aligned}
\Pr\left(\text{sign}\left(\mathbf{x}^\top\mathbf{y}\right) \neq \text{sign}\left(\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}}\right)\right) &= \Pr\left((\mathbf{x}^\top\mathbf{y} > 0) \wedge (\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} < 0)\right) + \Pr\left((\mathbf{x}^\top\mathbf{y} < 0) \wedge (\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} > 0)\right) \\
&\geq \Pr\left((\mathbf{x}^\top\mathbf{y} > 0) \wedge (\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} < 0)\right).
\end{aligned}$$

Furthermore, since the distribution of the normalized vectors $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ is symmetric, the same holds true for the dot product $\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}}$, which implies that

$$\begin{aligned}
\Pr\left((\mathbf{x}^\top\mathbf{y} > 0) \wedge (\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} < 0)\right) &= 1 - \Pr\left((\mathbf{x}^\top\mathbf{y} \leq 0) \vee (\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} \geq 0)\right) \geq 1 - \Pr\left(\mathbf{x}^\top\mathbf{y} \leq 0\right) - \Pr\left(\tilde{\mathbf{x}}^\top\tilde{\mathbf{y}} \geq 0\right) \\
&\geq 1 - 0.5 - \Pr\left(\mathbf{x}^\top\mathbf{y} \leq 0\right) = 0.5 - \Pr\left(\mathbf{x}^\top\mathbf{y} \leq 0\right).
\end{aligned}$$

It remains to show that

$$\Pr\left(\mathbf{x}^\top\mathbf{y} \leq 0\right) \leq 0.1. \tag{16}$$

To establish this, we see that

$$\Pr\left(\mathbf{x}^\top \mathbf{y} \leq 0\right) = \Pr\left(\mathbf{x}^\top \mathbf{y} - \mathbb{E}\left[\mathbf{x}^\top \mathbf{y}\right] \leq -\mathbb{E}\left[\mathbf{x}^\top \mathbf{y}\right]\right) = \Pr\left(\mathbf{x}^\top \mathbf{y} - \boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y \leq -\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)$$
$$= \Pr\left(\sum_{i=1}^{D} Z_i \leq -\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right)$$

with the random variables Z_i defined above. Using the triangle inequality of the ℓ_2-norm, it follows from equation 15 that

$$\sqrt{\sum_{i=1}^{D} \|Z_i\|_{\psi_1}^2} \leq 2\sqrt{6}\sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\sigma}_y^2} + \sqrt{\frac{8}{3\log 2}}\left(\sqrt{(\boldsymbol{\sigma}_y^2)^\top \boldsymbol{\mu}_x^2} + \sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\mu}_y^2}\right)$$

and that

$$\max_{i=1}^{D} \|Z_i\|_{\psi_1} \leq 2\sqrt{6}\|\boldsymbol{\sigma}_x \circ \boldsymbol{\sigma}_y\|_\infty + \sqrt{\frac{8}{3\log 2}}\left(\|\boldsymbol{\sigma}_y \circ |(\boldsymbol{\mu}_x)|\|_\infty + \|\boldsymbol{\sigma}_x \circ |(\boldsymbol{\mu}_y)|\|_\infty\right),$$

which implies that

$$\begin{aligned}
&\sqrt{2\sum_{i=1}^{D} \|Z_i\|_{\psi_1}^2}\sqrt{\log(10)} + \max_{i=1}^{D} \|Z_i\|_{\psi_1} \log(10) \\
&\leq 4\sqrt{3\log(10)}\sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\sigma}_y^2} + \frac{4\sqrt{\log(10)}}{\sqrt{3\log(2)}}\left(\sqrt{(\boldsymbol{\sigma}_y^2)^\top \boldsymbol{\mu}_x^2} + \sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\mu}_y^2}\right) \\
&+ 2\sqrt{6}\log(10)\|\boldsymbol{\sigma}_x \circ \boldsymbol{\sigma}_y\|_\infty + \sqrt{\frac{8}{3\log 2}}\log(10)\left(\|\boldsymbol{\sigma}_y \circ |(\boldsymbol{\mu}_x)|\|_\infty + \|\boldsymbol{\sigma}_x \circ |(\boldsymbol{\mu}_y)|\|_\infty\right) \\
&\leq 12\left(\sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\sigma}_y^2} + \|\boldsymbol{\sigma}_x \circ \boldsymbol{\sigma}_y\|_\infty\right) + 5\left(\sqrt{(\boldsymbol{\sigma}_y^2)^\top \boldsymbol{\mu}_x^2} + \sqrt{(\boldsymbol{\sigma}_x^2)^\top \boldsymbol{\mu}_y^2} + \|\boldsymbol{\sigma}_y \circ |(\boldsymbol{\mu}_x)|\|_\infty + \|\boldsymbol{\sigma}_x \circ |(\boldsymbol{\mu}_y)|\|_\infty\right) \\
&\leq |\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y|,
\end{aligned}$$

using the assumption equation 1 in the last inequality. With this inequality, we can use the fact that the Z_i are independent mean-zero sub-exponential random variables and Bernstein's inequality as stated in Lemma B.1 to conclude that

$$\Pr\left(\sum_{i=1}^{D} Z_i \leq -\boldsymbol{\mu}_x^\top \boldsymbol{\mu}_y\right) \leq \Pr\left(\sum_{i=1}^{D} Z_i \leq -\left(\sqrt{2\sum_{i=1}^{D} \|Z_i\|_{\psi_1}^2}\sqrt{\log(10)} + \max_{i=1}^{D} \|Z_i\|_{\psi_1} \log(10)\right)\right)$$
$$\leq \exp(-\log(10)) = 0.1.$$

This establishes equation 16, which concludes the proof. □

Lemma B.1 (Bernstein's Inequality, cf. Lemma 5.1 of Dirksen (2015) and the Bernstein inequality for sub-exponential random variables in Vershynin (2018))**.** *Let $Z_1, \ldots Z_D$ be independent mean-zero sub-exponential random variables. Then for every $t \geq 0$,*

$$\Pr\left(\sum_{i=1}^{D} Z_i \leq -\left(\sqrt{2\sum_{i=1}^{D} \|Z_i\|_{\psi_1}^2}\sqrt{t} + \max_{i=1}^{D} \|Z_i\|_{\psi_1} t\right)\right) \leq \exp(-t).$$

Lemma B.2 (Bounds on ψ_2-norm of Gaussians)**.** *1. If $X \sim \mathcal{N}(0, \sigma^2)$ is a centered Gaussian random variable with variance σ^2, then its ψ_2-norm equation 14 satisfies*

$$\|X\|_{\psi_2} \leq \sqrt{\frac{8}{3}}\sigma,$$

cf. Exercise 2.40(d) of Vershynin (2018).

2. If $X \sim \mathcal{N}(\mu, \sigma^2)$ is a Gaussian random variable with mean μ and variance σ^2, then its ψ_2-norm equation 14 satisfies

$$\|X\|_{\psi_2} \leq \max\left(2\sigma, \sqrt{\frac{\mu^2}{\log 2} + \sigma^2}\right),$$

which follows from the triangle inequality and the bound on the centered case.

B.2 Proof of Theorem 3.1

Proof of Theorem 3.1. Given the output of an affine transformation $\mathbf{x} = \mathbf{W}\mathbf{v} + \mathbf{b}$, where $\mathbf{W}$ and $\mathbf{b}$ are learnable parameters. If $\mathbf{x}' = (\alpha\mathbf{W})\mathbf{v} + (\alpha\mathbf{b})$, then the output of UnitNorm is unchanged, *i.e.*, $\tilde{\mathbf{x}}' = \tilde{\mathbf{x}}$, while the gradients to loss $\mathcal{L}$ are given as follows:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{W})} &= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{W}} &&= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J}\mathbf{v}^\top \\
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{b})} &= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{b}} &&= \frac{1}{\alpha} \cdot \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J} \\
\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} \cdot \frac{\partial \tilde{\mathbf{x}}'}{\partial \mathbf{v}} &= \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{v}} &&= \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}} \cdot \mathbf{J}\mathbf{W}^\top
\end{aligned} \tag{17}$$

Proof: First we will show $\tilde{\mathbf{x}}' = \tilde{\mathbf{x}}$, for which we have:

$$\begin{aligned}
\tilde{\mathbf{x}}' =& D^{\frac{k}{2}} \frac{\mathbf{x}'}{\|\mathbf{x}'\|} \\
=& D^{\frac{k}{2}} \frac{\alpha\mathbf{x}}{\alpha\|\mathbf{x}\|} \\
=& D^{\frac{k}{2}} \frac{\mathbf{x}}{\|\mathbf{x}\|} \\
=& \tilde{\mathbf{x}}
\end{aligned} \tag{18}$$

And thus for the gradients to loss $\mathcal{L}$, we have $\frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}'} = \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{x}}}$. Also, for the Jacobian matrix $\mathbf{J}$ of $\tilde{\mathbf{x}}$ *w.r.t.*$\mathbf{x}$, we have

$$\begin{aligned}
\mathbf{J} =& \frac{\partial D^{\frac{k}{2}} \frac{\mathbf{x}}{\|\mathbf{x}\|}}{\partial \mathbf{x}} \\
=& D^{\frac{k}{2}} \left(\frac{\mathbf{I}}{\|\mathbf{x}\|} - \frac{\mathbf{x}\mathbf{x}^\top}{\|\mathbf{x}\|^3} \right)
\end{aligned} \tag{19}$$

And the Jacobian matrix of $\tilde{\mathbf{x}}'$ *w.r.t.*$\mathbf{x}'$ is given as:

$$\begin{aligned}
\frac{\partial D^{\frac{k}{2}} \frac{\mathbf{x}'}{\|\mathbf{x}'\|}}{\partial \mathbf{x}'} =& D^{\frac{k}{2}} \left(\frac{\mathbf{I}}{\|\mathbf{x}'\|} - \frac{\mathbf{x}'\mathbf{x}'^\top}{\|\mathbf{x}'\|^3} \right) \\
=& D^{\frac{k}{2}} \left(\frac{\mathbf{I}}{\alpha\|\mathbf{x}\|} - \frac{\alpha^2\mathbf{x}\mathbf{x}^\top}{\alpha^3\|\mathbf{x}\|^3} \right) \\
=& \frac{1}{\alpha} D^{\frac{k}{2}} \left(\frac{\mathbf{I}}{\|\mathbf{x}\|} - \frac{\mathbf{x}\mathbf{x}^\top}{\|\mathbf{x}\|^3} \right) \\
=& \frac{1}{\alpha}\mathbf{J}
\end{aligned} \tag{20}$$

Then we have the gradient of loss $w.r.t.\mathbf{W}$ and $\alpha\mathbf{W}$:

$$\begin{aligned}
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{W}} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{x}}\cdot\frac{\partial\mathbf{x}}{\partial\mathbf{W}} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J}\mathbf{v}^\top \\
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{W})} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial\mathbf{x}'}\cdot\frac{\partial\mathbf{x}'}{\partial(\alpha\mathbf{W})} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{1}{\alpha}\mathbf{J}\mathbf{v}^\top \\
\Rightarrow \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{W})} =& \frac{1}{\alpha}\cdot\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{W}} = \frac{1}{\alpha}\cdot\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J}\mathbf{v}^\top
\end{aligned} \tag{21}$$

Similarly, for $\mathbf{b}$ and $\alpha\mathbf{b}$ we have:

$$\begin{aligned}
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{b}} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{x}}\cdot\frac{\partial\mathbf{x}}{\partial\mathbf{b}} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J} \\
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{b})} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial\mathbf{x}'}\cdot\frac{\partial\mathbf{x}'}{\partial(\alpha\mathbf{b})} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{1}{\alpha}\mathbf{J} \\
\Rightarrow \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial(\alpha\mathbf{b})} =& \frac{1}{\alpha}\cdot\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{b}} = \frac{1}{\alpha}\cdot\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J}
\end{aligned} \tag{22}$$

And for $\mathbf{v}$, we have:

$$\begin{aligned}
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{v}} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{x}}\cdot\frac{\partial\mathbf{x}}{\partial\mathbf{v}} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J}\mathbf{W}^\top \\
\frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial\mathbf{v}} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial\mathbf{x}'}\cdot\frac{\partial\mathbf{x}'}{\partial\mathbf{v}} \\
=& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{1}{\alpha}\mathbf{J}(\alpha\mathbf{W})^\top \\
\Rightarrow \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}'}\cdot\frac{\partial\tilde{\mathbf{x}}'}{\partial\mathbf{v}} =& \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\frac{\partial\tilde{\mathbf{x}}}{\partial\mathbf{v}} = \frac{\partial\mathcal{L}}{\partial\tilde{\mathbf{x}}}\cdot\mathbf{J}\mathbf{W}^\top
\end{aligned} \tag{23}$$

□

B.3 Proofs of Theorem 3.2 and Theorem 3.3

Proof of Theorem 3.2. Let $\mathbf{X}\in\mathbb{R}^{L\times D}$ be a single sequence of token vectors, and let $\tilde{\mathbf{X}}$ be the unit normalized output with modulus k, the entropy lower bound (ELB) of the attention scores is given by the following expression:

$$\begin{aligned}
\mathrm{ELB}(k;L,D) =& \min_{i=1}^{L} H(\mathbf{A}_i) \\
=& \min_{i=1}^{L}\left(-\sum_{j=1}^{L}\mathbf{A}_{i,j}\log\mathbf{A}_{i,j}\right) \\
=& \log\left(L-1+\exp\left(2D^{k-\frac{1}{2}}\right)\right) - \frac{2D^{k-\frac{1}{2}}\exp\left(2D^{k-\frac{1}{2}}\right)}{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)}
\end{aligned} \tag{24}$$

Proof: Let $\tilde{\mathbf{X}} = D^{\frac{k}{2}}\mathbf{e}$ where $\mathbf{e}$ are the vectors of unit norm. Without loss of generality, we can assume the ELB is achieved at anchor index i, where we can compute the attention scores as follows:

$$\begin{aligned}\mathbf{A}_i =& \text{softmax}\left(\frac{\tilde{\mathbf{X}}_i\tilde{\mathbf{X}}^\top}{\sqrt{D}}\right)\\ =& \text{softmax}\left(\frac{D^k\mathbf{e}_i\mathbf{e}^\top}{\sqrt{D}}\right)\\ =& \text{softmax}\left(D^{k-\frac{1}{2}}\mathbf{e}_i\mathbf{e}^\top\right)\end{aligned} \tag{25}$$

Since $\mathbf{e}_i\mathbf{e}_j^\top \in (-1,1)\,, \forall i,j = 1,2,\cdots,L$, the entropy of the attentions scores is lower bounded by the following expression when it satisfies that $\mathbf{e}_i\mathbf{e}_j = \begin{cases}1, & j=i\\ -1, & j\neq i\end{cases}$:

$$\begin{aligned}&H(\mathbf{A}_i)\\ =& -\sum_{j=1}^{L}\mathbf{A}_{i,j}\log\mathbf{A}_{i,j}\\ =& -(L-1)\cdot\frac{\exp\left(-D^{k-\frac{1}{2}}\right)}{(L-1)\exp\left(-D^{k-\frac{1}{2}}\right)+\exp\left(D^{k-\frac{1}{2}}\right)}\log\frac{\exp\left(-D^{k-\frac{1}{2}}\right)}{(L-1)\exp\left(-D^{k-\frac{1}{2}}\right)+\exp\left(D^{k-\frac{1}{2}}\right)}\\ &-\frac{\exp\left(D^{k-\frac{1}{2}}\right)}{(L-1)\exp\left(-D^{k-\frac{1}{2}}\right)+\exp\left(D^{k-\frac{1}{2}}\right)}\log\frac{\exp\left(D^{k-\frac{1}{2}}\right)}{(L-1)\exp\left(-D^{k-\frac{1}{2}}\right)+\exp\left(D^{k-\frac{1}{2}}\right)}\\ =&\frac{L-1}{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)}\log\left(L-1+\exp\left(2D^{k-\frac{1}{2}}\right)\right)\\ &+\frac{\exp\left(2D^{k-\frac{1}{2}}\right)}{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)}\log\frac{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)}{\exp\left(2D^{k-\frac{1}{2}}\right)}\\ =&\log\left(L-1+\exp\left(2D^{k-\frac{1}{2}}\right)\right)-\frac{2D^{k-\frac{1}{2}}\exp\left(2D^{k-\frac{1}{2}}\right)}{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)}\end{aligned} \tag{26}$$

Therefore, the entropy lower bound (ELB) for any L, D and k is:

$$\text{ELB}(k;L,D) = \log\left(L-1+\exp\left(2D^{k-\frac{1}{2}}\right)\right)-\frac{2D^{k-\frac{1}{2}}\exp\left(2D^{k-\frac{1}{2}}\right)}{L-1+\exp\left(2D^{k-\frac{1}{2}}\right)} \tag{27}$$

□

Proof of Theorem 3.3. The ELB is a monotonically decreasing function of k bounded between 0 and $\log L$.

Proof: Let $d = 2D^{k-\frac{1}{2}}$, then it is obvious that d is monotonically increasing with k, therefore we only need to prove that $\text{ELB}(k;L,D)$ is monotonically decreasing with d. The derivative of $\text{ELB}(k;L,D)$ with respect to

d is given as follows:

$$\begin{aligned}\frac{\partial \operatorname{ELB}(k;L,D)}{\partial d} =& \frac{e^d}{L-1+e^d} - \frac{\left(L-1+e^d\right)(d+1)e^d - (de^d)e^d}{\left(L-1+e^d\right)^2} \\ =& \frac{e^d}{\left(L-1+e^d\right)^2}\left(\left(L-1+e^d\right) - \left(L-1+e^d\right)(d+1) + de^d\right) \\ =& \frac{de^d}{\left(L-1+e^d\right)^2}(1-L) \\ (\forall L>1)\ <& 0\end{aligned} \tag{28}$$

Therefore, $\operatorname{ELB}(k;L,D)$ is monotonically decreasing with d and with k. If the limits of $\operatorname{ELB}(k;L,D)$ as $k \to -\infty$ and $k \to +\infty$ exist, then $\operatorname{ELB}(k;L,D)$ is bounded between these two limits. The limits are given as follows:

$$\begin{aligned}\lim_{k\to-\infty} \operatorname{ELB}(k;L,D) =& \lim_{d\to 0^+}\left(\log\left(L-1+e^d\right) - \frac{de^d}{L-1+e^d}\right) \\ =& \log\left(L-1+1\right) - \frac{0}{L-1+1} \\ =& \log L\end{aligned} \tag{29}$$

$$\begin{aligned}\lim_{k\to+\infty} \operatorname{ELB}(k;L,D) =& \lim_{d\to+\infty}\left(\log\left(L-1+e^d\right) - \frac{de^d}{L-1+e^d}\right) \\ =& \lim_{d\to+\infty}\log e^d - \lim_{d\to+\infty}\frac{d}{(L-1)e^{-d}+1} \\ =& d-d \\ =& 0\end{aligned} \tag{30}$$

Therefore, $\operatorname{ELB}(k;L,D)$ is bounded between 0 and $\log L$. □

C Discussion

C.1 Difference between the proposed normalization and the other normalization

BatchNorm and LayerNorm are all normalization methods that are widely used in deep learning. They share the same center-and-scale normalization paradigm by first subtracting the mean and then divide by standard deviation. The only difference between them in terms of computation is the dimensions of data used to compute these statistics, as shown in Table A1.

In terms of application, BatchNorm is often used in fully connected layers and convolution layers, while LayerNorm is often used in recurrent neural networks and Transformers. The subtle difference between LayerNorm (theory) and LayerNorm (practice) might be attributed to the fact that the sequence length L is often variable in Transformers, thus normalization within each token might be more stable. But this will require further investigation to come to a conclusion.

The proposed UnitNorm is a normalization method that is used to normalize the input data to have unit norm, which takes the same dimension for computation as LayerNorm, yet it distinguishes itself from LayerNorm by the fact that it does not subtract the mean and divide by standard deviation. Also, UnitNorm discard the center operation on the normalized output, as it will also cause the problem of token shift (Section 2.2).

C.2 Feasibility of switching the order of normalization and projection in theoretical analysis

Let $\mathbf{X} \in \mathbb{R}^{L\times D}$ be a single sequence of token vectors, and the normalization operation is given in the following form:

$$f : \mathbf{X} \mapsto \frac{\mathbf{X}-\boldsymbol{\mu}}{\boldsymbol{\sigma}} \equiv \mathbf{X}\mathbf{W} + \mathbf{b} \tag{31}$$

Table A1: Computation of the statistics for different normalization methods. Input data $\mathbf{X} \in \mathbb{R}^{N \times L \times D}$, where N is the batch size, L is the sequence length and D is the feature dimension. $\mathbf{X}_{n,l,d}$ denotes the d-th feature of the l-th token in the n-th sequence. Normalization is broadcasted over the same dimension as the statistics and mathematical operations are done element-wise. For BatchNorm, LayerNorm (theory) and LayerNorm (practice), $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ are optional learnable parameters that will re-scale and re-center the normalized output element-wise, which is enabled by default in the PyTorch's implementation.

Method	Statistics	Normalization
BatchNorm	$\mu_d = \frac{1}{NL}\sum_{n=1}^{N}\sum_{l=1}^{L}\mathbf{X}_{n,l,d}$ $\sigma_d^2 = \frac{1}{NL}\sum_{n=1}^{N}\sum_{l=1}^{L}(\mathbf{X}_{n,l,d}-\mu_d)^2$ $\boldsymbol{\mu} = [\mu_1 \quad \mu_2 \quad \cdots \quad \mu_D]^\top \in \mathbb{R}^{1\times1\times D}$ $\boldsymbol{\sigma}^2 = [\sigma_1^2 \quad \sigma_2^2 \quad \cdots \quad \sigma_D^2]^\top \in \mathbb{R}^{1\times1\times D}$	$\tilde{\mathbf{X}} = \frac{\mathbf{X}-\boldsymbol{\mu}}{\sqrt{\boldsymbol{\sigma}^2+\varepsilon}}$ $\mathbf{Y} = \tilde{\mathbf{X}} \odot \boldsymbol{\gamma} + \boldsymbol{\beta}$
LayerNorm (theory)	$\mu_n = \frac{1}{LD}\sum_{l=1}^{L}\sum_{d=1}^{D}\mathbf{X}_{n,l,d}$ $\sigma_n^2 = \frac{1}{LD}\sum_{l=1}^{L}\sum_{d=1}^{D}(\mathbf{X}_{n,l,d}-\mu_n)^2$ $\boldsymbol{\mu} = [\mu_1 \quad \mu_2 \quad \cdots \quad \mu_N]^\top \in \mathbb{R}^{N\times1\times1}$ $\boldsymbol{\sigma}^2 = [\sigma_1^2 \quad \sigma_2^2 \quad \cdots \quad \sigma_N^2]^\top \in \mathbb{R}^{N\times1\times1}$	
LayerNorm (practice)	$\mu_{n,l} = \frac{1}{D}\sum_{d=1}^{D}\mathbf{X}_{n,l,d}$ $\sigma_{n,l}^2 = \frac{1}{D}\sum_{d=1}^{D}(\mathbf{X}_{n,l,d}-\mu_n)^2$ $\boldsymbol{\mu} = \begin{bmatrix} \mu_{1,1} & \mu_{1,2} & \cdots & \mu_{1,L} \\ \mu_{2,1} & \mu_{2,2} & \cdots & \mu_{2,L} \\ \vdots & \vdots & \ddots & \vdots \\ \mu_{N,1} & \mu_{N,2} & \cdots & \mu_{N,L} \end{bmatrix}^\top \in \mathbb{R}^{N\times L\times1}$ $\boldsymbol{\sigma}^2 = \begin{bmatrix} \sigma_{1,1}^2 & \sigma_{1,2}^2 & \cdots & \sigma_{1,L}^2 \\ \sigma_{2,1}^2 & \sigma_{2,2}^2 & \cdots & \sigma_{2,L}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{N,1}^2 & \sigma_{N,2}^2 & \cdots & \sigma_{N,L}^2 \end{bmatrix}^\top \in \mathbb{R}^{N\times L\times1}$	
RMSNorm	$\lVert\mathbf{X}\rVert_{n,l} = \sqrt{\sum_{d=1}^{D}\mathbf{X}_{n,l,d}^2}$	$\tilde{\mathbf{X}} = \sqrt{D}\frac{\mathbf{X}}{\lVert\mathbf{X}\rVert}$ $\mathbf{Y} = \tilde{\mathbf{X}} \odot \boldsymbol{\gamma} + \boldsymbol{\beta}$
UnitNorm	$\lVert\mathbf{X}\rVert = \begin{bmatrix} \lVert\mathbf{X}\rVert_{1,1} & \lVert\mathbf{X}\rVert_{1,2} & \cdots & \lVert\mathbf{X}\rVert_{1,L} \\ \lVert\mathbf{X}\rVert_{2,1} & \lVert\mathbf{X}\rVert_{2,2} & \cdots & \lVert\mathbf{X}\rVert_{2,L} \\ \vdots & \vdots & \ddots & \vdots \\ \lVert\mathbf{X}\rVert_{N,1} & \lVert\mathbf{X}\rVert_{N,2} & \cdots & \lVert\mathbf{X}\rVert_{N,L} \end{bmatrix}^\top \in \mathbb{R}^{N\times L\times1}$	$\tilde{\mathbf{X}} = D^{\frac{k}{2}}\frac{\mathbf{X}}{\lVert\mathbf{X}\rVert}$ $\mathbf{Y} = \tilde{\mathbf{X}}$

where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the mean and standard deviation of the input vector $\mathbf{X}$, respectively, and $\mathbf{W} = \boldsymbol{\sigma}^{-1}$ and $\mathbf{b} = \boldsymbol{\mu}\boldsymbol{\sigma}^{-1}$. Depending on the normalization method, the mean and standard deviation can be computed over different dimensions.

The projection in the attention mechanism maps the input vectors to query, key and value vectors, and here we only consider the query and key vectors for this discussion, which are computed as follows:

$$\begin{aligned}\mathbf{Q} =& \mathbf{X}\mathbf{W}_Q + \mathbf{b}_Q \\ \mathbf{K} =& \mathbf{X}\mathbf{W}_K + \mathbf{b}_K\end{aligned} \tag{32}$$

where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{D\times D}$ are the projection matrices and $\mathbf{b}_Q, \mathbf{b}_K \in \mathbb{R}^D$ are the bias vectors for query and key, respectively. As the normalization and projection are both linear operations, we can combine them into a single linear operation as follows:

$$\begin{aligned}\mathbf{Y} =& \tilde{\mathbf{X}}\mathbf{W}_Y + \mathbf{b}_Y \\ =& \left(\mathbf{X}\mathbf{W} + \mathbf{b}\right)\mathbf{W}_Y + \mathbf{b}_Y \\ =& \mathbf{X}\left(\mathbf{W}\mathbf{W}_Y\right) + \left(\mathbf{b}\mathbf{W}_Y + \mathbf{b}_Y\right)\end{aligned} \tag{33}$$

for $Y \in \{Q, K\}$. Therefore, there must exist some $\mathbf{W}'$, $\mathbf{b}'$, $\mathbf{W}'_Y$ and $\mathbf{b}'_Y$ such that:

$$\begin{aligned}\mathbf{W}'_Y\mathbf{W}' =& \mathbf{W}\mathbf{W}_Y \\ \mathbf{b}_Y\mathbf{W}' + \mathbf{b}' =& \mathbf{b}\mathbf{W}_Y + \mathbf{b}_Y\end{aligned} \tag{34}$$

Therefore, the order of normalization and projection does not affect the theoretical analysis. And in favor of simplicity, we can assume the normalization is performed after the projection.

D Supplementary Figures

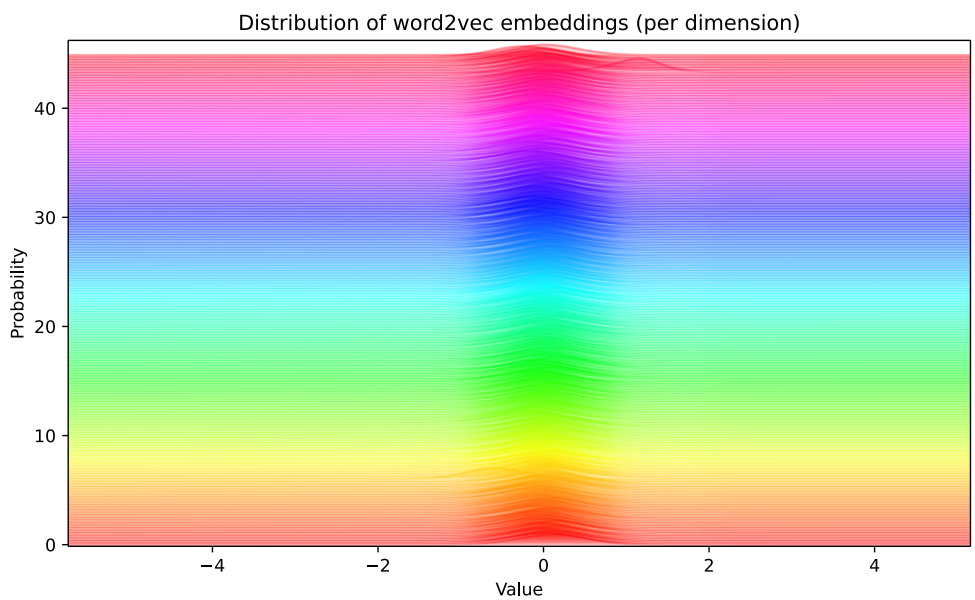

Figure A1: Distribution of values in each dimension of the word2vec embedding. The word2vec embedding is a 300-dimensional vector, and the distribution follows a normal distribution with means mostly around 0.

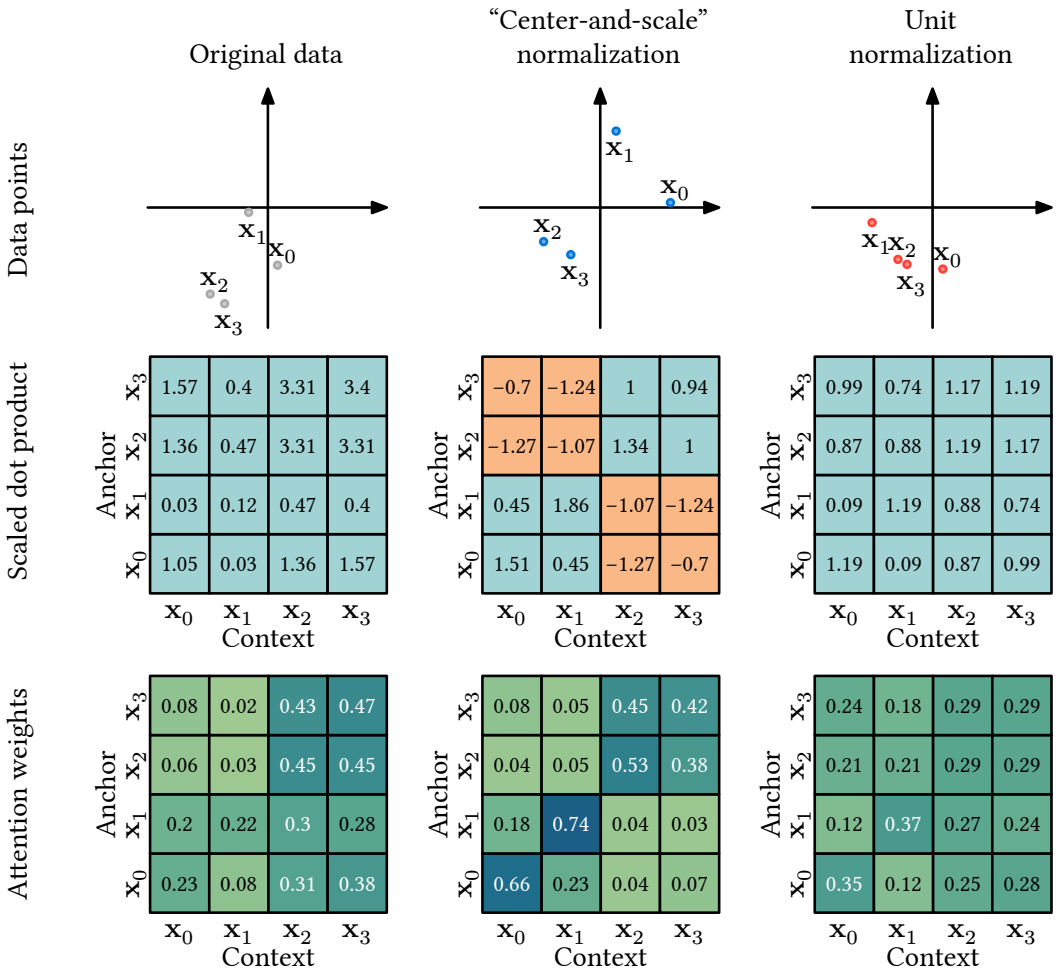

Figure A2: Demonstration of the token shift and attention shift problems using artificial data. The $\mathbf{x}_0$ and $\mathbf{x}_1$ exhibit typical token shift as shifting away from their original quadrants, resulting in sign flip in scaled dot product (marked in orange), and leading to less attention weights distributed to $\mathbf{x}_2$ and $\mathbf{x}_3$ than original. Attention shift and sparse attention problem can also be observed as the maximum attention weight is altered from $\mathbf{x}_2$ and $\mathbf{x}_3$ to nearly solely onto themselves.

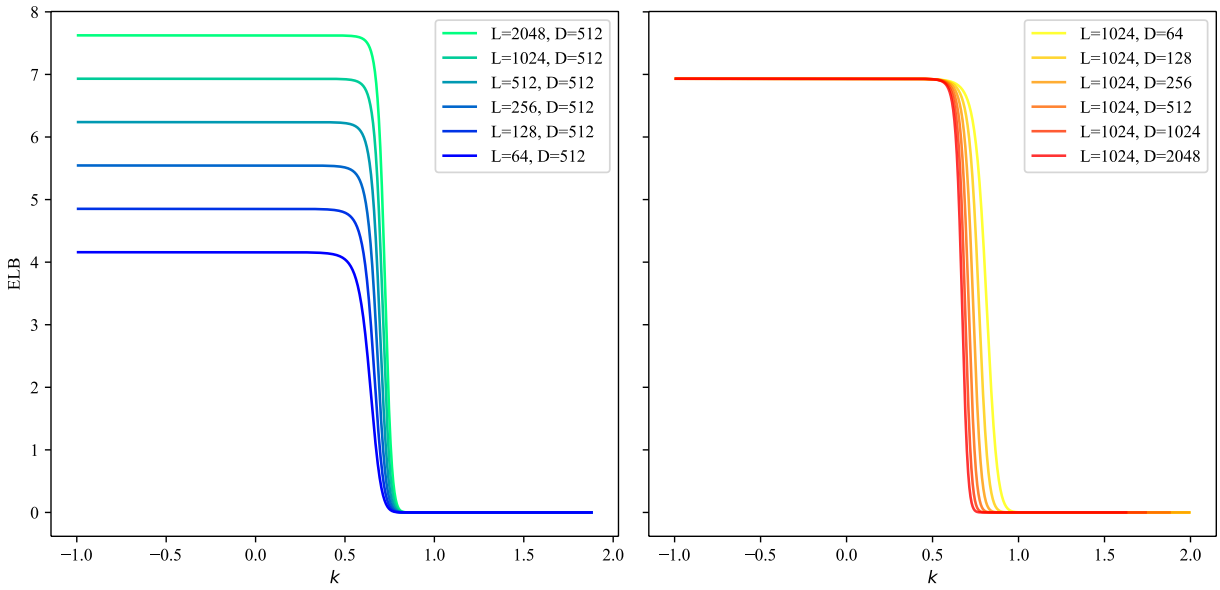

Figure A3: Entropy lower bound (ELB) against k for different L, D. The left figure shows the curve for fixed $D = 512$ and varying L, and the right figure shows the curve for fixed $L = 1024$ and varying D.

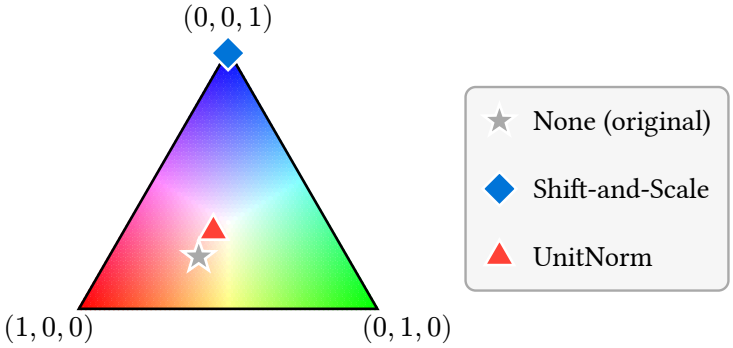

Figure A4: Graphical representation of attention weights showing a simple scenario of 3 tokens. Each corner represents a single-point distribution (red, blue and green) and the center representing a uniform distribution (white). Gray star, blue diamond, and red triangle mark the attention weights with no normalization, center-and-scale normalization, and UnitNorm, respectively.

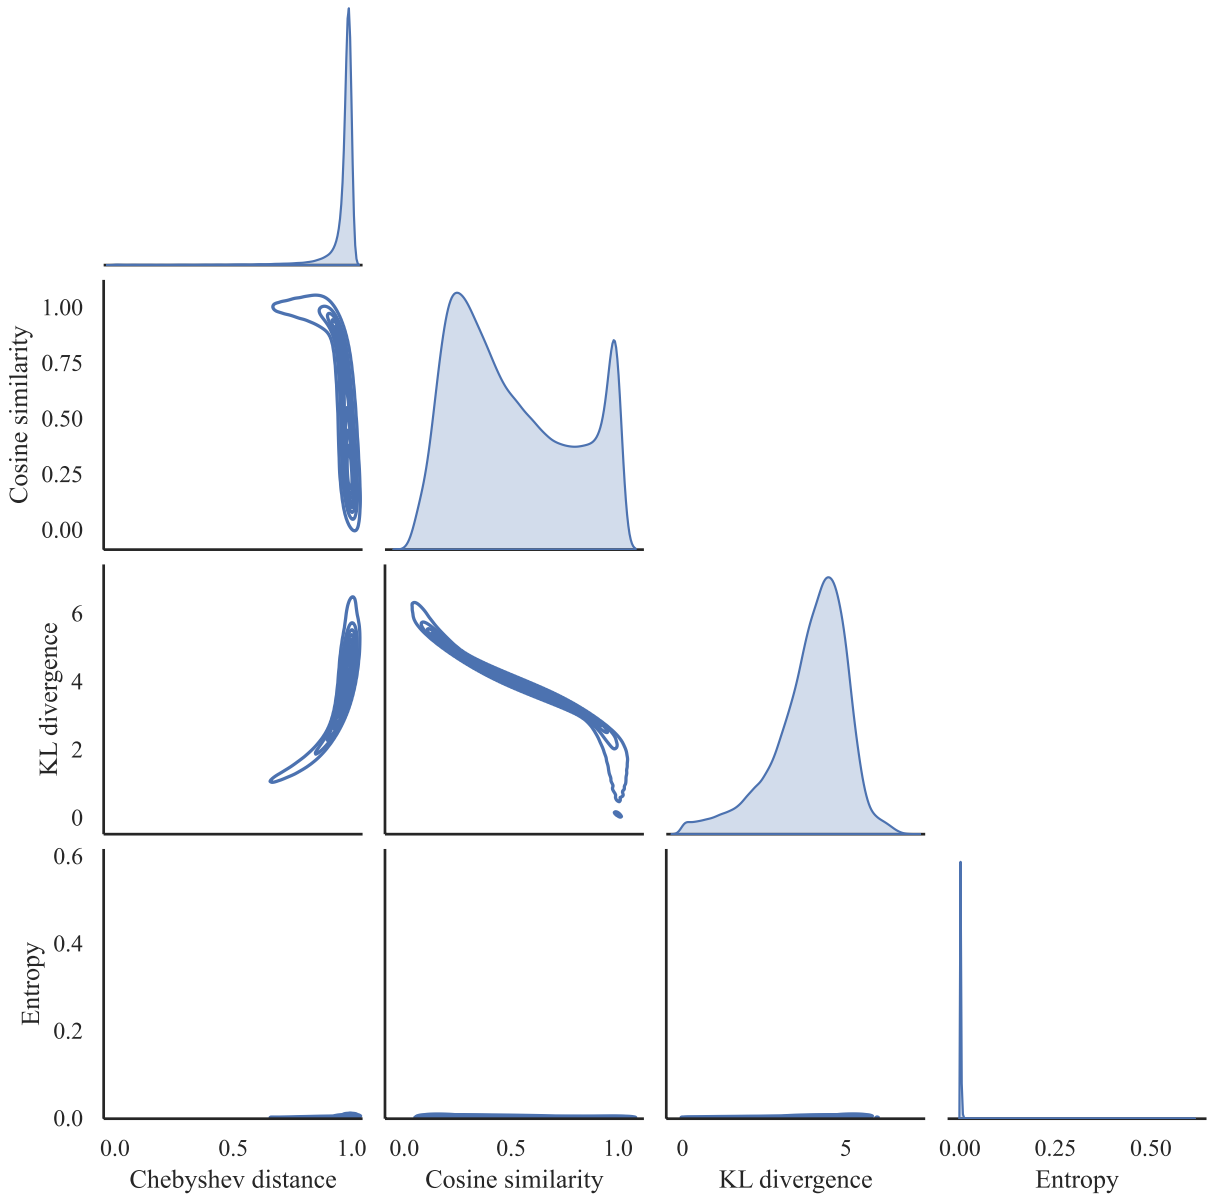

Figure A5: Joint distribution of metrics for LayerNorm (practice). Metrics used are defined as in Table A10.

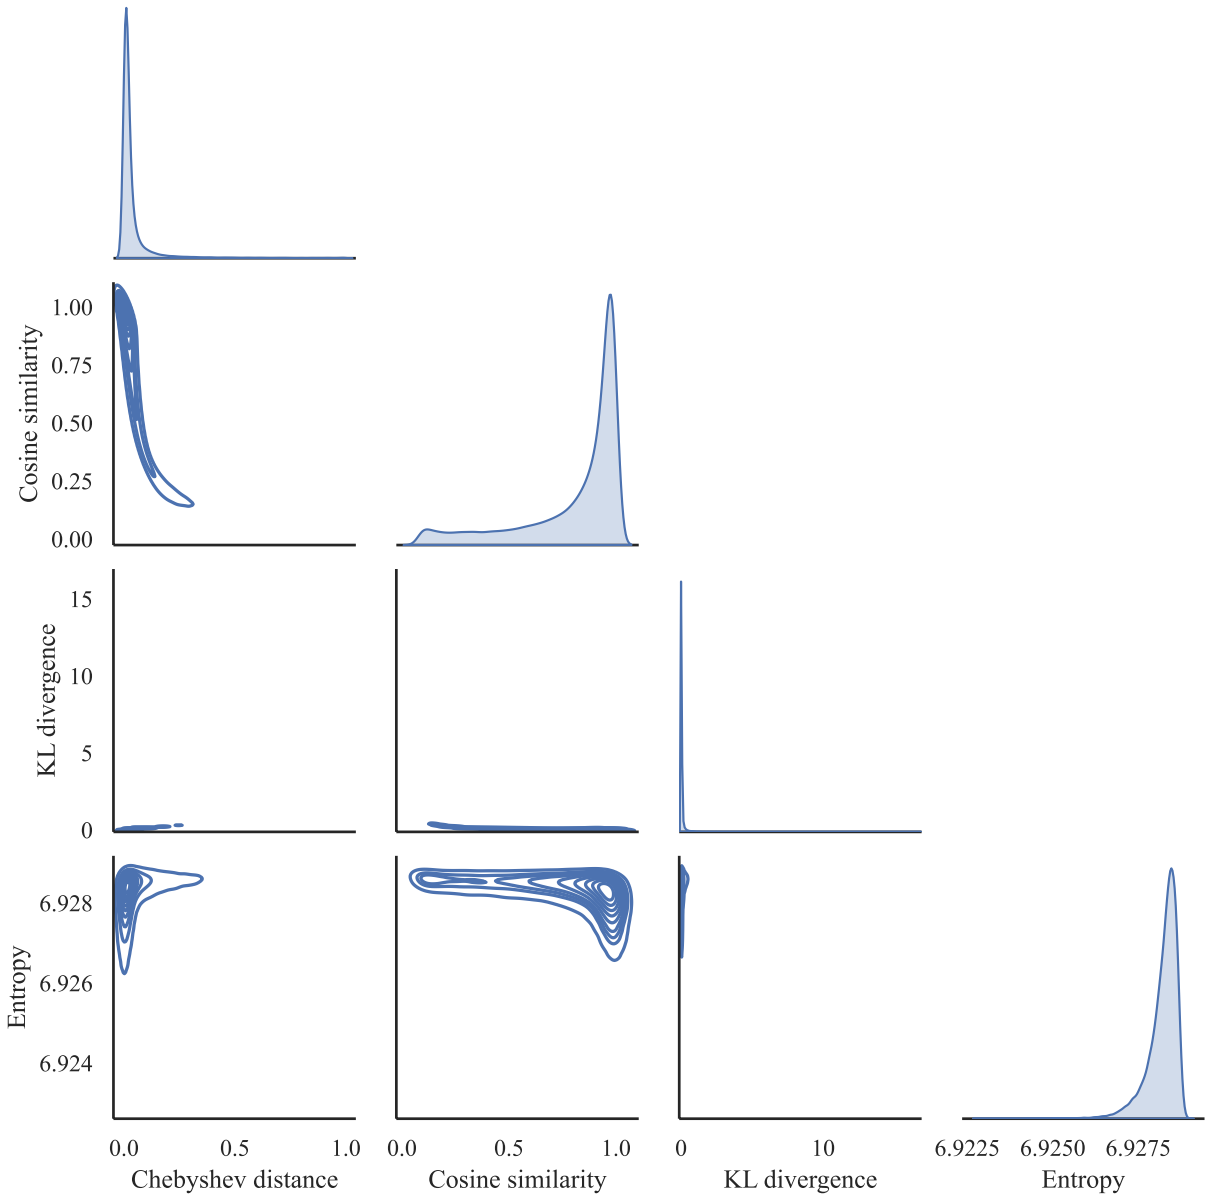

Figure A6: Joint distribution of metrics for UnitNorm. Metrics used are defined as in Table A10.

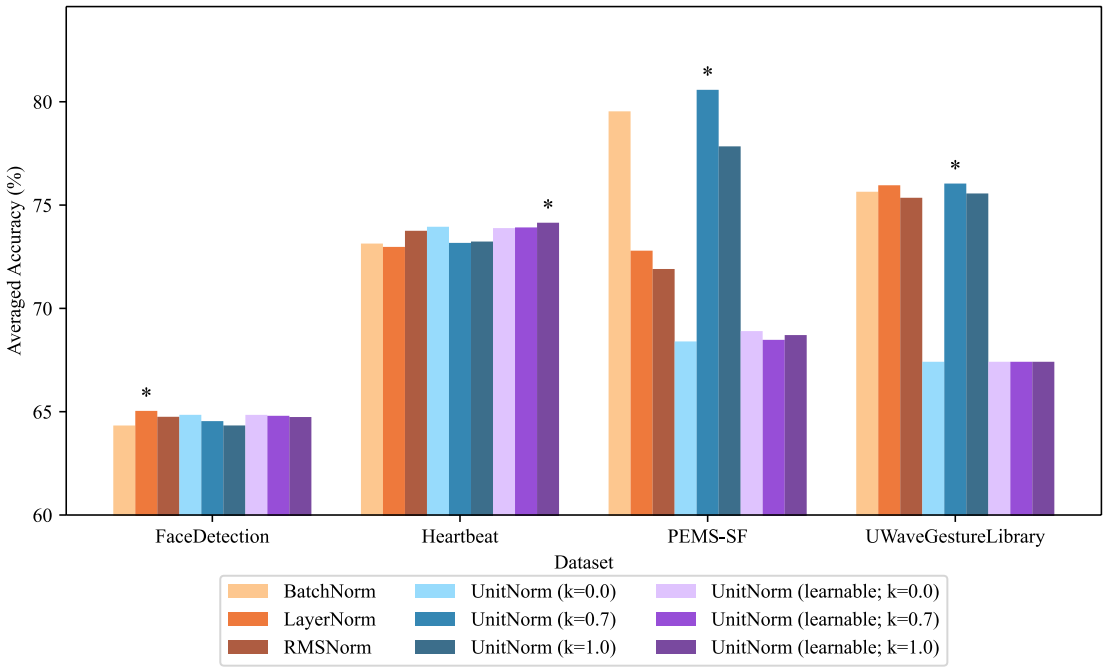

Figure A7: Average rank of normalization methods on the classification tasks. X-axis: Dataset with different normalization, Y-axis: average rank over models. Ranks are computed based on the accuracy of each model on each task with different normalization methods (lower is better). $*$ indicates the best performing normalization method(s) on each task. UnitNorm and UnitNorm (learnable) outperform other normalization methods on 3 out of 5 datasets, showing its potential in classification tasks.

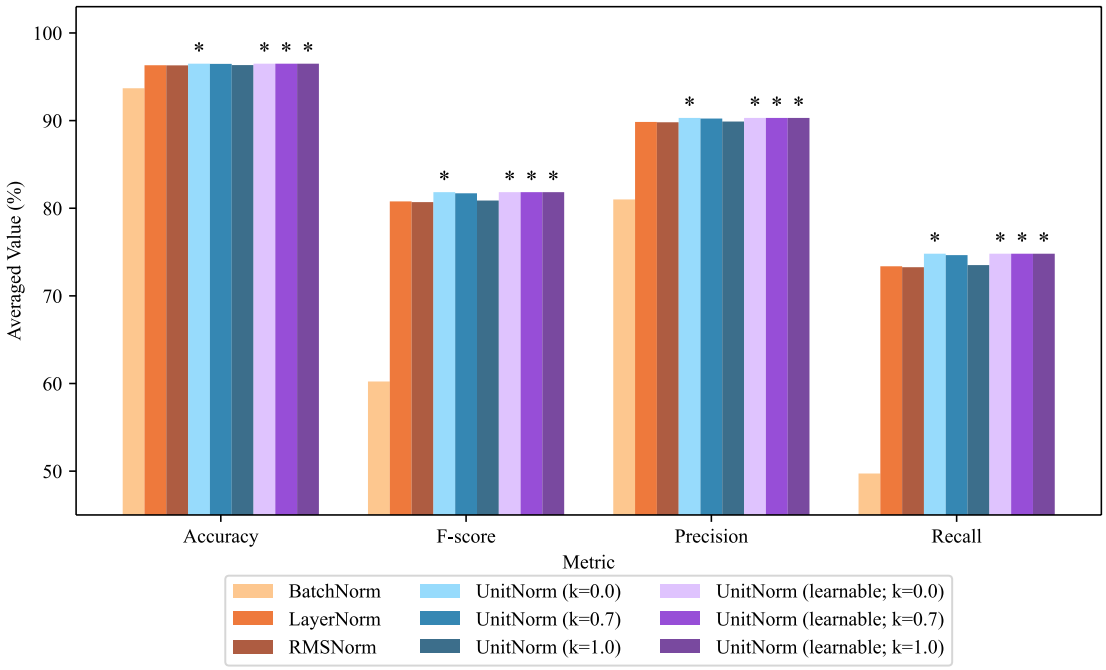

Figure A8: Average rank of normalization methods on the anomaly detection tasks. X-axis: Metrics under different normalization, Y-axis: average rank over models. Ranks are computed based on every metric of each model with different normalization methods (lower is better). $*$ indicates the best performing normalization method(s) on each metric. UnitNorm and UnitNorm (learnable) show a dominating performance gain over the other normalization methods in all metrics.

Table A2: Performance on a synthetic periodic dataset with two channels of sine waves (varying periods 2-5 Hz, amplitudes 0.5-2, plus Gaussian noise level = 0.1). Results shown for PatchTST model with different normalization methods, averaged over 3 random seeds with standard deviation. Best results are in **bold**.

Normalization	MSE	MAE
BatchNorm	2.721 ± 3.409	0.956 ± 0.550
LayerNorm	3.117 ± 3.078	1.054 ± 0.532
RMSNorm	3.114 ± 3.057	1.056 ± 0.528
UnitNorm ($k = 0.5$)	**1.127 ± 1.085**	**0.732 ± 0.269**
UnitNorm ($k = 0.7$)	1.600 ± 1.762	0.800 ± 0.361
UnitNorm ($k = 1.0$)	3.115 ± 3.059	1.056 ± 0.528

E Supplementary Tables

Table A3: Summary of long term forecasting benchmark settings. The sequence length is the number of historical time steps fed into the encoder, and the label length is the number of time steps fed into the decoder as the ground truth output of the decoder. The prediction length is the number of time steps to be predicted by the decoder.

Datasets	Feature number	Sequence length	Label length	Prediction length	Metrics	License
ETTh1, ETTh2 Zhou et al. (2021)	7	384	96	{96, 192, 384, 720}	MSE, MAE	CC BY-ND 4.0
ECL Trindade (2015)	321					CC BY 4.0
Exchange Lai et al. (2018)	8					N/A

Table A4: Summary of classification benchmark settings. All datasets are from UEA Archive Bagnall et al. (2018). The sequence length is the number of time steps in each sequence fed into the encoder, and the prediction is made on the flattened output of the encoder by a fully connected layer.

Datasets	Feature number	Class number	Sequence length	Metrics	License
FaceDetection	144	5890	96	Accuracy	N/A
Heartbeat	61	204			
PEMS-SF	963	173			
UWaveGestureLibrary	3	320			

Table A5: Summary of anomaly detection benchmark settings. The sequence length is the number of time steps in each sequence fed into the model for reconstruction using MSE as loss. The threshold is determined by the distribution of reconstruction error on the training set, and the metrics are computed on the test set based on this threshold.

Datasets	Feature number	Sequence length	Reconstruction error	Metrics	License
MSL Hundman et al. (2018)	55	100	MSE	Accuracy, F1-score, Precision, Recall	N/A

Table A6: Summary of compute resources used for the experiments. Depending on the dataset and model, the GPU memory usage varies from 4G to 64G.

CPU	Memory	GPU	GPU Memory
AMD Threadripper 3995WX	512G	4 × NVIDIA RTX A5000	4 × 24G

Table A7: Long term forecasting test losses on different datasets using different models and normalization methods. For each dataset, prediction length, metric and for each model, the best performing normalization method(s) are bolded, and the second best are underlined.

	dataset	ECL								ETTh1								ETTh2								Exchange							
	prediction length	96		192		336		720		96		192		336		720		96		192		336		720		96		192		336		720	
	metric	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
BatchNorm	Crossformer	0.408	0.278	0.435	0.311	**0.246**				0.814	**0.395**	**0.661**	0.422	**0.426**				1.093	0.346	1.414	0.395	**0.557**				0.608	0.204	**0.672**	0.297	**0.265**			
	FEDformer	0.310	0.184	0.357	0.198	**0.144**				1.008	**0.374**	**0.790**	0.389	0.417				1.852	0.292	3.319	0.352	**0.563**				0.665	0.088	**0.770**	0.169	**0.122**			
	Informer	**0.409**	0.285	0.448	**0.322**	0.262				0.805	0.429	0.792	0.449	0.507				1.258	0.399	2.105	**0.434**	0.878				0.921	0.300	0.861	**0.383**	**0.445**			
	PatchTST	**0.321**	0.191	0.364	**0.211**	0.162				1.014	0.426	1.015	0.433	0.531				2.488	0.376	6.878	**0.423**	1.355				1.458	0.180	1.308	0.276	**0.315**			
	Transformer	**0.463**	0.300	0.455	**0.342**	0.290				0.886	0.462	0.886	0.477	0.620				1.603	0.435	2.215	**0.473**	0.961				1.242	0.419	1.367	0.487	**0.646**			
		0.407	0.207	0.376	**0.234**	**0.189**				1.161	0.478	1.230	0.483	0.684				3.874	0.421	6.966	0.469	1.626				2.578	0.336	3.221	0.438	**0.658**			
		0.510	0.332	0.478	0.372	**0.338**				0.819	0.497	0.940	0.507	0.707				1.833	0.453	2.307	0.486	1.491				1.444	0.703	1.602	0.823	**0.846**			
		0.461	0.248	0.417	0.270	0.255				1.054	0.506	1.357	0.516	0.834				4.782	0.435	8.194	0.479	3.213				3.069	0.874	4.284	1.151	**1.086**			
LayerNorm	Crossformer	0.416	**0.270**	**0.410**	**0.309**	0.249				0.748	0.396	0.720	**0.420**	0.439				1.188	0.345	1.502	0.395	0.721				0.589	0.209	0.733	0.297	0.378			
	FEDformer	0.320	**0.181**	**0.325**	**0.195**	0.148				0.884	0.377	0.877	**0.379**	0.419				2.205	0.297	3.574	0.351	1.014				0.581	0.091	0.830	0.169	0.274			
	Informer	0.436	**0.276**	**0.436**	0.324	**0.262**				0.802	0.426	0.792	**0.444**	0.500				1.820	0.399	2.056	0.438	0.909				0.786	0.301	**0.836**	0.383	0.531			
	PatchTST	0.358	**0.187**	**0.355**	0.211	**0.162**				1.003	0.422	1.013	**0.420**	0.511				5.161	0.380	6.156	0.425	1.487				1.047	0.180	**1.081**	**0.276**	0.502			
	Transformer	0.499	**0.291**	**0.443**	0.347	0.291				0.834	**0.447**	0.842	**0.471**	**0.572**				1.699	0.431	1.864	0.475	1.331				1.051	0.412	1.073	0.488	0.890			
		0.458	**0.204**	**0.364**	0.236	0.192				1.081	**0.460**	1.148	0.464	**0.621**				4.549	0.417	4.990	**0.469**	2.686				1.745	0.324	1.817	0.440	1.288			
		0.542	**0.325**	**0.455**	0.366	0.344				0.833	0.492	0.855	0.507	0.667				1.390	0.447	1.652	0.488	1.506				1.212	0.713	1.409	0.824	1.047			
		0.514	**0.246**	**0.391**	0.262	0.266				1.078	0.511	1.174	0.507	0.765				2.908	0.429	3.943	0.481	3.057				2.297	0.901	2.967	1.153	1.701			
RMSNorm	Crossformer	0.403	0.271	0.417	0.309	0.249				0.745	0.396	0.709	0.421	0.431				1.211	0.345	1.573	0.393	0.709				0.580	0.209	0.730	0.297	0.371			
	FEDformer	**0.303**	0.182	0.334	0.195	0.150				0.882	0.377	0.864	0.380	0.409				2.276	0.297	4.058	0.350	0.963				**0.569**	0.091	0.841	0.169	0.261			
	Informer	0.427	0.277	0.444	0.324	0.262				0.802	0.425	0.781	0.445	0.495				1.815	0.399	2.126	0.437	0.904				**0.783**	0.300	0.866	0.383	0.522			
	PatchTST	0.343	0.188	0.366	0.211	0.162				0.984	**0.421**	0.995	0.420	0.508				5.030	0.379	6.495	0.425	1.451				**1.046**	0.179	1.158	0.277	0.485			
	Transformer	0.490	0.293	0.448	0.346	0.292				0.870	0.448	0.822	0.471	0.597				1.806	0.431	1.941	0.475	1.116				1.045	0.411	1.077	0.488	0.882			
		0.441	0.204	0.370	0.235	0.196				1.123	0.463	1.112	0.462	0.658				5.043	0.418	5.358	0.469	1.985				1.742	0.322	1.840	0.440	1.262			
		0.539	0.327	0.460	**0.365**	0.344				1.123	0.463	1.112	0.462	0.658				1.480	0.448	1.663	0.488	1.525				1.246	0.715	1.386	0.824	1.045			
		0.511	0.247	0.399	**0.261**	0.264				0.808	0.495	0.857	0.507	0.703				3.239	0.429	3.887	0.481	3.315				2.354	0.904	2.875	1.153	1.695			
UnitNorm (k=0.0)	Crossformer	0.445	0.280	0.499	0.325	0.282				0.800	0.401	0.816	0.440	0.591				0.999	**0.341**	1.155	0.399	0.576				1.023	0.206	1.016	0.298	0.688			
	FEDformer	0.359	0.189	0.457	0.211	0.184				0.990	0.384	1.104	0.410	0.686				1.518	**0.290**	2.060	0.355	0.694				1.535	0.088	1.497	0.170	0.913			
	Informer	0.460	0.284	0.520	0.345	0.295				0.791	0.431	0.814	0.457	0.607				1.189	0.394	1.258	0.441	**0.756**				1.121	**0.299**	1.066	0.385	0.772			
	PatchTST	0.392	0.193	0.485	0.238	0.200				0.979	0.431	1.111	0.440	0.705				2.231	0.369	2.582	0.430	**1.214**				1.848	**0.178**	1.646	0.279	1.057			
	Transformer	0.505	0.300	0.530	0.368	0.309				0.842	0.453	0.843	0.477	0.607				1.248	0.429	1.287	0.480	**0.811**				1.203	0.418	1.186	0.487	1.177			
		0.454	0.209	0.497	0.265	0.212				1.076	0.471	1.141	0.472	0.709				2.358	0.417	2.396	0.476	**1.297**				2.122	0.331	2.042	0.438	2.187			
		0.498	0.332	0.535	0.398	0.348				0.844	0.483	0.891	0.503	0.649				1.291	0.449	1.306	0.490	1.071				1.111	0.707	1.200	0.820	1.338			
		0.435	0.252	0.517	0.303	0.257				1.077	**0.486**	1.212	0.497	0.759				**2.379**	0.429	2.447	0.486	1.982				1.717	0.876	2.155	1.145	2.711			

Continued on next page

dataset	prediction length	metric	UnitNorm (learnable; k=1.0) Transformer	UnitNorm (learnable; k=1.0) PatchTST	UnitNorm (learnable; k=1.0) Informer	UnitNorm (learnable; k=1.0) FEDformer	UnitNorm (learnable; k=1.0) Crossformer	UnitNorm (learnable; k=0.7) Transformer	UnitNorm (learnable; k=0.7) PatchTST	UnitNorm (learnable; k=0.7) Informer	UnitNorm (learnable; k=0.7) FEDformer	UnitNorm (learnable; k=0.7) Crossformer	UnitNorm (learnable; k=0.0) Transformer	UnitNorm (learnable; k=0.0) PatchTST	UnitNorm (learnable; k=0.0) Informer	UnitNorm (learnable; k=0.0) FEDformer	UnitNorm (learnable; k=0.0) Crossformer	UnitNorm (k=1.0) Transformer	UnitNorm (k=1.0) PatchTST	UnitNorm (k=1.0) Informer	UnitNorm (k=1.0) FEDformer	UnitNorm (k=1.0) Crossformer	UnitNorm (k=0.5) Transformer	UnitNorm (k=0.5) PatchTST	UnitNorm (k=0.5) Informer	UnitNorm (k=0.5) FEDformer	UnitNorm (k=0.5) Crossformer
ECL	96	MAE	0.445	0.280	0.502	0.325	0.282	0.445	0.280	0.500	0.325	0.282	0.449	0.280	0.502	0.324	0.282	0.405	0.271	0.418	0.309	0.249	0.420	0.274	0.463	0.320	0.257
		MSE	0.359	0.189	0.462	0.211	0.184	0.359	0.189	0.461	0.211	0.184	0.361	0.189	0.461	0.211	0.184	0.304	0.182	0.335	0.195	0.149	0.331	0.183	0.401	0.204	0.156
	192	MAE	0.460	0.284	0.522	0.345	0.295	0.460	0.284	0.520	0.345	0.295	0.459	0.284	0.508	0.344	0.293	0.428	0.277	0.446	0.324	0.265	0.446	0.280	0.492	0.335	0.274
		MSE	0.392	0.193	0.490	0.238	0.200	0.392	0.193	0.486	0.238	0.200	0.388	0.193	0.460	0.237	0.198	0.346	0.188	0.368	0.211	0.165	0.371	0.189	0.434	0.222	0.176
	336	MAE	0.505	0.300	0.531	0.368	0.309	0.505	0.300	0.530	0.368	0.309	0.492	0.300	0.526	0.367	0.305	0.491	0.292	0.447	0.346	**0.290**	0.504	0.296	0.501	0.354	0.299
		MSE	0.454	0.209	0.499	0.265	0.212	0.454	0.209	0.497	0.265	0.212	0.436	0.209	0.491	0.265	0.208	0.442	0.204	0.368	0.235	0.190	0.460	0.205	0.441	0.242	0.204
	720	MAE	0.498	0.332	0.547	0.394	0.348	0.498	0.332	0.538	0.394	0.348	**0.493**	0.331	0.539	0.398	0.349	0.540	0.327	0.464	0.365	0.341	0.508	0.328	0.558	0.385	0.339
		MSE	0.435	0.252	0.531	0.300	0.257	0.435	0.252	0.519	0.300	0.257	**0.427**	0.251	0.523	0.304	0.258	0.511	0.247	0.402	0.261	0.254	0.459	0.247	0.543	0.283	**0.251**
ETTh1	96	MAE	0.813	0.401	0.817	0.440	0.583	0.813	0.401	0.815	0.440	0.583	0.780	0.400	0.850	0.440	0.593	0.745	0.396	0.713	0.421	0.431	**0.662**	0.397	0.748	0.428	0.436
		MSE	1.003	0.384	1.104	0.410	0.653	1.003	0.384	1.107	0.410	0.653	0.963	0.384	1.169	0.410	0.671	0.881	0.377	0.869	0.380	0.409	**0.751**	0.381	0.963	0.386	**0.408**
	192	MAE	0.828	0.431	0.810	0.456	0.605	0.828	0.431	0.807	0.456	0.605	0.808	0.430	0.820	0.457	0.605	0.802	**0.425**	0.780	0.445	**0.487**	**0.736**	0.428	**0.751**	0.448	0.498
		MSE	1.037	0.431	1.112	0.437	0.710	1.037	0.431	1.094	0.437	0.710	1.010	0.431	1.120	0.439	0.700	0.984	0.421	0.992	0.420	**0.492**	**0.917**	0.427	**0.981**	0.423	0.499
	336	MAE	0.849	0.453	0.846	0.480	0.604	0.849	0.453	0.845	0.480	0.604	0.849	0.453	0.844	0.477	0.611	0.871	0.448	0.823	0.471	0.606	**0.821**	0.451	**0.800**	0.471	0.600
		MSE	1.091	0.471	1.142	0.476	0.704	1.091	0.471	1.141	0.476	0.704	1.101	0.471	1.149	0.470	0.719	1.125	0.463	1.113	0.462	0.674	**1.073**	0.466	**1.061**	**0.462**	0.643
	720	MAE	0.854	0.483	0.888	0.509	**0.621**	0.854	0.483	0.887	0.509	**0.621**	0.817	**0.483**	0.894	0.508	0.651	**0.806**	0.495	**0.849**	**0.501**	0.735	0.874	0.484	0.871	0.502	0.673
		MSE	1.100	**0.486**	1.211	0.504	**0.723**	1.100	**0.486**	1.209	0.504	**0.723**	1.036	0.486	1.228	0.503	0.755	**1.026**	0.519	**1.160**	**0.496**	0.903	1.186	0.493	1.190	0.497	0.759
ETTh2	96	MAE	0.985	**0.341**	1.119	0.396	0.568	0.985	**0.341**	1.093	0.396	0.568	0.964	0.341	**1.059**	0.399	0.603	1.213	0.345	1.559	**0.393**	0.723	**0.914**	0.343	1.099	0.401	0.657
		MSE	1.495	**0.290**	1.917	0.355	0.673	1.495	**0.290**	1.861	0.355	0.673	1.393	0.291	**1.722**	0.355	0.742	2.283	0.297	3.951	**0.350**	0.993	**1.243**	0.291	1.884	0.356	0.863
	192	MAE	**1.160**	0.394	**1.233**	0.445	0.797	**1.160**	0.394	1.251	0.445	0.797	1.216	0.394	1.237	0.441	0.794	1.818	0.399	2.119	0.437	0.910	1.391	**0.393**	1.410	0.438	0.832
		MSE	**2.112**	0.369	2.415	0.434	1.268	**2.112**	0.369	**2.407**	0.434	1.268	2.312	0.370	2.430	0.430	1.324	5.050	0.380	6.448	0.425	1.467	3.050	**0.365**	3.014	0.427	1.326
	336	MAE	1.229	0.429	1.268	0.480	0.834	1.229	0.429	1.284	0.480	0.834	**1.228**	**0.428**	**1.259**	0.480	0.811	1.802	0.430	1.849	0.475	1.072	1.356	0.433	1.343	0.480	1.120
		MSE	2.378	0.417	2.314	0.475	1.379	2.378	0.417	2.388	0.475	1.379	**2.284**	**0.416**	**2.273**	0.476	1.316	5.019	0.418	4.921	0.469	1.886	2.854	0.420	2.696	0.475	2.004
	720	MAE	**1.290**	0.449	**1.305**	0.485	**1.041**	**1.290**	0.449	1.306	0.485	**1.041**	1.306	0.448	1.327	0.490	1.058	1.480	**0.446**	1.661	0.488	1.554	1.379	0.449	1.448	**0.483**	1.369
		MSE	2.399	0.429	2.436	0.480	**1.914**	2.399	0.429	**2.427**	0.480	**1.914**	2.386	0.428	2.491	0.486	1.940	3.242	**0.426**	3.883	0.481	3.409	2.669	0.431	2.927	**0.474**	2.851
Exchange	96	MAE	1.023	**0.202**	1.021	**0.293**	0.688	1.023	**0.202**	1.031	**0.293**	0.688	0.982	0.204	0.986	0.298	0.663	**0.580**	0.209	0.750	0.297	0.373	0.782	0.202	0.838	0.298	0.472
		MSE	1.535	**0.084**	1.527	**0.164**	0.913	1.535	**0.084**	1.557	**0.164**	0.913	1.426	0.086	1.438	0.170	0.841	0.569	0.091	0.877	0.169	0.263	0.965	0.085	1.068	0.170	0.430
	192	MAE	1.121	0.300	1.081	0.385	0.772	1.121	0.300	1.075	0.385	0.772	1.097	0.300	1.029	0.385	0.754	0.783	0.300	0.860	0.383	0.522	0.915	0.300	0.983	0.385	0.644
		MSE	1.848	0.180	1.725	0.280	1.057	1.848	0.180	1.705	0.280	1.057	1.772	0.179	1.541	0.279	1.000	1.046	0.179	1.137	0.277	0.486	1.328	0.178	1.457	0.279	0.702
	336	MAE	1.203	0.418	1.186	0.487	1.177	1.203	0.418	1.186	0.487	1.177	1.192	0.417	1.174	0.487	1.159	1.046	**0.411**	1.077	0.488	0.886	**0.993**	0.411	**1.028**	**0.487**	0.964
		MSE	2.122	0.331	2.042	0.438	2.187	2.122	0.331	2.042	0.438	2.187	2.084	0.330	2.001	0.438	2.132	1.742	0.322	1.841	0.440	1.270	**1.484**	**0.321**	**1.595**	**0.438**	1.526
	720	MAE	1.111	0.703	1.150	0.830	1.338	1.111	0.703	1.163	0.830	1.338	1.079	0.710	1.147	0.820	1.323	1.246	0.715	1.396	0.824	1.048	**0.911**	**0.698**	**1.023**	**0.820**	1.128
		MSE	1.717	0.869	2.031	1.167	2.711	1.717	0.869	2.083	1.167	2.711	1.618	0.882	1.963	1.145	2.659	2.354	0.904	2.936	1.153	1.702	**1.196**	**0.861**	**1.608**	**1.145**	1.971

Table A8: Classification accuracies of different datasets using different models and normalization methods. For each dataset and for each model, the best performing normalization method(s) are bolded, and the second best are underlined.

	dataset	Face Detection	Heartbeat	PEMS-SF	UWave Gesture Library
BatchNorm	Crossformer	50.435	**75.122**	**68.401**	**83.438**
	FEDformer	68.275	73.984	78.035	47.812
	Informer	68.606	73.984	**87.476**	82.083
	PatchTST	65.683	66.992	79.576	81.354
	Transformer	68.663	75.610	84.200	83.542
LayerNorm	Crossformer	**52.176**	73.008	26.397	82.708
	FEDformer	68.861	73.659	84.393	48.438
	Informer	68.076	73.821	85.742	81.979
	PatchTST	67.329	69.919	**84.971**	81.562
	Transformer	68.757	74.472	82.466	85.104
RMSNorm	Crossformer	51.693	73.659	23.699	81.042
	FEDformer	68.000	72.846	85.164	47.708
	Informer	68.275	75.447	84.586	83.125
	PatchTST	66.648	70.894	82.852	80.729
	Transformer	**69.154**	**75.935**	83.237	84.167
UnitNorm (k=0.0)	Crossformer	50.000	72.195	16.763	29.167
	FEDformer	**69.041**	**74.146**	83.430	**53.333**
	Informer	**69.088**	75.285	78.035	**85.000**
	PatchTST	67.546	72.195	81.118	82.396
	Transformer	68.568	**75.935**	82.659	**87.188**
UnitNorm (k=0.7)	Crossformer	50.236	73.008	65.896	81.667
	FEDformer	68.067	72.846	84.586	47.083
	Informer	68.142	72.846	83.237	83.958
	PatchTST	**67.641**	72.358	**84.971**	**82.604**
	Transformer	68.634	74.797	**84.200**	84.896
UnitNorm (k=1.0)	Crossformer	50.019	72.195	56.455	80.521
	FEDformer	67.357	73.984	**85.356**	48.125
	Informer	68.492	73.984	84.008	83.646
	PatchTST	67.452	72.033	83.044	81.146
	Transformer	68.350	73.984	80.347	84.375
UnitNorm (learnable; k=0.0)	Crossformer	50.000	72.195	16.763	29.167
	FEDformer	**69.041**	**74.146**	83.430	**53.333**
	Informer	69.079	74.959	80.539	**85.000**
	PatchTST	67.546	72.195	81.118	82.396
	Transformer	68.568	**75.935**	82.659	**87.188**
UnitNorm (learnable; k=0.7)	Crossformer	50.000	72.195	16.763	29.167
	FEDformer	**69.041**	**74.146**	83.430	**53.333**
	Informer	68.852	75.122	78.420	**85.000**
	PatchTST	67.546	72.195	81.118	82.396
	Transformer	68.568	**75.935**	82.659	**87.188**
UnitNorm (learnable; k=1.0)	Crossformer	50.000	72.195	16.763	29.167
	FEDformer	**69.041**	**74.146**	83.430	**53.333**
	Informer	68.558	**75.610**	79.576	**85.000**
	PatchTST	67.546	**72.846**	81.118	82.396
	Transformer	68.568	**75.935**	82.659	**87.188**

Table A9: Anomaly detection accuracies of MSL dataset using different models and normalization methods. For each metric and for each model, the best performing normalization method(s) are bolded, and the second best are underlined.

	metric	Accuracy	F-score	Precision	Recall
BatchNorm	Crossformer	93.507	59.110	81.410	47.903
	FEDformer	95.543	75.820	88.630	66.240
	Informer	93.040	56.927	81.760	43.733
	PatchTST	95.947	78.613	88.603	70.650
	Transformer	90.417	30.680	64.623	20.123
LayerNorm	Crossformer	96.313	80.640	90.330	72.823
	FEDformer	**96.603**	**82.427**	**90.697**	**75.537**
	Informer	96.390	81.193	90.120	73.877
	PatchTST	95.950	78.727	88.347	70.993
	Transformer	96.333	80.910	89.740	73.660
RMSNorm	Crossformer	96.307	80.613	90.323	72.790
	FEDformer	96.573	82.263	90.647	75.300
	Informer	96.373	81.067	90.097	73.680
	PatchTST	95.940	78.670	88.307	70.927
	Transformer	96.330	80.883	89.677	73.657
UnitNorm (k=0.0)	Crossformer	**96.533**	**82.000**	90.610	**74.880**
	FEDformer	96.547	82.097	90.653	75.013
	Informer	96.543	82.067	90.640	74.973
	PatchTST	**96.317**	**80.943**	**88.990**	**74.227**
	Transformer	96.540	**82.060**	**90.637**	**74.960**
UnitNorm (k=0.7)	Crossformer	96.523	81.943	**90.617**	74.783
	FEDformer	96.537	82.040	90.620	74.943
	Informer	**96.557**	**82.147**	**90.643**	**75.103**
	PatchTST	96.213	80.360	88.713	73.440
	Transformer	**96.540**	**82.043**	90.580	74.977
UnitNorm (k=1.0)	Crossformer	96.347	80.857	90.380	73.143
	FEDformer	96.587	82.350	90.670	75.433
	Informer	96.470	81.653	90.400	74.450
	PatchTST	95.933	78.607	88.273	70.847
	Transformer	96.337	80.933	89.753	73.687
UnitNorm (learnable; k=0.0)	Crossformer	**96.533**	**82.000**	90.610	**74.880**
	FEDformer	96.547	82.097	90.653	75.013
	Informer	96.543	82.067	90.640	74.973
	PatchTST	**96.317**	**80.943**	**88.990**	**74.227**
	Transformer	96.540	**82.060**	**90.637**	**74.960**
UnitNorm (learnable; k=0.7)	Crossformer	**96.533**	**82.000**	90.610	**74.880**
	FEDformer	96.547	82.097	90.653	75.013
	Informer	96.543	82.067	90.640	74.973
	PatchTST	**96.317**	**80.943**	**88.990**	**74.227**
	Transformer	96.540	**82.060**	**90.637**	**74.960**
UnitNorm (learnable; k=1.0)	Crossformer	**96.533**	**82.000**	90.610	**74.880**
	FEDformer	96.547	82.097	90.653	75.013
	Informer	96.543	82.067	90.640	74.973
	PatchTST	**96.317**	**80.943**	**88.990**	**74.227**
	Transformer	96.540	**82.060**	**90.637**	**74.960**

Table A10: Metrics used for characterizing the distribution of attention scores from the original and normalized data, denoted as $\mathbf{A}_{n,i}$ and $\tilde{\mathbf{A}}_{n,i}$, respectively, for the i-th sample in the n-th batch.

Metric	Definition	Evaluation
Chebyshev distance	$D_{\text{Chebyshev}}\left(\mathbf{A}_{n,i}, \tilde{\mathbf{A}}_{n,i}\right) = \max_{j=1}^{L} \left\lvert \mathbf{A}_{n,i,j} - \tilde{\mathbf{A}}_{n,i,j} \right\rvert$	Lower is Better
Cosine similarity	$D_{\text{Cosine}}\left(\mathbf{A}_{n,i}, \tilde{\mathbf{A}}_{n,i}\right) = \frac{\mathbf{A}_{n,i}^{\top} \tilde{\mathbf{A}}_{n,i}}{\lVert \mathbf{A}_{n,i} \rVert \lVert \tilde{\mathbf{A}}_{n,i} \rVert}$	Higher is Better
KL divergence	$D_{\text{KL}}\left(\mathbf{A}_{n,i} \lVert \tilde{\mathbf{A}}_{n,i}\right) = \sum_{j=1}^{L} \mathbf{A}_{n,i,j} \left(\log \mathbf{A}_{n,i,j} - \log \tilde{\mathbf{A}}_{n,i,j}\right)$	Lower is Better
Entropy	$E\left(\tilde{\mathbf{A}}_{n,i}\right) = -\sum_{j=1}^{L} \tilde{\mathbf{A}}_{n,i,j} \log \tilde{\mathbf{A}}_{n,i,j}$	Higher is Better

Table A11: Last-value predictor versus PatchTST with different normalization methods on Exchange dataset. Best results *aside from last-value predictor* are in **bold**.

Model & Metric\Prediction	96	192	336	720
Last-value predictor (MAE)	0.196	0.289	0.398	0.676
PatchTST + UnitNorm (MAE)	**0.202**	**0.300**	**0.411**	**0.698**
PatchTST + BatchNorm (MAE)	0.204	0.300	0.419	0.703
PatchTST + LayerNorm (MAE)	0.209	0.301	0.412	0.713
PatchTST + RMSNorm (MAE)	0.209	0.300	0.411	0.715
Last-value predictor (MSE)	0.081	0.167	0.306	0.810
PatchTST + UnitNorm (MSE)	**0.085**	**0.178**	**0.321**	**0.861**
PatchTST + BatchNorm (MSE)	0.088	0.180	0.336	0.874
PatchTST + LayerNorm (MSE)	0.091	0.180	0.324	0.901
PatchTST + RMSNorm (MSE)	0.091	0.179	0.322	0.904

F Comparison with Last-Value Predictor

To provide a more comprehensive evaluation of UnitNorm's effectiveness, we compared it against a simple but effective baseline: the last-value predictor. This baseline simply uses the last observed value as the prediction for all future time steps, which can be surprisingly effective for time series with slow-changing patterns or high levels of noise.

Tables A11 and A12 present the comparison between the last-value predictor and PatchTST model with different normalization methods on the Exchange and ETTh2 datasets, respectively.

On the Exchange dataset (Table A11), the last-value predictor outperforms all Transformer-based models. This is likely due to the highly stochastic nature of exchange rates, where the most recent value is often the best predictor for future values. However, among the Transformer models, UnitNorm consistently achieves the best performance across all prediction horizons for both MSE and MAE metrics. This demonstrates that even when sophisticated models struggle to outperform simple baselines on challenging datasets, UnitNorm still provides relative advantages over other normalization techniques.

On the ETTh2 dataset (Table A12), all Transformer models substantially outperform the last-value predictor, with UnitNorm showing the best overall performance. UnitNorm achieves the best MSE scores across all prediction horizons and the best MAE for horizons 96, 192, and 336. This confirms UnitNorm's ability to enable Transformer models to effectively capture periodic patterns present in this dataset, as indicated by the higher periodicity scores shown in Table A14.

These results highlight an important insight: UnitNorm's effectiveness is most pronounced on datasets with strong periodic components (like ETTh2), while its advantages may be less significant on datasets where complex models are less suitable overall (like Exchange). This aligns with our theoretical analysis that UnitNorm helps preserve attention distributions and mitigate token shift issues, which are particularly beneficial for modeling periodic patterns in time series data.

Table A12: Last-value predictor versus PatchTST with different normalization methods on ETTh2 dataset. Best results are in **bold**.

Model & Metric\Prediction	96	192	336	720
Last-value predictor (MAE)	0.422	0.473	0.511	0.519
PatchTST + UnitNorm (MAE)	**0.341**	**0.394**	**0.429**	0.449
PatchTST + BatchNorm (MAE)	0.346	0.399	0.435	0.453
PatchTST + LayerNorm (MAE)	0.345	0.399	0.431	**0.447**
PatchTST + RMSNorm (MAE)	0.345	0.399	0.431	0.448
Last-value predictor (MSE)	0.432	0.534	0.597	0.594
PatchTST + UnitNorm (MSE)	**0.290**	**0.369**	**0.417**	**0.429**
PatchTST + BatchNorm (MSE)	0.292	0.376	0.421	0.435
PatchTST + LayerNorm (MSE)	0.297	0.380	0.417	0.429
PatchTST + RMSNorm (MSE)	0.297	0.379	0.418	0.429

G Selection of hyperparameter k

The hyperparameter k in UnitNorm plays a crucial role in determining the sparsity of attention scores and consequently affects model performance. To empirically identify the optimal k values for time series applications, we conducted extensive experiments with various k values on multiple datasets.

G.1 Experimental setup

We performed a comprehensive sweep of fixed k values ($k = 0.1, 0.3, \ldots, 0.9, 1.1$) on the ETTh2 datasets using the PatchTST model. We trained the model with different prediction horizons (96, 192, 336, and 720 time steps) and measured the Mean Square Error (MSE). All experiments were conducted with 3 different random seeds (41, 42, and 43), and the reported results are averaged across these runs.

G.2 Results on ETTh2

Table A13 presents the MSE performance of PatchTST on the ETTh2 dataset with different fixed k values. The results show that the optimal k value depends on the prediction horizon, but generally falls within the range of 0.5 to 0.7. For the ETTh2 dataset, $k = 0.7$ achieves the best performance for prediction horizons of 96, 336, and 720, while $k = 0.5$ performs best for the horizon of 192.

G.3 Analysis

Our experiments reveal several important findings regarding the selection of k:

1. **Performance Robustness**: The performance of UnitNorm is generally not highly sensitive to small perturbations in the k value. This is evident from the relatively small percentage differences between adjacent k values, especially in the 0.5 to 0.7 range.

2. **Optimal Range**: The optimal k value consistently falls within the range of 0.5 to 0.7 across different prediction horizons and datasets. This range provides a good balance between attention diversity and focus.

Table A13: MSE of PatchTST on ETTh2 with different fixed k values. The relative difference from the best performing k value is shown in parentheses. The best performing k for each prediction length is highlighted in **bold**.

k \ prediction length	96	192	336	720
0.1	0.2897 (+0.633%)	0.3716 (+1.866%)	0.4148 (+0.3%)	0.4278 (+0.787%)
0.3	0.293 (+1.769%)	0.3668 (+0.535%)	0.414 (+0.088%)	0.4314 (+1.629%)
0.5	0.29 (+0.719%)	**0.3648 (+0.0%)**	0.4197 (+1.469%)	0.4252 (+0.168%)
0.7	**0.2879 (+0.0%)**	0.3717 (+1.896%)	**0.4136 (+0.0%)**	**0.4245 (+0.0%)**
0.9	0.2967 (+3.069%)	0.3764 (+3.187%)	0.4175 (+0.937%)	0.4283 (+0.904%)
1.1	0.2987 (+3.744%)	0.365 (+0.057%)	0.4188 (+1.265%)	0.435 (+2.472%)

3. **Performance Gain**: The appropriate selection of k can lead to moderate but meaningful performance improvements. For example, using $k = 0.7$ instead of $k = 1.1$ for the 720-step prediction horizon on ETTh2 results in a 2.5% reduction in MSE.

4. **Dataset Dependence**: The optimal k value is somewhat dataset-dependent, reflecting the varying degrees of periodicity and other temporal patterns across different time series data.

G.4 Practical Recommendations

Based on our findings, we offer the following practical recommendations for selecting k values:

- For datasets with strong periodicity, starting with $k = 0.7$ is recommended as it generally provides good performance across various prediction horizons.
- For datasets where the optimal k is uncertain, we recommend trying both $k = 0.5$ and $k = 0.7$ to determine which works better for the specific application.
- In scenarios requiring maximum flexibility, implementing UnitNorm with a learnable k parameter allows the model to adaptively determine the optimal attention sparsity during training.
- For comprehensive optimization, a validation-based approach can be employed where different fixed k values are evaluated on a validation set to select the best configuration.

H Extended Experimental Results

This section provides additional experimental results and analyses that supplement the main experiments presented in Section 4.

H.1 Periodicity Measurement Analysis

Time series periodicity is a crucial factor in forecasting performance, especially for long-term predictions. To quantify the periodic patterns present in our experimental datasets, we conducted a comprehensive periodicity analysis using normalized correlation between input and expected output series.

As shown in Table A14, the datasets exhibit varying degrees of periodicity across different prediction horizons. ETTh2 consistently shows the strongest periodic patterns (with scores around 0.41-0.42), followed by Exchange (0.33-0.40) and ETTh1 (0.34-0.35), while ECL demonstrates the weakest periodicity (0.16-0.17).

Interestingly, the Exchange dataset shows an increasing trend in periodicity as the prediction horizon extends, reaching its highest correlation value (0.396) at the 720-step horizon. This suggests that longer-term patterns become more apparent in financial exchange data when viewed over extended time frames.

Table A14: Periodicity measurement on datasets. Scores are given as the maximum normalized correlation between input and expected output series. Higher scores indicate stronger periodic patterns that may be leveraged by forecasting models.

Dataset\Prediction	96	192	336	720
ETTh1	0.347	0.351	0.353	0.338
ETTh2	0.404	0.415	0.420	0.413
ECL	0.156	0.162	0.165	0.158
Exchange	0.328	0.340	0.353	0.396

The ability of UnitNorm to maintain consistent attention distributions, as discussed in Section 3, makes it particularly well-suited for capturing these periodic patterns, especially in datasets with stronger periodicity. This helps explain UnitNorm's superior performance on datasets like ETTh2 and Exchange in the long-term forecasting experiments.

H.2 Performance on Large-Scale Dataset with Modern Architecture

To evaluate UnitNorm's effectiveness beyond standard benchmarks, we conducted additional experiments using the Pathformer model on the Solar dataset, which represents a significant increase in scale and complexity compared to our other test datasets.

Solar Dataset The Solar dataset is a large-scale time series collection containing data from 137 photovoltaic power plants across the United States, with approximately 52,000 samples. Its high dimensionality and real-world nature make it an excellent test case for evaluating normalization techniques in complex, practical scenarios.

Pathformer Model Pathformer is a state-of-the-art time series Transformer architecture that introduces innovations in handling multivariate time series through path-wise modeling. Unlike traditional Transformer architectures, Pathformer incorporates specialized mechanisms for capturing both temporal and cross-series dependencies.

Table A15: Pathformer performance on the Solar dataset (large-scale, 137 channels, 52K samples) with prediction horizon of 96. Results are averaged over 3 random seeds with standard deviation reported. Best results are in **bold**; second best are in *italic*. These results demonstrate UnitNorm's effectiveness even with modern architectures and large-scale datasets.

Normalization	MSE	MAE
BatchNorm	0.2215 ± 0.0080	0.2075 ± 0.0121
LayerNorm	0.2177 ± 0.0064	**0.1996 ± 0.0100**
RMSNorm	0.2225 ± 0.0031	0.2097 ± 0.0151
UnitNorm ($k = 0.0$)	0.2202 ± 0.0038	0.2062 ± 0.0082
UnitNorm ($k = 0.5$)	**0.2176 ± 0.0041**	*0.2053 ± 0.0181*
UnitNorm ($k = 0.7$)	*0.2177 ± 0.0091*	0.2074 ± 0.0133
UnitNorm ($k = 1.0$)	0.2267 ± 0.0092	0.2074 ± 0.0054

Result Analysis As shown in Table A15, UnitNorm with $k = 0.5$ achieves the best MSE (0.2176 ± 0.0041) and competitive MAE (0.2053 ± 0.0181, second only to LayerNorm). Several observations are worth noting:

- UnitNorm's performance remains strong even with more complex model architectures and larger datasets, demonstrating its generalizability.
- The optimal value of k for this dataset appears to be around 0.5-0.7, which aligns with our findings in the main experiments.
- While LayerNorm achieves the best MAE, UnitNorm provides more balanced performance across both MSE and MAE metrics.
- Compared to RMSNorm and BatchNorm, UnitNorm consistently delivers superior results, reinforcing its advantages over these traditional normalization methods.

These results further substantiate UnitNorm's potential as a broadly applicable normalization technique for time series analysis tasks across various model architectures and dataset scales. The consistent performance on the Solar dataset, with its high dimensionality and large sample size, suggests that UnitNorm's benefits extend to real-world, large-scale applications.

I Experimental Setup

We provide detailed experimental configurations to facilitate reproducibility. All experiments were conducted on 4× NVIDIA A5000 (24GB) GPUs using PyTorch Paszke et al.. The source code is available at `https://anonymous.4open.science/r/UnitNorm-5B84`.

I.1 Model Architectures

Table A16 summarizes the key architectural hyperparameters across all evaluated Transformer-based models. Unless otherwise specified, all models employ the post-normalization (Post-LN) architecture (Xiong et al., 2020), where normalization is applied after each residual connection. UnitNorm serves as a drop-in replacement for existing normalization layers; the models are otherwise identical to their published implementations.

Table A16: Model architectural hyperparameters. d_{model} denotes the hidden dimension, d_{ff} the feed-forward dimension, h the number of attention heads, N the number of encoder layers, and M the number of decoder layers. Models marked with * use only the encoder; their decoder layer count is shown as "—". Crossformer uses dataset-dependent configurations; values shown are for the ETT datasets.

Model	d_{model}	d_{ff}	h	N	M
PatchTST*	512	2048	8	2	—
Transformer	512	2048	8	2	1
FEDformer	512	2048	8	2	1
Crossformer	512	2048	8	2	1
Nonstationary-Transformer	512	2048	8	2	1
Informer	512	2048	8	2	1

For the classification and anomaly detection tasks, the encoder-only PatchTST model is used. For the long-term forecasting task, all six models in Table A16 are evaluated. The number of attention heads h is adjusted per prediction length to maintain appropriate representational capacity while preventing overfitting for shorter prediction horizons; exact per-configuration values are provided in the released experiment scripts. The hidden dimension d_{model} and feed-forward dimension d_{ff} follow the default settings from the respective model publications.

I.2 Training Configuration

Table A17 lists the training hyperparameters shared across all experiments. The Adam optimizer (Kingma & Ba) is used throughout with default moment parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$. A learning rate decay strategy (type "type1") is applied: the learning rate is reduced by half after each epoch following a predefined schedule. Early stopping with patience of 3 epochs is employed, monitoring validation loss. Each experiment is repeated 3 times (itr=3) with fixed random seed 2021, and the mean and standard deviation of the test metrics are reported.

For the long-term forecasting task, the input sequence length is fixed at 96, and prediction lengths are set to $\{96, 192, 336, 720\}$ for ETT, Weather, Electricity, and Exchange datasets, and $\{24, 36, 48, 60\}$ for the ILI dataset, following the standard evaluation protocol (Wu et al., 2023). The batch size is 32 for most datasets, and increased to 128 for large-scale datasets (Weather, Electricity, Traffic). Models use early stopping and typically converge within 3–10 epochs; exact per-dataset epoch budgets are provided in the released experiment scripts.

Table A17: Training configuration shared across all experiments.

Hyperparameter	Value
Optimizer	Adam (β_1=0.9, β_2=0.999)
Initial learning rate	10^{-4}
Learning rate schedule	halving per epoch (type1)
Loss function	MSE (forecasting, anomaly detection); Cross-Entropy (classification)
Batch size	32 (128 for Weather/Electricity/Traffic)
Early stopping patience	3 epochs
Maximum training epochs	10
Dropout	0.1
Activation	GeLU
Random seed	2021
GPU	4× NVIDIA A5000 (24GB)

I.3 Baseline Fairness

To ensure fair comparison, all normalization variants (LayerNorm, BatchNorm, RMSNorm, and UnitNorm with different k values) are evaluated under identical training configurations for each model–dataset pair. No model-specific hyperparameter tuning is performed beyond the standard settings prescribed by the original publications. The only varying factor is the normalization method, implemented as a drop-in replacement via the `-normalization` command-line argument. For UnitNorm variants, the k parameter is swept over $\{0.0, 0.5, 0.7, 1.0\}$ in preliminary experiments, and the best-performing fixed k values (typically $k = 0.5$ or 0.7) are reported in the main results. When k is learnable, it is initialized to 0.7 and optimized alongside model parameters using the same Adam optimizer.

For datasets referenced from prior work, the data splits and preprocessing pipelines are identical to those in Wu et al. (2023), ensuring direct comparability of results.

J Ablation: UnitNorm vs RMSNorm + Temperature Scaling

Setup. All experiments use the long-term forecasting setup (96→96) on PatchTST (Nie et al., 2022) and Crossformer (Zhang & Yan, 2022) with ETTh2 and Weather, following hyperparameters in Section I. The equivalence formula $\tau = D^{1-k}$ maps UnitNorm(k) to RMSNorm+τ at the attention-logit level; any remaining gap stems from V-path, FFN-path, and k's dual role (attention sparsity via ELB and signal magnitude via $D^{k/2}$), as discussed in Section 3.4.

Single-point comparison (k=0.7). Table A18 reports the $k = 0.7$ comparison over 3 random seeds across two architectures and datasets (24 runs total).

Table A18: Single-point comparison: UnitNorm($k = 0.7$) vs RMSNorm+τ ($\tau = D^{0.3}$). MSE (lower is better). Δ computed as $(\text{RMS} + \tau - \text{UN})/(\text{RMS} + \tau)$.

Model×Dataset	UnitNorm($k = 0.7$)	RMSNorm+τ	Δ
PatchTST + ETTh2	0.2957 ± 0.0035	0.2991 ± 0.0061	+1.15%
PatchTST + Weather	0.1746 ± 0.0005	0.1743 ± 0.0003	−0.19%
Crossformer + ETTh2	0.7987 ± 0.0357	0.8842 ± 0.2375	+9.67%
Crossformer + Weather	0.1760 ± 0.0002	0.1743 ± 0.0021	−0.99%
Combined	$\mathbf{0.3613 \pm 0.2692}$	$\mathbf{0.3830 \pm 0.3232}$	**+5.67%**

On ETTh2, where strong periodic structure makes attention modulation critical, UnitNorm consistently outperforms RMSNorm+τ across both architectures. On Weather, where MSE is an order of magnitude lower, the methods effectively tie.

k/τ sweep. We sweep $k \in \{0.5, 0.7, 1.0\}$ on PatchTST + ETTh2 with 3 seeds per condition.

Table A19: k/τ sweep: UnitNorm vs RMSNorm +τ on PatchTST + ETTh2. MSE (lower is better).

k	UnitNorm(k)	RMSNorm+τ	Δ
0.5	0.2927 ± 0.0074	0.2956 ± 0.0043	+0.99%
0.7	0.2957 ± 0.0035	0.2991 ± 0.0062	+1.15%
1.0	0.2981 ± 0.0029	0.2978 ± 0.0027	−0.11%

UnitNorm achieves lower MSE at $k = 0.5$ and $k = 0.7$; at $k = 1.0$ where UnitNorm reduces to RMSNorm the two methods tie, confirming internal consistency.

Partial replacement. We isolate the attention-path and FFN-path contributions of k in this Post-LN architecture (Figure 3), where norm_1 normalizes the FFN input while norm_2 normalizes the layer output, which becomes the next layer's attention input. All configurations use the same encoder depth ($e\_layers = 2$) as the main experiments. PatchTST + ETTh2, 3 seeds per condition (mean ± s.d.).

Attention path: $\text{norm}_2 = \text{UN}(k)$, $\text{norm}_1 = \text{RMS}$ (fixed). *FFN path:* $\text{norm}_1 = \text{UN}(k)$, $\text{norm}_2 = \text{RMS}$ (fixed). *Control:* both RMS.

Table A20: Path isolation: attention-path vs FFN-path contribution of k. MSE on PatchTST + ETTh2, 3 seeds per condition. Control is independent of k (multi-column).

	$k = 0.5$	$k = 0.7$	$k = 1.0$
Full UN (norm_1=UN, norm_2=UN)	0.2927 ± 0.0074	0.2957 ± 0.0035	0.2981 ± 0.0029
Attn path (norm_2=UN, norm_1=RMS)	0.2932 ± 0.0076	0.2937 ± 0.0057	0.2980 ± 0.0029
FFN path (norm_1=UN, norm_2=RMS)	0.2978 ± 0.0024	0.2977 ± 0.0024	0.2978 ± 0.0027
Control (norm_1=RMS, norm_2=RMS)		0.2978 ± 0.0027	

The attention path recovers most of UnitNorm's benefit across all k values, with MSE statistically indistinguishable from the full UN configuration. The FFN path is essentially flat across k (all values within ~ 0.0002 of the control), confirming that k's effect operates through the attention input path (via norm_2), consistent with Theorem 3.2, and that the ELB-driven attention sparsity modulation is the primary mechanism.