

Incorporating Exponential Smoothing into MLP: A Simple but Effective Sequence Model

Jiquan Chu
Peking University
chujiquan@pku.edu.cn

Zuoquan Lin *
Peking University
linzuoquan@pku.edu.cn

Abstract

Modeling long-range dependencies in sequential data is a crucial step in sequence learning. A recently developed model, the Structured State Space (S4), demonstrated significant effectiveness in modeling long-range sequences. However, it is unclear whether the success of S4 can be attributed to its intricate parameterization and HiPPO initialization or simply due to State Space Models (SSMs). To further investigate the potential of the deep SSMs, we start with exponential smoothing (ETS), a simple SSM, and propose a stacked architecture by directly incorporating it into an element-wise MLP. We augment simple ETS with additional parameters and complex field to reduce the inductive bias. Despite increasing less than 1% of parameters of element-wise MLP, our models achieve comparable results to S4 on the LRA benchmark.¹

1 Introduction

Transformer (Vaswani et al., 2017) and its variants have been the most successful architecture in various domains of deep learning. However, the self-attention layer, which plays a crucial role in contextualizing the input, poses a significant computational and memory burden with a complexity of $O(L^2)$. This limitation hinders the application of the transformers in modeling long sequences, particularly when operating under hardware constraints, which is a common scenario for large language models. To alleviate this issue, several models have been proposed to reduce the computational and memory requirements of the transformers (Beltagy et al., 2020; Choromanski et al., 2020; Kitaev et al., 2020; Wang et al., 2020; Guo et al., 2021; Kasai et al., 2021; Peng et al., 2021; Dao et al., 2022; Hua et al., 2022; Tay et al., 2022; Fournier et al., 2023; Zandieh et al., 2023). Despite these efforts, all the models are only partial

¹Our codes and scripts are available at <https://github.com/PKUAI-LINGroup/ETSMLP>.

modifications of the attention mechanism and struggle to perform well on long-range sequence benchmarks such as Long Range Arena (LRA) (Tay et al., 2020).

In a recent breakthrough result, (Gu et al., 2021) introduced a novel framework called the "structured state space sequence" (S4) that leveraged the State Space Models (SSMs). S4 builds upon continuous-time SSMs and addresses the computational bottleneck of previous approaches by introducing the Normal Plus Low-Rank (NPLR) decomposition of the state matrices. Additionally, the initialization of state matrices utilizes HiPPO matrices which have been demonstrated to be effective in sequence learning in (Gu et al., 2020). Notably, S4 exhibited exceptional performance across various sequential tasks, particularly in the LRA, where it outperformed the existing transformer variants by an impressive accuracy.

Despite the impressive performance of S4, its intricate parameterization and strict initialization schemes impede researchers from fully comprehending, implementing, and analyzing the model. Although there have been attempts to simplify the S4 framework by (Smith et al., 2022; Gupta et al., 2022), these models still required the HiPPO initialization process. Other studies have explored the relationship between SSMs and recurrent units or global convolutions and demonstrated strong performance on various tasks (Li et al., 2022; Orvieto et al., 2023). These works highlight the potential of SSMs and suggest that simpler yet effective SSM architectures may exist.

In our work, we deviate from the methodology proposed by S4, which begins with the continuous SSM and then simplifies the process. We initiate our approach with a discrete SSM, namely Exponential Smoothing (ETS), and introduce additional parameters to reduce the inductive bias. This alternative approach offers two notable advantages. Firstly, it circumvents the simplification of

the continuous SSMs that need sophisticated mathematical derivations and thus enhances accessibility and comprehensibility. Secondly, it explores the possibility of random initialization departing from HiPPO initialization for continuous SSMs. For the streamlining of the model, our architecture directly integrates a parameterized ETS into an element-wise Multi-Layer Perceptron (MLP). By incorporating less than 3% of the total parameters after the initial linear layer of the MLP, we successfully transform a channel-only MLP into a sequence learner.

We conducted experiments on multiple datasets, including the LRA and several Natural Language Understanding (NLU) datasets. Despite its simplicity, surprisingly, our model performs comparably to S4. In all six tasks in the LRA, our results slightly surpass the performance of S4 and DSS by 2.61 points on average and significantly outperform the transformer variants by about 20 points. In addition, we evaluated our model on seven NLU datasets and consistently achieved comparable performance with the transformer encoders. The findings of our work shed light on the potential of SSMs from a unique standpoint, where simply incorporating an ETS into an MLP can achieve a similar effect as the transformer model. A thorough examination of the proposed model was undertaken through an ablation study on the hyperparameters and an evaluation of the initialization method. Additional experiments were conducted to compare our model with the transformer model for efficiency and memory utilization, especially in handling lengthy texts. The results of these experiments provide evidence of the advantages of our model over the transformer model in terms of time and memory complexity.

In summary, our main contributions are as follows:

- We introduce the Exponential Smoothing Multi-Layer Perceptron (ETSMLP) model. We integrate the enhanced ETS module into an element-wise MLP to create an effective sequence model.
- We evaluate ETSMLP on the LRA and conduct comparative experiments with transformer encoders on various NLU datasets. The empirical results demonstrate the effective capacity in long-range sequence modeling.
- We conduct ablation studies on the proposed

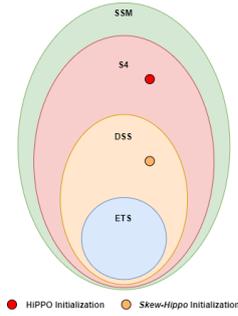


Figure 1: The relations among SSM, S4, DSS, and ETS. The HiPPO initialization is pointed out in red while the *Skew-Hippo* initialization is pointed out in orange.

parameters and initialization methods. Additionally, we emphasize the advantages of SSMs over the attention mechanism in speed and memory efficiency.

2 Preliminaries

We introduce basic notations and briefly review SSMs and ETS in this section. Focusing on time-invariant sequence models, we aim to transform a sequence of inputs $X = \{x_1, \dots, x_n\} \in \mathbb{R}^{n \times d}$ into a corresponding sequence of outputs $Y = \{y_1, \dots, y_n\} \in \mathbb{R}^{n \times d}$ with each output y_i is exclusively based on historical data x_1, \dots, x_i .

2.1 State space models

The continuous-time SSM is characterized by the differential equation (1), which establishes the relationship between a continuous-time scalar input $x(t)$ to a scalar output $y(t)$ with the state matrix $A \in \mathbb{R}^{d \times d}$ and vectors $B \in \mathbb{R}^{d \times 1}$, $C \in \mathbb{R}^{1 \times N}$:

$$\frac{dh}{dt}(t) = Ah(x) + Bx(t) \quad , \quad y(t) = Ch(t). \quad (1)$$

If we set a sample time interval $\Delta > 0$, and assume that the duration of sampling remains constant Δ , we convert the continuous-time SSM into a discrete-time one using a recursive equation in the following:

$$h_k = \bar{A}h_{k-1} + \bar{B}x_k \quad , \quad y_k = \bar{C}h_k, \quad (2)$$

where $\bar{A} = e^{A\Delta}$, $\bar{B} = (\bar{A} - I)A^{-1}B$ and $\bar{C} = C$. With $x_0 = 0$, we unroll this recursion explicitly as the equation (3) which can be vectorized into a convolution in the equation (5) with the SSM convolution kernel defined in the equation (4) as follows:

$$y_k = \sum_{j=0}^k \bar{C} \bar{A}^j \bar{B} x_{k-j}, \quad (3)$$

$$\bar{K} \in \mathbb{R}^L = (\bar{C} \bar{B}, \dots, \bar{C} \bar{A}^{L-1} \bar{B}), \quad (4)$$

$$y = \bar{K} * x. \quad (5)$$

If we obtain the kernel \bar{K} , the convolution function aforementioned can be efficiently computed with Fast Fourier Transform (FFT) in $O(L \log(L))$ (Cormen et al., 2022). Nevertheless, the main challenge in the computation of SSMs is how to efficiently compute \bar{K} from the matrices \bar{A} , \bar{B} and \bar{C} . S4 proposes an effective parameterization through decomposing matrix \bar{A} to the NPLR matrices (Gu et al., 2021), and diagonal state spaces (DSS) only consider the circumstances when A is diagonalizable over \mathbb{C} (Gupta et al., 2022). Both methods involve intricate mathematics, sophisticated parameterization, and strict initialization, all of which are indispensable for achieving excellent performance. Our method will start from a special SSM, namely ETS, which gives a new insight into this problem, and requires fewer mathematical operations, fewer parameters, and more flexible initialization.

2.2 Exponential smoothing

ETS is a time series forecasting method that utilizes a weighted average of past observations to predict future values (Winters, 1960; Hunter, 1986; Hyndman et al., 2008). The fundamental idea behind ETS is to give more weight to recent observations and less to older ones, with the weights decreasing exponentially as the observations get older. The core recursive equation for this method is the equation (6) with the smoothing factor λ in the range $(0, 1)$:

$$y_t = \lambda x_t + (1 - \lambda)y_{t-1}. \quad (6)$$

ETS is a special SSM, with the substitution $\bar{A} = 1 - \lambda$, $\bar{B} = \lambda$, $\bar{C} = 1$. However, compared with SSMs, ETS cannot capture sequential information effectively. Figure 1 illustrates the relationship among SSMs, S4, DSS, and ETS. S4 and DSS are derived from the continuous-time SSMs with the difference that S4 decomposes the matrix A into an NPLR matrix, while DSS assumes A to be diagonalizable. As a result, the HiPPO initialization in S4 cannot directly adapt to DSS (Gu et al., 2020). The initialization in DSS is *skew-Hippo* initialization which is the normal part of the HiPPO matrix.

ETS serves as a special case within the realm of discrete-time SSM. In our approach, we incorporate parameters directly from ETS, distinguishing ours from S4 and DSS methods that simplify equations based on continuous-time SSMs.

3 Exponential Smoothing Multi-layer Perceptrons

In this section, we present our ETSMLP. We first introduce a complex exponential smoothing module which is the pivotal component of our architecture. We then describe the full architecture with two proposed versions, ESMLP and ESMLP-Gate.

3.1 Complex exponential smoothing module

Learnable damped factors. Damped factors are a commonly used technique of ETS for attenuating the influence of specific factors (Gardner Jr, 1985; McKenzie and Gardner Jr, 2010). We introduce two learnable damped factors α and β into simple ETS in equation (7). The factor α controls the learning of λ in an exponential scalar. A small α close to zero amplifies the impact of λ and results in λ^α approximating 1 while a large α diminishes its impact, driving the combination closer to 0. The factor β serves as a multiplicative factor that controls the influence of the current input x_t . The recursion can be unrolled in equation (8) with the kernel defined by the equation (9) as follows:

$$y_t = (1 - \lambda^\alpha)\beta x_t + (\lambda^\alpha)y_{t-1}, \quad (7)$$

$$y_t = \sum_{i=0}^t (\lambda^\alpha)^i (1 - \lambda^\alpha)\beta x_{t-i}, \quad (8)$$

$$K = ((1 - \lambda^\alpha)\beta, \dots, (\lambda^\alpha)^{L-1}(1 - \lambda^\alpha)\beta). \quad (9)$$

Complex parameters. Complex parameters in ETS have been demonstrated to capture both level and trend information in forecasting (Svetunkov et al., 2022). By extending the learning capacity and enlarging the parameter space, the transformation from real to complex numbers is beneficial. Therefore, we treat α , λ , β as complex numbers, and keep the input x_t and the output y_t real. Consequently, only the real part of the kernel coefficients is utilized, and the corresponding computation formula is as follows:

$$y_t = \sum_{i=0}^t \Re((\lambda^\alpha)^i (1 - \lambda^\alpha)\beta) x_{t-i}. \quad (10)$$

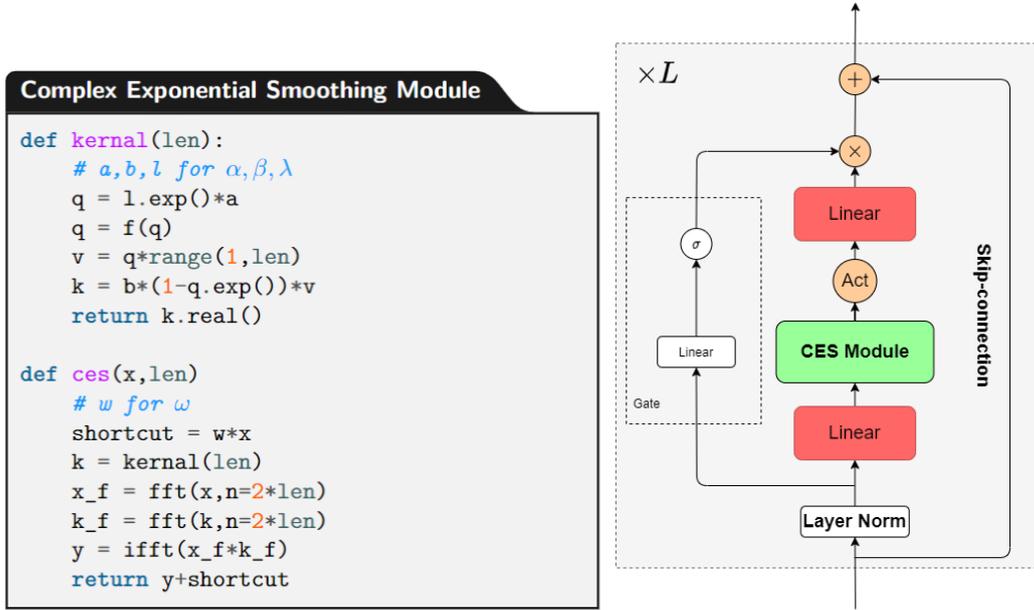


Figure 2: Overview of the ETSMLP architecture. The left is the pseudo-code of the complex exponential smoothing (CES) module. The right is the entire architecture with a gate mechanism.

Exponential parameterization. Directly training the parameters λ is infeasible, due to the rapid explosion of the gradient λ within a few steps. This challenge becomes evident upon inspecting the equation (6). Given the gradients of $y_t \frac{dL}{dy_t}$, the gradient of $\lambda \frac{dL}{d\lambda}$ could be derived from the formulas (11). This reveals that the gradients of λ are proportional to $\frac{1}{1-\lambda}$. Consequently, as λ approaches 1, the gradients of λ will explode.

$$\begin{aligned}
\frac{dL}{d\lambda} &= \sum_{t=1}^N \frac{dL}{dy_t} \sum_{j=1}^t \left[\frac{(t-j)}{\lambda} - (t-j+1) \right] \lambda^{t-j} x_j \\
&\approx \sum_{t=1}^N \frac{dL}{dy_t} \sum_{j=1}^t (-1 * \lambda^{t-j} x_j) = - \frac{\sum_{t=1}^N (\frac{dL}{dy_t} y_t)}{1-\lambda} \quad (11)
\end{aligned}$$

To address this issue, we propose an exponential transformation of the parameters. We train the parameter $\lambda' = \log \log \lambda$ instead of the λ . We prove that the stability of learning λ' constrains the gradients of λ' within a specified range, as described in Proposition 1.

Proposition 1. *Let $\lambda \in \mathbb{C}$ be within the interior of the hollow unit disc $D^\circ(0, 1) = \{z \mid |z| < 1\} \setminus \{(0, 0)\}$. We define $\lambda' = \log \log \lambda$ which substitutes λ in the equation (10). If the gradients of y_t satisfy $\sum_0^L \frac{dL}{dy_t} y_t < \infty$, the gradients of the real and imaginary parts of λ' are bounded for all $\lambda \in D^\circ(0, 1)$.*

The proof of Proposition 1 is elementary and is provided in the appendix A. This proposition proves

that the exponential parameterization λ' gradients are stable in the feasible region of λ .

Constraint function and shortcut. In addition to the settings as aforementioned, λ^α must lie within the feasible field $D^\circ(0, 1)$. To address this, we introduce a constraint function to enforce the validity of these parameters, which can be formulated in equation (12):

$$f(\lambda) = \begin{cases} \lambda, & \text{if } |\lambda| < \max_\lambda; \\ \frac{\max_\lambda}{|\lambda|} \lambda, & \text{if } |\lambda| \geq \max_\lambda. \end{cases} \quad (12)$$

Although this solution is simple, it yields remarkably effective results. We also explored an alternative approach inspired by the separation of real and imaginary parts, as discussed in (Gu et al., 2022; Orvieto et al., 2023). Unfortunately, its performance is unsatisfactory, because the gradients of imaginary parts appear unstable and may explode in a few steps.

Moreover, we introduce a parameter ω to establish a shortcut from input to output, a commonly used technique in deep learning. This parameter serves as a gating unit that regulates the incoming input. The final output of our model can be described with the sigmoid function σ as follows:

$$o = \sigma(\omega)x + y. \quad (13)$$

Bidirectional. We describe a bidirectional model incorporating backward recursion as $y_t = (1 -$

$\lambda_2)x_t + \lambda_2y_{t+1}$. By employing a bidirectional model, the influence of tokens is determined by both preceding and succeeding tokens, resulting in a wide-ranging receptive field. The kernel function is formed by combining the forward and backward kernels in equation (14). We employ the circular convolution to compute the output with the input being zero-padded on the right side to twice its length.

$$K = (1 - \lambda_1, \dots, (1 - \lambda_1)\lambda_1^{L-1}, \\ (1 - \lambda_2)\lambda_2^{L-1}, \dots, 1 - \lambda_2). \quad (14)$$

A sketch of the full computation is presented on the left of Figure 2. Initially, we calculate the kernel corresponding to the sequence length and subsequently apply FFT to compute the convolution of the inputs and the kernel. The Complex Exponential Smoothing (CES) module produces the final results by combining the shortcut and the convolution outputs. Although the current code is designed for a unidirectional kernel, a bidirectional kernel can be easily achieved by connecting two unidirectional kernels using the equations as aforementioned.

3.2 ETSMLP blocks

We incorporate the CES module into the element-wise MLP to learn token-level information. The CES module facilitates the mix of input information at the token level, resulting in a mixed output containing sequence information. We integrated the CES module just before the activation function into the MLP in the full architecture, depicted in Figure 2. The functions are described as follows:

$$X = \text{LayerNorm}(X_i) \in \mathbb{R}^{L \times d}, \\ H = W_1 X \in \mathbb{R}^{L \times D}, \\ Y = \sigma(\text{CES}(H)) \in \mathbb{R}^{L \times D}, \\ Z = W_2 Y \in \mathbb{R}^{L \times d}.$$

Compared to standard MLP, the increased parameters constitute only $\frac{3}{d}$ of the original MLP, where d is the embedding dimension. For a typical model with $d = 512$, a modest increase 0.58% parameters enables channel-only MLP to learn sequence information, which is previously unattainable. Moreover, the computational and memory complexity is lower than that of the self-attention, as detailed in Section 4.3.3

Gated architecture. To further enhance the expressive capacity of our model, we add a gate mechanism like (Cho et al., 2014; Shazeer, 2020; Hua et al., 2022). This gate unit controls the output of each block. After obtaining the representation after layernorm, we pass it through a linear layer, derive the score using the sigmoid function, multiply it with the output from the preceding module, and obtain the output of our layer through a residual connection. As in Figure 2, we express these processes as follows:

$$G = \text{sigmoid}(W_g X) \in \mathbb{R}^{L \times d}, \\ O = G \otimes Z \in \mathbb{R}^{L \times d}, \\ X_{i+1} = X_i + O \in \mathbb{R}^{L \times d}.$$

4 Experiments

We present an empirical comparison between our ETSMLP and other baseline models. Our experiments encompass a set of sequence modeling tasks, including LRA, MNLI, IMDB, etc. The main experiment results are divided into two subsections: LRA and NLU benchmarks. Furthermore, we conduct an ablation study to examine the influence of hyperparameters. Additional information about the experimental details and datasets can be found in Appendix B.

4.1 LRA

The LRA benchmarks are a collection of long sequence modeling tasks ranging from 1024 to over 16000 (Tay et al., 2020). In Table 1, we compare our models to several variants of SSM and Transformer. We observe that our model outperforms all the Transformer variants and achieves the comparable performance of S4 on average which is 83.09 vs 80.48. Although we don't gain the highest average scores among all concurrent works, it still produces comparable results without relying on the attention in MEGA (Ma et al., 2022), or Hippo initialization in S5 (Smith et al., 2022). When comparing the individual tasks horizontally, we observe that our model performs significantly better in text tasks such as *ListOps* and *Text*, while slightly underperforming on image tasks like *Image*. This discrepancy may be attributed to the weight decaying exponentially with distance, which is unsuitable for flattened patches.

models	<i>ListOps</i>	<i>Text</i>	<i>Retrieval</i>	<i>Image</i>	<i>Pathfinder</i>	<i>Path-X</i>	Avg.
Transformer	36.37	64.27	57.46	42.44	71.40	-	53.66
Reformer	37.27	56.10	53.40	38.07	68.50	-	50.36
Linear Trans	16.13	65.90	53.09	42.34	75.30	-	50.46
Performer	18.01	65.40	53.82	42.77	77.05	-	51.18
S4	58.35	76.02	87.09	87.26	86.05	88.10	80.48
$DSS_{SOFTMAX}$	60.6	84.8	87.8	85.7	84.6	87.8	81.88
DSS_{EXP}	59.7	84.6	87.6	84.9	84.7	85.6	81.18
S5	62.15	89.02	<u>91.4</u>	<u>88.0</u>	<u>95.33</u>	<u>98.58</u>	<u>87.46</u>
Mega-chunk	58.76	<u>90.19</u>	90.97	85.80	94.41	93.81	85.66
ETSMLP	61.35	87.2	85.78	78.14	85.86	87.21	80.92
ETSMLP-Gate	62.55	88.49	86.72	75.34	91.66	93.78	83.09

Table 1: Performance on the LRA benchmark tasks. We follow the procedure reported in (Ma et al., 2022), and report means across three seeds for our methods. The Bold scores indicate the best performance between S4, DSS, and our models. We also include and underline the state-of-art results of concurrent methods such as Mega and S5.

models	Classification			Similarity		Inference	
	CoLA	SST-2	IMDB	QQP	MRPC	MNLI	QNLI
transformer	69.2	81.7	88.2	80.6	71.1	58.7	61.2
ESMLP	69.2	81.3	87.1	81.6	71.6	60.6	64.5
ESMLP-Gate	69.3	81.2	87.1	82.3	70.3	61.3	64.8

Table 2: Performance on the several NLU tasks. We report accuracy scores averaged across three seeds for all the datasets. All models are trained from scratch and are of a fairly similar size. The bold scores indicate the highest performance of each dataset.

We provide the hyperparameters used in our experiments in Appendix B.

4.2 NLU

The LRA results demonstrate the benefits of our method in sequential text tasks. Furthermore, we conduct experiments on various NLU tasks and compare our models with a transformer encoder architecture trained from scratch. Our experimental evaluations were divided into three categories: sentence classification, including *CoLA* (Warstadt et al., 2019), *SST-2* (Socher et al., 2013), and *IMDB* (Maas et al., 2011); sentence similarity, including *QQP*,² and *MRPC* (Dolan and Brockett, 2005); and natural language inference, including *MNLI* (Williams et al., 2018) and *QNLI* (Rocktäschel et al., 2015). We present the experiment results in Table 2, which reveal that our architecture can achieve comparable or even superior performance to transformers on all the datasets. Considering the simple computation and slight increase in param-

eters on MLP, these results suggest that the ETS has tremendous potential in sequence learning tasks.

4.3 Analysis

4.3.1 Role of damped factors and fields

The experimental results presented above are encouraging and demonstrate the effectiveness of the ETS for sequence modeling tasks. It is proved empirically that even the simplest SSM like ETS can achieve a competitive result compared with other state space model variants. To further consider if we would simplify the ETS in fewer parameters, we conducted ablation studies on the damped factors and the number fields. Table 3 shows the accuracy results of *Listops* if we remove α , β , or ω or change all parameters from complex to real fields. We could observe that whether to remove the α or β or ω or the complex field, the performance of our method drops significantly, especially α , ω and complex field. These experiments illustrate the necessity of our architecture in sequence modeling tasks.

²<https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>

model	Arch	Acc
ETSMLP	Real	40.5
	No α	41.5
	No ω	40.75
	No β	56.05
	-	61.35

Table 3: Ablation analysis of the learnable parameters α, β, ω on *ListOps*. "-" means the keeping all the parameters of our methods.

model	Initialization	Acc
ETSMLP	Stable(0.3,0)	52.4
	Stable(0.7,0)	52.4
	Stable(0.5,0)	52.95
	Ring(0.1,0.6)	60.55
	Ring(0.6,0.9)	58.95
	Ring(0.1,0.9)	61.35

Table 4: Performance of *ListOps* with different initialization. "Ring" means uniform sampling on a ring $\{r_{min} \leq |z| \leq r_{max} | z \in \mathbb{C}\}$. "Stable" means initialization λ' on the same point.

4.3.2 Initialization of parameters

S4 and its variants conducted several experiments on *HiPPO* initialization and concluded that random initialization may hurt performance (Gu et al., 2021, 2022). Because of the different computation processes, *HiPPO* initialization doesn't work in our models. Therefore, we consider the ring initialization method, which involves uniform sampling on a ring defined by the range $\{r_{min} \leq |z| \leq r_{max} | z \in \mathbb{C}\}$. By predefining values for r_{min} and r_{max} , we uniformly sample λ along the ring, between circles with radii r_{min} and r_{max} . In addition to examining the effects of different initializations, we conducted experiments using fixed-value initialization operations. Our experimental results on *listops* are displayed in Figure 4. It can be observed that our model exhibits consistent performance across rings of varying sizes. However, when dealing with fixed points, the effectiveness diminishes significantly.

4.3.3 Efficiency and memory analysis

To assess the speed and memory efficiency across different lengths, we performed experiments using a synthetic dataset that combines multiple sentences to achieve sufficient length. Our chosen task is language modeling, as it allows us to seg-

ment sentences into desired lengths. The maximum length of our synthetic dataset is 8192. We adjusted the sample length within each batch to compare the words per second (WPS) and memory usage (in GB) between the transformer, S4, and our model at comparable sizes. The batch size was uniformly set to 1 to ensure accurate memory usage. All training procedures are carried out on an NVIDIA GeForce GTX 2080 GPU.

The comparison results are presented in Figure 3. Notice that our approach consistently achieves the highest WPS for all the sequence lengths. The slower performance of S4 can be attributed to its complex calculations on the NPLR. Both our model and S4 share a common characteristic: the WPS remains constant as the sequence length increases, while the transformer shows a decrease. Furthermore, the memory requirements of the transformer exhibit an almost quadratic growth, whereas our model and S4 demonstrate linear growth, with our model having a lower slope. For sequence lengths below 3072, there is minimal difference between our model and the transformer. However, as the training length increases, the undesirable quadratic growth in computation and memory complexity becomes apparent in the transformer, whereas our method avoids this issue.

5 Related Works

Since the Transformer was introduced, the quadratic time cost of the attention operation has been numerously researched. Optimizing this operation can improve the efficiency when training and inferencing long context for large language models (Xiao et al., 2023). Recently, many transformer variants have been introduced to reduce the complexity of attentions (Tay et al., 2022), including sparse attention (Beltagy et al., 2020; Kitaev et al., 2020; Guo et al., 2021), kernel-based methods (Choromanski et al., 2020; Kasai et al., 2021; Peng et al., 2021), chunked attention with gating (Hua et al., 2022; Ma et al., 2022) and other efficient methods (Wang et al., 2020; Dao et al., 2022). Another line of research tries to replace the attention mechanism with other modules for long sequences and avoid quadratic time costs. A dizzying number of attention-free models have emerged, where SSMs are becoming one of the most promising models among them.

SSMs. S4 first investigated the SSM for long sequence modeling (Gu et al., 2021). They showed

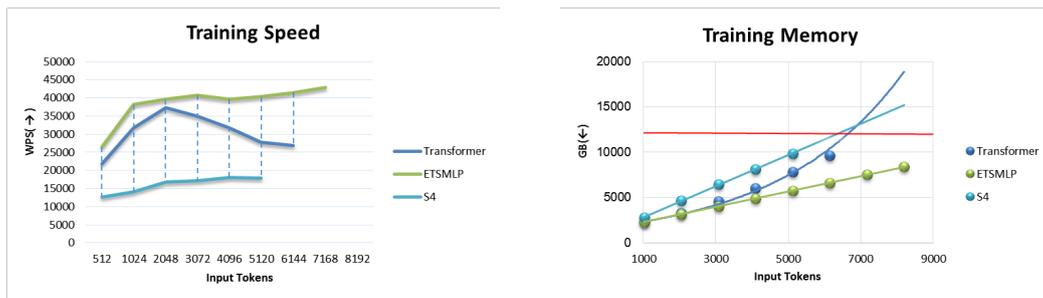


Figure 3: A training speed and memory comparison between the transformer and ETSMLP. Both models have approximately 30M parameters, and the batch size remains constant at 1 under all circumstances.

that naive instantiations of the SSM did not perform well but HiPPO-LegS matrix did (Gu et al., 2020), and hence introduced the DPLR that efficiently computed the complex diagonal plus low-rank matrix. DSS observed that a fully diagonal matrix could preserve the performance of the original S4 (Gupta et al., 2022), and S4D (Gu et al., 2022) then showed that the initialization is critical for DSS. Inspired by S4, many SSM variants emerged recently. S5 replaced single-input, single-output (SISO) SSMs in S4 with multi-input, multi-output (MIMO) (Smith et al., 2022). SGConv viewed the SSM as a global convolution model and suggested that the convolution kernel’s sub-linear decay in sequence length is indispensable (Li et al., 2022). Linear Recurrent Unit (LRU) explored the relationship between the SSM and linear RNN and showed the importance of initialization, exponential parameterization, and normalization for SSMs (Orvieto et al., 2023). MEGA was the most similar approach to ours and plugged exponential moving average into the attention mechanism to improve position-wise local dependency (Ma et al., 2022). Our CES mechanism only considered a position-aware but representation-agnostic dependency which is completely different from the attention mechanism but matches the performance of the transformer.

Other attention free models. MLP-Mixer (Tolstikhin et al., 2021), and its variants proposed to replace the attention with MLPs in computer vision task (Touvron et al., 2022; Yu et al., 2022; Tatsunami and Taki, 2022). Another MLP-based model gMLP showed the potential of MLPs to model sequence dependency and achieved comparable results in pretraining and downstream NLP tasks (Liu et al., 2021). The Attention Free Transformer (AFT) replaced the attention mechanism with an element-wise multiplication and avoided the quadratic computation burden of the attention

matrix (Zhai et al., 2021). Recurrent Memory Transformer (RMT) added a special cache token and utilized the recursive components to increase the context length in the transformer (Bulatov et al., 2022, 2023). Receptance Weighted Key Value (RWKV) leveraged token shift for parallel training a simple linear RNN (Peng et al., 2023). Our models do not conflict with those models in spirit. Our CES modules can be integrated into their models to improve their capabilities of sequence learning.

6 Conclusion

We proposed the ETSMLP model for long-range sequence modeling. Our approach began with a special SSM, namely ETS, and incorporated additional hyperparameters. Moreover, we proposed an exponential parameterization and a constraint function essential for stable training. The experimental results demonstrated the effectiveness of the ETSMLP in sequence learning. Our proposed module could become a plug-in module on other models to enhance their sequence learning capabilities. We hope our research could provide valuable insights into the application of SSMs and encourage further exploration in this area.

7 Limitations

Our approach focuses on evaluating datasets containing fewer than 100,000 samples, where the influence of prior knowledge on performance is substantial. In the next phase, we aim to conduct experiments on pretraining. The considerable disparity between pretraining and training from scratch requires meticulous adjustment of exponential smoothing and ingenious design of the architecture, something like Mamba (Gu and Dao, 2023).

Another limitation of our approach lies in the empirical design of the constraint function. This arises from the potential for lambda to surpass the

precision of 32-bit floating point numbers if its training range is not restricted, and leads to NaN results during backpropagation. We believe that a low granularity parameterization can effectively mitigate this concern. Our future work will prioritize establishing a smooth training process on the parameter space.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under grant number 62076009.

References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Aydar Bulatov, Yuri Kuratov, and Mikhail S Burtsev. 2023. Scaling transformer to 1m tokens and beyond with rmt. *arXiv preprint arXiv:2304.11062*.
- Aydar Bulatov, Yury Kuratov, and Mikhail Burtsev. 2022. Recurrent memory transformer. *Advances in Neural Information Processing Systems*, 35:11079–11091.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the properties of neural machine translation: Encoder–decoder approaches](#). In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarrlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*.
- Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to algorithms*. MIT press.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Quentin Fournier, Gaétan Marceau Caron, and Daniel Aloiise. 2023. [A practical survey on faster and lighter transformers](#). *ACM Comput. Surv.*, 55(14s).
- Everette S Gardner Jr. 1985. Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28.
- Albert Gu and Tri Dao. 2023. [Mamba: Linear-time sequence modeling with selective state spaces](#).
- Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. 2020. Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33:1474–1487.
- Albert Gu, Karan Goel, Ankit Gupta, and Christopher Ré. 2022. On the parameterization and initialization of diagonal state space models. *Advances in Neural Information Processing Systems*, 35:35971–35983.
- Albert Gu, Karan Goel, and Christopher Ré. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2021. Longt5: Efficient text-to-text transformer for long sequences. *arXiv preprint arXiv:2112.07916*.
- Ankit Gupta, Albert Gu, and Jonathan Berant. 2022. Diagonal state spaces are as effective as structured state spaces. *Advances in Neural Information Processing Systems*, 35:22982–22994.
- Weizhe Hua, Zihang Dai, Hanxiao Liu, and Quoc Le. 2022. Transformer quality in linear time. In *International Conference on Machine Learning*, pages 9099–9117. PMLR.
- J Stuart Hunter. 1986. The exponentially weighted moving average. *Journal of quality technology*, 18(4):203–210.
- Rob Hyndman, Anne B Koehler, J Keith Ord, and Ralph D Snyder. 2008. *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.
- Jungo Kasai, Hao Peng, Yizhe Zhang, Dani Yogatama, Gabriel Ilharco, Nikolaos Pappas, Yi Mao, Weizhu Chen, and Noah A Smith. 2021. Finetuning pre-trained transformers into rnns. *arXiv preprint arXiv:2103.13076*.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#).
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*.
- Yuhong Li, Tianle Cai, Yi Zhang, Deming Chen, and Debadepta Dey. 2022. What makes convolutional models great on long sequence modeling? *arXiv preprint arXiv:2210.09298*.
- Hanxiao Liu, Zihang Dai, David So, and Quoc V Le. 2021. Pay attention to mlps. *Advances in Neural Information Processing Systems*, 34:9204–9215.

- Xuezhe Ma, Chunting Zhou, Xiang Kong, Junxian He, Liangke Gui, Graham Neubig, Jonathan May, and Luke Zettlemoyer. 2022. Mega: moving average equipped gated attention. *arXiv preprint arXiv:2209.10655*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Eddie McKenzie and Everette S Gardner Jr. 2010. Damped trend exponential smoothing: a modelling viewpoint. *International Journal of Forecasting*, 26(4):661–665.
- Antonio Orvieto, Samuel L Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pascanu, and Soham De. 2023. Resurrecting recurrent neural networks for long sequences. *arXiv preprint arXiv:2303.06349*.
- Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, Kranthi Kiran GV, et al. 2023. Rwkv: Reinventing rnns for the transformer era. *arXiv preprint arXiv:2305.13048*.
- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A Smith, and Lingpeng Kong. 2021. Random feature attention. *arXiv preprint arXiv:2103.02143*.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2015. Reasoning about entailment with neural attention. *arXiv preprint arXiv:1509.06664*.
- Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Jimmy TH Smith, Andrew Warrington, and Scott W Linderman. 2022. Simplified state space layers for sequence modeling. *arXiv preprint arXiv:2208.04933*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Ivan Svetunkov, Nikolaos Kourentzes, and John Keith Ord. 2022. Complex exponential smoothing. *Naval Research Logistics (NRL)*, 69(8):1108–1123.
- Yuki Tatsunami and Masato Taki. 2022. Raftmlp: How much can be done without attention and with less spatial locality? In *Proceedings of the Asian Conference on Computer Vision*, pages 3172–3188.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. 2020. Long range arena: A benchmark for efficient transformers. *arXiv preprint arXiv:2011.04006*.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2022. [Efficient transformers: A survey](#). *ACM Comput. Surv.*, 55(6).
- Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. 2021. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34:24261–24272.
- Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Gautier Izacard, Armand Joulin, Gabriel Synnaeve, Jakob Verbeek, et al. 2022. Resmlp: Feedforward networks for image classification with data-efficient training. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):5314–5321.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Peter R Winters. 1960. Forecasting sales by exponentially weighted moving averages. *Management science*, 6(3):324–342.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. 2022. S2-mlp: Spatial-shift mlp architecture for vision. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 297–306.

Amir Zandieh, Insu Han, Majid Daliri, and Amin Karbasi. 2023. [Kdeformer: Accelerating transformers via kernel density estimation](#).

Shuangfei Zhai, Walter Talbott, Nitish Srivastava, Chen Huang, Hanlin Goh, Ruixiang Zhang, and Josh Susskind. 2021. An attention free transformer. *arXiv preprint arXiv:2105.14103*.

A Proof of Proposition 1

We first restate Proposition 1 for the sake of convenience.

Proposition. *Let $\lambda \in \mathbb{C}$ be within the interior of the hollow unit disc $D^\circ(0, 1) = \{z \mid |z| < 1\} \setminus \{(0, 0)\}$. We define $\lambda' = \log \log \lambda$ which substitutes λ in the equation (10). If the gradients of y_t satisfy $\sum_0^L \frac{dL}{dy_t} y_t < \infty$, the gradients of the real and imaginary parts of λ' are bounded for all $\lambda \in D^\circ(0, 1)$.*

Proof. Let the real and imaginary parts of λ' be a and b , thus $\lambda' = a + bi$.

We define intermediate variables $y'_t \in \mathbb{C}$ for t in range $\{0, 1, \dots, L\}$ and

$$y'_t = \sum_{j=0}^{t-1} (\lambda^\alpha)^j (1 - \lambda^\alpha) \beta x_{t-j}.$$

Compared with the equation (10), we find that $y_t = \Re(y'_t)$.

We have that if $\sum_0^L \frac{dL}{dy_t} y_t < \infty$, then $\sum_0^L \frac{dL}{dy_t} \Im(y'_t) < \infty$. This is proved by the following:

$$\begin{aligned} \frac{\Re(y'_t)}{\Im(y'_t)} &= \frac{\sum_{j=0}^t \Re(\beta(1 - \lambda^\alpha) \lambda^{\alpha j} x_{t-i})}{\sum_{j=0}^t \Im(\beta(1 - \lambda^\alpha) \lambda^{\alpha j} x_{t-i})} \\ &= \frac{\sum_{j=0}^t \Re((\beta_1 + i\beta_2)((\lambda^\alpha)^j - (\lambda^\alpha)^{j+1})) x_{t-i}}{\sum_{j=0}^t \Im((\beta_1 + i\beta_2)((\lambda^\alpha)^j - (\lambda^\alpha)^{j+1})) x_{t-i}} \\ &= \frac{\sum_{j=0}^t |\lambda^\alpha|^j (\beta_1 (\cos(j\theta) - |\lambda^\alpha| \cos((j+1)\theta)) - \beta_2 (\sin(j\theta) - |\lambda^\alpha| \sin((j+1)\theta))) x_{t-i}}{\sum_{j=0}^t |\lambda^\alpha|^j (\beta_2 (\cos(j\theta) - |\lambda^\alpha| \cos((j+1)\theta)) + \beta_2 (\sin(j\theta) - |\lambda^\alpha| \sin((j+1)\theta))) x_{t-i}}. \end{aligned}$$

It is obvious that the ratio $\frac{\Re(y'_t)}{\Im(y'_t)}$ is bounded for the finite summation $\sum_0^L \frac{dL}{dy_t} \Im(y'_t)$ is bounded too.

To compute the gradients of a and b , we consider the gradients of λ' for y'_t . The function of y'_t is holomorphic function for λ' thus the gradients is:

$$\begin{aligned} \frac{dy'_t}{d\lambda'} &= \sum_{j=0}^t \beta x_{t-j} (\alpha j \lambda^{\alpha j-1} - \alpha(j+1) \lambda^{\alpha(j+1)-1}) e^{e^{\lambda'}} e^{\lambda'} \\ &= \sum_{j=0}^t \beta x_{t-j} \alpha \lambda^{\alpha j} (j - (j+1) \lambda^\alpha) \log(\lambda) \\ &= \sum_{j=0}^t \beta x_{t-j} \alpha \lambda^{\alpha j} (j - (j+1) \lambda^\alpha) \log(\lambda). \end{aligned}$$

As y'_t is a holomorphic function for λ' and $y_t = \Re(y'_t)$, we thus have:

$$\frac{dy'_t}{d\lambda'} = \frac{dy_t}{da} - \frac{dy_t}{db} i.$$

Thus, the gradients of a and b is computed by the chain rule in the following:

$$\begin{aligned} \frac{dL}{da} &= \sum_{t=1}^L \frac{dL}{dy_t} \frac{dy_t}{da} \\ &= \sum_{t=1}^L \frac{dL}{dy_t} \Re\left(\sum_{j=0}^t \beta x_{t-j} \alpha \lambda^{\alpha j} (j - (j+1) \lambda^\alpha) \log(\lambda)\right); \end{aligned}$$

$$\begin{aligned} \frac{dL}{db} &= \sum_{t=1}^L \frac{dL}{dy_t} \frac{dy_t}{db} \\ &= - \sum_{t=1}^L \frac{dL}{dy_t} \Im\left(\sum_{j=0}^t \beta x_{t-j} \alpha \lambda^{\alpha j} (j - (j+1) \lambda^\alpha) \log(\lambda)\right). \end{aligned}$$

Obviously, $\frac{dL}{da}$ and $\frac{dL}{db}$ are continue on all $\lambda \in D^\circ(0, 1)$.

We consider the boundary of $D^\circ(0, 1)$. We first take a look at the zero point. For $\lim_{\lambda \rightarrow \infty} \lambda^{\alpha j} \log(\lambda) = 0$, we can easily compute the limitation :

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \frac{dL}{da} &= \sum_{t=1}^L \frac{dL}{dy_t} \Re\left(\sum_{j=0}^t \beta x_{t-j} \alpha j * (\lambda^{\alpha j} \log(\lambda))\right) \\ &= 0. \end{aligned}$$

For the point λ_0 on the unit cycle, we can find a constant $C(\lambda_0) \in D(0, 2L)$ which satisfies $\Re(\beta \alpha \lambda_0^{\alpha j} (j - (j+1) \lambda_0^\alpha)) \leq \Re(\beta \alpha \lambda_0^{\alpha j} * C(\lambda_0))$ for all $j \in \{1, 2, \dots, L\}$. Therefore,

$$\begin{aligned} \frac{dL}{da} &= \sum_{t=1}^L \frac{dL}{dy_t} \frac{dy_t}{da} \\ &\leq \sum_{t=1}^L \frac{dL}{dy_t} \Re\left(\sum_{j=0}^t \beta x_{t-j} \alpha \lambda^{\alpha j} C(\lambda_0) \log(\lambda)\right) \\ &= \sum_{t=1}^L \frac{dL}{dy_t} \Re\left(\frac{y'_t}{1 - \lambda^\alpha} \alpha C(\lambda_0) \log(\lambda)\right). \end{aligned}$$

Task	D	L	H	$Norm$	LR	WD	DP
<i>ListOps</i>	160	12	160	layer	0.01	0.01	0.0
<i>Text</i>	160	4	160	layer	0.005	0.01	0.1
<i>Retrieval</i>	160	6	160	layer	0.005	0.01	0.1
<i>Image</i>	160	12	320	batch	0.01	0.01	0.0
<i>Pathfinder</i>	128	6	256	batch	0.01	0.01	0.0
<i>Path-X</i>	128	6	256	batch	0.05	0.01	0.0
<i>CoLA</i>	512	3	512	layer	1e-5	0.1	0.1
<i>SST-2</i>	512	12	512	layer	1e-5	0.1	0.1
<i>iMDB</i>	512	4	512	layer	1e-5	0.1	0.3
<i>QQP</i>	512	6	512	layer	1e-5	0.1	0.1
<i>MRPC</i>	512	6	512	layer	1e-5	0.1	0.1
<i>MNLI</i>	512	6	512	layer	1e-5	0.1	0.1
<i>QNLI</i>	512	6	512	layer	1e-5	0.1	0.1

Table 5: The hyperparameters of the ESMLP on LRA and NLU tasks. D is the embedding size, H is the hidden features, and L is the number of layers. LR is the learning rate, WD is weight decay and DP is dropout. BN and LN in the column $Norm$ refer to Batch Normalization and Layer Normalization. For NLU tasks, the small and base models have different model scales.

We know that $\frac{\log(\lambda)\alpha}{1-\lambda^\alpha}$ is finite on the unit cycle except for 1 and the $\lim_{\lambda \rightarrow 1} \frac{\log(\lambda)\alpha}{1-\lambda^\alpha} = -1$ is also finite. As a result, we can find a constant D which satisfies $\Re(y'_t \frac{\alpha \log(\lambda_0)}{1-\lambda_0^\alpha} C(\lambda_0)) \leq \Re(y'_t)D$ for all $t = \{1, \dots, L\}$ and λ_0 . Thus, for all λ_0 on the unit cycle we have:

$$\frac{dL}{da} \leq \sum_{t=1}^L \frac{dL}{dy_t} \Re(y'_t)D < \infty.$$

Similarly, we have:

$$\frac{dL}{da} \leq \sum_{t=1}^L \frac{dL}{dy_t} \Im(y'_t)D < \infty.$$

As the function $\frac{dL}{da}$ and $\frac{dL}{db}$ are continues and the boundaries are finite, by the boundedness theorem, we conclude that the gradients of a and b are bounded for all $\lambda \in D^\circ(0, 1)$. \square

B Experimental setup

Architecture. We present an overview of our architecture in Figure 2. The ETSMLP architecture contains L blocks and each block contains a normalization, skip connection, and an MLP plus CES. We use the ReLU activation function in MLPs. For ETSMLP-Gate architecture, an extra gate mechanism is added parallel to the original architecture.

For the sake of performance, we add extra normalization like LRU (Orvieto et al., 2023). We use bidirectional models for all datasets.

Experimental details. We use the Adam optimizer (Kingma and Ba, 2017), with the hyperparameter $\beta_1 = 0.9$, $\beta_2 = 0.98$. We use warmup for the learning rate LR that we start from a value of 10^{-7} and increase the learning rate linearly up a specified value for the first 10% of training. Then a linear annealing schedule is conducted for the rest of the training. All experiments except for Path-X were carried out on an NVIDIA GeForce GTX 2080 GPU, while Path-X requires 8 NVIDIA GeForce GTX 2080 GPUs.

Hyperparameters. We follow the general optimization approach used by Mega (Ma et al., 2022). Table 5 presents the main hyperparameters for each experiment. For all the experiments, we tune the embedding size D , the number of layers L , and the hidden features H . We also tune the learning rate LR and weight decay WD for all the datasets. Besides, the max_λ of the constraint function is all set to 0.9999.