

Unlocking the Power of LSTM for Long Term Time Series Forecasting

Yaxuan Kong^{1*}, Zepu Wang^{2,5*}, Yuqi Nie³, Tian Zhou⁴, Stefan Zohren¹, Yuxuan Liang^{2†},
Peng Sun^{5†}, Qingsong Wen⁶

¹University of Oxford

²The Hong Kong University of Science and Technology (Guangzhou)

³Princeton University

⁴Alibaba Group

⁵Duke Kunshan University

⁶Squirrel AI

yaxuan.kong@eng.ox.ac.uk, zw163@duke.edu

Abstract

Traditional recurrent neural network architectures, such as long short-term memory neural networks (LSTM), have historically held a prominent role in time series forecasting (TSF) tasks. While the recently introduced sLSTM for Natural Language Processing (NLP) introduces exponential gating and memory mixing that are beneficial for long term sequential learning, its potential short memory issue is a barrier to applying sLSTM directly in TSF. To address this, we propose a simple yet efficient algorithm named P-sLSTM, which is built upon sLSTM by incorporating patching and channel independence. These modifications substantially enhance sLSTM’s performance in TSF, achieving state-of-the-art results. Furthermore, we provide theoretical justifications for our design, and conduct extensive comparative and analytical experiments to fully validate the efficiency and superior performance of our model.

Code — <https://github.com/Eleanorkong/P-sLSTM>

Introduction

Time Series Forecasting (TSF) is a prominent area of research in statistics and machine learning with diverse applications ranging from financial forecasting to traffic and human trajectory prediction (Wang et al. 2022, 2023; Wang, Sun, and Boukerche 2022; Nie et al. 2024; Zhao et al. 2024; Guo et al. 2024; Ruan and Di 2024). Historically, traditional Recurrent Neural Networks (RNNs) have been the cornerstone for modeling sequential data due to their inherent ability to capture temporal dependencies. Among different RNN architectures, Long Short-Term Memory (LSTM) networks are among the most popular and successful ones, as they solve the traditional vanishing and exploding gradient problems of RNNs and exhibit better performance with long sequences. Numerous researchers have adopted LSTM networks to predict diverse time series data, demonstrating their effectiveness across various domains (Hochreiter

and Schmidhuber 1997; Gers, Schmidhuber, and Cummins 1999; Siami-Namini, Tavakoli, and Namin 2018).

However, LSTM has several limitations in TSF. Firstly, although LSTM is designed to capture longer sequential correlations compared to traditional RNNs, it fails to memorize long sequential information statistically and experimentally (Zhao et al. 2020). Additionally, LSTMs lack the ability to revise storage decisions dynamically with their historical forget gates, which restricts their flexibility to adapt to changes in data (Oreshkin et al. 2020; Beck et al. 2024; Zhao et al. 2020). Consequently, traditional RNNs have gradually lost their dominant position in time series tasks, with the leaderboard for TSF now being dominated by alternative architectures such as Transformers (Wen et al. 2023; Zhou et al. 2021) and Multi-Layer Perceptrons (MLPs) (Wang et al. 2023; Zeng et al. 2023).

Method	Type	Time Complexity	Memory Complexity
LSTM	RNN	$O(L)$	$O(L)$
DLinear	MLP	$O(L)$	$O(L)$
Informer	Transformer	$O(L \log L)$	$O(L \log L)$
Crossformer	Transformer	$O(L^2)$	$O(L^2)$

Table 1: Time and Space Complexity Comparison of TSF Models. L is the input sequence length/look-back window. LSTM (Hochreiter and Schmidhuber 1997); DLinear (Zeng et al. 2023); Informer (Zhou et al. 2021); Crossformer (Zhang and Yan 2023).

Despite these limitations, LSTM and, more broadly, RNNs remain powerful tools and a promising research direction due to their advantages (Hewamalage, Bergmeir, and Bandara 2021). Firstly, compared to most transformer-based models, RNNs have much lower time complexity and memory complexity, making them more efficient for certain applications, as shown in Table 1. Secondly, RNNs have a clear temporal flow, which makes it easier to interpret their decisions and understand how information flows through the sequence compared to transformer-based structures and MLP-based structures (Graves 2013; Hou and Zhou 2020). At the same time, state-space models (SSMs) have received wide attention in TSF and broader deep learning research

*These authors contributed equally.

†Corresponding authors.

areas. Compared to other black-box deep learning models, SSMs can be explained from statistical and physical perspectives, potentially offering better interpretability (Gu et al. 2022; Wang et al. 2024). Researchers have shown that LSTM/RNNs can be viewed as a special case of SSMs, where the recurrent process illustrates the flow of information (Kim and Kim 1994; Durbin and Koopman 2012a; Zhao et al. 2020). Therefore, LSTM/RNNs are also a reasonable research direction, through the lens of those findings.

Recently, an advanced version of the extended LSTM, named sLSTM (Beck et al. 2024), was introduced showing that one can not only revise the memory storage decisions but also improve its memory capacities, leading to very competitive performance in various NLP tasks (Beck et al. 2024). Given the success of advanced LSTMs in NLP, can we unlock the power of LSTMs for time series forecasting? In this paper, we give an affirmative answer by repropoing sLSTM for multivariate TSF, leading to our new method called P-sLSTM. Our main contributions are as follows.

- We explain why the framework of sLSTM improves memory capacity and revises memory storage, which is suitable for TSF;
- We rigorously show that sLSTM cannot be guaranteed to have a long memory to capture long-term dependencies. Based on the previous limitation, we apply the patching technique to solve this problem and develop an LSTM-based structure, P-sLSTM, for TSF;
- Extensive evaluations on various datasets shows that P-sLSTM outperforms the original LSTM by 20% of accuracy and reaches comparable performance against state-of-the-art (SOTA) models.

Related Work

TSF has seen significant advancements with the introduction of various model architectures such as Transformer (Lin et al. 2023), and MLP (Lu et al. 2024). Transformer-based models have gained popularity in TSF due to their ability to capture long-range dependencies and handle complex temporal patterns (Wen et al. 2023; Lim et al. 2021). For instance, the Informer model by Zhou et al. (2021) improves the efficiency and scalability of Transformers by using a sparse self-attention mechanism and a generative style decoder, allowing it to handle long sequence forecasting more effectively. Additionally, many other Transformer-based models are developed for TSF, such as Autoformer (Wu et al. 2021), Fedformer (Zhou et al. 2022), iTransformer (Liu et al. 2024).

In comparison, MLP-based models typically involve feeding lagged time series data into a fully connected neural network to predict future values (Zeng et al. 2023). For instance, DLinear combines decomposition techniques to get the seasonal information and trend information of the time series data and uses MLP layers to get a good prediction accuracy (Zeng et al. 2023). Another example is Rlinear, which applies linear mapping to the data before linear layers (Li et al. 2023).

State space models (SSMs) have gained prominence due to their inherent flexibility and interpretability in

TSF (Durbin and Koopman 2012b). SSMs capture temporal dynamics by representing time series as a combination of observable and hidden states that evolve over time through specified transition equations (Gu et al. 2022; Wang et al. 2024; Hu et al. 2024). One notable framework leveraging SSMs is the Mamba model, which employs a hierarchical Bayesian approach to capture intricate temporal patterns and adapt to changing dynamics over time (Gu and Dao 2023).

sLSTM for Time Series Forecasting

In this section, we define the TSF problem and revisit the general concept of LSTM. We then briefly explain why the new sLSTM framework can revise historical memory storage and why it is suitable for time series data.

Problem Formulation

The objective of multivariate TSF is to learn a mapping function f that takes historical data X with M channels as input, denoted as $X \in \mathbb{R}^{L \times M}$, where L represents the look-back window. The goal is to predict future data for a time interval of length T , denoted as $\hat{X} \in \mathbb{R}^{T \times M}$.

Recap of LSTM Structure

The LSTM (Hochreiter and Schmidhuber 1997; Hochreiter et al. 2001) was original introduced to avoid the vanishing gradient problem encountered in traditional RNNs. The LSTM memory cell contains three gates: input, output, and forget gates. The LSTM memory cell update rules at time step t are as follows:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{z}_t \quad \text{cell state} \quad (1)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tilde{\mathbf{h}}_t, \quad \tilde{\mathbf{h}}_t = \psi(\mathbf{c}_t) \quad \text{hidden state} \quad (2)$$

$$\mathbf{z}_t = \varphi(\tilde{\mathbf{z}}_t), \quad \tilde{\mathbf{z}}_t = \mathbf{W}_z \mathbf{x}_t + \mathbf{R}_z \mathbf{h}_{t-1} + \mathbf{b}_z \quad \text{cell input} \quad (3)$$

$$\mathbf{i}_t = \sigma_{lstm}(\tilde{\mathbf{i}}_t), \quad \tilde{\mathbf{i}}_t = \mathbf{W}_i \mathbf{x}_t + \mathbf{R}_i \mathbf{h}_{t-1} + \mathbf{b}_i \quad \text{input gate} \quad (4)$$

$$\mathbf{f}_t = \sigma_{lstm}(\tilde{\mathbf{f}}_t), \quad \tilde{\mathbf{f}}_t = \mathbf{W}_f \mathbf{x}_t + \mathbf{R}_f \mathbf{h}_{t-1} + \mathbf{b}_f \quad \text{forget gate} \quad (5)$$

$$\mathbf{o}_t = \sigma_{lstm}(\tilde{\mathbf{o}}_t), \quad \tilde{\mathbf{o}}_t = \mathbf{W}_o \mathbf{x}_t + \mathbf{R}_o \mathbf{h}_{t-1} + \mathbf{b}_o \quad \text{output gate} \quad (6)$$

In this structure, the input weight matrices ($\mathbf{W}_z, \mathbf{W}_i, \mathbf{W}_f$, and \mathbf{W}_o) link the input \mathbf{x}_t to the cell input and the three gates (input, forget, and output). The recurrent weight matrices ($\mathbf{R}_z, \mathbf{R}_i, \mathbf{R}_f$, and \mathbf{R}_o) connect the previous hidden state \mathbf{h}_{t-1} to the cell input and the same three gates. Each component also has its associated bias vector ($\mathbf{b}_z, \mathbf{b}_i, \mathbf{b}_f$, and \mathbf{b}_o). The functions φ and ψ are activation functions for the cell input and hidden state, respectively. Typically, both use the hyperbolic tangent (tanh) function. The purpose of ψ is to constrain the cell state within a specific range, preventing it from growing without bounds. In the standard LSTM network, the activation function for all gates is the sigmoid function.

sLSTM

The sLSTM (scalar LSTM) has recently been introduced as part of the xLSTM (extend LSTM) family (Beck et al. 2024). It is designed to enhance the LSTM architecture by incorporating exponential gating and a new memory mixing mechanism. Compared to the traditional LSTM architecture, the sLSTM introduces several key differences:

- Forget Gate: The sLSTM allows for either sigmoid or exponential activation function in the forget gate, whereas traditional LSTM uses sigmoid. This is represented as:

$$\mathbf{f}_t = \exp(\tilde{\mathbf{f}}_t) \quad (7)$$

- Input Gate: The sLSTM uses an exponential function for the input gate, replacing the sigmoid function used in traditional LSTM units:

$$\mathbf{i}_t = \exp(\tilde{\mathbf{i}}_t) \quad (8)$$

- Normalizer State and Hidden State: The sLSTM introduces a normalizer state \mathbf{n}_t and modifies the hidden state calculation:

$$\mathbf{h}_t = \mathbf{o}_t \odot \tilde{\mathbf{h}}_t, \quad \tilde{\mathbf{h}}_t = \mathbf{c}_t \odot \mathbf{n}_t^{-1} \quad (9)$$

where the normalizer state is updated as:

$$\mathbf{n}_t = \mathbf{f}_t \odot \mathbf{n}_{t-1} + \mathbf{i}_t \quad (10)$$

These modifications on the exponential gating in both the forget and input gates provides more flexibility in controlling information flow, while the normalizer state helps in stabilizing the hidden state calculations over long sequences.

The sLSTM architecture further introduces a multi-head structure and a new memory mixing mechanism. Each sLSTM block can have multiple heads, with memory mixing occurring within each head but not across heads. This is achieved through block-diagonal recurrent weight matrices, allowing for efficient parameter sharing while maintaining separate information flows. The memory mixing process can be represented as $\mathbf{h}_t = f(\mathbf{x}_t, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}, \mathbf{n}_{t-1})$, where f encapsulates all sLSTM operations including the block-diagonal connections and exponential gating. Multiple sLSTM blocks can be stacked together to form deeper networks, enhancing the model’s capacity to capture complex temporal dependencies in sequential data (see sLSTM structure in Figure 1).

Why Exponential Gating?

The introduction of exponential gating in sLSTM enables the model to revise memory storage decisions compared to the standard LSTM. Unlike the sigmoid activation function, exponential gating allows for a broader range of values, providing the forget gate with more diverse outputs in different situations. This is beneficial in TSF, where the importance of past information can change dynamically, and the model could revise stored values when encountering more relevant inputs. Moreover, exponential gating helps mitigate the vanishing gradient problem often encountered in sigmoid-based gates. Since the derivative of the exponential function is the function itself, gradients can flow more easily during back-propagation, even in long sequences.

Why Memory Mixing?

The sLSTM model also introduces a memory mixing mechanism to incorporate multiple heads, which are diagonal blocks within the recurrent gate pre-activations, and exponential gating. Both of them establish a new approach to

memory mixing. This feature enables the model to dynamically combine and integrate memories from different time steps, allowing it to retain and utilize important information over extended periods more effectively. The relevance of historical data can vary, and memory mixing provides an advantage by dynamically adjusting the importance of past information. This flexibility allows the model to adapt to changing patterns and trends in the data and capture complex temporal dependencies. Memory mixing ensures that relevant historical data, regardless of its temporal distance, can be effectively integrated into the forecasting process. This leads to more informed predictions that take into account both short-term fluctuations and long-term trends.

Methodology

In this section, we will begin by providing a theoretical foundation for memory property of sLSTM. Next, we will define sLSTM Markov Chain process and provide detailed theoretical analysis on memory property of sLSTM. Then, we design a simple yet efficient algorithm named P-sLSTM, built upon sLSTM with patching and channel independence.

Theoretical Preparation

Let us first recall some basic property results related to RNN/LSTM (Zhao et al. 2020).

Assumption 1. (i) The joint density function of $\epsilon^{(t)}$ is continuous and positive everywhere. Note that $\epsilon^{(t)}$ is a sequence of independent and identically distributed random vectors, where the target sequence $y^{(t)} =$ the output term $e^{(t)} +$ the error term $\epsilon^{(t)}$. (ii) For some $\kappa \geq 2$, $\mathbb{E}\|\epsilon^{(t)}\|^\kappa < \infty$.

Definition 1. A process $\{X_t\}$ is geometrically ergodic if it converges exponentially fast over time. From Harris Theorem (Meyn and Tweedie 2012), $\gamma_X(k) \sim \rho^k$ as $k \rightarrow \infty$, indicating short memory. This means the influence of past observations on future distributions diminishes exponentially fast.

Theorem 1. Under Assumptions 1, if there exist real numbers $0 < a < 1$ and b such that $\|\mathcal{M}(x)\| \leq a\|x\| + b$, then recurrent network process with transition function $\mathcal{M}(x)$ is geometrically ergodic, and hence has short memory.¹

sLSTM Process

Many time series models, including LSTM and general RNNs, can be viewed as state-space models and transformed into Markov Chain processes. Zhao et al. (2020) reformulated the LSTM as a Markov Chain process and proved, under mild conditions, that it exhibits geometric ergodicity and short-term memory characteristics. To examine the memory properties of sLSTM, we will define an sLSTM Markov Chain process. We adopt a framework similar to that outlined by (Zhao et al. 2020) for this definition.

Consider a many-to-many sLSTM structure for TSF problems, using square loss as an example:

¹We formally define the transition function $\mathcal{M}(x)$ in subsection sLSTM Process.

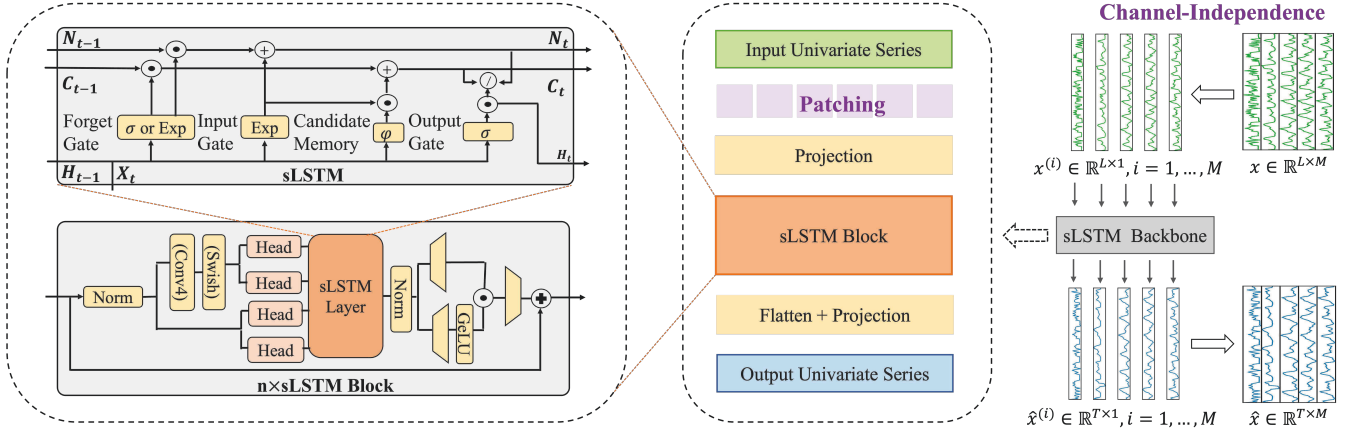


Figure 1: Overview of P-sLSTM Architecture (Top Left: sLSTM structure; Bottom Left: sLSTM block (Beck et al. 2024)).

$$\begin{cases} l^{(t)} = \|y^{(t)} - e^{(t)}\|^2 & \text{loss} \\ e^{(t)} = g(W_{eh}h^{(t)} + b_e) & \text{output} \\ h^{(t)} = o^{(t)} \odot \left(\frac{c^{(t)}}{n^{(t)}}\right) & \text{hidden state} \end{cases} \quad (11)$$

where the hidden unit can be calculated using the sLSTM gating operations. Note here that as mentioned in Zhao et al. (2020), including long-range dependencies in an external input could interfere with our following analysis of the memory characteristics of $\{y^{(t)}\}$. Thus, for simplicity, we also assume no external input, meaning $x^{(t)} = y^{(t-1)}$. The process is defined for $t \in \{1, \dots, T\}$, where $y^{(0)} = 0$, $h^{(0)} = 0$, $y^{(t)}, e^{(t)} \in \mathbb{R}^p$, $h^{(t)}, c^{(t)}, z^{(t)}, f^{(t)}, i^{(t)}, o^{(t)} \in \mathbb{R}^q$, g is an elementwise output function, and \odot denotes elementwise product.

For sLSTM process, we define the hidden states $s^{(t)} \in \mathbb{R}^q$ to include the cell state $c^{(t)}$, the normalizer state $n^{(t)}$, and the hidden state $h^{(t)}$. The sLSTM process can then be represented as a homogeneous Markov chain with the transition function $\mathcal{M}_{\text{sLSTM}}$:

$$\begin{pmatrix} y^{(t)} \\ h^{(t)} \\ c^{(t)} \\ n^{(t)} \end{pmatrix} = \mathcal{M}_{\text{sLSTM}}(y^{(t-1)}, h^{(t-1)}, c^{(t-1)}, n^{(t-1)}) + \begin{pmatrix} \epsilon^{(t)} \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (12)$$

In the above formulation, the transition function $\mathcal{M}_{\text{sLSTM}}$ maps $(\mathbb{R}^p \times \mathbb{R}^q \times \mathbb{R}^q \times \mathbb{R}^q)$ to $(\mathbb{R}^p \times \mathbb{R}^q \times \mathbb{R}^q \times \mathbb{R}^q)$, where q is the dimension of the hidden state vectors: $h^{(t)}, c^{(t)}, n^{(t)}$.

Memory Property of sLSTM Process

Zhao et al. (2020) analysis revealed that the forget gate plays a crucial role in controlling the output range of the forget gate and determining the memory properties of LSTM. Notably, the primary distinction between LSTM and sLSTM lies in the activation function of the forget gate: LSTM typically employs a sigmoid function, while sLSTM uses the exponential function. Intuitively, in sLSTM, the range of the

forget gate with an exponential activation function is unbounded (please refer to Case 2). Consequently, an exponentially large output from the forget gate suggests a greater retention of past information. This characteristic allows the model to better capture long-term dependencies in time series data, thereby enhancing its memory capacity.

However, does this mean that sLSTM is guaranteed to have a long memory? Building upon Zhao et al. (2020)'s work, we extended the analysis from LSTM to sLSTM and introduced an additional corollary based on Assumption 1 and Theorem 1. We prove that under mild conditions, the memory loss property is preserved as the sLSTM process also maintains geometric ergodicity, indicating that it potentially has a short memory. Specifically, this short memory property holds under the condition that $\|f(u, v)\|_\infty \leq a$ for some $a < 1$. For cases where this condition exceeds 1, we examined the long-term behavior of the forget gate to understand its implications for memory retention.

Case 1: When $\|f(u, v)\|_\infty \leq a$ for some $a < 1$, we observe that the sLSTM network exhibits geometric ergodicity, and hence short memory.

Corollary. *The input series features $\{y^{(t-1)}\}$ are scaled to the range $[-1, 1]$. Suppose that $M := \sup_{\mathbf{x} \in B_\infty^q} \|g(W_{eh}\mathbf{x} + b_e)\|_{l_1} < \infty$ and $\|f(u, v)\|_{l_\infty} = \|\exp(W_f u + R_f v + b_f)\|_{l_\infty} \leq a$ for some $a < 1$, where B_∞^q is the q -dimensional l_∞ -ball and $\|W\|_{l_\infty} = \max_{1 \leq i \leq m} \sum_{j=1}^n |w_{ij}|$ is the matrix l_∞ -norm. If Assumption 1 holds, then the sLSTM process is geometrically ergodic and has short memory.*

Proof: Let $a = a_0 \in (0, 1)$ and $b = M + 3q$, then:

$$\begin{aligned} & \|\mathcal{M}_{\text{sLSTM}}(u, v, w, n)\|_{l_1} - a_0 \|(u', v', w', n')\|_{l_1} \\ & \leq \|g(W_{eh}x + b_e)\|_{l_1} + \|x\|_{l_1} + \|\mathbf{1}_q + f(u, v) \odot w\|_{l_1} \\ & \quad + \|f(u, v) \odot n + \mathbf{1}_q\|_{l_1} - a_0 \|(u', v', w', n')\|_{l_1} \\ & \leq M + q + q + \|f(u, v)\|_{l_\infty} \|w\|_{l_1} + \|f(u, v)\|_{l_\infty} \|n\|_{l_1} + q \\ & \quad - a_0 \|u\|_{l_1} - a_0 \|v\|_{l_1} - a_0 \|w\|_{l_1} - a_0 \|n\|_{l_1} \end{aligned}$$

$$\leq M + 3q - a_0 \|u\|_{l_1} - a_0 \|v\|_{l_1} + (\|f(u, v)\|_{l_\infty} - a_0) \|w\|_{l_1} \\ + (\|f(u, v)\|_{l_\infty} - a_0) \|n\|_{l_1} \leq b = M + 3q,$$

where $x = o(u, v) \odot [i(u, v) \odot \tanh(R_z v + W_z u + b_z) + f(u, v) \odot w] \odot [f(u, v) \odot n + i(u, v)]^{-1} \in B_\infty^q$, $v = h_t \in B_\infty^q$, and $u = y_{t-1} \in B_\infty^p$. The second inequality holds due to the definition of M and $x \in B_\infty^q$. The fourth inequality holds due to $\|f(u, v)\|_{l_\infty} = \|\exp(W_f u + R_f v + b_f)\|_{l_\infty} \leq a_0$. By Assumption 1 and Theorem 1 in Zhao et al. (2020), the sLSTM model (12) is geometrically ergodic and has short memory if $\|f(u, v)\|_{l_\infty} \leq a$ for some $a < 1$. \square

The above proof aligns with Zhao’s (2020) observations on LSTM, where we extend the analysis to sLSTM, showing that the forget gate primarily influences the memory property of sLSTM. As the geometric ergodicity of sLSTM only holds when $\|f(u, v)\|_{l_\infty} = \|\exp(W_f u + R_f v + b_f)\|_{l_\infty} \leq a$ for some $a < 1$, this implies that $W_f u + R_f v + b_f$ need to be negative (since the exponential function of a negative value is less than 1). This condition is analogous to the use of a sigmoid activation function on the forget gate in standard LSTM.

Case 2: When $\|f(u, v)\|_{l_\infty} > 1$, we examine the long term behavior of the exponential forget gate of sLSTM.

It is also important to consider the case when $\|f(u, v)\|_{l_\infty}$ exceeds 1. This is when $W_f u + R_f v + b_f$ is positive, leading to an exponential value greater than 1.

In this scenario, the forget gate $f_t = \exp(\tilde{f}_t)$ tends to amplify the cell state c_t over time instead of allowing it to decay. This behavior seems to revise the memory storage and assign larger values to new information that is important; however, it may diminish the model’s ability to effectively integrate and prioritize new information over time.

The cell state update equation $c_t = f_t \odot c_{t-1} + i_t \odot z_t$ shows that when $f_t > 1$ and grows larger, the past state’s contribution ($f_t \odot c_{t-1}$) becomes larger than the new input information ($i_t \odot z_t$). This leads to exponential growth of both c_t and the normalizer state n_t ($n_t = f_t \odot n_{t-1} + i_t$) over time, potentially resulting in unbounded values.

As f_t becomes larger, the impact of the input gate $i_t \odot z_t$ on the cell state and normalizer state diminishes, as the contribution from the past state dominates. In this case, we can approximate that $c_t \approx f_t \odot c_{t-1}$ and $n_t \approx f_t \odot n_{t-1}$.

Consequently, the model’s capacity to learn is diminished. It may struggle to integrate new information effectively as the cell state becomes dominated by amplified past states. Furthermore, the hidden state calculation $h_t = o_t \cdot \frac{c_t}{n_t}$ may face precision problems due to the exponential growth of both c_t and n_t .

Even if c_t and n_t grow at the same rate, and their ratio $\frac{c_t}{n_t}$ ($\frac{c_t}{n_t} \approx \frac{f_t \odot c_{t-1}}{f_t \odot n_{t-1}} = \frac{c_{t-1}}{n_{t-1}}$) might remain stable, the absolute values can become very large, leading to overflow issues.

The above analysis shows that the memory properties

of sLSTM are closely tied to the behavior of its forget gate. When the forget gate’s output is bounded below 1, sLSTM exhibits geometric ergodicity and short memory, similar to standard LSTM. However, when the forget gate’s output exceeds 1, it can lead to exponential growth in the cell state, potentially causing computational issues and diminishing the model’s ability to effectively integrate new information.

General Discussion of the RNN memory property

Is it possible to design an LSTM or a general RNN framework that has a long memory? Theoretically, we might be able to. However, in practice, it is quite hard to ensure that we can use neural networks to reach this goal.

Using LSTM as an example, to maintain long-term memory, we need to ensure that the forget gate is activated with a high value. Without specialized training techniques, this can easily lead to overflow and instability in neural networks. To stabilize the networks, we typically need to restrict gradient changes to prevent gradient explosion, often through methods like gradient clipping (Arjovsky, Shah, and Bengio 2016; Jing et al. 2017). However, by limiting gradient changes within certain training epochs, we inadvertently place an invisible bound on the forget gate’s value. Consequently, neural network may still exhibit short-term memory behavior in practice. To address the problem of short memory in general RNNs for TSF, additional techniques are often necessary. In this paper, we propose to utilize the patching method to overcome this problem as follows.

Overall Structure of P-sLSTM

Figure 1 illustrates the overall structure of the designed P-sLSTM, where multivariate time series data is divided into different channels that share the same backbone but are processed independently. Each channel’s univariate series is segmented into patches, processed by a linear layer, and after several system blocks, another linear layer produces the final prediction. Specifically, the batch (B) of samples $x \in \mathbb{R}^{L \times M}$ with size $B \times L \times M$ is initially transformed to $B \times M \times L$ and undergoes channel independence, resulting in $(B \cdot M) \times L$. This data is then passed through a patching operator, producing $(B \cdot M) \times N \times P$, where N is the number of patches and P is the patch size. It is projected to $(B \cdot M) \times N \times \text{Embedding}$ by a linear layer and processed by the xLSTM block. After flattening, the data becomes $(B \cdot M) \times (N \cdot \text{Embedding})$ and is finally projected to $(B \cdot M) \times T$ by a linear layer, with the output reshaped to $B \times T \times M$.

Patching As we discussed before, we show that although the activation function of sLSTM is adjusted, the model could not have a long memory all the time; hence, it still cannot capture long-term dependencies very successfully. Therefore, inspired by the success of patches in transformer architectures (Nie et al. 2023), we use patches to artificially divide original time series into components and the sLSTM can extract different shorter information; and finally, we can combine them to get global information by linear layers.

Channel Independence (CI) Channel Independence (CI) has been proved to avoid overfitting problems and improve

the computational efficiency in Transformer-based models (Nie et al. 2023; Zhou et al. 2023) and MLP-based models in time series analysis (Zeng et al. 2023; Wang et al. 2023). To our best knowledge, it is the first time to apply CI in RNNs-based models for TSF.

Experiments

Dataset, Baselines, Metrics and Settings

We evaluate the performance of proposed P-sLSTM model on 5 popular multivariate datasets, Weather, Electricity, Solar, ETtm1 and PEMS03. These datasets have been extensively used in the previous literature (Wu et al. 2023; Liu et al. 2024). As baselines, we select SOTA and representative models in the TSF domain, including the following categories: (i) RNNs: sLSTM (without patch), LSTM; (ii) Transformers: iTransformer (Liu et al. 2024), FEDformer (Zhou et al. 2022), Crossformer (Zhang and Yan 2023); (iii) MLPs: RLinear (Li et al. 2023), DLinear (Zeng et al. 2023); (iv) convolutional neural networks (CNNs): TimesNet (Wu et al. 2023); (vi) SSMS: Mamba (Gu and Dao 2023). Mean Squared Error (MSE) and Mean Absolute Error (MAE) are used to assess the performance of TSF models.

We follow the same data loading parameters (Ex. train/val/test split ratio) from (Wu et al. 2023; Liu et al. 2024). All models use the same experimental setup with prediction length $T \in \{12, 24, 48, 96\}$ and look-back window $L = 96$ for PEMS dataset, following the setting of (Liu et al. 2024), and $T \in \{96, 192, 336, 720\}$ and look-back window $L = 336$ for other datasets, following (Nie et al. 2023). All experiments are conducted on RTX 4090 GPU.

Results and Analysis

Main Forecasting Results The quantitative results for TSF with various baselines are presented in Table 2. P-sLSTM demonstrates exceptional performance across multiple datasets and prediction length settings, securing 23 first-place and 10 second-place rankings out of 40 settings.

Notably, among all experimental results, P-sLSTM outperforms sLSTM by 90% of them and outperforms LSTM in 95% of the settings, indicating the enhancement of our method over conventional LSTM methods. We also found that P-sLSTM does not work perfectly in PEMS03 dataset. A potential explanation is that PEMS03 is a very noisy series and P-sLSTM does not include any denoising mechanisms for this situation. However, it still has a better performance than many other SOTA methods.

Influence of Patch Size As you can see from Figure 2, as patch size increases, the prediction accuracy will increase at first, reach an experimental optimal solution, and then the prediction accuracy will decrease as patch size further increases. Generally, patching will divide the time series sequence into multiple pieces, and each piece contains parts of the sequential information. A small patch will destroy the original temporal sequential information, such that sLSTM cannot process the sequential information. In comparison, a big patch will contain too much information, causing previous information to diminish the model’s ability to effectively integrate new information. This observation aligns

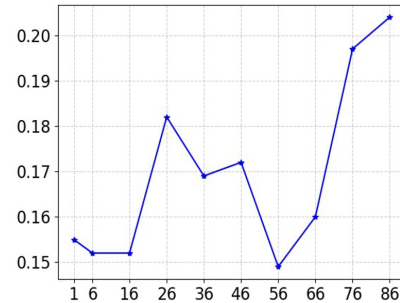


Figure 2: Exploration of the patch size on MSE results of P-sLSTM on Weather dataset.

with our analysis of the sLSTM’s memory properties. However, based on our experiments, for different datasets, there are different optimal patches. Hence, there might be a need to find optimal patches based on trial and error.

Impact of Look-back Window Sizes As we can see in Figure 3, compared to LSTM and sLSTM, the mechanism of patching can help models capture long-term dependencies. Hence, the prediction accuracy will increase as the look-back window increases.

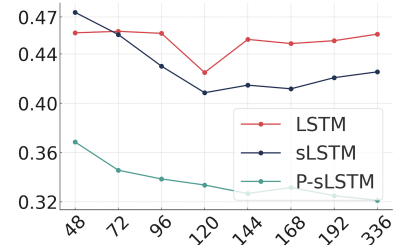


Figure 3: The MSE results (Y-axis) of models with different look-back window sizes (X-axis) of long-term forecasting (T=720) on Weather dataset.

Ablation study of Memory Mixing To test the effectiveness of memory mixing of sLSTM in TSF, we conducted an ablation study to remove memory mixing as in normal LSTM. The results of the weather dataset with a horizon of 96 with length 336 are shown in Table 3. As we can see, the memory mixing strategy slightly helps us select important previous temporal information, resulting in an increase in prediction accuracy. However, its improvement is limited.

Importance of Channel Independence (CI) To test the importance of CI, we prepared a variant of P-sLSTM with channel mixing (CM) and tested its performance. For CI P-sLSTM, the input dimension block of sLSTM is $(B \cdot M) \times N \times \text{Embedding}$ instead of $B \times N \times (\text{Embedding} \cdot M)$. The final results using the solar energy dataset are shown in Table 4. As we can see, CI sLSTM has higher training errors but lower validation errors and testing errors. Hence, we can conclude that CI can significantly prevent overfitting. Therefore, through these experiments, we can conclude that

Models		P-sLSTM		sLSTM		LSTM		iTransformer		Fedformer		Crossformer		RLinear		DLinear		TimesNet		Mamba	
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Weather	96	0.149	0.208	0.172	0.261	0.236	0.304	0.162	<u>0.210</u>	0.217	0.296	<u>0.160</u>	0.229	0.175	0.225	0.176	0.237	0.172	0.220	0.219	0.276
	192	0.197	<u>0.256</u>	0.253	0.337	0.262	0.330	<u>0.205</u>	0.253	0.276	0.336	0.229	0.304	0.218	0.260	0.220	0.282	0.219	0.261	0.290	0.319
	336	0.249	<u>0.297</u>	0.338	0.394	0.341	0.387	<u>0.265</u>	0.294	0.339	0.380	0.267	0.327	<u>0.265</u>	0.294	0.265	0.319	0.280	0.306	0.321	0.340
	720	0.320	<u>0.350</u>	0.426	0.443	0.456	0.457	0.329	0.339	0.403	0.428	0.354	0.394	0.329	0.339	<u>0.323</u>	0.362	0.365	0.359	0.381	0.377
Electricity	96	0.130	0.226	0.281	0.376	0.350	0.424	<u>0.132</u>	<u>0.228</u>	0.193	0.308	0.161	0.229	0.140	0.235	0.140	0.237	0.168	0.272	0.190	0.300
	192	0.148	0.243	0.280	0.376	0.354	0.429	0.154	0.249	0.201	0.315	0.229	0.304	0.154	<u>0.248</u>	<u>0.153</u>	0.249	0.184	0.289	0.201	0.309
	336	0.165	0.262	0.290	0.384	0.357	0.431	0.170	0.265	0.214	0.329	<u>0.169</u>	0.270	0.171	<u>0.264</u>	<u>0.169</u>	0.267	0.198	0.300	0.211	0.319
	720	<u>0.199</u>	<u>0.293</u>	0.297	0.383	0.368	0.438	0.195	0.288	0.246	0.355	0.248	0.340	0.209	0.297	0.203	0.301	0.220	0.320	0.219	0.323
Solar	96	0.167	0.232	0.210	0.271	0.212	0.293	0.217	0.248	0.287	0.383	<u>0.176</u>	<u>0.236</u>	0.238	0.268	0.221	0.292	0.238	0.295	0.206	0.253
	192	0.180	0.241	0.242	0.295	0.218	0.301	0.216	0.264	0.278	0.364	<u>0.212</u>	<u>0.259</u>	0.263	0.283	0.249	0.308	0.223	0.285	0.232	0.292
	336	0.190	0.248	0.244	0.308	0.272	0.335	0.220	0.271	0.319	0.397	<u>0.218</u>	<u>0.263</u>	0.284	0.295	0.268	0.324	0.238	0.295	0.232	0.292
	720	0.196	0.249	0.223	0.295	0.283	0.337	<u>0.221</u>	<u>0.274</u>	0.319	0.402	0.232	0.285	0.282	0.294	0.2708	0.3274	0.245	0.290	0.234	0.294
ETTm1	96	0.292	<u>0.343</u>	0.631	0.552	0.771	0.659	0.302	0.365	0.326	0.390	0.344	0.384	<u>0.301</u>	0.342	0.299	<u>0.343</u>	0.341	0.378	0.458	0.453
	192	0.329	0.369	0.702	0.598	0.820	0.685	0.344	0.382	0.365	0.415	0.352	0.387	<u>0.335</u>	0.363	<u>0.335</u>	<u>0.365</u>	0.462	0.429	0.504	0.484
	336	0.362	0.391	0.826	0.681	0.927	0.743	0.378	0.403	0.392	0.425	0.509	0.516	0.370	0.383	<u>0.369</u>	<u>0.386</u>	0.433	0.436	0.591	0.555
	720	0.421	0.424	0.937	0.733	0.960	0.772	0.438	0.437	0.446	0.458	0.717	0.646	<u>0.425</u>	0.414	<u>0.425</u>	<u>0.421</u>	0.454	0.451	0.587	0.553
PEMS03	12	<u>0.077</u>	<u>0.183</u>	0.103	0.203	0.111	0.214	0.071	0.174	0.126	0.251	0.090	0.203	0.126	0.236	0.122	0.243	0.085	0.192	0.116	0.232
	24	<u>0.109</u>	<u>0.219</u>	0.115	0.217	0.123	0.227	0.093	0.201	0.149	0.275	0.121	0.240	0.246	0.334	0.201	0.317	0.118	0.223	0.162	0.268
	48	0.163	0.270	<u>0.138</u>	<u>0.241</u>	0.176	0.270	0.125	0.236	0.227	0.348	0.202	0.317	0.551	0.529	0.333	0.425	0.155	0.260	0.247	0.334
	96	0.209	0.309	0.163	0.267	0.196	0.291	<u>0.164</u>	<u>0.275</u>	0.348	0.434	0.262	0.367	1.057	0.787	0.457	0.515	0.228	0.317	0.360	0.412

Table 2: Multivariate long-term forecasting errors in terms of MSE and MAE. The lower the better. PEMS03 dataset is with forecasting horizon $T \in \{12, 24, 48, 96\}$. For the others, $T \in \{96, 192, 336, 720\}$. The best results are highlighted in **bold** and the second best results are highlighted with an underline.

Length	Accuracy with M.M.	Accuracy without M.M.
96	0.149 / 0.208	0.151 / 0.209
192	0.197 / 0.256	0.201 / 0.263
336	0.249 / 0.297	0.255 / 0.306
720	0.320 / 0.350	0.320 / 0.352

Table 3: Ablation Study of Memory Mixing (M.M.) on Weather dataset. Values represent MSE/MAE.

Dataset	Length	CI P-sLSTM			CM P-sLSTM		
		Train	Val	Test	Train	Val	Test
Weather	96	0.151	0.138	0.167	0.049	0.215	0.227
	192	0.163	0.150	0.180	0.052	0.238	0.234
	336	0.170	0.158	0.190	0.054	0.246	0.265
	720	0.173	0.155	0.196	0.057	0.206	0.260

Table 4: CI vs CM strategies on P-sLSTM.

Dataset	P-sLSTM (seconds)	iTransformer (seconds)
Weather	52.54	79.12
ETTm1	93.13	114.25

Table 5: Comparison of Training Time for P-sLSTM and iTransformer across Weather and ETTm1 Datasets.

channel independence is not a unique idea in Transformer-

based models and MLP-based models. It has the potential to be expanded to more neural network structures.

Time Efficiency Study We compared the computational efficiency of P-sLSTM with iTransformer (the second best model) for Weather and ETTm1 in Table 5. P-sLSTM has the lowest computational cost, indicating its potential.

Conclusion

In this paper, we develop an LSTM-based method, P-sLSTM, for long term TSF task. We utilize the framework of the sLSTM from NLP in combination with patching to solve the potential short memory problem of original LSTMs or RNNs in general, as well as channel independence techniques to avoid overfitting problems. We anticipate that this work will inspire a renewed exploration of RNN-based/LSTM-based models in TSF tasks, offering valuable insights into RNN structures and applications.

In future work, we might consider more complex patching techniques to preserve as much of the original periodicity of time series as possible. Additionally, there are still known limitations of LSTM/RNNs, such as that they cannot be computed in parallel. To help models perform parallel computing, we can consider adding mLSTM, another LSTM structure that can perform parallel computing.

Acknowledgments

This work is partially supported by NSFC grant (62250410368), NSFC grant (62402414) and KGR fund (24KKSGR024).

References

- Arjovsky, M.; Shah, A.; and Bengio, Y. 2016. Unitary evolution recurrent neural networks. In *International conference on machine learning*, 1120–1128. PMLR.
- Beck, M.; Pöppel, K.; Spanring, M.; Auer, A.; Prudnikova, O.; Kopp, M.; Klambauer, G.; Brandstetter, J.; and Hochreiter, S. 2024. xLSTM: Extended Long Short-Term Memory. *arXiv preprint arXiv:2405.04517*.
- Durbin, J.; and Koopman, S. J. 2012a. *Time series analysis by state space methods*, volume 38. Oxford University Press.
- Durbin, J.; and Koopman, S. J. 2012b. *Time Series Analysis by State Space Methods*. Oxford University Press.
- Gers, F. A.; Schmidhuber, J.; and Cummins, F. 1999. Learning to forget: Continual prediction with LSTM. *Neural Computation*, 12(10): 2451–2471.
- Graves, A. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; Gupta, A.; and Ré, C. 2022. On the parameterization and initialization of diagonal state space models. *Advances in Neural Information Processing Systems*, 35: 35971–35983.
- Guo, F.; Mo, H.; Wu, J.; Pan, L.; Zhou, H.; Zhang, Z.; Li, L.; and Huang, F. 2024. A Hybrid Stacking Model for Enhanced Short-Term Load Forecasting. *Electronics*, 13(14): 2719.
- Hewamalage, H.; Bergmeir, C.; and Bandara, K. 2021. Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1): 388–427.
- Hochreiter, S.; Bengio, Y.; Frasconi, P.; Schmidhuber, J.; et al. 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In *A field guide to dynamical recurrent neural networks*. IEEE Press In.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hou, B.-J.; and Zhou, Z.-H. 2020. Learning with interpretable structure from gated RNN. *IEEE transactions on neural networks and learning systems*, 31(7): 2267–2279.
- Hu, J.; Lan, D.; Zhou, Z.; Wen, Q.; and Liang, Y. 2024. Time-SSM: Simplifying and Unifying State Space Models for Time Series Forecasting. *arXiv preprint arXiv:2405.16312*.
- Jing, L.; Shen, Y.; Dubcek, T.; Peurifoy, J.; Skirlo, S.; LeCun, Y.; Tegmark, M.; and Soljačić, M. 2017. Tunable efficient unitary neural networks (eunn) and their application to rnns. In *International Conference on Machine Learning*, 1733–1741. PMLR.
- Kim, H.-S.; and Kim, J.-G. 1994. Nonlinear state space model approach to time series analysis. *Journal of Time Series Analysis*, 15(6): 545–553.
- Li, Z.; Qi, S.; Li, Y.; and Xu, Z. 2023. Revisiting long-term time series forecasting: An investigation on linear mapping. *arXiv preprint arXiv:2305.10721*.
- Lim, B.; Arık, S. Ö.; Loeff, N.; and Pfister, T. 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4): 1748–1764.
- Lin, F.; Yuan, X.; Zhang, Y.; Sigdel, P.; Chen, L.; Peng, L.; and Tzeng, N.-F. 2023. Comprehensive transformer-based model architecture for real-world storm prediction. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 54–71. Springer.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024. itransformer: Inverted transformers are effective for time series forecasting. In *International Conference on Learning Representations*.
- Lu, J.; Han, X.; Sun, Y.; and Yang, S. 2024. CATS: Enhancing Multivariate Time Series Forecasting by Constructing Auxiliary Time Series as Exogenous Variables. *arXiv preprint arXiv:2403.01673*.
- Meyn, S. P.; and Tweedie, R. L. 2012. *Markov chains and stochastic stability*. Springer Science & Business Media.
- Nie, Y.; Kong, Y.; Dong, X.; Mulvey, J. M.; Poor, H. V.; Wen, Q.; and Zohren, S. 2024. A Survey of Large Language Models for Financial Applications: Progress, Prospects and Challenges. *arXiv preprint arXiv:2406.11903*.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A time series is worth 64 words: Long-term forecasting with transformers. *International Conference on Learning Representations*.
- Oreshkin, B. N.; Carpov, D.; Chapados, N.; and Bengio, Y. 2020. N-BEATS: Neural basis expansion analysis for time series forecasting. *International Conference on Learning Representations*.
- Ruan, K.; and Di, X. 2024. InfoSTGCAN: An Information-Maximizing Spatial-Temporal Graph Convolutional Attention Network for Heterogeneous Human Trajectory Prediction. *Computers*, 13(6): 151.
- Siami-Namini, S.; Tavakoli, N.; and Namin, A. S. 2018. A comparison of ARIMA and LSTM in forecasting time series. In *2018 17th IEEE international conference on machine learning and applications (ICMLA)*, 1394–1401. Ieee.
- Wang, Z.; Kong, F.; Feng, S.; Wang, M.; Zhao, H.; Wang, D.; and Zhang, Y. 2024. Is Mamba Effective for Time Series Forecasting? *arXiv preprint arXiv:2403.11144*.
- Wang, Z.; Nie, Y.; Sun, P.; Nguyen, N. H.; Mulvey, J.; and Poor, H. V. 2023. St-mlp: A cascaded spatio-temporal linear framework with channel-independence strategy for traffic forecasting. *arXiv preprint arXiv:2308.07496*.
- Wang, Z.; Sun, P.; and Boukerche, A. 2022. A novel time efficient machine learning-based traffic flow prediction method for large scale road network. In *ICC 2022-IEEE International Conference on Communications*, 3532–3537. IEEE.
- Wang, Z.; Sun, P.; Hu, Y.; and Boukerche, A. 2022. A novel mixed method of machine learning based models in vehicular traffic flow prediction. In *Proceedings of the 25th International ACM Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems*, 95–101.

Wen, Q.; Zhou, T.; Zhang, C.; Chen, W.; Ma, Z.; Yan, J.; and Sun, L. 2023. Transformers in time series: a survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 6778–6786.

Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *International Conference on Learning Representations*.

Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. *Advances in Neural Information Processing Systems*, 34: 22419–22430.

Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11121–11128.

Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *International Conference on Learning Representations*.

Zhao, J.; Huang, F.; Lv, J.; Duan, Y.; Qin, Z.; Li, G.; and Tian, G. 2020. Do RNN and LSTM have long memory? In *International Conference on Machine Learning*, 11365–11375. PMLR.

Zhao, S.; Dong, Z.; Cao, Z.; and Douady, R. 2024. Hedge Fund Portfolio Construction Using PolyModel Theory and iTransformer. *arXiv preprint arXiv:2408.03320*.

Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; and Xiong, H. 2021. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12): 11106–11115.

Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, 27268–27286. PMLR.

Zhou, T.; Niu, P.; Sun, L.; Jin, R.; et al. 2023. One fits all: Power general time series analysis by pretrained lm. *Advances in neural information processing systems*, 36: 43322–43355.