Understanding Transformers in Time Seires Forecasting: A Case Study on MOIRAI

Anonymous authorsPaper under double-blind review

ABSTRACT

We give a comprehensive theoretical analysis of transformers as time series prediction models, with a focus on MOIRAI (Woo et al., 2024). We study its approximation and generalization capabilities. First, we demonstrate that there exist transformers that fit an autoregressive model on input univariate time series via gradient descent. We then analyze MOIRAI, one of the state-of-the-art multivariate time series prediction models capable of modeling arbitrary number of covariates. We prove that MOIRAI is capable of automatically fitting autoregressive models with an arbitrary number of covariates, offering insights into its design and empirical success. For generalization, we establish learning bounds for pretraining when the data satisfies Dobrushin's condition. Experiments support our theoretical findings, highlighting the efficacy of using transformers for time series forecasting.

1 Introduction

Recent advancement of transformers is reshaping the field of time series forecasting. Numerous studies have demonstrated transformers are an effective architectures for various time series analysis tasks such as forecasting (Woo et al., 2024; Ansari et al., 2024; Liang et al., 2024; Das et al., 2023), anomaly detection (Zhang and Luo, 2025; Wen et al., 2025) and more. While recent works have made efforts to enhance transformers for times series forecasting by modifying model architecture (Zhang and Yan, 2023; Liu et al., 2023; Wu et al., 2021) or data processing (Reneau et al., 2023; Nie et al., 2022), the research community has yet to understand how transformers perform so well on forecasting tasks even on the simplest settings. Moreover, existing methods (Zhang and Yan, 2023; Wu et al., 2021; Woo et al., 2024; Liu et al., 2023) predominantly depend on heuristic reasoning and are notably deficient in rigorous theoretical analysis. A deeper understanding of transformers in time series forecasting will guide the design of more effective architectures by future practitioners.

A critical challenge for transformers to handle time series data is to handle arbitrary number of covariates as the architecture proposed in (Vaswani et al., 2017) is designed to handle a fixed vocabulary size. Therefore, recent studies have developed several ways to address this issue. MOIRAI (Woo et al., 2024) propose an unified way to concatenate all covariates into a single long univariate time series (any-variate encoding). MOIRAI also propose a novel any-variate mechanism for disambiguating different covariates. iTransformer (Liu et al., 2023) propose to use a pooling technique to reshape arbitrary number of covariates into a unified size. Other models choose to discard covariate features by only considering univariate time series (Ansari et al., 2024; Rasul et al., 2023). In this paper, we start by theoretically analyzing MOIRAI and its novel mechanisms. The main goal of this paper is to address the central question: What characteristics of architectures in MOIRAI contribute to their strong performance on time series data?

To understand MOIRAI in time series learning, we begin by examining one of the most widely used algorithms for time series regression: the auto-regressive (AR) regression (Hamilton, 2020). **1.** We show that transformers are indeed capable of performing AR regression on univariate time series, indicating they can handle time series in a principled, algorithmic fashion. **2.** We apply our theoretical results on a state-of-the-art time series model, MOIRAI (Woo et al., 2024), a transformer-based model that can handle arbitrary number of covariates in various time series tasks. While several novel designs in MOIRAI are for engineering purpose and heuristic, we show that these unique designs enable it to perform AR regression on arbitrary number of covariates. **3.** Our theoretical results contribute to understand the effectiveness of recent time series pretraining approaches (Ansari et al., 2024; Woo et al., 2024; Jin et al., 2023; Das et al., 2023; Rasul et al., 2023; Liang et al., 2024). These methods involve pretraining transformers on large collections of time series, often spanning diverse

domains, and have demonstrated strong empirical performance on a wide range of domains. To analyze this phenomenon, we derive generalization error bounds for pretrained transformer models under a formal statistical framework. The contributions of this paper is threefold:

- From an algorithmic approximation view, we prove the existence of a transformer capable of fitting an AR model on any given univariate time series via gradient descent. When extending this to the multi-variate setting, we theoretically verify that a MOIRAI transformer can automatically adjust the dimensionality of the AR model to fit time series with an arbitrary number of covariates. Our approximation results not only explain the strong performance of modern models across diverse datasets but also justify the design of MOIRAI.
- We present the first pretraining generalization bound for MOIRAI on time series learning. We show that when the pretraining data satisfies Dobrushin's condition, the test error can be effectively bounded even when the data does not satisfy i.i.d. Specifically, when pretraining MOIRAI on n multi-variate time series with length T, the test error decays by a rate of $1/\sqrt{nT}$.
- Our experimental results on both synthetic and real-world data match our theories by showing that
 the prediction error of transformers reduces as the input time series length increases, corresponding
 to our approximation result. This empirical finding not only verifies our theoretical results, but also
 demonstrates transformers are capable of acquiring algorithmic capability through training.

Organizations. Section 2 introduces our problem setup and preliminaries. Section 3 describes how transformers simulate different AR regression algorithms and a case study on MOIRAI. Section 4 analyzes the pretraining guarantee for learning transformers over different time series. Section 5 conducts experiments to verify our theoretical results on synthetic and real-world datasets. Section 6 discusses how our findings generalize to other models and limitations. Additional Related works are discussed in the Appendix C.

Notations. We use the following notation conventions. The vector-valued variable is given by boldfaced characters. We denote $[n] := \{1, \dots, n\}$ and $[i:j] := \{i, i+1, \dots, j\}$ for i < j. The universal constants are given by C and are ad hoc. Considering a sequence of vectors $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_T)$, we use \boldsymbol{x} without index to represent the whole sequence, and $\boldsymbol{x}_{i:j}$ represents $(\boldsymbol{x}_i, \dots, \boldsymbol{x}_T)$, for i < j. We impose periodic boundary conditions for the negative index, i.e., $\boldsymbol{x}_{-1} = \boldsymbol{x}_T$. For a vector \boldsymbol{v} we denote $\|\boldsymbol{v}\|_2$ as its L_2 norm. For a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ we denote its operator norm as $\|\boldsymbol{A}\|_2 := \sup_{\boldsymbol{v} \in \mathbb{S}^{n-1}} \|\boldsymbol{A}\boldsymbol{v}\|_2$. Random variables are given by calligraphic characters \mathcal{X} , and elements from a domain set are given by normal font \boldsymbol{x} . For more details, see Table 1.

2 PROBLEM SETUP

This section presents backgrounds and formula definitions of the transformer model, and then introduce the auto-regressive model.

Transformers. We consider a sequence of N input vectors $\{h_i\}_{i=1}^N \subset \mathbb{R}^D$. We introduce $\boldsymbol{H} \coloneqq [h_1, \cdots, h_N] \in \mathbb{R}^{D \times N}$. Given any $\boldsymbol{H} \in \mathbb{R}^{D \times N}$, we define the attention layer as follows.

Definition 2.1 (Attention layer). A self-attention layer with M heads is denoted as $\operatorname{Attn}_{\theta_0}^{\dagger}(\cdot)$ with parameters $\theta_0 = \{(V_m), (Q_m), (K_m)\}_{m \in [M]} \subset \mathbb{R}^{D \times D}$. The self-attention layer processes any given input sequence $H \in \mathbb{R}^{D \times N}$ as

$$\mathit{Attn}_{\boldsymbol{\theta}_0}^{\dagger}\left(\boldsymbol{H}\right) \coloneqq \boldsymbol{H} + \frac{1}{N} \sum_{m=1}^{M} (\boldsymbol{V}_{m}\boldsymbol{H}) \times \sigma\left(\left(\boldsymbol{Q}_{m}\boldsymbol{H}\right)^{\top}\left(\boldsymbol{K}_{m}\boldsymbol{H}\right)\right),$$

where $\sigma(t) := ReLU(t)/N$ is the ReLU function normalized by N.

Next, we introduce the any-variate attention, where Woo et al. (2024) uses it to replace the standard attention in transformers. The any-variate attention introduces two learnable variables: Attention Bias $u_1, u_2 \in \mathbb{R}$, for disambiguation between variates. In (Woo et al., 2024), they use this mechanism to allow transformers to distinguish different covariates.

Definition 2.2 (Any-variate Attention.). An any-variate attention layer with M heads is denoted as $Attn_{\theta_1}(\cdot)$ with parameters $\theta_1 = \{(V_m), (Q_m), (K_m), (u_m^1), (u_m^2)\}_{m \in [M]}$. With any input $H \in \mathbb{R}^{D \times N}$, we have

$$\textit{Attn}_{\boldsymbol{\theta}_1}\left(\boldsymbol{H}\right) \coloneqq \boldsymbol{H} + \frac{1}{N} \sum_{m=1}^{M} (\boldsymbol{V}_m \boldsymbol{H}) \times \sigma\left((\boldsymbol{Q}_m \boldsymbol{H})^\top \left(\boldsymbol{K}_m \boldsymbol{H}\right) + u_m^1 * \boldsymbol{U} + u_m^2 * \bar{\boldsymbol{U}} \right),$$

where $\sigma(t) = ReLU(t)/N$. $U \in \mathbb{R}^{N \times N}$ is a block diagonal matrix with block size $T \in \mathbb{N}^+$, such that each block consists of 1s, $\bar{U} = I - U$, and * denotes element-wise multiplication.

Definition 2.3 (MLP Layer). We denote an MLP layer with hidden state dimension D' as $\text{MLP}_{\theta}(\cdot)$ with parameters $\theta_2 = (W_1, W_2) \in \mathbb{R}^{D' \times D} \times \mathbb{R}^{D \times D'}$. The MLP layer processes any given input sequence $H \in \mathbb{R}^{D \times N}$ as $MLP_{\theta_2}(H) := H + W_2\sigma(W_1H)$.

Finally, we define a *L*-layer transformer.

Definition 2.4 (MOIRAI Transformer). We define the L-layer MOIRAI transformer (Woo et al., 2024), $TF_{\theta}(\cdot)$, as

$$\mathit{TF}_{\boldsymbol{\theta}}(\boldsymbol{H}) = \mathit{MLP}_{\boldsymbol{\theta}_2^L} \left(\mathit{Attn}_{\boldsymbol{\theta}_1^L} \left(\cdots \mathit{MLP}_{\boldsymbol{\theta}_2^1} \left(\mathit{Attn}_{\boldsymbol{\theta}_1^1}(\boldsymbol{H}) \right) \right) \right).$$

Note that this transformer is equipped with any-variate attention instead of the standard attention. For transformers with standard attention, we denote it as $TF_{\theta}^{\dagger}(\cdot)$.

We use θ to denote the vectorization of all parameters in a transformer and super-index ℓ to denote the parameter of the ℓ -th layer. Thus, the parameter of a transformer is defined by

$$\boldsymbol{\theta} = \left\{ \left\{ \left(\{\boldsymbol{Q}_m^{\ell}, \boldsymbol{K}_m^{\ell}, \boldsymbol{V}_m^{\ell}, u_m^{1,\ell}, u_m^{2,\ell} \}_{m \in [M]}, \boldsymbol{W}_1^{\ell}, \boldsymbol{W}_2^{\ell} \right) \right\}_{\ell \in [L]} \right\}.$$

We denote the "attention-only" transformers with $W_1^{(\ell)}, W_2^{(\ell)} = 0$, as $TF_{\theta}^0(\cdot)$ for shorthand. We define the following norm of a MOIRAI transformer as

$$\|\boldsymbol{\theta}\|_{op} \coloneqq \max_{\ell \in [L]} \Big\{ \max_{m \in [M^{\ell}]} \Big\{ \|\boldsymbol{Q}_m^{\ell}\|_2, \|\boldsymbol{K}_m^{\ell}\|_2, |u_m^{1\ell}|, |u_m^{2\ell}| \Big\} + \sum_{m=1}^{M^{\ell}} \|\boldsymbol{V}_m^{\ell}\|_2 + \|\boldsymbol{W}_1^{\ell}\|_2 + \|\boldsymbol{W}_2^{\ell}\|_2 \Big\},$$

where M^{ℓ} is the number of heads of the ℓ -th Attention layer.

Auto-regressive Model. Here, we first consider the case where we aim to find a multi-layered transformer that performs regression via In-context learning (ICL). Specifically, we assume our data is generated from an autoregressive process $\mathsf{AR}_d(q)$ as follows, where q,d denotes the steps of lag and number of covariates, respectively. Consider a sequence of data $x \in \mathbb{R}^{d \times T} := (x_1, \dots, x_T)$, where $x_t = (x_t^1, \dots, x_t^d) \in \mathbb{R}^d$. Assuming our target (variate of interest) is univariate, we define the $\mathsf{AR}_d(q)$ process generates x_t^1 as follows:

$$x_t^1 = \sum_{i=1}^q \sum_{j=1}^d a_i^j \cdot x_{t-i}^j + \epsilon_t = \sum_{j=1}^d \langle \boldsymbol{w}^j, \boldsymbol{x}_{t-q:t-1}^j \rangle + \epsilon_t,$$
 (2.1)

where $\epsilon_t \sim N(0,1)$, $a_i^j \in \mathbb{R}^1$. We denote the concatenation of all weights $\boldsymbol{w}^\star = (\boldsymbol{w}_1, \cdots, \boldsymbol{w}^j) \in \mathbb{R}^{qd}$. We assume bounded features $||\boldsymbol{x}_{t-q:t-1}||_2 \leq B_x$, for all $t \in [T]$. The first equation writes the AR process in scalar form, and the second writes it in vector form. In the following chapters, we will start by considering the uni-variate case (AR₁) and then move on to the multi-variate case (AR_d).

3 TRANSFORMERS SIMULATE AUTO-REGRESSIVE REGRESSION

We first investigate how MOIRAI Next, we will move on to multi-variate AR regression and analyze how MOIRAI's unique design and pre-processing methods enable its advantages.

3.1 WARM UP: UNIVARIATE AUTOREGRESSIVE REGRESSION

We start our analysis with a warm-up example on the $AR_1(q)$ model. We show that standard transformers are capable of performing gradient descent via in-context learning on autoregressive data. The results in this subsection apply to univariate time series models such as (Ansari et al., 2024). Here, we consider an input sequence with the following form

$$\boldsymbol{H} \coloneqq \begin{bmatrix} x_1 & x_2 & \dots & x_T & 0 \\ \boldsymbol{p}_1 & \boldsymbol{p}_2 & \dots & \boldsymbol{p}_T & \boldsymbol{p}_{T+1} \end{bmatrix} \in \mathbb{R}^{D \times (T+1)}, \quad \boldsymbol{p}_i \coloneqq \begin{bmatrix} \boldsymbol{0}_{d'} \\ \boldsymbol{e}_i \\ 1 \\ 1\{i < T\} \end{bmatrix} \in \mathbb{R}^{d' + T + 3}, \quad (3.1)$$

where e_i is an one-hot vector with 1 at the *i*-th entry, and d' + T + 3 = D. Here, our goal is to predict x_{T+1} .

Remark 3.1. Most in-context learning studies (Akyürek et al., 2023; Bai et al., 2024; Li et al., 2023) make an assumption on the input data, where they assume it is formatted with features and labels in the same column, i.e., $\begin{bmatrix} x_1 & x_2 & \dots & x_N \\ y_1 & y_2 & \dots & y_N \end{bmatrix}$. In contrast, we adopt a natural approach that leverages the raw structure of the data, particularly for the $AR_d(q)$ process. In this setting, each time step's label also serves as a feature for future steps. Further, the unknown value of q complicates the task of achieving such a format in Remark 3.1.

Our next lemma shows that transformers are indeed capable of reformatting H into the form in Remark 3.1. Notably, we relax the assumption in Remark 3.1 of previous studies as well.

Lemma 3.2. Given a sequence of token H in the form of Equation 3.1, there exists a one-layer, q_{max} head attention layer, such that for any $q \leq q_{\text{max}}$, the columns of $Attn^{\dagger}_{\theta}(H)$ has the following form:

$$Attn_{\boldsymbol{\theta}_1}(\boldsymbol{H})_i := \begin{bmatrix} x_i & x_{i-1} & \cdots & x_{i-q} & \boldsymbol{p}_i' \end{bmatrix}^\top, \quad \boldsymbol{p}_i' := \begin{bmatrix} \mathbf{0}_{d'-q} \\ \boldsymbol{e}_i \\ 1 \\ 1\{i < T\} \end{bmatrix}. \tag{3.2}$$

The proof is in Appendix E.1. Lemma 3.2 is crucial in our analysis as it connects the theoretical results in ICL (Bai et al., 2024) to univariate time series forecasting. When data formats in the form of Remark 3.1, Bai et al. (2024) show that there exists a multi-layer transformer that performs linear regression via gradient descent on the first N-1 data points and evaluates the N-th one. Thus, Lemma 3.2 implies transformers are also capable of performing linear regression on time series data.

This lemma applies to both any-variate attention and standard attention, as the latter can be viewed as a special case of any-variate attention by setting $u^1, u^2 = 0$. Additionally, the construction of a single layer with q heads is not a strict requirement; the lemma also holds for c layers of q/c head attention, for any c satisfies $q/c \geq 2$.

With Lemma 3.2, we are able to apply the in-context learning results in (Bai et al., 2024) on the $AR_1(q)$ case. Consider the data generated by the AR process in Equation 2.1. Given an input time series $x \in \mathbb{R}^{d \times T}$, we define the least squares estimator as the empirical risk minimizer over the time series, i.e.,

$$egin{aligned} \ell_{ ext{reg}}(oldsymbol{w}, oldsymbol{x}_{t-1:t-q}) &\coloneqq rac{\left[\langle oldsymbol{w}, [oldsymbol{x}_{t-1:t-q}^1; \ldots; oldsymbol{x}_{t-1:t-q}^d]
angle - x_t^1
ight]^2}{2} \ L_{ ext{reg}}(oldsymbol{w}, oldsymbol{x}) &\coloneqq rac{1}{T-1} \sum_{t=1}^{T-1} \ell_{ ext{reg}}\left(oldsymbol{w}, oldsymbol{x}_{t-1:t-q}
ight), \quad \widehat{oldsymbol{w}}_{ ext{ERM}} \coloneqq \operatornamewithlimits{argmin}_{oldsymbol{w} \in \mathbb{R}^{dq}} L_{ ext{reg}}\left(oldsymbol{w}, oldsymbol{x}
ight), \end{aligned}$$

where [v;u] denotes the concatenation between vectors, as $[x_{t-1:t-q}^1;x_{t-1:t-q}^2]=(x_{t-1}^1,x_{t-2}^1,\cdots,x_{t-q+1}^2,x_{t-q}^2)\in\mathbb{R}^{2q},\ \widetilde{x}$ denotes masking out the last time step of the target variate, and L_{reg} is a loss, which is α -strongly convex, and β -smooth over \mathbb{R}^{dq} . We make the following assumption and then present our first result on univariate time series (d=1).

Proposition 3.3 (Univariate Autoregressive Regression via MOIRAI). Assume the regression problem is well-defined and has a bounded solution and fix a $q_{\max} > 0$. For any $0 \le \alpha \le \beta$ with $\kappa := \frac{\beta}{\alpha}$, $B_w > 0$, and $\epsilon < B_x B_w / 2$, there exists a L-layer MOIRAI transformer TF_{θ}^0 (·), with

$$L = L_1 + L_2, \ L_1 = \lceil 2\kappa \log(\frac{B_x B_w}{2\epsilon}) \rceil, \ L_2 = \lceil \frac{q_{\max}}{3} \rceil, \ _{max\ell \in [L]} M^{(\ell)} \le 3, \ ||\boldsymbol{\theta}||_{op} \le |4R + 8\beta^{-1}|,$$

 $(R \coloneqq \max\{B_xB_w,B_x,1\})$, the following holds. On any input data x generated by any $\mathtt{AR}_1(q)$ process such that $0 < q \leq q_{\max}, \ ||\widehat{\boldsymbol{w}}_{\mathit{ERM}}||_2 \leq {}^{B_w}/2$, we have

$$\|\widehat{\boldsymbol{x}}_T - \left\langle \widehat{\boldsymbol{w}}_{\text{ERM}}, [\boldsymbol{x}_{t-1:t-q}^1; \dots; \boldsymbol{x}_{t-1:t-q}^d] \right\rangle\| \le \epsilon, \tag{3.3}$$

where $\widehat{x}_T = \text{read}(TF_{\theta}^0(H))$. read(H) operation reads out the first entry of T-th column of H.

This proposition follows immediately from Lemma 3.2 and (Bai et al., 2024, Theorem 4). The above result applies for standard transformers as well as in our construction u_m^1 , $u_m^2 = 0$ in all heads and

layers. Further, one can replace the least squares ERM with lasso or ridge ERM and obtain a similar result by applying Theorem 4, 7, and 13 of (Bai et al., 2024).

So far, we show that transformers are capable of solving univariate AR regression with, at best, one additional layer compared to the results in (Bai et al., 2024). The result above provides insights on transformer-based univariate time series foundation models (Ansari et al., 2024; Rasul et al., 2023; Das et al., 2023). To study MOIRAI, we then include two ingredients into our analysis: the any-variate encoding and the covariates in the following chapters.

3.2 CASE STUDY: MOIRAI TRANSFORMER

In this subsection, we extend our results to the multivariate autoregressive process (d > 1) and the encoding method of MOIRAI. Note that in the multi-variate case, we only focus on MOIRAI as it is the only transformer-based model that is compatible with arbitrary number of covariates. We start by introducing the any-variate encoding.

Any-Variate Encoding. Woo et al. (2024) propose to flatten a d-dimensional time series, $\boldsymbol{x} \in \mathbb{R}^{d \times T}$, into a 1-dimensional sequence, i.e., $\boldsymbol{x}' \in \mathbb{R}^{1 \times Td}$. This operation transforms time series with arbitrary number of covariates (d), into a long sequence with fixed dimension, enabling consistent input dimension for transformers. Following the flattening operation, Woo et al. (2024) also proposes to add two types of indices into the input sequence: the time and variate ID. We term the above operations as the any-variate encoding, which transforms a multivariate sequence $\boldsymbol{x} \in \mathbb{R}^{d \times T}$, as follows:

$$\begin{bmatrix} x_1^1 & \cdots & x_T^1 \\ x_1^2 & \cdots & x_T^2 \\ \vdots & \vdots & \vdots \\ x_1^d & \cdots & x_T^d \end{bmatrix} \rightarrow \begin{bmatrix} x_1^1 & \cdots & x_T^1 & \cdots & x_1^d & \cdots & x_T^d \\ \mathbf{p}_1 & \cdots & \mathbf{p}_T & \cdots & \mathbf{p}_1 & \cdots & \mathbf{p}_T \\ \mathbf{e}_1 & \cdots & \mathbf{e}_1 & \cdots & \mathbf{e}_d & \cdots & \mathbf{e}_d \end{bmatrix}, \tag{3.4}$$

where e_i is the variate index, a one-hot vector with i-th entry being 1, and p_i is the time index, which is defined the same as Equation (3.1). This is without loss of generality because the discrete-time and variate ID used in (Woo et al., 2024) can be easily transformed into a high-dimensional vector with the embedding layer. Note that only the target variate has length T, we highlight x_T^1 as it is our prediction target and will be masked as 0.

Now we define the history matrix $\mathbf{A}_i(q) \in \mathbb{R}^{q+1 \times T}$ for the *i*-th covariates (x_1^i, \cdots, x_T^i) , with order q, such that

$$\mathbf{A}_i(q)_{\mu,\nu} := x_{\nu-\mu+1}^i, \quad \text{ for } \mu \in [d], \nu \in [q],$$

where in the j-th column of $A_i(q)$, it contains historical values of x_i^i with lag q > 0.

Lemma 3.4. Fix q_{\max} , $D \in \mathbb{N}^+$. Given any T > 0, d' > q > 0, d > 0 such that T > q, $q_{\max} \ge q$. For any input matrix \boldsymbol{H} in the form of any-variate encoding in Equation 3.4, such that $\boldsymbol{H} \in \mathbb{R}^{D \times dT}$. There exists a one layer, q_{\max} head **any-variate attention** that performs the following operation.

$$\begin{bmatrix} x_1^1 & \cdots & x_T^1 & x_1^2 & \cdots & x_T^2 & \cdots & x_1^d & \cdots & x_T^d \\ \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \cdots & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T \\ \boldsymbol{e}_1 & \cdots & \boldsymbol{e}_1 & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_d & \cdots & \boldsymbol{e}_d \end{bmatrix} \mapsto \begin{bmatrix} \mathbf{A}_1(q) & \mathbf{A}_2(q) & \cdots & \mathbf{A}_d(q) \\ \mathbf{0}_{d'' \times T} & \mathbf{0}_{d'' \times T} & \cdots & \mathbf{0}_{d'' \times T} \\ \vdots & \vdots & \ddots & \ddots & \vdots \end{bmatrix},$$

where $d^{"}=d'-q_{\max}$.

The proof is in Appendix E.2. Intuitively, the above operation performs the same operation in Lemma 3.2 but in a variate-wise fashion. Lemma 3.4 is crucial to our understanding to MOIRAI as it shows that any-variate attention is capable of organizing the history of each variate efficiently and in a parallel way. To again achieve the format in Remark 3.1, one has to stack all $A_i(q)$ in the same columns, which can be easily done by a single layer of attention via Lemma 3.2 and (Bai et al., 2024, Proposition A.5) (details in Appendix E.2). This lemma serves as a foundation for MOIRAI to handle multi-variate time series with in-context learning which we present as the theorem below.

Remark 3.5. Comparing to Lemma 3.2, Lemma 3.4 is specifically for any-variate attention in our construction, where we demonstrate that several special mechanisms in any-variate attention enables variate-wise operations in parallel.

Theorem 3.6 (Any-variate Autoregressive Regression via MOIRAI). Let $L_1 = \lceil \frac{q_{\max}}{3} \rceil + 1$, $L_2 = \lceil 2\kappa \log \frac{B_x B_w}{2\epsilon} \rceil$, with some $d_{\max}, q_{\max} > 0$. For any $0 \le \alpha \le \beta$ with $\kappa := \frac{\beta}{\alpha}$, $B_w > 0$, and

 $\epsilon < B_x B_w/2$, there exists an $(L_1 + L_2)$ -layer of MOIRAI transformer equipped with any-variate Attention, satisfies the following

$$\max_{\ell \in [L_1 + 1, L_2]} M^{(\ell)} \le 3, \ \|\boldsymbol{\theta}\| \le |4R + 8\beta^{-1}|, \quad \sum_{\ell = 1}^{L_1} M^{(\ell)} = d_{\max} + q_{\max}, \quad D = (q_{\max} + 1)d_{\max} + T + 2,$$

such that for any input time series x with length T generated from an $AR_d(q)$ process, where $x \in \mathbb{R}^{d \times T}$, $q \leq q_{\max}$, $d \leq d_{\max}$, it satisfies

$$\|\widehat{\boldsymbol{x}}_{T}^{1} - \left\langle \boldsymbol{w}_{i}^{\star}, [\boldsymbol{x}_{T-1:T-q}^{1}; \dots; \boldsymbol{x}_{T-1:T-q}^{d}] \right\rangle \| \leq \epsilon, \tag{3.5}$$

where $\widehat{x}_T^1 = \text{read}(TF_{\boldsymbol{\theta}}^0(\boldsymbol{H}))$, and $\boldsymbol{H} \in \mathbb{R}^{D \times N}$ is the any-variate encoding of \boldsymbol{x} .

Remark 3.7. Theorem 3.6 indicates there exists a MOIRAI transformer that fits an autoregressive model on time series as long as the number of covariates no greater than d_{\max} and lags no greater than q_{\max} . This shows its ability to infer the underlying AR model in a principled way and provides a possible explanation for its zero-shot performance on a wide range of datasets. To modify the setting from multi-variate prediction, one can easily apply Bai et al. (2024, proposition A.5), where the hidden dimension scales linearly w.r.t. the number of variates to predict.

The proof is in Appendix D. Observe that there exist two trade-offs in Theorem 3.6. First, $q_{\rm max}d_{\rm max}$ is upper bounded by the D (up to constant), which is a natural trade-off in our construction. Second, the approximation error is roughly $O(e^{-L})$, suppressed exponentially by the number of layers, as in our analysis, each layer of MOIRAI performs a single step of gradient descent on $L_{\rm reg}$.

4 Analysis on Pretraining

Let π be a meta distribution, and each distribution drawn from it $P^{(T)} \sim \pi$, satisfies Dobrushin's condition (Dobrushin, 1968) (which we will introduce shortly). For pretraining data, we first sample n distributions $P_j^{(T)}$ i.i.d. from π , and for each distribution, we sample a time series $(\boldsymbol{x}_{1j}, \cdots, \boldsymbol{x}_{Tj})$, for $j \in [n]$, and each of them contains no more than d covariates and with lag step no more than q.

For each time series, we encode it with any-variate encoding into an input matrix denoted as $\mathbf{H} \in \mathbb{R}^{D \times N}$, ¹ We define each pretraining sample as $\mathbf{z}_j \coloneqq (\mathbf{H}_j, y_j)$, where $y_j = \mathbf{z}_{Tj}^1$. We consider the squared loss between model prediction and the label, i.e.

$$\ell(\boldsymbol{z}_{t},\boldsymbol{\theta})\coloneqq\frac{1}{2}\Big[y_{t}-\mathtt{Clip}_{B_{x}}\Big(\mathtt{read}_{y}\Big(\mathtt{TF}_{\boldsymbol{\theta}}^{R}\left(\boldsymbol{H}\right)\Big)\Big)\Big]^{2},$$

where $\mathrm{Clip}_{B_x}(t) \coloneqq \max\{\min\{t, B_x\}, -B_x\}$, and $\mathrm{TF}_{\pmb{\theta}}^R$ is the MOIRAI transformer defined in Definition 2.4 with $\mathrm{Clip}(\cdot)$ applied after each layer. The pretraining loss and test loss is defined as the following:

$$\widehat{L}(\boldsymbol{\theta}) := \frac{1}{nT} \sum_{t=1}^{T} \sum_{j=1}^{n} \ell(\boldsymbol{\theta}, \boldsymbol{z}_{jt}), \ L(\boldsymbol{\theta}) := \mathbb{E}_{\boldsymbol{z}, P^{(T)}} \left[\ell(\boldsymbol{\theta}, \boldsymbol{z}) \right]. \tag{4.1}$$

The goal of our pretraining algorithm is to find an empirical risk minimizer (ERM) over MOIRAI transformers with L layers, M heads, and norm bounded by B:

$$\widehat{\boldsymbol{\theta}} \coloneqq \underset{\boldsymbol{\theta} \in \Theta_{L,M,D',B}}{\operatorname{argmin}} \widehat{L}(\boldsymbol{\theta}),
\Theta_{L,M,D',B} \coloneqq \left\{ \boldsymbol{\theta} = \left(\boldsymbol{\theta}_1^{(1:L)}, \boldsymbol{\theta}_2^{(1:L)} \right) : \max_{\ell \in [L]} M^{(\ell)} \le M, \max_{\ell \in [L]} D^{(\ell)} \le D', \ \|\boldsymbol{\theta}\|_{op} \le B \right\}.$$
(4.2)

4.1 Non IID Time Series

Here, we consider the training data x to be drawn from a distribution P satisfying Dobrushin's condition. Under this condition, we are able to present several generalization bounds on pretraining. **Definition 4.1** (Dobrushin's Uniqueness Condition). Let $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_T)$ be a sequence of random variables over $\mathcal{D}_{\mathcal{X}}^T$. The Dobrushin coefficient α of \mathcal{X} and the influence \mathbf{I} of variable \mathcal{X}_j on variable \mathcal{X}_i are defined as

$$\alpha(\mathcal{X}) \coloneqq \max_{1 \leq i \leq T} \sum_{j \neq i} \mathbf{I}_{j \rightarrow i}(\mathcal{X}), \quad \mathbf{I}_{j \rightarrow i}(\mathcal{X}) \coloneqq \max_{x_{-i-j}, x_j, x_j'} \left| \left| P_{\mathcal{X}_i \mid \mathcal{X}_{-i}} \left(\cdot | x_{-i-j}, x_j \right), P_{\mathcal{X}_i \mid \mathcal{X}_{-i}} \left(\cdot | x_{-i-j}, x_j' \right) \right| \right|_{TV},$$

¹Due to any-variate encoding, N = dT.

where $x_{-i-j} \in \mathcal{D}_{\mathcal{X}}^{T-2}, x_j, x_j' \in \mathcal{D}_{\mathcal{X}}, ||\cdot||_{TV}$ denotes the total variation distance, and x_{-i} represents the vector \mathbf{x} after omitting the i-th element. We say the variable satisfies Dobrushin's uniqueness condition if $\alpha(\mathcal{X}) < 1$. For a distribution P, we denote $\alpha(P) = \sup_{\mathcal{X} \sim P} \alpha(\mathcal{X})$.

Definition 4.2 (Log Dobrushin's Coefficients). Let $\mathcal{X} = (\mathcal{X}_1, \cdots, \mathcal{X}_T)$ be a random variable over $\mathcal{D}_{\mathcal{X}}^T$ and let P_z denote its density. Assume that $P_z > 0$ on all Ω^T . For any $i \neq j \in [T]$, the log influence between j and i is defined as:

$$I_{j,i}^{\log}(\mathcal{X}) = \frac{1}{4} \sup \log \frac{P[x_i, x_j, x_{-i-j}] P[x_i', x_j', x_{-i-j}]}{P[x_i', x_j, x_{-i-j}] P[x_i, x_j', x_{-i-j}]},$$

where the sup is taken over $x_{-i-j}, x_i, x_i', x_j, x_j'$, and the log-coefficient of \mathcal{X} is defined as $\alpha_{\log}(\mathcal{X}) = \max_{i \in [T]} \sum_{j \neq i} I_{j,i}^{\log}(\mathcal{X})$. Note that $\alpha(\cdot)$ has a natural bound in [0, T-1], with $\alpha=0$ implies i.i.d.

Remark 4.3 (On the choice of Dobrushin's condition). In general, time series analysis (non IID) requires certain level of regularity condition on data, to derive meaningful learning bounds and exclude extream cases. For example, mixing properties (Mohri and Rostamizadeh, 2010; Steinwart and Christmann, 2009; Kuznetsov and Mohri, 2014), hypercontractivity (Ziemann and Tu, 2022), or Dobrushin's condition (Dagan et al., 2019). Dobrushin's condition is not only a commonly used regularity condition for time series data, it also fits several common statistical models, which makes our derivation suitable for further analysis on different models. An example is shown in our section 4.3. One more benefit of having this condition instead of mixing properties is that we do not need to assume stationarity of the data, which can easily be violated in practice. We highlight that verifying all those conditions is difficult in practice, as the true distribution of data is intractable.

4.2 GENERALIZATION BOUNDS OF MOIRAI

Theorem 4.4 (Pretraining Generalization Bound). Let $\Theta_{L,M,D',B}$ be the parameter space defined in Equation 4.2. Assume $\alpha_{\log}(P^{(T)}) < 1/2$. Then with probability at least $1 - \varepsilon$, ERM $\widehat{\theta}$ satisfies the following:

$$L(\widehat{\boldsymbol{\theta}}) \leq \inf_{\boldsymbol{\theta} \in \Theta_{L,M,D',B}} L(\boldsymbol{\theta}) + O\left(\frac{B_x^2}{1 - \alpha^n(\mathbf{P}^{(T)})} \sqrt{\frac{L(MD^2 + DD')\zeta + \log(1/\varepsilon)}{nT}}\right),$$

where C is an universal constant, and $\zeta = O(\log(2 + \max\{B, R, B_x, T, d\}))$.

The proof is in Appendix E.4. Here we observe that the bound mainly depends on two terms, the Dobrushin's coefficient, and the demoninator nT. Under Dobrushin's condition, increasing n alleviates the performance degradation of $1/1-\alpha$ exponentially in n. Further, increasing n also tightens the bound with the order $1/\sqrt{n}$. Further, in Theorem 4.4, we do not assume our data is generated from the AR process, only its Dobrushin coefficient. When the data is generated by the AR process, we are able to give a more explicit bound on the same test loss as described below.

Proposition 4.5 (Test Error Bound). Following the setup in Theorem 4.4, if pretraining samples are generated by some $AR_d(q)$ process with noise sampled from $N(0, \sigma_{\epsilon}^2)^2$, then with probability $\Delta(1-\varepsilon)$, ERM $\hat{\theta}$ satisfies the following:

$$L(\widehat{\boldsymbol{\theta}}) \leq O\Bigg(B_x B_w \exp\bigg(\frac{-L}{\kappa}\bigg) + \frac{B_x^2}{1 - \alpha^n(\mathbf{P}^{(T)})} \sqrt{\frac{L(MD^2 + DD')\zeta + \log(1/\varepsilon)}{nT}}\Bigg).$$

where $\Delta = O\left(1 - \left(\frac{\sigma_{\epsilon}}{B_x B_w e^{-L/2\kappa}}\right)^2\right)$, C is an universal constant, and $\zeta = O(\log(2 + \max\{B, B_x, T, d\}))$.

Considering the model parameters (M,D,D',d) are of constant level, one is able to further optimize the bound to $L(\widehat{\boldsymbol{\theta}}) \lesssim (nT)^{-1/2}$ by selecting L appropriately. Next we provide an example of the application of Proposition 4.5 on AR(1) process with the following form

$$\boldsymbol{x}_{t+1} = \langle \boldsymbol{w}, \boldsymbol{x}_t \rangle + \epsilon_t, \quad \epsilon \sim N(0, \sigma_{\epsilon}^2), \quad \text{with } \boldsymbol{x}_t \in \mathbb{R}^d, \ \boldsymbol{w} \in \mathbb{R}^d, \ \epsilon \in \mathbb{R}, \ \boldsymbol{y}_{t+1} = \boldsymbol{x}_{t+1}^1.$$

To satisfy the condition of $\alpha(P) < \frac{1}{2}$, we assume $B_x^2 < \ln 1/2 + (\sigma_\epsilon^2)$, $||\boldsymbol{w}||_\infty < 1$. The first condition comes from the fact that we require the pair-wise potential of this time series to be less than 1/2 (For more details, see Appendix E.5). The second condition comes from the requirement of it being stationary.

 $^{^2}$ Here we assume fixed d,q across all samples as one can describe a lower dimension/order AR process with zero coefficients.

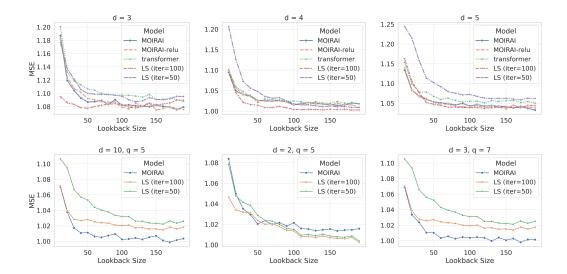


Figure 1: Top: Model performance on data with different number of covariates. For both MOIRAI-based models, we observe their performance behave like least squares. As in our construction, the longer the lookback size is, the more examples available for transformers to fit an AR. Note that our test data has variance $\sigma^2=1$, thus the MSE for both models are expected to converge to 1 as the q increases. Bottom: Generalization to unseen values of d, q. From left to right, we have MOIRAI's generalization power (pretrained on $d \in \{4,5\}, q \in \{4,5\}$) on high dimensional (d=10), low dimensional data (d=2) and high lag step + low dimensional data (d=3, q=7). Note that high and low is compared with pretraining data. We observe that even when MOIRAI did not learn from any time series with d=10, it is still able to generalize well and shows even better sample complexity than LSR. Finally, even when both q, d are unseen, it does not impact MOIRAI's ability to make accuracy predictions. The lines are the average over 5 runs.

Corollary 4.6 (Generalization Bound for MOIRAI on AR(1)). Considering n AR(1) processes with each of its Dobrushin's coefficient bounded by 1/2. With probability at least $\delta(1-\varepsilon)$, ERM $\widehat{\theta}$ satisfies the following:

$$L(\widehat{\boldsymbol{\theta}}) = O\Bigg(\frac{\sigma_{\epsilon}}{\sqrt{1-\delta}} + \frac{\sigma_{\epsilon}^2}{B_x} \exp\bigg(\frac{-L}{\kappa}\bigg) + \frac{\sigma_{\epsilon}^2}{1-\alpha^n(\mathtt{AR}(1))} \sqrt{\frac{L(MD^2 + DD')\zeta + \log(1/\varepsilon)}{nT}}\Bigg),$$

with

$$\alpha(\mathtt{AR}(1)) \; \leq \; \min \left\{ \frac{1}{2}, \; \frac{2\sqrt{2/\pi} \, B_x \|w\|_*}{\sigma_\epsilon} \right\},$$

and $\zeta = O(\log(2 + \max\{B, \mathtt{R}, B_x, d\}))$. If we further optimize the bound by viewing the hyperparameters as constants, the test error obeys $O(e^{-L} + \sqrt{\frac{L}{nT}})$ w.h.p. whenever σ_ϵ is small.

Remark 4.7. Corollary 4.6 provides an application of Dobrushin's condition on common statistical models, indicating that Dobrushin's condition can be further used to certify stability and derive generalization bounds for a broad class of time series processes beyond the i.i.d. setting.

5 EXPERIMENTS

To verify our analysis, we first train transformers on synthetic datasets generated from AR process with different parameters. The goal of this experiment is to verify the existence of a transformer that performs least squares regression (LSR) on input time series with bounded lags and number of covariates. Next we study whether a pretrained transformer is capable of generalize such an ability to unseen values of d, q. More empirical results such as on *real-world* data are in Appendix F.5.

Setup. We use MOIRAI-base, a 12 layer MOIRAI transformer. The hyperparameters of this experiment are in Table 2. We use MSE loss for pretraining instead of NLL loss for simplicity. For pretraining, we follow the setup in Woo et al. (2024) but set the patch size as 1 to minimize the impact of patch embedding. During pretraining, each time series is randomly sampled. We evaluate the pretrained model on our test data with $d = \{3, 4, 5\}$, q = 5 and $\sigma^2 = 1$ with different input length.

We compare MOIRAI with LSR performing different gradient descent steps. For LSR, we assume q is known. When MOIRAI takes a T length input, LSR is trained on T-1 samples with each having dq features. A detailed example on our implementation is in Appendix F.4. We also include the standard transformers and MOIRAI with ReLU replacing Softmax, which we term it as MOIRAI-relu. For standard transformers, we keep the any-variate encoding but use standard attention. ³Details of data generation can be found in Appendix F.3.

Results. Since our test data generation process obeys noise variance = 1, when fitting a linear model, the expected MSE will converge towards 1 as lookback size increases. Based on Theorem 3.6, the length of input time series also corresponds to the number of examples model perform least squares on via gradient descent. We observe that as the input length increases, the predictive error of MOIRAI decreases similarly to least squares, which verifies Theorem 3.6. Next, when pretrained on diverse dataset, pretrained MOIRAI is able to adapt to different number of covariates and perform least squares accordingly. Further, when replace softmax with ReLU, the performance gap is negligible. For standard transformer, while it also behaves similar to other models, it does present higher error comparing to other baselines, indicating the advantages of using any-variate attention.

Next we are interested in whether a pretrained transformer can generalize to unseen values of d and q. Thus, we train transformers (MOIRAI) on synthetic data generated with AR with $d \in \{4,5\}$, and $q \in \{4,5\}$. In our construction, pretrained transformer is compatible with lower order and dimension AR data. We evaluate the trained model on data with unseen values of d. We select d=2, d=10, to represent the scenario when the number of covariates is lower and higher than pretraining data.

Results. We observe that even when facing data with unseen number of covariates, MOIRAI is still capable of performing AR regression effectively. Note that for d=10, LS require higher sample complexity to obtain similar performance to d=5 case. However, the pretrained MOIRAI is able to outperform it from such an aspect. For d=2 all models perform well, again verifies our theoretical results. Finally, when facing data with unseen both d and q, it is still capable of performing well.

6 DISCUSSION

Here we discuss the implications of our results. First, we demonstrate how transformers perform univariate AR regression, a classic time series prediction algorithm. Next, we demonstrate that the unique any-variate attention proposed by Woo et al. (2024), indeed provides benefits for transformers to process each covariate in a parallel fashion (Lemma 3.4). This observation not only leads to our next result, showing MOIRAI can perform AR regression on arbitrary number of covariates, also provides a justification on the design of any-variate encoding (Woo et al., 2024). The above approximation results provide profound insight on designing transformer-based time series models. Next, for pretraining, we derive generalization bounds when data is under Dobrushin's condition, where the bound scales roughly by $\frac{1}{\sqrt{nT}}$. It is worth investigating how to further relax this assumption, which we leave to future work. For predicting multiple time steps or multiple variates, one can easily apply Bai et al. (2024, Proposition A.5) to obtain the same result with hidden dimension D scales linearly. More discussions on the limitations of this work is in Appendix B.

Which transformers do we cover? Our results in Section 3.1 also covers notable SOTA models including classic transformers (Vaswani et al., 2017), Chronos (Ansari et al., 2024) (Theorem E.5), MOMENT (Goswami et al., 2024). Architectures using transformer-decoders are also compatible, such as TimeGPT (Garza and Mergenthaler-Canseco, 2023), TimesFM (Das et al., 2023), see Appendix B in (Bai et al., 2024). Furthermore, the use of patch embedding Nie et al. (2022) does not affect our theoretical result as an AR sequence remains an AR sequence after patch embedding transformation (linear), thus solvable by our AR regression framework. One simply needs to use a weight-tying trick (for the last layer embedding), to obtain the prediction of the original input. Univariate models with patch embedding also fits Section 3.1.

Notable transformer-based SOTA models we do not cover include iTransformer (Liu et al., 2023), CrossFormer (Zhang and Yan, 2023), as they also applied attention mechanism on the covariate dimension. This novel technique requires further analysis thus we leave to future work. Autoformer (Wu et al., 2021) also utilizes a technique named auto-correlation, which requires different theoretical tool to study. Similarly, Time-LLM (Jin et al., 2023) and Lag-Llama (Rasul et al., 2023) utilize either natural language inputs, or the lag features, therefore cannot be covered by our theoretical analysis.

³In (Woo et al., 2024), without any-variate attention, the error of MOIRAI-small increases roughly 40%.

REFERENCES

- Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. Transformers learn to implement preconditioned gradient descent for in-context learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algorithm is in-context learning? investigations with linear models. In *The Eleventh International Conference on Learning Representations*, 2023.
- Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C Maddix, Syama Rangapuram, David Salinas, Jasper Schulz, et al. Gluonts: Probabilistic and neural time series modeling in python. *Journal of Machine Learning Research*, 21(116):1–6, 2020.
- Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.
- Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. *Advances in neural information processing systems*, 36, 2024.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Yuval Dagan, Constantinos Daskalakis, Nishanth Dikkala, and Siddhartha Jayanti. Learning from weakly dependent data under dobrushin's condition. In *Conference on Learning Theory*, pages 914–928. PMLR, 2019.
- Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. *arXiv preprint arXiv:2310.10688*, 2023.
- P. L. Dobrushin. The description of a random field by means of conditional probabilities and conditions of its regularity. *Theory of Probability and Its Applications*, 13:197–224, 1968.
- Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598, 2022.
- Azul Garza and Max Mergenthaler-Canseco. Timegpt-1, 2023.
- Rakshitha Godahewa, Christoph Bergmeir, Geoffrey I Webb, Rob J Hyndman, and Pablo Montero-Manso. Monash time series forecasting archive. *arXiv preprint arXiv:2105.06643*, 2021.
- Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski. Moment: A family of open time-series foundation models. In *International Conference on Machine Learning*, 2024.
- James D Hamilton. *Time series analysis*. Princeton university press, 2020.
- Yihan He, Yuan Cao, Hong-Yu Chen, Dennis Wu, Jianqing Fan, and Han Liu. Learning spectral methods by transformers. *arXiv preprint arXiv:2501.01312*, 2025.
- Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yux-uan Liang, Yuan-Fang Li, Shirui Pan, et al. Time-llm: Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728*, 2023.
- Christof Külske. Concentration inequalities for functions of gibbs fields with application to diffraction and random gibbs measures. *Communications in mathematical physics*, 239:29–51, 2003.
- Vitaly Kuznetsov and Mehryar Mohri. Generalization bounds for time series prediction with non-stationary processes. In *International conference on algorithmic learning theory*, pages 260–274. Springer, 2014.

- Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 95–104, 2018.
 - Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. Transformers as algorithms: Generalization and stability in in-context learning. In *International Conference on Machine Learning*, pages 19565–19594. PMLR, 2023.
 - Yuxuan Liang, Haomin Wen, Yuqi Nie, Yushan Jiang, Ming Jin, Dongjin Song, Shirui Pan, and Qingsong Wen. Foundation models for time series analysis: A tutorial and survey. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pages 6555–6565, 2024.
 - Licong Lin, Yu Bai, and Song Mei. Transformers as decision makers: Provable in-context reinforcement learning via supervised pretraining. *arXiv preprint arXiv:2310.08566*, 2023.
 - Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. *arXiv preprint arXiv:2310.06625*, 2023.
 - Arvind Mahankali, Tatsunori B Hashimoto, and Tengyu Ma. One step of gradient descent is provably the optimal in-context learner with one layer of linear self-attention. *arXiv* preprint arXiv:2307.03576, 2023.
 - Mehryar Mohri and Afshin Rostamizadeh. Stability bounds for stationary φ -mixing and β -mixing processes. *Journal of Machine Learning Research*, 11(2), 2010.
 - Eshaan Nichani, Alex Damian, and Jason D Lee. How transformers learn causal structure with gradient descent. *arXiv preprint arXiv:2402.14735*, 2024.
 - Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.
 - Arushi Rai, Kyle Buettner, and Adriana Kovashka. Strategies to leverage foundational model knowledge in object affordance grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1714–1723, 2024.
 - Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Arian Khorasani, George Adamopoulos, Rishika Bhagwatkar, Marin Biloš, Hena Ghonia, Nadhir Hassen, Anderson Schneider, et al. Lag-llama: Towards foundation models for time series forecasting. In *R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*, 2023.
 - Alex Daniel Reneau, Jerry Yao-Chieh Hu, Ammar Gilani, and Han Liu. Feature programming for multivariate time series prediction. In *International Conference on Machine Learning*, pages 29009–29029. PMLR, 2023.
 - Michael E Sander, Raja Giryes, Taiji Suzuki, Mathieu Blondel, and Gabriel Peyré. How do transformers perform in-context autoregressive learning? *arXiv preprint arXiv:2402.05787*, 2024.
 - Kai Shen, Junliang Guo, Xu Tan, Siliang Tang, Rui Wang, and Jiang Bian. A study on relu and softmax in transformer. *arXiv preprint arXiv:2302.06461*, 2023.
 - Ingo Steinwart and Andreas Christmann. Fast learning from non-iid observations. *Advances in neural information processing systems*, 22, 2009.
 - Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
 - Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pages 35151–35174. PMLR, 2023.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Mi Wen, ZheHui Chen, Yun Xiong, and YiChuan Zhang. Lgat: A novel model for multivariate time series anomaly detection with improved anomaly transformer and learning graph structures. *Neurocomputing*, 617:129024, 2025.
- Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Unified training of universal time series forecasting transformers. *arXiv preprint arXiv:2402.02592*, 2024.
- Mitchell Wortsman, Jaehoon Lee, Justin Gilmer, and Simon Kornblith. Replacing softmax with relu in vision transformers. *arXiv preprint arXiv:2309.08586*, 2023.
- Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.
- Biao Zhang, Ivan Titov, and Rico Sennrich. Sparse attention with linear units. *arXiv preprint* arXiv:2104.07012, 2021.
- Wenxin Zhang and Cuicui Luo. Decomposition-based multi-scale transformer framework for time series anomaly detection. *Neural Networks*, 187:107399, 2025.
- Yunhao Zhang and Junchi Yan. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=vSVLM2j9eie.
- Ingvar Ziemann and Stephen Tu. Learning with little mixing. *Advances in Neural Information Processing Systems*, 35:4626–4637, 2022.

648			SUPPLEMENTARY MATERIAL	
649 650		TD-1-1 6 NI - 4 - 4*	_	12
651	A	Table of Notation	S	13
652	В	Limitations and I	mpact	14
653 654	C	Related Works		14
655	D	Additional Theorem	atical Packground	15
656	ע	Auditional Theor	etical Background	13
657	\mathbf{E}	Proofs		17
658			nma 3.2	17
659			orem 3.6	17
660			Lipschitzness of Any-Variate Transformers	22
661			orem 4.4	25
662		•	Corollary 4.6	30 30
663		E.0 Additional L	etans	30
664	F	Experimental De	tails	31
665				. 31
666		F.2 Model Archi	tecture	. 31
667			ta Generation	
668				. 31
669			xperiments	32
670		F.6 Evaluation o	n Real-World Datasets	32
671				
672	A	TABLE OF NO	OTATIONS	
673				
674			Table 1: Mathematical Notations and Symbols	
675				
676		Symbol	Description	
677		$oldsymbol{x}_i$	The i -th component of vector $oldsymbol{x}$	
678		$\langle oldsymbol{a}, oldsymbol{b} angle$	Inner product for vectors $oldsymbol{a}, oldsymbol{b} \in \mathbb{R}^d$	
679		[I]	Index set $\{1, \dots, I\}$, where $I \in \mathbb{N}^+$	
680		•[Spectral norm, equivalent to the l_2 -norm when applied to a vector	
681		$\left \left \cdot\right \right _{2,\infty}$	The largest L2 norm of column vectors of a matrix	
682		$oldsymbol{A_{ij}}$	The element on the i -th row and j -th column of matrix A	
683		$oldsymbol{x}_{i:j}$	The sub-sequence of sequence x from coordinate i to j	
684		\oplus	Concatenation between column vectors $\boldsymbol{v} \oplus \boldsymbol{u} \mapsto (\boldsymbol{v}^\top, \boldsymbol{u}^\top)^\top$	
685		$[\boldsymbol{u};\boldsymbol{v}]$	Concatenation between two row vectors	
686		$\overline{}$	Length of a transformer input sequence	
687		$\overset{r}{T}$	Number of time steps of a time series	
688		$\stackrel{ au}{M}$	Number of attention heads.	
689		q	Lag of an AR process.	
690		$\overset{1}{d}$	The number of covariates in an AR process	
691		\overline{v}	Vector (bold lower)	
692		$oldsymbol{a}$	Matrix (bold upper)	
693		$\overset{\mathbf{A}}{\mathcal{X}}$	random variable (calligraphic)	
694		X	element from a domain set	
695		$\mathfrak{D}_{\mathcal{X}}$	Domain of random variable X	
		$oldsymbol{e}_i$	one-hot vector with its <i>i</i> -th entry as 1	
696		· · · · · · · · · · · · · · · · · · ·		
697		$P_{\mathcal{X}}$	Probability distribution of \mathcal{X} The probability $P[x = x \mid x = y]$	
698		$P_{\boldsymbol{z}\mid\boldsymbol{w}}(z\mid w)$	The probability $P[z = z \mid w = w]$	

B LIMITATIONS AND IMPACT

Limitations. One limitation in our analysis is that we consider ReLU instead of softmax in attention mechanisms. While the same approach also is in theoretical (5; 22; 15) and empirical works (40; 42; 32), one might obtain a different approximation bound comparing to Theorem 3.6. However, in our generalization analysis, the difference is small as softmax does not affect the model complexity too much. Another aspect is that we mainly focus on AR processes. While in the appendix, we do show the approximation result for non-linear AR processes generated by a ReLU network, to achieve universal forecasting, a more general assumption on data is required. Another limitation is we consider MSE loss instead of NLL loss used in (39) for pretraining.

Impact Statement. This paper studies the theoretical aspect of transformers as time series foundation models. No negative societal impacts that the authors feel should be specifically highlighted here.

C RELATED WORKS

Transformers in Time Series Forecasting. The recent progress in foundation models (34; 6; 28) has begun to reshape the field of time series forecasting, a critical task of predicting the future based on history (14). However, there are two major challenges in building a time series foundation model: (a) the model must be able to handle an arbitrary number of covariates, and (b) the model must generalize to unseen time series domains. To circumvent (a), several studies simplify the task by considering only univariate time series (4; 29; 8). (8) propose a decoder-only transformer pretrained on both real and synthetic datasets. (29) incorporate lag features and the Llama architecture to pretrain a large univariate time series foundation model. (4) leverage the power of large language models (LLMs) by using pretrained LLMs backbones.

Recently, (39) proposed MOIRAI, the first time series foundation model capable of handling an arbitrary number of covariates. It addresses (a) by concatenating all covariates into a uni-dimensional sequence, ensuring a consistent input dimension across datasets. It addresses (b) by pretraining on a large collection of time series datasets (12; 3; 41; 19) spanning domains such as weather, traffic, electricity, and industry. MOIRAI not only generalizes across a wide range of domains, but its *zero-shot* performance also surpasses several strong supervised learning baselines (23; 27; 44). However, the machine learning community has yet to provide a suitable explanation for MOIRAI's impressive performance. Therefore, this paper is the first to offer theoretical guarantees for MOIRAI as a time series foundation model.

In-Context Learning. In-context learning (ICL) is an emerging capability of large foundation models, enabling them to learn diverse and unseen tasks from given examples. (6) first provide empirical evidence of ICL in large language models (LLMs); by presenting several examples of (x, y) pairs, GPT-3 effectively infers the relationship between x and y. (10) then conduct quantitative experiments on simple function classes, such as linear regression. Their results demonstrate that large foundation models can learn the parameters of these function classes. Subsequently, several theoretical studies (5; 36; 1; 2) have proven that different types of transformers can implement algorithms such as gradient descent. This discovery provides a theoretical foundation for the empirical findings in (10).

The closest studies to this paper are (26; 31). However, (31) examines ICL in the context of next-token prediction using a linear transformer. While their theoretical results relate to in-context learning on sequential data, they are insufficient to explain transformers' success in time series forecasting. (26) explores another case where the data is modeled as a Markov chain generated by a transition matrix. They demonstrate the existence of induction heads that enable transformers to perform next-token prediction. However, their scenario does not align with multivariate time series, which is where our main contribution lies.

The biggest different between our work to existing ICL studies (5; 26; 36; 24) is we do not consider structured inputs as in practice, time series data is considered a long sequence instead of many feature-label pairs. Next, in our case study, we show that MOIRAI can perform AR regression on different dimensions of feature, while ICL studies mostly focused on a fixed dimension of features. Last, our pretraining analysis considers the case where data is non IID, which is critical in time series data due to the dependent nature. Therefore, this paper indeed provides new insights on using transformers for time series learning.

Learning Theory on Time Series Data. Efforts in time series learning theory have developed frameworks for non-i.i.d. data using mixing or ergodic assumptions (7; 45; 25; 33). Yet these results are largely agnostic to model architecture and thus offer limited insight into the specific advantages of transformers. In short, a clear theoretical understanding of how transformer architectures contributes to time series forecasting remains elusive.

ADDITIONAL THEORETICAL BACKGROUND D

Here, we include several technical lemme that are intensively used throughout our paper. The Lipschitzness of an MLP layer is obtained in (5, Lemma J.1), which we restate it below

Lemma D.1 ((5)). For a single MLP layer, $\theta_2 = (W_1, W_2)$, we introduce its norm

$$|||\boldsymbol{\theta}_2||| = ||\boldsymbol{W}_1||_{op} + ||\boldsymbol{W}_2||_{op}.$$

For any fixed hidden dimension D', we consider

$$\Theta_{2,B} := \{ \boldsymbol{\theta}_2 : |||\boldsymbol{\theta}_2||| \le B \}.$$

Then for $H \in \mathcal{H}_R$, $\theta_2 \in \Theta_{2,B}$, the function $(\theta_2, H) \mapsto MLP_{\theta_2}$ is (BR)-Lipschitz w.r.t. θ_2 and $(1+B^2)$ -Lipschitz w.r.t. \mathbf{H} .

The following lemma shows any-variate attention is capable of performing variate-wise operation on arbitrary number of covariates under any-variate encoding.

Lemma D.2 (Group-Wise Operation via Any-Variate Attention). Let $||\mathbf{H}||_{2,p} := (\sum_{i=1}^{N} ||\mathbf{h}_i||_2^p)^{1/p}$ denote the column-wise (2,p)-norm of \mathbf{H} . For any input matrix $\mathbf{H} = (\mathbf{h}_1, \cdots, \mathbf{h}_T)$ such that $||\mathbf{H}||_{2,\infty} \leq \mathbb{R}$, suppose $\psi(\cdot) : \mathbb{R}^{D \times T} \to \mathbb{R}^{D \times T}$ is a sequence-to-sequence function implemented by a single layer standard transformer (TF_{θ}^{\dagger}) such that

$$TF_{\boldsymbol{\theta}}^{\dagger}(\boldsymbol{H}) \coloneqq \psi(\boldsymbol{H}).$$

Then there exists a single layer MOIRAI transformer $TF_{\theta}(\cdot)$ such that for any input

$$\boldsymbol{H}^{\star} = [\boldsymbol{H}_1 \quad \boldsymbol{H}_2 \quad \cdots \quad \boldsymbol{H}_K],$$

where $\mathbf{H}_k \in \mathbb{R}^{D \times T}$. $TF_{\boldsymbol{\theta}}(\cdot)$ performs

$$TF_{\theta}(\mathbf{H}^{\star}) = [\psi(\mathbf{H}_1) \quad \psi(\mathbf{H}_2) \quad \cdots \quad \psi(\mathbf{H}_K)].$$

Proof of Lemma D.2. We start by showing the case of a single-head, single-layer standard transformer. Let

$$\begin{split} \mathsf{MLP}_{\boldsymbol{\theta}_2} \circ \mathsf{Attn}_{\boldsymbol{\theta}}^{\dagger}(\boldsymbol{H}) &= \mathsf{MLP}_{\boldsymbol{\theta}_2} \circ \boldsymbol{V} \boldsymbol{H} \sigma(\langle \boldsymbol{Q}\boldsymbol{H}, \boldsymbol{K}\boldsymbol{H} \rangle) = \boldsymbol{V} \boldsymbol{H} \boldsymbol{A}_{\boldsymbol{H}}, \\ \text{where } \boldsymbol{A}_{\boldsymbol{H}} &= \sigma(\langle \boldsymbol{Q}\boldsymbol{H}, \boldsymbol{K}\boldsymbol{H} \rangle). \end{split}$$

Let $\psi_1(H) := VHA_H$, to apply group-wise operation of $\psi_1(\cdot)$ on some input such that

$$\psi_1(\mathbf{H}^{\star}) = [\psi_1(\mathbf{H}_1) \quad \psi_1(\mathbf{H}_2) \quad \cdots \quad \psi_1(\mathbf{H}_K)].$$

Let $\mathbf{0} \in \mathbb{R}^{T \times T}$ be a zero matrix, and $\mathbf{1} \in \mathbb{R}^{T \times T}$ be a 1s matrix, for for any input $||\mathbf{H}^{\star}||_{2,\infty} \leq \mathbb{R}$, one can find some $u^2 < 0$ to decompose $\psi_1(\cdot)$ into the following form.

$$\begin{split} \psi_{1}(\boldsymbol{H}^{\star}) &= \boldsymbol{V} \left[\boldsymbol{H}_{1}\boldsymbol{A}_{\boldsymbol{H}_{1}} \quad \boldsymbol{H}_{2}\boldsymbol{A}_{\boldsymbol{H}_{2}} \quad \cdots \quad \boldsymbol{H}_{K}\boldsymbol{A}_{\boldsymbol{H}_{K}} \right] \\ &= \boldsymbol{V}\boldsymbol{H}^{\star} \begin{bmatrix} \boldsymbol{A}_{\boldsymbol{H}_{1}} & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{A}_{\boldsymbol{H}_{2}} & \cdots & \boldsymbol{0} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{A}_{\boldsymbol{H}_{K}} \end{bmatrix} \\ &= \boldsymbol{V}\boldsymbol{H}^{\star} \times \\ &\sigma \begin{pmatrix} \left\langle \boldsymbol{Q}\boldsymbol{H}_{1}, \boldsymbol{K}\boldsymbol{H}_{1} \right\rangle & \left\langle \boldsymbol{Q}\boldsymbol{H}_{1}, \boldsymbol{K}\boldsymbol{H}_{2} \right\rangle & \cdots & \left\langle \boldsymbol{Q}\boldsymbol{H}_{1}, \boldsymbol{K}\boldsymbol{H}_{K} \right\rangle \\ \left\langle \boldsymbol{Q}\boldsymbol{H}_{2}, \boldsymbol{K}\boldsymbol{H}_{1} \right\rangle & \left\langle \boldsymbol{Q}\boldsymbol{H}_{2}, \boldsymbol{K}\boldsymbol{H}_{2} \right\rangle & \cdots & \left\langle \boldsymbol{Q}\boldsymbol{H}_{2}, \boldsymbol{K}\boldsymbol{H}_{K} \right\rangle \\ &\vdots & \vdots & \ddots & \vdots \\ \left\langle \boldsymbol{Q}\boldsymbol{H}_{K}, \boldsymbol{K}\boldsymbol{H}_{1} \right\rangle & \left\langle \boldsymbol{Q}\boldsymbol{H}_{K}, \boldsymbol{K}\boldsymbol{H}_{2} \right\rangle & \cdots & \left\langle \boldsymbol{Q}\boldsymbol{H}_{K}, \boldsymbol{K}\boldsymbol{H}_{K} \right\rangle \end{pmatrix} + \begin{pmatrix} \boldsymbol{0} & u^{2} \cdot \boldsymbol{1} & \cdots & u^{2} \cdot \boldsymbol{1} \\ u^{2} \cdot \boldsymbol{1} & \boldsymbol{0} & \cdots & u^{2} \cdot \boldsymbol{1} \\ \vdots & \vdots & \ddots & \vdots \\ u^{2} \cdot \boldsymbol{1} & u^{2} \cdot \boldsymbol{1} & \cdots & \boldsymbol{0} \end{bmatrix} \end{pmatrix} \end{split}$$

Further, observe that operations in an MLP layer are either left multiplication or element-wise operations, which implies group-wise as well. We then finish the proof by setting $u^1 = 0$.

Theorem D.3 ((37, Section 5.6)). Suppose $\psi: [0, +\infty_{\to}[0, +\infty)$ is a convex, non-decreasing function satisfying $\psi(x+y) \geq \psi(x)\psi(y)$. For any random variable X, we consider the Orlicz norm induced by $\psi: ||X||_{\psi} \coloneqq \inf\{K>0: \mathbb{E}_{\psi}(|X|/K)\} \leq 1$. Suppose that $\{X_{\theta}\}$ is a zero-mean random process indexed by $\theta \in \Theta$ such that $||X_{\theta} - X_{\theta'}|| \leq \rho(\theta, \theta')$ for some metric ρ on Θ . Then the following holds

$$P\left(\sup_{\theta,\theta'\in\Theta}|X_{\theta}-X_{\theta'}|\leq 8(J+t)\right)\leq \frac{1}{\psi(t/D)}, \quad \textit{for all } t\geq 0,$$

where D is the diameter of the metric space (Θ, ρ) , and the generalized Dudley entropy integral J is given by

824 given l

$$J:=\int_0^D \psi^{-1}(N(\delta;\Theta,\rho))d\delta,$$

where $N(\delta; \Theta, \rho)$ is the δ -covering number of (Θ, ρ) .

The next technical lemma is in (5). Let $\mathsf{B}_\infty^k(R) = [-R,R]^k$ denotes the standard ℓ_∞ ball in \mathbb{R}^k with radius R>0.

Definition D.4 (Sufficiently smooth k-variable function). We say a function $g: \mathbb{R}^k \to \mathbb{R}$ is (R, C_ℓ) -smooth if for $s = \lceil (k-1)/2 \rceil + 2$, g is a C^s function on $\mathsf{B}^k_\infty(R)$, and

$$\sup_{\boldsymbol{z} \in \mathsf{B}_{\infty}^k(R)} \left| \left| \nabla^i g(\boldsymbol{z}) \right| \right|_{\infty} = \sup_{\boldsymbol{z} \in \mathsf{B}_{\infty}^k(R)} \sup_{j_1, \cdots, j_i \in [k]} \left| \partial_{x_{j1} \cdots x_{ji}} g(\boldsymbol{x}) \leq L_i$$

for all $i = 0, 1, \dots, s$, with $\max_{0 \le i \le s} L_i R^i \le C_{\ell}$.

Lemma D.5 (Approximating smooth k-variable functions). For any $\varepsilon_{approx} > 0$, $R \ge 1$, $C_{\ell} > 0$, we have the following: Any (R, C_{ℓ}) -smooth function $g : \mathbb{R}^k \mapsto \mathbb{R}$ is $(\varepsilon_{approx}, R, M, C)$ -approximable by sum of relus with $M \le C(k)C_{\ell}^2 \log(1 + C_{\ell}/\varepsilon_{approx}^2)$ and $C \le C(k)C_{\ell}$, where C(k) > 0 is a constant that depends only on k, i.e.,

$$f(\boldsymbol{z}) = \sum_{m=1}^{M} c_m \sigma(\boldsymbol{a}_m^{\top}[\boldsymbol{z};1]) \quad \textit{with } \sum_{m=1}^{M} |c_m| \leq C, \quad \max_{m \in [M]} ||\boldsymbol{a}_m||_1 \leq,$$

such that $\sup_{z \in [-R,R]^k} |f(z) - g(z)| \le \varepsilon_{approx}$.

E PROOFS

E.1 PROOF OF LEMMA 3.2

Here we prove a slightly simpler result with the positional encoding containing only zero vectors and a one-hot vector. One can easily extend the proof by padding the weight matrices.

$$\boldsymbol{H} \coloneqq \begin{bmatrix} x_1 & x_2 & \dots & x_T & 0 \\ \boldsymbol{p}_1 & \boldsymbol{p}_2 & \dots & \boldsymbol{p}_T & \boldsymbol{p}_{T+1} \end{bmatrix} \in \mathbb{R}^{D \times (T+1)}, \quad \boldsymbol{p}_i \coloneqq \begin{bmatrix} \boldsymbol{0}_{d'} \\ \boldsymbol{e}_i \end{bmatrix} \in \mathbb{R}^{d'+T}, \tag{E.1}$$

Lemma E.1 (Lemma 3.2 Restate). Given a sequence of token \mathbf{H} in the form of Equation E.1, there exists a one-layer, q-1 head ReLU attention layer, such that the columns of $Attn_{\theta}(\mathbf{H})$ has the following form:

$$Attn_{\boldsymbol{\theta}_{1}}^{\dagger}(\boldsymbol{H})_{i} := \begin{bmatrix} x_{i} \\ x_{i-1} \\ \vdots \\ x_{i-q} \\ \boldsymbol{p}_{i}' \end{bmatrix}, \quad where \ \boldsymbol{p}_{i}' := \begin{bmatrix} \mathbf{0}_{d'-q} \\ 1 \\ 1\{i < T+1\} \end{bmatrix} \in \mathbb{R}^{d'-q+2}. \tag{E.2}$$

Proof. Consider an input of the following form

$$m{x} = egin{bmatrix} x_1 & x_2 & \cdots & x_T \ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \ e_1 & e_2 & \cdots & e_T \end{bmatrix},$$

where $x_t \in \mathbb{R}^d$, $\mathbf{p}_t \in \mathbb{R}^T$, for all $t = 1, \dots, T$. We construct weights of the m-th head \mathbf{W}_Q^m , \mathbf{W}_Q^m as following,

$$oldsymbol{W}_K^m = egin{bmatrix} \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_1^ op \ \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_2^ op \ dots & dots & dots \ \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_T^ op \end{bmatrix}, \quad oldsymbol{W}_Q^m = egin{bmatrix} \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_{1-m}^ op \ \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_{2-m}^ op \ dots & dots & dots \ \mathbf{0}^ op & \mathbf{0}^ op & \mathbf{e}_{T-m}^ op \end{bmatrix},$$

where we define the negative index as rotational index, i.e., $e_{-1} = e_T$, $e_{-2} = e_{T-1}$. We have

$$egin{aligned} \left(oldsymbol{W}_{K}^{m}\mathbf{X}
ight)^{ op}\left(oldsymbol{W}_{Q}^{m}\mathbf{X}
ight) = egin{bmatrix} egin{bmatrix} egin{matrix} egin{$$

Note that the result of $\sigma\left((\boldsymbol{W}_K^m\mathbf{X})^\top\left(\boldsymbol{W}_Q^m\mathbf{X}\right)\right)$ is a rotation matrix, where right multiplication on \mathbf{X} will rotate the columns of \mathbf{X} . Therefore, we have \boldsymbol{W}_V^m that performs row-wise shifting and the attention matrix $\sigma\left((\boldsymbol{W}_K^m\mathbf{X})^\top\left(\boldsymbol{W}_Q^m\mathbf{X}\right)\right)$ performs column-wise shifting.

E.2 PROOF OF THEOREM 3.6

Autoregressive Linear Regression under Any-Variate Encoding. The ultimate goal of this setup is to perform the following mechanism. Let x be the target variate we wish to predict, z^j be the j-th covariate of x, for $j \in [M]$. We denote the lookback window size as q, and each covariate has length T (T-time steps.). We denote the time encoding as p_i for $i \in [T]$, and the variate encoding as q_j for $j \in [M]$. Finally, our goal is to predict x_T .

$$\begin{bmatrix} x_1^1 & \cdots & x_T^1 & x_1^2 & \cdots & x_T^2 & \cdots & x_1^d & \cdots & x_T^d \\ \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \cdots & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T \\ \boldsymbol{e}_1 & \cdots & \boldsymbol{e}_1 & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_d & \cdots & \boldsymbol{e}_d \end{bmatrix} \mapsto \begin{bmatrix} \mathsf{A}_1(q) & \cdots \\ \mathsf{A}_2(q) & \cdots \\ \vdots & & \\ \mathsf{A}_d(q) & \cdots \\ \vdots & & \\ \end{bmatrix}$$

Here, different colors represent different covariates. The motivation for performing such an operation is to apply the in-context learning property of transformers proved in (5).

Lemma E.2 (Lemma 3.4 Restate). Define the matrix $A_i(q)$ for the *i*-th covariates (x_1^i, \dots, x_T^i) , with order q, such that

$$\mathbf{A}_i(q) \coloneqq \begin{bmatrix} x_1^i & x_2^i & \cdots & x_t^i & x_{t+1}^i & x_{t+2}^i & \cdots \\ x_T^i & x_{T-1}^i & \cdots & x_{t-1}^i & x_t^i & x_{t+1}^i & \cdots \\ x_{T-1}^i & x_{T-2}^i & \cdots & x_{t-2}^i & x_{t-1}^i & x_t^i & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \\ x_{T-q}^i & x_{T-q+1}^i & \cdots & x_{t-q}^i & x_{t-q+1}^i & x_{t-q+2}^i & \cdots \end{bmatrix},$$

where in the j-th column of $A_i(q)$, it contains historical values of x_j^i with lag q.

Given fixed $D, T \in \mathbb{N}^+$, where T > q. For any input matrix \mathbf{H} in the form of Any-Variate Encoding in Equation 3.4, such that $\mathbf{H} \in \mathbb{R}^{D' \times dT'}$, and $D' \leq D$, T' < T. There exists a 1-layer, q head Any-Variate Attention that performs the following operation.

$$\begin{bmatrix} x_1^1 & \cdots & x_T^1 & x_1^2 & \cdots & x_T^2 & \cdots & x_1^d & \cdots & x_T^d \\ \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T & \cdots & \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_T \\ \boldsymbol{e}_1 & \cdots & \boldsymbol{e}_1 & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_2 & \cdots & \boldsymbol{e}_d & \cdots & \boldsymbol{e}_d \end{bmatrix} \mapsto \begin{bmatrix} \mathtt{A}_1(q) & \mathtt{A}_2(q) & \cdots & \mathtt{A}_d(q) \\ \vdots & \vdots & \ddots & \ddots & \ddots \end{bmatrix}$$

Proof. The proof comes as a direct corollary of Lemma D.2 and (5, Proposition A.5). By Lemma 3.2, there exists a single layer standard transformer layer (with W_1, W_2 being 0s) that generates $A_i(q)$ for each univariate (covariate). It then left applying Lemma D.2 for variate-wise operation and applying (5, Proposition A.5) to keep the time indices p_t unchanged.

Corollary E.3. There exists a d_{max} head standard attention layer that performs the following

$$\begin{bmatrix} \mathtt{A}_1(q) & \mathtt{A}_2(q) & \cdots & \mathtt{A}_d(q) \\ \vdots & \vdots & \ddots & \cdots & \ddots \end{bmatrix} \mapsto \begin{bmatrix} \mathtt{A}_1(q) & \cdots \\ \widetilde{\mathtt{A}}_2(q) & \cdots \\ \vdots & & \\ \widetilde{\mathtt{A}}_d(q) & \cdots \\ \vdots & & \ddots & \ddots \end{bmatrix}, \quad \textit{for any } d \leq d_{\max},$$

where $\widetilde{A}_i(q)$ is $A_i(q)$ without the first row.

Proof. Note that this operation in Corollary E.3 is straightforward with Lemma 3.2 and (5, Proposition A.5). As for each $i \in [d]$, $i \neq 1$, the attention layer performs two operations to each element of $A_i(q)$:

```
\begin{cases} iT \text{ columns to the left} & \text{right multiplication} \\ q_{\max} \text{ rows below} & \text{left multiplication} \\ \text{zero out} & \text{if in first row (left multiplication)} \end{cases}
```

Note that one can simply construct weight matrices to perform the above permutations and masking. In total, we need d_{\max} heads to perform such operations for each $A_i(q)$, for any $d \leq d_{\max}$. For

 $q < q_{\rm max}$, the remaining entries will be zero padded. Finally, with at best 2 layers of $d_{\rm max}$ head any-variate attention, we then obtain

$$\widetilde{H}^{(2)} \coloneqq egin{bmatrix} imes_{1}^{A_{1}(q)} & \cdots \ ilde{A}_{2}(q) & \cdots \ ilde{A}_{2}(q) & \cdots \ ilde{P} \ ilde{e} \end{bmatrix} = egin{bmatrix} \cdots & x_{T-1}^{1} & x_{T-1}^{1} & \cdots \ vx_{T-2}^{1} & x_{T-1}^{1} & \cdots \ vx_{T-1}^{1} & x_{T-q}^{1} & \cdots \ vx_{T-1}^{1} & x_{T-1}^{1} & \cdots \ vx_{T-2}^{1} & x_{T-2}^{1} & \cdots \ vx_{T-2}^{1} & x_{T-2}^{1} & \cdots \ vx_{T-1}^{1} & vx_{T-2}^{1} & \cdots \ vx_{T-1}^{1} & vx_{T-1}^{1} & vx_{T-1}^{1} & \cdots \ vx_{T-1}^{1} & vx_{T-1}^{1}$$

where p is the matrix of (p_1, \dots, p_T) , e is the matrix of (e_1, \dots, e_1) .

assume the autoregressive process generates x_t^1 as follows:

Note that x_T^1 in red is the target we wish to predict (masked as 0 initially), and the entries in blue is considered the input feature of our AR model (a linear regression model in this case), and we are able to directly apply several theoretical results in (5) with input $\widetilde{\boldsymbol{H}}^{(2)}$. Specifically, for Theorem 3.6, it follows directly from (5, Theorem 4) by setting $\lambda=0$.

Next, we present several approximation results from (5), which our approximation results follows immediately from. Considering the general form of autoregressive data: $\boldsymbol{x} \in \mathbb{R}^{d \times T} := (\boldsymbol{x}_1, \dots, \boldsymbol{x}_T)$, where $\boldsymbol{x}_t = (x_1^t, \dots, x_t^d) \in \mathbb{R}^d$. Assuming our target (variate of interest) is in dimension 1, we

$$x_t^1 = f(x_{t-a;t-1}^{1:d}) + \epsilon_t,$$
 (E.3)

where $\epsilon_t \sim N(0, \sigma^2)$, $a_i^j \in \mathbb{R}^1$, and f is a function of interest. We then present several results when f varies.

Non-Linear AR. Here we analyze that when the autoregressive process is generated by a 2 layer ReLU network with look back window size q. Suppose the prediction function $\operatorname{pred}(\boldsymbol{x}, \boldsymbol{w}) := \sum_{k=1}^K u_k r(\boldsymbol{v}_k^\top \boldsymbol{x})$ is given by a two-layer neural network, parameterized by $\boldsymbol{w} = [\boldsymbol{w}_k, u_k]_{k \in [K]} \in \mathbb{R}^{K(d+1)}$. Consider the ERM problem:

$$\min_{\boldsymbol{w} \in \mathcal{W}} \widehat{L}_N(\boldsymbol{w}) \coloneqq \frac{1}{2N} \sum_{i=1}^N \ell(\operatorname{pred}(\boldsymbol{x}_i, \boldsymbol{w}), y_i) = \frac{1}{2N} \sum_{i=1}^N \ell\left(\sum_{k=1}^K u_k r(\boldsymbol{v}_k^\top \widetilde{\boldsymbol{x}}_i), x_T^1\right),$$

where W is a bounded domain and $\tilde{x}_i \in \mathbb{R}^{qd}$ is a flatten version of $x_{t-q:t-q} \in \mathbb{R}^{d \times q}$.

Proposition E.4. Fix any B_w , $B_u > 0$, $L \ge 3$, $\nu > 0$, and $\varepsilon > 0$. Suppose that

- 1. Both the activation function r and the loss function ℓ is C^4 -smooth.
- 2. W is a closed domain such that $W \subset \{w = [v_k; u_k]_{k \in [K]} \in \mathbb{R}^{K(d+1)} : ||v_k||_2 \leq B_v, |u_K| \leq B_u \}$, and $Proj_{\mathcal{W}} = MLP_{\theta_2}$ for some MLP layer with hidden dimension D_w and $||\theta_2||_{op} \leq C_w$.

Then there exists a $(L_1 + 2L_2)$ -layer MOIRAI transformer with

$$\max_{\ell \in [L_1+1, 2L_2]} M^{(\ell)} \le \widetilde{O}(\varepsilon^{-2}), \quad \max_{\ell \in [L_1+1, 2L_2]} D^{(\ell)} \le \widetilde{O}(\varepsilon^{-2}) + D_w,$$
$$||\boldsymbol{\theta}||_{op} \le O(1+\eta) + C_w, \quad \sum_{\ell=1}^{L_1} M^{(\ell)} = d_{\max} + q_{\max}.$$

where we hide the constants K, B_x, B_u, B_v, C^4 , satisfies the following

$$||\widehat{\boldsymbol{w}} - \boldsymbol{w}_{GD}^L||_2 \le L_f^{-1} (1 + \eta L_f)^L \varepsilon,$$

where $L_f = \sup_{w \in \mathcal{W}} \left\| \nabla^2 \widehat{L}_N(w) \right\|_2$.

Maximum Likelihood Estimation (Gaussian) via Transformers. The next result shows that MOIRAI transformers are also capable of performing maximum likelihood estimation on any input multi-variate time series. Given a data generated by some $AR_d(q)$ process with parameter $(\boldsymbol{w}_1, \dots, \boldsymbol{w}_q) \subset \mathbb{R}^d$: $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_T) \subset \mathbb{R}^d$, the conditional likelihood $f(\cdot)$ of observing \boldsymbol{x}_t is

$$f(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1}, \cdots, \boldsymbol{x}_{t-q}) \coloneqq \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(\boldsymbol{x}_t - \sum_{i=1}^q \langle \boldsymbol{w}_i, \boldsymbol{x}_{t-i} \rangle)^2}{2\sigma^2}\right).$$

The goal is to estimate the mean vector $(\boldsymbol{w}_1,\cdots,\boldsymbol{w}_q)$ and the variance σ^2 by minimizing the negative log-likelihood loss. Note that with $n\geq d$, the loss is strongly convex. The optimization over the NLL Loss has two steps: estimating the mean vector: $\hat{\boldsymbol{w}}$, and then derive the variance $\hat{\sigma}^2$ with the following closed-form solution:

$$\sigma^2 = \frac{1}{T} \sum_{t=1}^{T} \left(\boldsymbol{x}_t - \sum_{i=1}^{q} \langle \widehat{\boldsymbol{w}}_i, \boldsymbol{x}_{t-i} \rangle \right)^2.$$

Theorem E.5. Given a set of input data generated by some $AR_d(q)$ process: $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\mathbf{Y} \in \mathbb{R}^n$, considering the following negative log-likelihood loss, the goal is to find a set of parameters $\mathbf{w} \in \mathbb{R}^d$, $\sigma^2 \in \mathbb{R}^+$ to minimize the following loss

$$L_{\mathit{NLL}}(oldsymbol{w}, \sigma) \coloneqq rac{n}{2} \log(2\pi\sigma^2) + rac{1}{2\sigma^2} \sum_{t=q+1}^T \left(oldsymbol{x}_t - \sum_{i=1}^q \langle oldsymbol{w}_i, oldsymbol{x}_{t-i}
angle
ight)^2$$

We denote \mathbf{w}^{\star} , σ^{\star} as the ERM satisfying the NLL Loss. There exists a $(L_1 + L_2 + 2)$ -layer MOIRAI Transformer such that its first $L_1 + L_2$ layers follow the same requirement in Theorem 3.6, and the last two layers each has two and one heads, it estimates \mathbf{w} , σ with bounded error:

$$||\widehat{\boldsymbol{w}} - \boldsymbol{w}^{\star}|| \leq \varepsilon,$$

and the estimated variance is bounded by

$$\left|\widehat{\sigma}^2 - {\sigma^{\star}}^2\right| \le 2EB_x\varepsilon + B_x^2\varepsilon = \widetilde{O}(\varepsilon + \varepsilon^2),$$

where $E \leq B(1+B_w)$, and \widetilde{O} hides the values dependent on B_x, B_w .

Proof of Theorem E.5.

$$L_{\text{NLL}}(\boldsymbol{w}, \sigma) \coloneqq \frac{n}{2} \log(2\pi\sigma^2) + \frac{1}{2\sigma^2} \sum_{t=1}^{T} \left(\boldsymbol{x}_t - \sum_{i=1}^{q} \langle \boldsymbol{w}_i, \boldsymbol{x}_{t-i} \rangle \right)^2.$$

Following Theorem 3.6, the first $L_1 + L_2$ layers of MOIRAI obtains \hat{w} such that the $L_1 + L_2 + 1$ -th layer takes the following as input

$$\widetilde{\boldsymbol{h}}_{i}^{(L_{1}+L_{2})} = \left[x_{i}^{1}; \boldsymbol{x}_{i-1:i-q}^{1}; \boldsymbol{x}_{i-1:i-q}^{2}; \cdots; \boldsymbol{x}_{i-1:i-q}^{d}; \boldsymbol{w}^{\star} + \boldsymbol{\varepsilon}; \boldsymbol{0}; 1; t_{i}\right],$$

where $\boldsymbol{w}^* + \boldsymbol{\varepsilon} \in \mathbb{R}^{qd}$ is the flatten mean vectors. For the simplicity of notations, for the *i*-th column, we denote x_i^1 with $\widetilde{\boldsymbol{y}}_i$, and denote $[\boldsymbol{x}_{i-1:i-q}^1; \boldsymbol{x}_{i-1:i-q}^2; \cdots; \boldsymbol{x}_{i-1:i-q}^d]$ as $\widetilde{\boldsymbol{x}}_i \in \mathbb{R}^{qd}$, as they correspond to the label and feature of our AR model, respectively. $\boldsymbol{\varepsilon} \in \mathbb{R}^{dq}$ satisfies

$$||\boldsymbol{\varepsilon}|| \leq \varepsilon \cdot (\eta B_x).$$

Now we start to construct the $(L_1 + L_2 + 1)$ -th layer. One can then construct

$$\begin{aligned} & \boldsymbol{Q}_{1}^{L+1}\boldsymbol{h}_{i}^{L} = [\boldsymbol{0}; \widetilde{\boldsymbol{x}}_{i}; \boldsymbol{0}], \quad \boldsymbol{K}_{1}^{L+1}\boldsymbol{h}_{j}^{L} = [\boldsymbol{0}; \widehat{\boldsymbol{w}}; \boldsymbol{0}], \quad \boldsymbol{V}_{1}^{L+1}\boldsymbol{h}_{k}^{L} = [\boldsymbol{0}; 1; \boldsymbol{0}] \\ & \boldsymbol{Q}_{2}^{L+1}\boldsymbol{h}_{i}^{L} = [\boldsymbol{0}; \widetilde{\boldsymbol{x}}_{i}; \boldsymbol{0}], \quad \boldsymbol{K}_{2}^{L+1}\boldsymbol{h}_{j}^{L} = [\boldsymbol{0}; -\widehat{\boldsymbol{w}}; \boldsymbol{0}], \quad \boldsymbol{V}_{2}^{L+1}\boldsymbol{h}_{k}^{L} = [\boldsymbol{0}; -1; \boldsymbol{0}]. \end{aligned}$$

The above construction gives us

$$\begin{aligned} \boldsymbol{h}_{i}^{L+1} &= \boldsymbol{h}_{i}^{L} + \frac{1}{n} \sum_{j=1}^{n} \sum_{m=1}^{2} \sigma\left(\langle \boldsymbol{Q}_{m}^{L+1} \boldsymbol{h}_{n+1}^{L}, \boldsymbol{K}_{m}^{L+1} \boldsymbol{h}_{j}^{L} \rangle\right) \boldsymbol{V}_{m}^{L+1} \boldsymbol{h}_{j}^{L} \\ &= \left[\widetilde{\boldsymbol{y}}_{i}; \widetilde{\boldsymbol{x}}_{i}; \widehat{\boldsymbol{w}}; \boldsymbol{0}; 1; t_{i}\right] + \left(\sigma(\langle \widehat{\boldsymbol{w}}, \boldsymbol{x}_{i} \rangle) - \sigma(-\langle \widehat{\boldsymbol{w}}, \boldsymbol{x}_{i} \rangle)\right) \cdot \left[\boldsymbol{0}; 1; \boldsymbol{0}\right] \\ &= \left[\widetilde{\boldsymbol{y}}_{i}; \widetilde{\boldsymbol{x}}_{i}; \widehat{\boldsymbol{w}}; \langle \widehat{\boldsymbol{w}}, \widetilde{\boldsymbol{x}}_{i} \rangle, \boldsymbol{0}; 1; t_{i}\right]. \end{aligned}$$

Next, we construct the last layer as

$$\boldsymbol{Q}_{1}^{L+1}\boldsymbol{h}_{i}^{L}=[...;\widetilde{\boldsymbol{y}}_{i}-\langle\widehat{\boldsymbol{w}},\widetilde{\boldsymbol{x}}_{i}\rangle;...],\quad\boldsymbol{K}_{1}^{L+1}\boldsymbol{h}_{j}^{L}=[...;\widetilde{\boldsymbol{y}}_{j}-\langle\widehat{\boldsymbol{w}},\widetilde{\boldsymbol{x}}_{j}\rangle;...],\quad\boldsymbol{V}_{1}^{L+1}\boldsymbol{h}_{k}^{L}=[\boldsymbol{0};1;\boldsymbol{0}]$$

Finally, the result becomes

$$oldsymbol{h}_i = rac{1}{n} [...; \sum_{\mu=1}^n (oldsymbol{y}_\mu - \langle oldsymbol{x}_\mu, \widehat{oldsymbol{w}}
angle)^2; ...] = [...; \widehat{\sigma^2}; ...].$$

Thus, we complete the proof.

E.3 PROOF OF THE LIPSCHITZNESS OF ANY-VARIATE TRANSFORMERS

We first show the Lipschitzness of each component in an Any-Variate Transformer. For any $p \in [1, \infty]$, let $||\boldsymbol{H}||_{2,p} := (\sum_{i=1}^N ||\boldsymbol{h}_i||_2^p)^{1/p}$ denote the column-wise (2,p)-norm of \boldsymbol{H} . For any radius $\mathbb{R} > 0$, we denote $\mathcal{H}_\mathbb{R} := \{\boldsymbol{H} : ||\boldsymbol{H}||_{2,\infty} \le \mathbb{R}\}$ be the ball of radius \mathbb{R} under norm $||\cdot||_{2,\infty}$.

Lemma E.6. For a single Any-Variate attention layer, $\theta_1 = \{(V_m, Q_m, K_m, u_m^1, u_m^2)\}_{m \in [M]}$, we introduce its norm

$$\|m{ heta}_1\| \coloneqq \max_{m \in [M]} \{||m{Q}_m||_{op}, ||m{K}_m||_{op}, |u_m^1|, |u_m^2|\} + \sum_{m=1}^M ||m{V}_m||_{op}$$

For any fixed hidden dimension D', we consider

$$\Theta_{1,B} := \{ \boldsymbol{\theta}_1 : |||\boldsymbol{\theta}_1||| \le B \}.$$

Then for $\mathbf{H} \in \mathcal{H}_{\mathbb{R}}$, $\mathbf{\theta}_1 \in \Theta_{1,B}$, the function $(\mathbf{\theta}_1, \mathbf{H}) \mapsto Attn_{\mathbf{\theta}_1}$ is $(1 + \iota)$ -Lipschitz w.r.t. $\mathbf{\theta}_1$, where $\iota = \max\{B^2\mathbb{R}^2 + T + (T-1)d, B(T-1)d\}$, and $(1 + B^3\mathbb{R}^2)$ -Lipschitz w.r.t. \mathbf{H} .

Proof. Given some $\epsilon > 0$, some set X and a function class \mathcal{F} . If \mathcal{F} is L-Lipschitzness, i.e.,

$$||f(x_1) - f(x_2)|| \le L ||x_1 - x_2||, \text{ for all } f \in \mathcal{F}.$$

Then, the following holds

$$N(\epsilon, \mathcal{F}, ||\cdot||) \le N(\epsilon/L, X, ||\cdot||).$$

Define

$$\Theta_{\mathsf{attn},B} := \{\theta_{\mathsf{attn}} : |\|\theta_{\mathsf{attn}}\|| \le B\}.$$

The output of the Any-Variate Attention $[\widetilde{h}_i]$ is given by

$$\widetilde{h}_i = h_i + \sum_{m=1}^M \frac{1}{N} \sum_{j=1}^N \sigma\left(\langle \mathbf{Q}_m h_i, \mathbf{K}_m h_j \rangle \cdot \mathbf{V}_m h_j + u_m^1 \star \mathbf{U} + u_m^2 \star \overline{\mathbf{U}}\right).$$

We also define $\theta'_{\text{attn}} = \{\mathbf{V}'_m, \mathbf{Q}'_m, \mathbf{K}'_m, u_m^{1\prime}, u_m^{2\prime}\}_{m \in [M]}$. \widetilde{h}'_i as

$$\widetilde{h}'_{i} = h_{i} + \sum_{m=1}^{M} \frac{1}{N} \sum_{j=1}^{N} \sigma \left(\langle \mathbf{Q}'_{m} h_{i}, \mathbf{K}'_{m} h_{j} \rangle \cdot \mathbf{V}'_{m} h_{j} + u_{m}^{1\prime} \star \mathbf{U} + u_{m}^{2\prime} \star \overline{\mathbf{U}} \right).$$

Now we bound $\left|\left|\operatorname{Attn}_{\boldsymbol{\theta}_1}(\boldsymbol{H}) - \operatorname{Attn}_{\boldsymbol{\theta}_1'}(\boldsymbol{H})\right|\right|_{2,\infty} = \max_i \left|\left|\widetilde{\boldsymbol{h}}_i - \widetilde{\boldsymbol{h}}_i'\right|\right|_2$ as follows

$$\left\| \tilde{\boldsymbol{h}}_{i} - \tilde{\boldsymbol{h}}_{i}' \right\|_{2} = \left\| \sum_{m=1}^{M} \frac{1}{N} \left[\sum_{j=1}^{N} \sigma \left(\langle \mathbf{Q}_{m} h_{i}, \mathbf{K}_{m} h_{j} \rangle + u_{m}^{1} \star \mathbf{U} + u_{m}^{2} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m} h_{j} - \sum_{j=1}^{N} \sigma \left(\langle \mathbf{Q}_{m}' h_{i}, \mathbf{K}_{m}' h_{j} \rangle + u_{m}^{1\prime} \star \mathbf{U} + u_{m}^{2\prime} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m}' h_{j} \right] \right\|_{2}$$

$$\leq \sum_{m=1}^{M} \frac{1}{N} \sum_{j=1}^{N} \left\| \sigma \left(\langle \mathbf{Q}_{m} h_{i}, \mathbf{K}_{m} h_{j} \rangle + u_{m}^{1} \star \mathbf{U} + u_{m}^{2} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m}' h_{j} \right\|_{2}$$

$$- \sigma \left(\langle \mathbf{Q}_{m}' h_{i}, \mathbf{K}_{m}' h_{j} \rangle + u_{m}^{1\prime} \star \mathbf{U} + u_{m}^{2\prime} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m}' h_{j} \right\|_{2}$$

$$\leq \sum_{m=1}^{M} \frac{1}{N} \sum_{j=1}^{N} \left\| h_{j} \right\|_{2} \left\| \sigma \left(\langle \mathbf{Q}_{m} h_{i}, \mathbf{K}_{m} h_{j} \rangle + u_{m}^{1} \star \mathbf{U} + u_{m}^{2} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m}' \right\|_{2}$$

$$- \sigma \left(\langle \mathbf{Q}_{m}' h_{i}, \mathbf{K}_{m}' h_{j} \rangle + u_{m}^{1\prime} \star \mathbf{U} + u_{m}^{2\prime} \star \bar{\mathbf{U}} \right) \mathbf{V}_{m}' \right\|_{2}.$$

1189 Let

$$A = \langle \mathbf{Q}_m h_i, \mathbf{K}_m h_j \rangle + u_m^1 \star \mathbf{U} + u_m^2 \star \bar{\mathbf{U}}$$
$$B = \langle \mathbf{Q}'_m h_i, \mathbf{K}'_m h_j \rangle + u_m^{1\prime} \star \mathbf{U} + u_m^{2\prime} \star \bar{\mathbf{U}}.$$

By triangle inequality, we have

$$||\sigma(A)V_m - \sigma(B)V_m'|| \leq ||\sigma(A)||_{\operatorname{op}} ||\boldsymbol{V}_m - \boldsymbol{V}_m'||_{\operatorname{op}} + ||\sigma(A) - \sigma(B)||_{\operatorname{op}} ||V_m'||_{\operatorname{op}} \,.$$

Note that $\sigma(\cdot)$ is 1-Lipschitz, we get

$$\begin{aligned} ||\sigma(A) - \sigma(B)||_{\text{op}} &\leq ||A - B||_{\text{op}} \\ &= \left| \left| \langle \mathbf{Q}_m h_i, \mathbf{K}_m h_j \rangle - \langle \mathbf{Q}_m' h_i, \mathbf{K}_m' h_j \rangle + (u_m^1 - u_m^{1\prime}) \boldsymbol{U} + (u_m^2 - u_m^{2\prime}) \widehat{\boldsymbol{U}} \right| \right|_{\text{op}} \\ &\leq \left| \left| \left| \langle \mathbf{Q}_m h_i, \mathbf{K}_m h_j \rangle - \langle \mathbf{Q}_m' h_i, \mathbf{K}_m' h_j \rangle \right| + \left| \left| \left| \left| (u_m^1 - u_m^{1\prime}) \star \boldsymbol{U} \right| \right| + \left| \left| \left| \left| (u_m^2 - u_m^{2\prime}) \star \bar{\boldsymbol{U}} \right| \right| \right| \right| \right|. \end{aligned}$$

For the first term in the last inequality, we have

$$\begin{split} \langle \mathbf{Q}_{m}h_{i}, \mathbf{K}_{m}h_{j} \rangle - \langle \mathbf{Q}'_{m}h_{i}, \mathbf{K}'_{m}h_{j} \rangle &\leq ||\mathbf{Q}_{m} - \mathbf{Q}'_{m}|| \, ||h_{i}|| \, ||h_{j}|| \, ||\mathbf{K}_{m}|| + ||\mathbf{K}_{m} - \mathbf{K}'_{m}|| \, ||h_{i}|| \, ||h_{j}|| \, ||\mathbf{Q}_{m}|| \\ &= \mathbb{R}^{2}B \left(||\mathbf{Q}_{m} - \mathbf{Q}'_{m}|| + ||\mathbf{K}_{m} - \mathbf{K}'_{m}|| \right). \end{split}$$

Further, we have

$$||(u^1 - u^{1})||$$

$$\left| \left| (u_m^1 - u_m^{1\prime}) \star \boldsymbol{U} \right| \right| \le |u_m^1 - u_m^{1\prime}| \, ||\boldsymbol{U}|| \le T|u_m^1 - u_m^{1\prime}|,$$

where T is the length of each variate (lookback window size).

$$\left|\left|(u_m^2-u_m^{2\prime})\star\bar{\boldsymbol{U}}\right|\right|\leq |u_m^2-u_m^{2\prime}|\left|\left|\bar{\boldsymbol{U}}\right|\right|\leq (T-1)d|u_m^2-u_m^{2\prime}|,$$

where d is the number of variates.

Thus, we have

$$||\sigma(A) - \sigma(B)||_{\text{op}} ||V'_m||_{\text{op}} \le B \left(\mathbb{R}^2 B \left(||Q_m - Q'_m|| + ||K_m - K'_m|| \right) + T \left(|u_m^1 - u_m^{1\prime}| \right) + (T - 1)d \left(|u_m^2 - u_m^{2\prime}| \right) \right)$$

$$\le B \cdot \max \{ \mathbb{R}^2 B, (T - 1)d \} \cdot \left(||Q_m - Q'_m|| + ||K_m - K'_m|| + |u_m^1 - u_m^{1\prime}| + |u_m^2 - u_m^{2\prime}| \right).$$

Next, we bound

$$||\sigma(A)||_{\operatorname{op}} \leq ||A||_{\operatorname{op}} \leq B^2 \mathtt{R}^2 + (T + (T-1)d),$$

due to the fact that

$$||A|| \leq ||\boldsymbol{Q}_m h_i|| \, ||\boldsymbol{K}_m h_j|| \, \left| \left| u_m^1 \boldsymbol{U} \right| \right| \, \left| \left| u_m^2 \bar{\boldsymbol{U}} \right| \right|.$$

Overall, the Any-Variate Attention is $\max\{B^2\mathtt{R}^2+T+(T-1)d,B(T-1)d\}$ -Lipshcitz in $\pmb{\theta}_1$. \square

Proof. We start by considering $H' = [h'_i]$ and

$$\widetilde{\boldsymbol{h}}_{i}' = \boldsymbol{h}_{i}' + \sum_{m=1}^{M} \frac{1}{N} \sum_{j=1}^{N} \sigma \left(\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}', \mathbf{K}_{m} \boldsymbol{h}_{j}' \rangle + u_{m}^{1} \cdot \mathbf{U} + u_{m}^{2} \cdot \bar{\mathbf{U}} \right) \cdot \mathbf{V}_{m} \boldsymbol{h}_{j}'.$$

We then bound $\left|\left|(\widetilde{m{h}}_i'-m{h}_i')-(\widetilde{m{h}}_i-m{h}_i)\right|\right|$ $= \left\| \sum_{m=1}^{M} \frac{1}{N} \sum_{i=1}^{N} \left[\sigma \left(\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}^{\prime}, \mathbf{K}_{m} \boldsymbol{h}_{j}^{\prime} \rangle + u_{m}^{1} \cdot \mathbf{U} + u_{m}^{2} \cdot \bar{\mathbf{U}} \right) \mathbf{V}_{m} \boldsymbol{h}_{j} - \left(\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}^{\prime}, \mathbf{K}_{m} \boldsymbol{h}_{j}^{\prime} \rangle + u_{m}^{1} \cdot \mathbf{U} + u_{m}^{2} \cdot \bar{\mathbf{U}} \right) \mathbf{V}_{m} \boldsymbol{h}_{j}^{\prime} \right] \right\|$ $\leq \sum_{m=1}^{M} \frac{1}{N} \sum_{i=1}^{N} \left| \left| \mathbf{V}_{m} \right| \right|_{\text{op}} \left| \left| \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}^{\prime}, \mathbf{K}_{m} \boldsymbol{h}_{j}^{\prime} \right\rangle + u_{m}^{1} \cdot \mathbf{U} + u_{m}^{2} \cdot \bar{\mathbf{U}} \right) \boldsymbol{h}_{j} - \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}^{\prime}, \mathbf{K}_{m} \boldsymbol{h}_{j}^{\prime} \right\rangle + u_{m}^{1} \cdot \mathbf{U} + u_{m}^{2} \cdot \bar{\mathbf{U}} \right) \boldsymbol{h}_{j}^{\prime} \right| \right|_{2}$ $\leq \sum_{m=1}^{M} \frac{1}{N} \sum_{m=1}^{N} \left| \left| \mathbf{V}_{m} \right| \right|_{\text{op}} \left\{ \left| \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}, \mathbf{K}_{m} \boldsymbol{h}_{j} \right\rangle + u_{m}^{1} \mathbf{U} + u_{m}^{2} \bar{\mathbf{U}} \right) \right| \cdot \left| \left| \boldsymbol{h}_{j} - \boldsymbol{h}_{j}' \right| \right|_{2}$ + $|\sigma\left(\langle \mathbf{Q}_m \mathbf{h}_i, \mathbf{K}_m \mathbf{h}_i \rangle + u_m^1 \mathbf{U} + u_m^2 \bar{\mathbf{U}}\right) - \sigma\left(\langle \mathbf{Q}_m \mathbf{h}_i', \mathbf{K}_m \mathbf{h}_i \rangle + u_m^1 \mathbf{U} + u_m^2 \bar{\mathbf{U}}\right) |\cdot||\mathbf{h}_i'||_{\mathbf{Q}}$

$$+ \left| \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}, \mathbf{K}_{m} \boldsymbol{h}_{j} \right\rangle + u_{m}^{1} \mathbf{U} + u_{m}^{2} \mathbf{U} \right) - \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}', \mathbf{K}_{m} \boldsymbol{h}_{j} \right\rangle + u_{m}^{1} \mathbf{U} + u_{m}^{2} \mathbf{U} \right) \left| \cdot \left| \left| \boldsymbol{h}_{j}' \right| \right|_{2} \right.$$

$$+ \left| \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}, \mathbf{K}_{m} \boldsymbol{h}_{j} \right\rangle + u_{m}^{1} \mathbf{U} + u_{m}^{2} \bar{\mathbf{U}} \right) - \sigma \left(\left\langle \mathbf{Q}_{m} \boldsymbol{h}_{i}', \mathbf{K}_{m} \boldsymbol{h}_{j}' \right\rangle + u_{m}^{1} \mathbf{U} + u_{m}^{2} \bar{\mathbf{U}} \right) \left| \left| \left| \boldsymbol{h}_{j}' \right| \right|_{2} \right. \right\}$$

$$\leq \sum_{m=1}^{M} \frac{1}{N} \sum_{j=1}^{N} \left|\left|\mathbf{V}_{m}\right|\right|_{\mathsf{op}} \cdot 3 \left|\left|\mathbf{Q}_{m}\right|\right|_{\mathsf{op}} \left|\left|\mathbf{K}_{m}\right|\right|_{\mathsf{op}} \mathtt{R}^{2} \left|\left|\boldsymbol{h}_{j}-\boldsymbol{h}_{j}^{\prime}\right|\right|_{2}$$

$$\leq B^3 \mathbb{R}^2 || \boldsymbol{H} - \boldsymbol{H}' ||_{2,\infty}$$
.

 Where the third inequality comes from the fact that ReLU is 1-Lipschitzness, and the fourth and fifth inequality comes from the AM-GM inequality. For more details, refer (5, Section J.2)

Corollary E.7 (Lipschitz Constant of Single Layer Moirai Transformer). For a fixed number of heads M and hidden dimension D', we consider

$$\Theta_{TF,1,B} = \{ \boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \} : M \text{ heads, hidden dimension } D', ||\boldsymbol{\theta}||_{op} \leq B.$$

Then for the function $TF^{\mathbb{R}}$ given by

$$\mathit{TF}^{\mathtt{R}}: (\boldsymbol{\theta}, \boldsymbol{H}) \mapsto \mathit{clip}_{\mathtt{R}}(\mathit{MLP}_{\boldsymbol{\theta}_2}(\mathit{Attn}_{\boldsymbol{\theta}_1}(\boldsymbol{H}))), \quad \boldsymbol{\theta} \in \Theta_{\mathit{TF},1,B}, \boldsymbol{H} \in \mathcal{H}_{\mathtt{R}}.$$

 $TF^{\mathbb{R}}$ is B_{Θ} -Lipschitz w.r.t. $\boldsymbol{\theta}$ and B_H -Lipschitz w.r.t. \boldsymbol{H} , where $B_{\Theta} = (1+B^2)(1+\iota) + B\mathbb{R}(1+B^3\mathbb{R}^2)$ and $B_H = (1+B^2)(1+B^3\mathbb{R}^2)$.

Proposition E.8 (Lipschitz Constant of Moirai Transformer). For a fixed number of heads M and hidden dimension D', we consider

$$\Theta_{TF,L,B} = \{ \boldsymbol{\theta} = (\boldsymbol{\theta}_1^{(1:L)}, \boldsymbol{\theta}_2^{(1:L)}) \} : M^{(\ell)} = M, D^{(\ell)} = D', ||\boldsymbol{\theta}||_{op} \leq B.$$

Then for the function $TF^{\mathbb{R}}$ is $(LB_H^{L-1}B_{\Theta})$ -Lipschitz in $\theta \in \Theta_{TF,L,B}$ for any fixed H.

Proof. For any $\theta=(\theta_1,\theta_2)$, $H\in\mathcal{H}_R$, and $\theta'=(\theta_1',\theta_2')$, we have

$$\begin{split} ||\mathrm{TF}_{\boldsymbol{\theta}}(\boldsymbol{H}) - \mathrm{TF}_{\boldsymbol{\theta}'}(\boldsymbol{H})||_{2,\infty} &\leq \left|\left|\mathrm{MLP}_{\boldsymbol{\theta}_2}(\mathrm{Attn}_{\boldsymbol{\theta}_1}(\boldsymbol{H})) - \mathrm{MLP}_{\boldsymbol{\theta}_2}(\mathrm{Attn}_{\boldsymbol{\theta}'_1}(\boldsymbol{H}))\right|\right|_{2,\infty} + \\ &\left|\left|\mathrm{MLP}_{\boldsymbol{\theta}_2}(\mathrm{Attn}_{\boldsymbol{\theta}'_1}(\boldsymbol{H})) - \mathrm{MLP}_{\boldsymbol{\theta}'_2}(\mathrm{Attn}_{\boldsymbol{\theta}'_1}(\boldsymbol{H}))\right|\right|_{2,\infty} \\ &\leq (1+B^2)\left|\left|\mathrm{Attn}_{\boldsymbol{\theta}'_1}(\boldsymbol{H}) - \mathrm{Attn}_{\boldsymbol{\theta}'_1}(\boldsymbol{H})\right|\right|_{2,\infty} + B\bar{\mathbf{R}} \left|\left|\boldsymbol{\theta}_2 - \boldsymbol{\theta}'_2\right|\right|_{\mathrm{op}} \\ &\leq (1+B^2)(1+\iota)\left|\left|\boldsymbol{\theta}_1 - \boldsymbol{\theta}'_1\right|\right|_{\mathrm{op}} + B\bar{\mathbf{R}} \left|\left|\boldsymbol{\theta}_2 - \boldsymbol{\theta}'_2\right|\right|_{\mathrm{op}} \leq B_{\boldsymbol{\Theta}} \left|\left|\boldsymbol{\theta} - \boldsymbol{\theta}'\right|\right|_{\mathrm{op}}, \end{split}$$

where $\bar{\mathbf{R}} = \mathbf{R} + B^3\mathbf{R}^3$, $\iota = \max\{B^2\mathbf{R}^2 + T + (T-1)d, B(T-1)d\}$. The second inequality comes from the fact $||\mathbf{Atm}_{\boldsymbol{\theta}}(\boldsymbol{H})|| \leq \mathbf{R} + B^3\mathbf{R}^3$.

Further, for $H' \in \mathcal{H}_R$, we have

$$||\operatorname{TF}_{\boldsymbol{\theta}}(\boldsymbol{H}) - \operatorname{TF}_{\boldsymbol{\theta}}(\boldsymbol{H}')||_{2,\infty} \le (1 + B^2) ||\operatorname{Attn}_{\boldsymbol{\theta}_1}(\boldsymbol{H}) - \operatorname{Attn}_{\boldsymbol{\theta}_1}(\boldsymbol{H}')||$$

$$\le (1 + B^2)(1 + B^3 R^2) ||\boldsymbol{H} - \boldsymbol{H}'||_{2,\infty}.$$

For the multi-layer case, one can simply follow (5, Proposition J.1) to conclude the proof.

E.4 PROOF OF THEOREM 4.4

We first describe our proving technique.

- (i) We write out the pretraining data generation and objective.
- (ii) We show that the objective of pretraining on n time series satisfying Dobrushins condition with length T, is equivalent to pretraining on a single time series with length nT under Dobrushin's condition with the same coefficient.
- (iii) We bound the complexity of transformers, and apply learning bounds from (7).

Let π be a meta distribution, and each distribution drawn from $P^{(T)} \sim \pi$ satisfies the Dobrushin's condition with max coefficient as α . We then define the single-path average loss as

$$Y_{\theta, \mathbf{P}^{(T)}} \coloneqq \frac{1}{T} \sum_{t=1}^{T} \ell(\theta, \boldsymbol{z}_t) - \mathbb{E}_{\boldsymbol{z} \sim \mathbf{P}^{(T)}} \left[\ell(\theta, \boldsymbol{z}) \right].$$

Now, we assume our pretraining data is generated by the following

- 1. Sample n distributions from π i.i.d. to get $P_j^{(T)}$, for $j=1,\cdots,n$
- 2. For each distribution $P_j^{(T)}$, we sample $(z_{j,1}, \dots, z_{j,T})$

Assumption E.9. We assume that for each $j \in [n]$, $(z_{j,t})$ has marginals equal to some distribution D for $t = 1, \dots, T$.

We first present several lemma and theorems that will be used later.

Lemma E.10 ((37, Example 5.8)). Given any well-defined norm $||\cdot||'$. Let \mathbb{B} be the \mathbb{R}^d unit-ball in $||\cdot||'$, i.e. $\mathbb{B} = \{\theta \in \mathbb{R}^d \mid ||\theta||' \leq 1\}$, we have

$$\log N(\delta, \mathbb{B}, ||\cdot||') \le d \log \left(1 + \frac{2}{\delta}\right).$$

Theorem E.11 ((7, Theorem 5.3)). Given a function class \mathcal{F} , such that $|f| \leq B$, for all $f \in \mathcal{F}$. Let $P^{(T)}$ be a distribution over some domain $Z^{(T)}$, assuming Assumption E.9 holds and $\alpha_{\log}(P^{(T)}) < 1/2$. Then for all t > 0,

$$P_{\boldsymbol{z} \sim \mathbf{P}^{(T)}} \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{T} \sum_{i=1}^{T} f(z_i) - \mathbb{E}_z[f(z)] \right| > C \left(\mathfrak{G}_{\mathbf{P}^{(T)}}(\mathcal{F}) + \frac{Bt}{\sqrt{T}} \right) \right) \le e^{-t^2/2},$$

for some universal constant whenever $1/2 - \alpha_{\log}(P^{(T)})$ is bounded away from zero.

The following theorem is from (7; 17).

Theorem E.12. Let $P_{\mathbf{z}}^{(T)}$ be a distribution satisfying the Dobrushin's condition with coefficient $\alpha(P_{\mathbf{z}}^{(T)})$. Let $(\mathbf{z}_1, \dots, \mathbf{z}_T) \sim P^{(T)}$, and let $f : \mathbf{Z}^{(T)} \to \mathbb{R}$ be a real-valued function with the following bounded difference property, with parameters $\lambda_1, \dots, \lambda_T \geq 0$:

$$|f(\boldsymbol{z}) - f(\boldsymbol{z}')| \leq \sum_{t=1}^{T} \mathbb{1}_{\boldsymbol{z}_t
eq \boldsymbol{z}_t'} \lambda_t.$$

Then for all t > 0,

$$P(|f(z) - \mathbb{E}[f(z)] \ge t) \le 2 \exp\left(-\frac{(1-\alpha)t^2}{2\sum_t \lambda_t^2}\right)$$

The following corollary directly follows from the above result

Corollary E.13. Following Theorem E.12, let

$$\ell(\boldsymbol{z}) \coloneqq \frac{1}{T} \sum_{t=1}^{T} \ell(\boldsymbol{z}_t),$$

where $0 \le \ell(z_t) \le B$ for all $t = 1, \dots, T$ and all $z \sim P_z$. Then the variance of $\ell(\cdot)$ is bounded by $|\ell(z) - \ell(z')| \le B$.

Then, the following holds

$$P\left(|\ell(\boldsymbol{z}) - \mathtt{E}[\ell(\boldsymbol{z})]| \geq t\right) \leq 2 \exp\left(\frac{-(1-\alpha)t^2}{2\sum_t B^2}\right).$$

Direct Application of Theorem E.11. By Theorem E.11, if Assumption E.9 holds, with probability over $1 - e^{-t^2/2}$, for any $\theta \in \Theta$, $\alpha_{\log}(\mathsf{P}_i^{(T)}) < 1/2$ we have

$$\sup_{\theta \in \Theta} \left| Y_{\theta, \mathbf{P}^{(T)}} \right| \leq C \left[\mathfrak{G}_{\mathbf{P}_j^{(T)}}(\ell(\Theta)) + \frac{2tB_x^2}{\sqrt{T}} \right],$$

where $\ell(\Theta)$ denotes the function class of $\ell(\theta,)$, for all $\theta \in \Theta$, and C > 0 is an universal constant. Note that the above bound presents the naive learning bound for learning a single time series, which is a direct result from (7).

Here we show properties of a time series generated by concatenating n series under Dobrushin's condition with bounded coefficient.

Lemma E.14 (Dobrushin submultiplicativity). *Let* P, Q *be Markov kernels on a measurable space* (X, \mathcal{F}) *with Dobrushin coefficients*

$$\alpha(P) := \sup_{x,x'} \|P(x,\cdot) - P(x',\cdot)\|_{\mathrm{TV}}, \qquad \alpha(Q) := \sup_{x,x'} \|Q(x,\cdot) - Q(x',\cdot)\|_{\mathrm{TV}}.$$

Then $\alpha(PQ) \leq \alpha(P) \alpha(Q)$.

Proof. For any probability measures μ, ν on X,

$$\|\mu PQ - \nu PQ\|_{\mathrm{TV}} = \|(\mu P - \nu P)Q\|_{\mathrm{TV}} \le \alpha(Q) \|\mu P - \nu P\|_{\mathrm{TV}} \le \alpha(Q) \alpha(P) \|\mu - \nu\|_{\mathrm{TV}}.$$
 Taking $\sup_{\mu,\nu}$ over Dirac measures yields $\alpha(PQ) \le \alpha(P)\alpha(Q)$.

Lemma E.15 (Concatenation preserves Dobrushin). Let $\{P_t\}_{t=1}^T$ be kernels with $\alpha(P_t) \leq \kappa_t < 1$. For the T-step kernel $P_{1:T} := P_1 P_2 \cdots P_T$,

$$\alpha(P_{1:T}) \le \prod_{t=1}^{T} \kappa_t.$$

In particular, if $\kappa_t \leq \kappa < 1$ for all t, then $\alpha(P_{1:T}) \leq \kappa^T$.

Proof. Apply the previous lemma repeatedly:

$$\alpha(P_{1:T}) \le \prod_{t=1}^{T} \alpha(P_t) \le \prod_{t=1}^{T} \kappa_t.$$

П

Corollary E.16 (Concatenating n Dobrushin blocks). Let $\{P_t^{(j)}\}_{t=1}^T$ for $j=1,\ldots,n$ be kernels for n blocks, and suppose $\alpha(P_t^{(j)}) \leq \kappa_t^{(j)} < 1$. For the concatenated nT-step chain

$$\underbrace{P_{1:T}^{(1)}}_{block 1}\underbrace{P_{1:T}^{(2)}}_{block 2} \cdots \underbrace{P_{1:T}^{(n)}}_{block n},$$

its per-step coefficients are bounded by $\kappa_* := \max_{i,t} \kappa_t^{(j)} < 1$, hence the entire nT-step kernel has

$$\alpha \left(\prod_{j=1}^{n} \prod_{t=1}^{T} P_t^{(j)} \right) \leq \prod_{j=1}^{n} \prod_{t=1}^{T} \kappa_t^{(j)} \leq \kappa_*^{nT}.$$

If all blocks share a common bound $\kappa < 1$, then the concatenated process is Dobrushin with the same per-step bound κ and nT-step contraction $\leq \kappa^{nT}$.

Proof. A direct result after applying Lemma E.15.

Remark E.17. A "reset" step (next state independent of current) has $\alpha = 0$ and only improves the bounds. The above results tell us that by concatenating n independent length T Dobrushin time series, we end up having a length n Dobrushin time series.

Concatenation with resets. For block $j \in [n]$, let $(P_t^{(j)})_{t=1}^T$ be the Markov kernels and let $\mu^{(j)}$ be its prescribed initial law. Define the *reset* kernel $R^{(j)}(x,\cdot) \equiv \mu^{(j)}(\cdot)$. The concatenated nT-step chain has kernels

$$\underbrace{P_1^{(1)}, \dots, P_T^{(1)}}_{\text{block } 1}, \ R^{(1)}, \ \underbrace{P_1^{(2)}, \dots, P_T^{(2)}}_{\text{block } 2}, \ R^{(2)}, \ \dots, \ \underbrace{P_1^{(n)}, \dots, P_T^{(n)}}_{\text{block } n}.$$

Lemma E.18. For any probability measure μ , the reset kernel $R(x,\cdot) \equiv \mu(\cdot)$ satisfies $\alpha(R) = 0$. Consequently, if $\alpha(P_t^{(j)}) \leq \kappa_t^{(j)} < 1$ for all blocks and steps, then the concatenated chain with resets has

$$\alpha \left(\prod_{j=1}^{n} \left(P_1^{(j)} \cdots P_T^{(j)} R^{(j)} \right) \right) \leq \prod_{j=1}^{n} \prod_{t=1}^{T} \kappa_t^{(j)}.$$

Lemma E.19 (Reduction to a single chain of length nT). Let $\widetilde{\mathsf{P}}^{(nT)}$ be the concatenation with resets. Then: (a) $\alpha_{\log}(\widetilde{P}_s) \leq \bar{\alpha} < \frac{1}{2}$ for all $s \leq nT$; (b) the marginal at every time under $\widetilde{\mathsf{P}}^{(nT)}$ equals D; (c) for every θ ,

$$\mathbb{E}_{\widetilde{\mathbf{p}}(n^T)} \Big[\frac{1}{nT} \sum_{t=1}^{nT} \ell(\boldsymbol{\theta}, \boldsymbol{z}_t) \Big] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{p}_j^{(T)}} \Big[\frac{1}{T} \sum_{t=1}^T \ell(\boldsymbol{\theta}, \boldsymbol{z}_{j,t}) \Big].$$

Proof. (a) Direct result from the submultiplicative property of Dobrushin series; resets satisfy $\alpha_{\log}(R^{(j)}) = 0$, hence $\sup_s \alpha_{\log}(\widetilde{P}_s) \leq \bar{\alpha}$. (b) By $DP_t^{(j)} = D$ and $R^{(j)}$ having output D, induction over time yields marginal D at all positions. (c) Linearity and the blocking t = (j-1)T + s. \square

Lemma E.20. Let

$$X_{\theta} \coloneqq \frac{1}{n} \sum_{j=1}^{n} \left(\frac{1}{T} \sum_{t=1}^{T} \ell(\theta, \boldsymbol{z}_{j,t}) - \mathbb{E}_{\mathsf{P}_{j}^{(T)}} \left[\frac{1}{T} \sum_{t=1}^{T} \ell(\theta, \boldsymbol{z}_{j,t}) \right] \right)$$

and

$$\widetilde{Y}_{\theta} \coloneqq \frac{1}{nT} \sum_{t=1}^{nT} \ell(\theta, \boldsymbol{z}_t) - \mathbb{E}_{\widetilde{\mathsf{P}}(nT)}[\frac{1}{nT} \sum_{t=1}^{nT} \ell(\theta, \boldsymbol{z}_t)].$$

Then $\{X_{\theta}\}_{\theta}$ and $\{\widetilde{Y}_{\theta}\}_{\theta}$ have the same expectation, i.e.

$$\mathbb{E}_{\widetilde{\mathbf{p}}(nT)} \Big[\frac{1}{nT} \sum_{t=1}^{nT} g(z_t) \Big] = \frac{1}{n} \sum_{j=1}^{n} \mathbb{E}_{\mathbf{p}_{j}^{(T)}} \Big[\frac{1}{T} \sum_{t=1}^{T} g(z_{j,t}) \Big].$$

Proof of Lemma E.20. For all $x, x', R(x, \cdot) = R(x', \cdot) = \mu$, hence $\|R(x, \cdot) - R(x', \cdot)\|_{\mathrm{TV}} = 0$, so $\alpha(R) = 0$. Submultiplicativity gives $\alpha(KL) \leq \alpha(K)\alpha(L)$ and $\alpha(R^{(j)}) = 0$. Use linearity for the expectation identity.

Remark E.21 (Equality in law of objectives). Under Lemma E.20, for each fixed θ , X_{θ} and \widetilde{Y}_{θ} have the same mean (0) and the same Lipschitz modulus in the sample coordinates. In particular, it suffices to bound $\sup_{\theta} |\widetilde{Y}_{\theta}|$.

Lemma E.22 (Increment control in θ). For any (z_1, \ldots, z_{nT}) ,

$$\left| \frac{1}{nT} \sum_{t=1}^{nT} f_{\theta}(z_t) - \frac{1}{nT} \sum_{t=1}^{nT} f_{\theta'}(z_t) \right| \le L_{\Theta} \rho(\theta, \theta').$$

Hence $\widetilde{Y}_{\theta} - \widetilde{Y}_{\theta'}$ is subgaussian with variance proxy $\frac{L_{\Theta}^2 \rho(\theta, \theta')^2}{(1 - \kappa_*) nT}$.

Proof. Direct application of He et al. (15, Theorem 2.3).

Lemma E.23 (Suprema coincide under concatenation). Let X_{θ} and \widetilde{Y}_{θ} be

$$X_{\theta} = \frac{1}{n} \sum_{j=1}^{n} \left(\frac{1}{T} \sum_{t=1}^{T} \ell(\theta, z_{j,t}) - \mathbb{E}_{\mathsf{P}_{j}^{(T)}} \left[\frac{1}{T} \sum_{t=1}^{T} \ell(\theta, z_{j,t}) \right] \right),$$

$$\widetilde{Y}_{\theta} = \frac{1}{nT} \sum_{t=1}^{nT} \ell(\theta, z_t) - \mathbb{E}_{\widetilde{\mathsf{P}}^{(nT)}} \left[\frac{1}{nT} \sum_{t=1}^{nT} \ell(\theta, z_t) \right],$$

where (z_1, \ldots, z_{nT}) is the concatenation (with resets) of the blocks $\{(z_{j,1}, \ldots, z_{j,T})\}_{j=1}^n$. Then for every realization of the data and for every θ ,

$$X_{\theta} = \widetilde{Y}_{\theta}, \quad \textit{hence} \quad \sup_{\theta \in \Theta} |X_{\theta}| = \sup_{\theta \in \Theta} |\widetilde{Y}_{\theta}|.$$

Proof. Write

$$X_{\theta} = \frac{1}{nT} \sum_{i=1}^{n} \sum_{t=1}^{T} \ell(\theta, z_{j,t}) - \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{\mathsf{p}_{j}^{(T)}} \left[\frac{1}{T} \sum_{t=1}^{T} \ell(\theta, z_{j,t}) \right].$$

By construction (z_1, \ldots, z_{nT}) is just the stacked sequence, so $\frac{1}{nT} \sum_{t=1}^{nT} \ell(\theta, z_t) = \frac{1}{nT} \sum_{j=1}^{n} \sum_{t=1}^{T} \ell(\theta, z_{j,t})$. Expectation alignment (concatenation with resets) gives

$$\mathbb{E}_{\widetilde{\mathsf{P}}^{(nT)}}\Big[\frac{1}{nT}\sum_{t=1}^{nT}\ell(\theta,z_t)\Big] = \frac{1}{n}\sum_{j=1}^{n}\mathbb{E}_{\mathsf{P}_{j}^{(T)}}\Big[\frac{1}{T}\sum_{t=1}^{T}\ell(\theta,z_{j,t})\Big].$$

Substitute these two equalities into the definition of \widetilde{Y}_{θ} to obtain $X_{\theta} = \widetilde{Y}_{\theta}$ pointwise, hence the equality of suprema.

Proof of Theorem 4.4. Recall from Lemma E.23, it is sufficient to bound $\sup_{\theta} |\widetilde{Y}_{\theta}|$.

Now, to take supremum over \widetilde{Y}_{θ} , we get

$$\sup_{\theta \in \Theta} X_{\theta} = \sup_{\theta \in \Theta} \frac{1}{n} \sum_{j=1}^{n} Y_{\theta, \mathbf{P}_{j}^{(T)}}$$
$$\leq \frac{1}{n} \sum_{j=1}^{n} \sup_{\theta \in \Theta} Y_{\theta, \mathbf{P}_{j}^{(T)}}.$$

To upper bound $\sup |\widetilde{Y}_{\theta}|$, we take a similar approach to (5, Proposition A.4).

Assuming the index set Θ is equipped with a distance metric ρ and diameter D. We assume that for any ball Θ' of radius r in Θ , there exists some constant C_1 such that the covering number admits upper bound

$$\log N(\delta, \Theta', \rho) \le d \log(2Ar/\delta),$$

for all $0 < \delta \le 2r$.

Now we select Θ_0 such that it is a $(D_0/2)$ -covering of Θ . The above assumption guarantees us that we can have a Θ_0 such that $\log |\Theta_0| \leq d \log(2AD/D_0)$. By Corollary E.13, \widetilde{Y}_{θ} is a ${}^{2B_x^2/(1-\alpha)}$ -subgaussian $(\alpha = \alpha(\mathsf{P}^{(nT)}))$. Then, with probability at least $1 - \delta/2$,

$$\sup_{\theta \in \Theta_0} |\widetilde{Y}_{\theta}| \le C \frac{2B_x^2}{(1-\alpha)} \sqrt{d \log(2AD/D_0) + \log(2/\delta)}.$$

Note that the uniform bound for independent subgaussian random variables still applies here as for each θ , we are re-sampling a new chain from a new distribution sampled from π .

Assume that $\Theta_0 = \{\theta_1, \cdots, \theta_n\}$. Now for each $j \in [m]$, we consider Θ_j is the ball centered at θ_j of radius D_0 in (Θ, ρ) . With Theorem D.3, for each process $\{\widetilde{Y}_\theta\}_{\theta \in \Theta_j}$, then

$$\psi = \psi_2, \quad \left| \left| \widetilde{Y}_{\theta} - \widetilde{Y}_{\theta'} \right| \right|_{\psi} \le \frac{B^1}{\sqrt{nT}} \rho(\theta, \theta'),$$

1518 where $\ell(\theta, z) - \ell(\theta', z)$ is a $B^1 \rho(\theta, \theta')$ -subgaussian random variable.

We then get

$$P\left(\sup_{\theta,\theta'\in\Theta_j}|\widetilde{Y}_{\theta}-\widetilde{Y}_{\theta'}|\leq C'B^1D_0\left(\sqrt{\frac{d\log(2A)}{nT}}+t\right)\right)\leq 2\exp(-nTt^2),\quad \text{ for all } t\geq 0.$$

If we further take $t \leq \sqrt{\log(2m/\delta)/Tn}$, then with probability at least $1 - \delta/2$, it holds that for all $j \in [m]$,

$$\sup_{\theta,\theta'\in\Theta_j} |\widetilde{Y}_{\theta} - \widetilde{Y}_{\theta'}| \le C' B^1 D_0 \sqrt{\frac{2d \log(2AD/D_0) + \log(4/\delta)}{nT}}.$$

By chaining, we have

$$|\widetilde{Y}_{\theta}| \leq |\widetilde{Y}_{\theta_i}| + |\widetilde{Y}_{\theta} - \widetilde{Y}_{\theta_i}|.$$

Hence with probability at least $1 - \delta$, it holds that

$$\sup_{\theta \in \Theta} |\widetilde{Y}_{\theta}| \leq \sup_{\theta \in \Theta_0} |\widetilde{Y}_{\theta}| + \sup_{j} \sup_{\theta \in \Theta_j} |\widetilde{Y}_{\theta} - \widetilde{Y}_{\theta_j}| \leq C''(\frac{2B_x^2}{(1-\alpha)} + B^1D_0)\sqrt{\frac{d \log(2AD/D_0) + \log(2/\delta)}{nT}}.$$

Next by taking $D_0 = D/\kappa, \kappa = 1 + B^1 D \frac{(1-\alpha)}{2B_x^2}$, we get

$$\sup_{\theta \in \Theta} |\widetilde{Y}_{\theta}| \leq C''(\frac{2B_x^2}{(1-\alpha)} + B^1 D\kappa) \sqrt{\frac{d\log(2A\kappa) + \log(2/\delta)}{nT}}.$$

Last, we check whether the assumptions we make above hold for our function class ℓ_{Θ} . Below, we slightly abuse our notation by using D as the dimension for weight matrices in TF_{θ} . By Lemma E.10, it holds that

$$\log N(\delta, B_{||\cdot||_{op}}(r), ||\cdot||_{op}) \le L(3MD^2 + DD' + 2)\log(1 + 2r/\delta),$$

where $B_{||\cdot||_{\mathrm{op}}}(r)$ is a ball of radius r under norm $||\cdot||_{\mathrm{op}}$.

We check that

$$||\ell(\theta, \boldsymbol{z}) - \ell(\theta', \boldsymbol{z})|| \le B_x (LB_H^{L-1}B_{\Theta}) ||\theta - \theta'||_{\text{op}},$$

where it is a direct result from Proposition E.8. By plugging all the parameters, we get

$$\sup_{\theta \in \Omega} |\widetilde{Y}_{\theta}| \le C(\frac{B_x^2}{(1-\alpha)}) \sqrt{\frac{L(3MD^2 + DD' + 2)\iota + \log(2/\delta)}{nT}},$$

where $\iota = \log(2 + 2(LB_H^{L-1}B_{\Theta})B\frac{1-\alpha}{B_x})$

By plugging the ERM $\hat{\theta}$, we get

$$L(\widehat{\boldsymbol{\theta}}) \le \inf_{\theta} L(\theta) + 2 \sup_{\theta} |\widetilde{Y}_{\theta}|.$$

Finally, apply Lemma E.15, we have $\alpha \leftarrow \alpha^n$.

E.5 ANALYSIS OF COROLLARY 4.6

Definition E.24 (Markov Random Field (MRF) with pairwise potentials). The random vector $\mathcal{Z} = (\mathcal{Z}_1, \dots, \mathcal{Z}_d)$ over Z^d is an MRF with pariwise potentials if there exist functions $\psi_i : Z \to \mathbb{R}$ and $\varphi_{ij} : Z^2 \to \mathbb{R}$ for $i \neq j \in \{1, \dots, d\}$ such that for all $z \in Z^d$,

$$\mathbb{P}_{z \sim \mathbf{P}^d} \left[\mathcal{Z} = z \right] = \prod_{i=1}^d e^{\psi(\mathcal{Z}_i)} \prod_{1 \leq i < j \leq d} e^{\varphi_{ij}(\mathcal{Z}_i, \mathcal{Z}_j)}$$

The functions ψ_i are called as element-wise potentials and ϕ_{ij} are pairwise potentials.

Definition E.25. Given an MRF \mathcal{Z} with potentials $\{\varphi_i\}$ and $\{\psi_{ij}\}$, we define

$$\beta_{i,j}(\mathcal{Z}) \coloneqq \sup_{\mathcal{Z}_i,\mathcal{Z}_j \in Z} \lvert \varphi_{ij}(\mathcal{Z}_i,\mathcal{Z}_j) \rvert; \quad \beta(\mathcal{Z}) \coloneqq \max_{1 \leq i \leq d} \sum_{j \neq i} \beta_{ij}(\mathbf{P}^d).$$

Lemma E.26. Given an MRF z with pairwise potentials, for any $i \neq j$, $I_{j \to i}(z) \leq \beta_{j,i}(z)$. $I_{j \to i}(\mathcal{Z}) \leq I_{j,i}^{\log}(\mathcal{Z}) \leq \beta_{j,i}(\mathcal{Z})$

Lemma E.26 implies that to satisfy the condition $\alpha^{\log}(\cdot) < 1/2$, it is sufficient to show that $\beta(\cdot) < 1/2$, leading to the following condition.

$$\langle \boldsymbol{w}\boldsymbol{x}_{t},\boldsymbol{x}_{t+1}\rangle < \ln\frac{1}{2} + (\sigma_{\epsilon}^{2}).$$
 (E.4)

Observe that

$$\langle \boldsymbol{w}\boldsymbol{x}_{t}, \boldsymbol{x}_{t+1} \rangle \leq ||\boldsymbol{w}|| \cdot \max_{t} ||\boldsymbol{x}_{t}||$$

= $B_{w}B_{x}$
 $< \ln \frac{1}{2} + (\sigma_{\epsilon}^{2}) \sim 0.3.$

E.6 ADDITIONAL DETAILS

The History Matrix. The matrix form of $A_i(q)$ is presented below

$$\mathbf{A}_{i}(q) \coloneqq \begin{bmatrix} x_{1}^{i} & x_{2}^{i} & \cdots & x_{t}^{i} & x_{t+1}^{i} & x_{t+2}^{i} & \cdots \\ x_{T}^{i} & x_{T-1}^{i} & \cdots & x_{t-1}^{i} & x_{t}^{i} & x_{t+1}^{i} & \cdots \\ x_{T-1}^{i} & x_{T-2}^{i} & \cdots & x_{t-2}^{i} & x_{t-1}^{i} & x_{t}^{i} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \\ x_{T-q}^{i} & x_{T-q+1}^{i} & \cdots & x_{t-q}^{i} & x_{t-q+1}^{i} & x_{t-q+2}^{i} & \cdots \end{bmatrix}$$
(E.5)

F EXPERIMENTAL DETAILS

F.1 ENVIRONMENT

We mostly train our model on NVIDIA-H100 GPUs with 2 cores each with 128GB RAM. 2 GPUs are sufficient for all of our experiments. We use PyTorch 2.1 and our code is based on the open source published by (39). Training and evaluate takes roughly 12 hours for one run.

F.2 MODEL ARCHITECTURE

For most of our experiments, we use MOIRAI-base model. The hyperparameters are listed in Table 2.

Table 2: Hyperparameters

parameter	values
batch size	64
initial learning rate	1e-3
learning rate decay	cosine annealing
hidden dimension	768
MLP dimension	3072
number of heads	12
training steps	20k
max sequence length	512
optimizer	AdamW
beta (β_1, β_2)	(0.9, 0.98)
weight decay	1e-1
warm up steps (linear)	10k

F.3 SYNTHETIC DATA GENERATION

We generate the AR synthetic data similar to Equation (2.1) but use normalization to stabilize the values. The parameters of synthetic data are in Table 3. Consider a sequence of data $\boldsymbol{x} \in \mathbb{R}^{d \times T} \coloneqq (\boldsymbol{x}_1, \dots, \boldsymbol{x}_T)$, where $\boldsymbol{x}_t = (x_t^1, \dots, x_t^d) \in \mathbb{R}^d$. Assuming our target (variate of interest) is in dimension 1, we assume the $\mathsf{AR}_d(q)$ process generates x_t^1 as follows:

$$x_t^1 = \frac{1}{qd} \sum_{i=1}^q \sum_{j=1}^d a_i^j \cdot x_{t-i}^j + \epsilon_t,$$
 (F.1)

where $\epsilon_t \sim N(0,1)$, $a_i^j \sim N(0,1) \in \mathbb{R}$. After recursively generating the time series, we remove its first 50 time steps as burnout. Each AR time series has a number of covariates between 1 to 5. For training data, we sampled 100 different time series, each with 20k time steps. For test data, we randomly generate one time series with time step 5k, and evaluate our model on all time steps. We set $q, d \leq 5$ in our experiments.

Seasonality. We also conduct experiments on datasets with seasonality information. Specifically, we consider monthly information. After generating a multi-variate time series with T time steps $x \in \mathbb{R}^{d \times T}$, we then add the seasonality information. For each time step t, its seasonal information is

$$a \cdot \sin \frac{2\pi T}{f} \in \mathbb{R},$$

where $a \in \mathbb{R}$ is the amplitude, $f \in \mathbb{N}^+$ is the frequency which is 30 for monthly information. The whole seasonal information will be added to the time series.

F.4 BASELINES

Least Squares Regression. Consider MOIRAI taking an input AR sequence $x \in \mathbb{R}^{d \times T}$, to match our theoretical results (Theorem 3.6), we transform x into the following input-label pairs

$$egin{aligned} \widetilde{m{x}}_1 &= ((m{x}_1, \cdots m{x}_q), m{x}_{q+1}) \ \widetilde{m{x}}_2 &= ((m{x}_2, \cdots m{x}_{q+1}), m{x}_{q+2}) \dots \end{aligned}$$

Table 3: Parameter of Synthetic Data

parameter	values
lag size	$\{1, 2, 3, 4, 5\}$
variance	$\operatorname{unif}(0.1, 1)$
length (T)	20k
number of covariates (d)	$\{1, 2, 3, 4, 5\}$
amplitude	unif(0, 1.5)
frequency	30

After fitting least squares on this transformed dataset with T-q samples, it predicts the T+1-th time step with the following input

$$\widetilde{\boldsymbol{x}}_{\text{test}} = (\boldsymbol{x}_{T-q+1}, \cdots \boldsymbol{x}_T)$$
.

For least squares, we use learning rate as 0.1, and perform full gradient descent with 50, 100 iterations.

F.5 ADDITIONAL EXPERIMENTS

Seasonality Data. Here we present the experimental results on training transformers on seasonality data. The data generation is the same as described above. We use the same setup for seasonality data, where our training data comes from time series with $d \in \{1, 2, 3, 4, 5\}$, and $q = \{1, 2, 3, 4, 5\}$. The evaluation results on seasonality data is in Figure 2. We observe that transformers are capable of inferring data with seasonality. Note that transformers are capable of achieving nearly optimal performance, while least squares regression fails, indicating that transformers are capable of fitting a more complicated model than AR on a given time series.

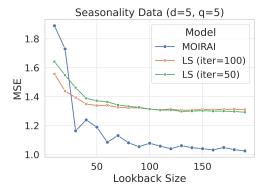


Figure 2: We observe that when least squares regression fails to obtain the optimal error rate for prediction, transformers are capable of having their MSE converge towards 1 as the lookback size increases. This indicates that these models are capable of fitting a more complex model other than linear regression on a given time series.

F.6 EVALUATION ON REAL-WORLD DATASETS.

Here we conduct a similar experiment to Section 5 to real-world datasets. We first pretrain a MOIRAI transformer to simulate AR regression, and then we test it on real-world datasets to see if it can handle real-world dataset in a principled way. Specifically, we use the same checkpoint of transformers in Section 5, and we pick two common datasets: ETTh1 and ETTm1 to test transformers. The results are in Figure 3.

We compare the pretrained transformer with AR models with different lag size varying from 3 to 20. Each point in the figure is the average over 5 runs, the variance is small thus we omit it for better visualization.

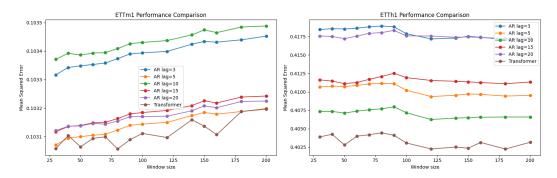


Figure 3: Comparison between pretrained transformer and AR models. While different lag size shows different performance, we are able to observe that transformers can indeed outperform all other AR models. Indicating transformers' predictive power might be more than just simulating the AR algorithm.

Datasets. ETTh1 and ETTm1 are datasets recording the electricity usage through out time. ETTh1 records the usage every hour, and ETTm1 records the usage every 15 minutes. The two datasets are in the same domain, but with different frequency. We test all models on the whole dataset.