

# COSA: CONTEXT-AWARE OUTPUT-SPACE ADAPTER FOR TEST-TIME ADAPTATION IN TIME SERIES FORECASTING

**Jeonghwan Im**

Department of Data Science  
Seoul National University of Science and Technology  
24520018@seoultech.ac.kr

**Hyuk-Yoon Kwon\***

Department of Data Science  
Seoul National University of Science and Technology  
hyukyoon.kwon@seoultech.ac.kr

## ABSTRACT

Deployed time-series forecasters suffer performance degradation under non-stationarity and distribution shifts. Test-time adaptation (TTA) for time-series forecasting differs from vision TTA because ground truth becomes observable shortly after prediction. Existing time-series TTA methods typically employ dual input/output adapters that indirectly modify data distributions, making their effect on the frozen model difficult to analyze. We introduce the *Context-aware Output-Space Adapter* (COSA), a minimal, plug-and-play adapter that directly corrects predictions of a frozen base model. COSA performs residual correction modulated by gating, utilizing the original prediction and a lightweight context vector that summarizes statistics from recently observed ground truth. At test time, only the adapter parameters (linear layer and gating) are updated under a leakage-free protocol, using observed ground truth with an adaptive learning rate schedule for faster adaptation. Across diverse scenarios, COSA demonstrates substantial performance gains versus baselines without TTA (13.91~17.03%) and SOTA TTA methods (10.48~13.05%), with particularly large improvements at long horizons, while adding a reasonable level of parameters and negligible computational overhead. The simplicity of COSA makes it architecture-agnostic and deployment-friendly. Source code: [https://github.com/bigbases/COSA\\_ICLR2026](https://github.com/bigbases/COSA_ICLR2026)

## 1 INTRODUCTION

Time-series forecasting serves as the foundation for critical decision-making across diverse domains, including finance (Chen et al., 2023), supply chain management (Aamer et al., 2020), energy grids (Di Piazza et al., 2021), and predictive maintenance (Makridis et al., 2020). Modern forecasting models, including Transformer-based architectures (Zhou et al., 2021; Liu et al., 2023; 2022), typically achieve high accuracy. However, they suffer performance degradation in real deployment settings due to non-stationarity and distribution shifts (Du et al., 2021; Chen et al., 2024a). Time series exhibit inherent non-stationarity, with changing temporal patterns and statistical characteristics over time, resulting in distributions at training that typically differ from those encountered after deployment.

To address this challenge, various approaches have been proposed, including online learning, continual learning, and domain adaptation. Online and continual learning methods adapt by updating model parameters directly to streaming data (Du et al., 2021; Zhang et al., 2024; Kirkpatrick et al., 2017; Rolnick et al., 2019; Giannini et al., 2023; Pham et al., 2022), but these approaches incur additional computational costs, memory requirements, catastrophic forgetting issues, and plasticity. Furthermore, these methods typically require labeled data or explicit knowledge of task boundaries, making them unsuitable for scenarios where only unlabeled streaming data is available during deployment. Domain adaptation methods learn robust representations by reducing source-target dis-

\*Corresponding author: [hyukyoon.kwon@seoultech.ac.kr](mailto:hyukyoon.kwon@seoultech.ac.kr)

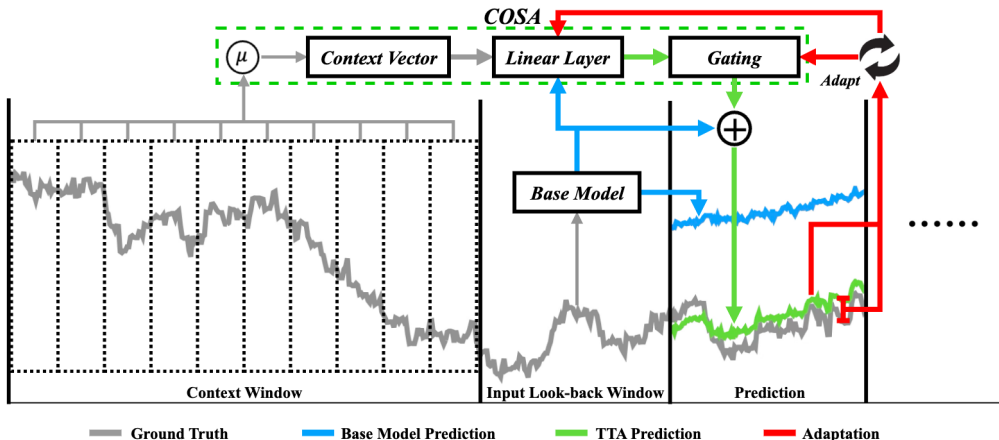


Figure 1: Overview of COSA operation showing the context-aware gated linear adapter architecture with input processing, linear transformation, gating mechanism, and output correction for test-time adaptation.

tribution differences (Wilson et al., 2020; Jin et al., 2022), but they rely on explicit target domain data and boundary definitions.

Test-time adaptation (TTA) offers an alternative approach that adapts to distribution changes by updating only lightweight modules using unlabeled test streams after deployment. TTA methods have evolved mainly in the vision domain through batch normalization coefficient optimization and entropy minimization (Wang et al., 2020), self-supervised/contrastive learning combined with pseudo-labeling (Liang et al., 2021; Chen et al., 2022; Gong et al., 2025), single-sample multi-augmentation-based adaptation (Zhang et al., 2022), and long-term adaptation stabilization (Wang et al., 2022).

Unlike vision tasks, time-series forecasting has unique characteristics that distinguish it from vision tasks: 1) it employs normalization methods different from vision tasks to preserve periodicity and level information, and 2) ground truth becomes sequentially observable after prediction with short delays, enabling the use of direct losses such as Mean Squared Error (MSE).

Time-series forecasting TTA is a recently evolving topic; to the best of our knowledge, only few methods (Kim et al., 2025; Medeiros et al., 2025; Grover & Etemad, 2025) have been proposed. All of them adopted dual-adapter architectures that place calibration modules at both input and output ends of the base model. They map inputs to domains that the base model can handle more easily and restore outputs to the original domain, controlling adaptation intensity through gating. However, these indirect distribution calibration methods involve design complexity and create uncertainty about the impact of input transformations on internal model representations.

To this end, we propose Context-aware Output-Space Adapter (COSA), which offers a direct output-space correction approach that operates with minimal computational overhead. Figure 1 presents the overview of COSA. COSA takes the predictions from a frozen base model and a lightweight context vector, summarizes recent observation statistics as input, computes residuals through linear correction, and controls correction strength using gating. At deployment, we freeze the base forecaster and update only a lightweight output adapter (i.e., linear correction with a learnable gate) under a leakage-free streaming protocol: after each prediction, adaptation uses only previously revealed ground truth, never current or future labels. COSA is architecture-agnostic and demonstrates consistent performance improvements over existing state-of-the-art time-series forecasting TTA methods across various predictors and horizons.

The main contributions of this study are summarized as follows:

1. **Architecture-agnostic output adapter.** Unlike existing time-series TTA methods that adopt dual input-output adapters, COSA consists of a single output adapter. COSA operates independently in the output space, correcting predictions from any base model without changes to training pipelines or internal parameters. COSA also shows compatibility with SOTA normalizers, consistently reducing prediction error.

2. **Context-aware linear residual with gating.** A linear correction uses the base prediction and a lightweight context vector that summarizes statistics of recent observed ground truth, and a learnable gate modulates correction strength.
3. **Consistent accuracy gains.** Across six benchmarks, four horizons, and six baseline architectures, COSA improves test MSE over baselines (13.91~17.03%) and SOTA TTA methods (10.48~13.05%), in particular, with the largest gains at longer horizons.
4. **Fast, efficient TTA.** Adaptive learning rate enables faster convergence of COSA, leading to higher accuracy within a few adaptation steps. Specifically, COSA enables 88.59~90.10% faster inference time against prior SOTA TTA methods.

## 2 RELATED WORK

### 2.1 TIME-SERIES FORECASTING

To handle non-stationarity in time-series forecasting, existing methods typically employ 1) on-line learning, 2) continual learning, and 3) domain adaptation. Representative online learning, D3A (Zhang et al., 2024) narrows source–target gaps through z-score monitoring of loss distributions and Gaussian noise injection, whereas Adarnn (Du et al., 2021) reduces temporal distribution shifts using temporal distribution characterization and distribution matching. In continual learning, cPNN (Giannini et al., 2023) grows temporal columns and transfers knowledge via lateral connections, and FSNet (Pham et al., 2022) separates per-layer adapters for rapid adaptation from associative memory for long-term retention to balance plasticity and stability. For domain adaptation, CoDATS (Wilson et al., 2020) learns domain-invariant features adversarially, and DAF (Jin et al., 2022) shares attention with domain-invariant queries/keys and domain-specific values. These families generally update the base model during training or online operation, differing from TTA, which adapts lightweight modules on unlabeled test streams while keeping the base model frozen.

### 2.2 TEST-TIME ADAPTATION

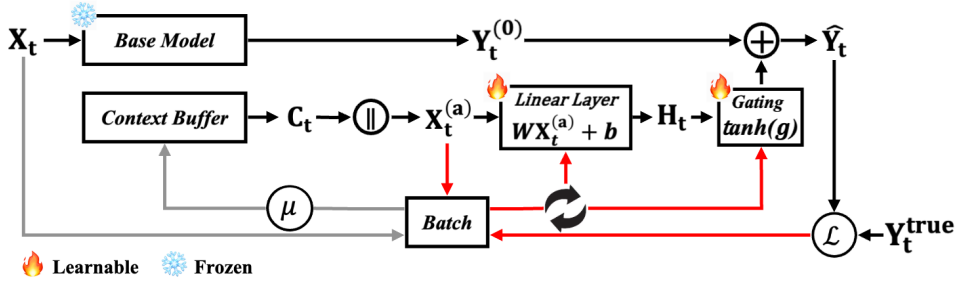
Tent (Wang et al., 2020) optimizes only batch-normalization affine parameters under entropy minimization, and SHOT (Liang et al., 2021) combines information maximization with self-supervised objectives to transfer source hypotheses to the target. AdaContrast (Chen et al., 2022) constructs pseudo-labels via contrastive learning with a dynamic memory bank for gradual adaptation, while MEMO (Zhang et al., 2022) applies multi-augmentation to a single test example to minimize marginal output entropy, updating all weights. CoTTA (Wang et al., 2022) limits error accumulation via weight and stochastic restoration, and ACCUP (Gong et al., 2025) integrates adaptive clustering with pseudo-labeling. However, they are proposed for vision tasks. TTA for time-series forecasting requires different approaches from those for vision tasks due to its own characteristics.

### 2.3 TEST-TIME ADAPTATION FOR TIME-SERIES FORECASTING

Time-series forecasting TTA methods typically employ dual adapters that calibrate distributions at both input and output. TAFAS (Kim et al., 2025) couples a calibration module to map inputs to a model-friendly domain and restores outputs to the original domain. It uses gating to modulate the calibration strength and utilizes Periodicity-Aware Adaptive Scheduling (PAAS) to adjust adaptation frequency using frequency patterns based on inputs. PETSAs (Medeiros et al., 2025) factorizes the calibration module with a low-rank structure and adopts a combined loss for stable adaptation with fewer parameters. DynaTTA (Grover & Etemad, 2025) adjusts the dynamic learning rate, based on local distribution shift, global distribution shift, loss z-score. Existing time-series forecasting TTA methods employ an indirect approach that bidirectionally calibrates distributions at the input and output sides of the base model. They entailed design complexity due to indirect calibration and difficulty in predicting the impact of input transformations on internal representations. In contrast, we aim to utilize a single output-space adapter that directly corrects predictions without requiring input calibration or bidirectional transformation, resulting in a simpler design and more predictable adaptation behavior.

Table 1: Adapter-specific notation. Basic sizes/indices are defined as  $(W, L, K, B; t, i, k)$ .

Symbol	Meaning (shape)
$\mathbf{Y}_t^{(0)}$	Base (frozen) $L$ -step prediction at time $t$ ( $\mathbb{R}^L$ ).
$\mathbf{Y}_t^{\text{true}}$	True $L$ -step target revealed after $t$ ( $\mathbb{R}^L$ ).
$\mathbf{C}_t$	Context vector from revealed batch statistics $[\mu_t - K, \dots, \mu_t - 1]^\top$ ( $\mathbb{R}^K$ ).
$\mathbf{X}_t$	Input look-back window ( $\mathbb{R}^W$ ).
$\mathbf{X}_t^{(a)}$	Adapter input $[\mathbf{Y}_t^{(0)} \parallel \mathbf{C}_t]$ ( $\mathbb{R}^{L+K}$ ).
$\mathbf{H}_t$	Linear residual $\mathbf{W}\mathbf{X}_t^{(a)} + \mathbf{b}$ ( $\mathbb{R}^L$ ).
$\hat{\mathbf{Y}}_t$	Corrected output $\mathbf{Y}_t^{(0)} + \alpha \mathbf{H}_t$ with $\alpha = \tanh(g) \in [-1, 1]$ ( $\mathbb{R}^L$ ).
$\mathbf{W}, \mathbf{b}, g$	Adapter weights ( $\mathbb{R}^L \times (L + K)$ ), bias ( $\mathbb{R}^L$ ), and gate parameter ( $\mathbb{R}$ ).
<i>Operators:</i> concatenation $[\mathbf{a} \parallel \mathbf{b}]$ ; $\ \cdot\ _2$ vector norm; $\ \cdot\ _F$ Frobenius.	

Figure 2: Detailed architecture of COSA illustrating the linear correction layer (weight matrix  $\mathbf{W}$  and bias  $\mathbf{b}$ ), learnable gating parameter ( $g$ ), and context vector ( $\mathbf{C}$ ) integration for output-space correction.

### 3 COSA:CONTEXT-AWARE OUTPUT-SPACE ADAPTER

#### 3.1 NOTATION AND PROBLEM FORMULATION

Table 1 shows the symbols necessary for COSA and their meanings.

This study targets univariate time-series forecasting, following the existing SOTA time-series forecasting TTA methods (Kim et al., 2025; Medeiros et al., 2025; Grover & Etemad, 2025). For multivariate time-series inputs, we decompose them into per-variable univariate forecasting tasks and perform the task iteratively for each variable. At time  $t$ , base model generates  $L$ -step original predictions  $\mathbf{Y}_t^{(0)} \in \mathbb{R}^L$  from input  $\mathbf{X}_t \in \mathbb{R}^W$ , where  $W$  denotes the input look-back window length. COSA generates corrected predictions  $\hat{\mathbf{Y}}_t \in \mathbb{R}^L$  from input  $\mathbf{X}_t^{(a)} \in \mathbb{R}^{L+K}$ , where  $K$  denotes the length of the context vector. After making predictions, the ground truth for that interval becomes sequentially observable following a short delay. Like other TTA approaches, we keep the base model completely frozen and perform only adapter adaptation at test time. Adaptation is performed by collecting the most recent  $B$  prediction, ground truth pairs (batch index  $i \in \{1, \dots, B\}$ ) and context index  $k \in \{1, \dots, K\}$ ).

#### 3.2 OVERALL ARCHITECTURE

Figure 2 illustrates the overall operation of COSA. COSA consists of a single output adapter that directly corrects the predictions. The key components are: 1) a linear layer composed of weight matrix  $\mathbf{W}$  and bias variable  $\mathbf{b}$  that computes correction values  $\mathbf{H}$ , 2) learnable gating  $g$  that controls correction strength, and 3) a context vector  $\mathbf{C}$  that summarizes and stores recent trend information.

We choose a single linear layer for two key reasons: 1) **Efficiency**: Linear operations provide lower latency and higher throughput compared to nonlinear modules, making them suitable for fast adaptation. We confirmed that a single-layer adapter shows 34.95% faster wall-clock time on average than a 2-layer MLP adapter. 2) **Simplicity-Performance balance**: As reported in LTSF-Linear (Zeng

et al., 2023), a linear layer sufficiently performs well in time-series forecasting, despite its simplicity. We also verified that a single linear layer adapter showed 5.71% even better performance on average against a 2-layer MLP adapter. These characteristics make the linear layer beneficial for TTA. Detailed results are provided in Appendix G.3.

The streaming protocol for leakage prevention is as follows (let the last adaptation was performed in  $t-1$ ):

1. **Prediction:** At time  $t$ , base model generates prediction  $\mathbf{Y}_t^{(0)}$  from input  $\mathbf{X}_t$ .
2. **Correction:** Feed  $\mathbf{Y}_t^{(0)}$  and context  $\mathbf{C}_t$  into COSA to generate the corrected prediction  $\hat{\mathbf{Y}}_t$ .
3. **Observation:** After delay  $\Delta \geq 0$ , values of ground truth of the prediction horizon  $\mathbf{Y}_t^{true}$  are sequentially observed.
4. **Adaptation:** Collect the most recent  $B$  prediction, ground truth pairs  $\{\hat{\mathbf{Y}}_{t+i-1}, \mathbf{Y}_{t+i-1}^{true}\}$ , and perform adaptation that updates COSA parameters  $\{\mathbf{W}, \mathbf{b}, \mathbf{g}\}$ .

### 3.3 OUTPUT-SPACE RESIDUAL CORRECTION

For time  $t$ , we concatenate the original prediction of base model and context vector to create the adapter input:

$$\mathbf{X}_t^{(a)} = [\mathbf{Y}_t^{(0)} \parallel \mathbf{C}_t].$$

The residual is computed using a linear transformation:

$$\mathbf{H}_t = \mathbf{W} \mathbf{X}_t^{(a)} + \mathbf{b}.$$

The correction magnitude is controlled through gating to compose the final output:

$$\hat{\mathbf{Y}}_t = \mathbf{Y}_t^{(0)} + \tanh(\mathbf{g}) \mathbf{H}_t.$$

The tanh activation stabilizes the correction magnitude.

### 3.4 CONTEXT CONSTRUCTION

To prevent information leakage, the context summarizes previously observed ground truth information. For time  $t$ , we compute batch-wise aggregation as:

$$\mu_t = \text{agg}\{y_{t-(kB)+i}^{true} : 1 \leq i \leq B\}, \quad 1 \leq k \leq K.$$

where the aggregation function **agg** can use statistics such as mean, median, etc. We construct the context vector by stacking the most recent  $K$  aggregated values:

$$\mathbf{C}_t = [\mu_1, \mu_2, \dots, \mu_K]^\top.$$

This context vector summarizes level/scale changes and gradual drift patterns to help interpret the relative magnitude of the base prediction  $\mathbf{Y}_t^{(0)}$  (reducing to single time-series values when  $B=1$ ).

### 3.5 ADAPTATION OBJECTIVE AND SCHEDULING

Because targets arrive with a delay, we employ a direct objective with weight decay:

$$\mathcal{L} = \sum_{i=1}^B \|\hat{\mathbf{Y}}_{t-i-1} - \mathbf{Y}_{t-i-1}^{true}\|_2^2 + \lambda(\|\mathbf{W}\|_F^2 + \|\mathbf{b}\|_2^2 + \|\mathbf{g}\|_2^2). \quad (1)$$

When  $B$  forecast–target pairs have been enqueued, we run  $S$  gradient steps on the adapter parameters using a cosine–adaptive learning-rate schedule, simply *CALR*. We apply cosine annealing within the  $S$  steps,

$$\eta^{(s+1)} = \eta_{\min} + \frac{1}{2}(\eta^{(s)} - \eta_{\min}) \left(1 + \cos \frac{s\pi}{S}\right). \quad (2)$$

and then adjust  $\eta$  online, based on short-horizon loss trends to encourage fast but stable convergence (decrease  $\eta$  on loss upticks; mildly increase on plateaus). When a new batch arrives, it is always

initialized with the same learning rate, and thereafter the learning rate for the next step within the batch is determined through Equation 2 according to the loss. Early stopping and gradient clipping are also implemented. The threshold values for learning rate adjustment are stability-induced by balancing adaptation speed against stability. Conservative thresholds ensure convergence while aggressive values enable faster response to distribution shifts. Full pseudocode and thresholds are given in Algorithm 1 in Appendix A.

COSA targets TSF-TTA under non-stationary environments in which the distribution of time-series data changes over time. In such environments, the classical notion of convergence toward a fixed optimal point is not well-defined. Instead, stable learning within each adaptation window is critical. CALR guarantees uniformly bounded step-wise updates through the following four mechanisms, which structurally prevent error amplification and thus ensure stability during adaptation.

1. **Upper-bounded learning rate:** The learning rate is constrained by  $\eta \leq \eta_{\max}$ , limiting the maximum magnitude of a single-step update.
2. **Gradient clipping:** At Line 18 of Algorithm 1, the gradient norm is adaptively bounded as  $\|\mathbf{g}_\phi\| \leftarrow \min(\|\mathbf{g}_\phi\|, \max(c, \mathcal{L}))$ .
3. **L2 regularization:** The weight-decay term in Equation 1,  $\lambda(\|\mathbf{W}\|_F^2 + \|\mathbf{b}\|_2^2 + \|g\|_2^2)$ , constrains parameter magnitude.
4. **Bounded gating:** Because  $\alpha = \tanh(g) \in [-1, 1]$ , the correction magnitude is structurally limited.

For every new batch, the learning rate is reinitialized to  $\eta_{\max}$  (Line 2 of Algorithm 1), giving each batch an equal opportunity for adaptation. The learning rate is then adapted according to the batch’s loss behavior. When the loss spikes, we reduce the learning rate as  $\eta \leftarrow \max(0.5\eta, \eta_{\min})$ , temporarily lowering update intensity. When the loss decreases stably, we increase it as  $\eta \leftarrow \min(1.1\eta, \eta_{\max})$ , strengthening adaptation. This enables stable learning even when short-term perturbations or anomalies appear in the input data, allowing rapid recovery.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETTINGS

We evaluate COSA on six benchmark datasets (ETTh1/2, ETTm1/2, Exchange Rate, and Weather) with a fixed look-back window ( $W = 96$ ) and four prediction horizons ( $L \in \{96, 192, 336, 720\}$ ). We used six representative base models spanning different architectures: Transformer-based (iTransformer (Liu et al., 2023), PatchTST (Nie et al., 2023)), linear-based (DLinear (Zeng et al., 2023), OLS (Toner & Darlow, 2024)), and MLP-based (FreTS (Yi et al., 2023), MICN (Wang et al., 2023)). By default, all input time series are treated as variable-wise univariate forecasting tasks, standard normalization is applied, and MSE serves as the performance comparison metric.

We compare *COSA* (our method) with *Baseline* (without TTA), *TAFAS* (Kim et al., 2025), and *PETSA* (Medeiros et al., 2025). All experiments were conducted according to the official benchmark library (Wang et al., 2024)<sup>1</sup>. The train:validation:test ratio is 7:1:2 for all datasets.

Unless otherwise noted, we fix the adapter hyperparameters to  $K=10$  and  $S=3$ , enabled *CALR*. We utilize the average as **agg**. Ablation studies for the variations of **agg** are provided in Appendix G.1.

We report two variants for COSA: COSA-F, which uses a fixed  $B=48$  (half of the look-back), and COSA-P, which sets  $B$  online following the PAAS in TAFAS (Kim et al., 2025). Hyperparameters for comparative methods follow the settings reported in the original papers or official code defaults. In tables, the best score is shown in bold and the second best is underlined.

We utilize Xavier uniform initializer (Glorot & Bengio, 2010) with  $\text{gain} = 0.1$  for parameters of the weight matrix  $\mathbf{W}$ . The bias  $\mathbf{b}$  and gating  $g$  are initialized to 0, and Adam optimizer is utilized.

<sup>1</sup>In the case of *DynaTTA* (Grover & Etemad, 2025), there were reproducibility issues when we used the officially released source code. Therefore, we report the comparison results with them in the Appendix F.5 with the used detailed hyperparameters.

Table 2: Prediction accuracy comparison. Standard deviations less than 0.001 are omitted.

	Transformer-based					Linear-based					MLP-based					
	iTransformer					DLinear					FreTS					
	Baseline	TAFAS	PETSA	COSA-F	COSA-P	Baseline	TAFAS	PETSA	COSA-F	COSA-P	Baseline	TAFAS	PETSA	COSA-F	COSA-P	
ETTh	96	.4507	.4411	.4393	<u>.4368</u>	<b>.4363</b>	.4695	.4618	.4594	<u>.4574</u>	<b>.4482</b>	.4462	.4403	.4387	<u>.4384</u>	<b>.4371</b>
	192	.5078	.4928	.4949	<u>.4961</u>	<b>.4919</b>	.5213	.5117	.5118	<u>.5066</u>	<b>.5050</b>	.5022	.4954	<u>.4942</u>	.4951	<b>.4940</b>
	336	.5658	.5629	.5640	<u>.5651</u>	<b>.5300</b>	.5659	.5604	.5617	<u>.5528</u>	<b>.5456</b>	.5544	.5521	.5527	<b>.5467</b>	<u>.5351</u>
	720	.7038	.6612	.6596	<u>.5958</u>	<b>.5638</b>	.7117	.6820	.6743	<u>.6107</u>	<b>.5896</b>	.7182	.6852	.6846	<u>.6259</u>	<b>.5959</b>
ETTh2	96	.2577	.2549	.2551	<u>.2504</u>	<b>.2493</b>	.2323	.2303	.2306	<u>.2300</u>	<b>.2281</b>	.2384	.2367	<u>.2364</u>	.2367	<b>.2350</b>
	192	.3161	.3010	.3006	<u>.2983</u>	<b>.2947</b>	.2862	.2842	.2876	<u>.2827</u>	<b>.2819</b>	.2866	.2824	.2832	<b>.2816</b>	<u>.2824</u>
	336	.3545	.3352	.3348	<u>.3241</u>	<b>.3339</b>	.3252	.3185	.3184	<b>.3050</b>	<u>.3083</u>	.3317	.3229	.3233	<b>.3031</b>	<u>.3153</u>
	720	.4276	.4023	.4043	<b>.3487</b>	<u>.3591</u>	.4087	.3873	.3853	<b>.3062</b>	<u>.3477</u>	.4119	.3857	.3860	<b>.3169</b>	<u>.3399</u>
ETTh1	96	.3823	.3558	.3570	<b>.3447</b>	<u>.3455</u>	.3715	.3497	.3524	<b>.3456</b>	<u>.3475</u>	.3675	.3582	.3583	<b>.3520</b>	<u>.3525</u>
	192	.4423	.4146	.4142	<b>.4124</b>	<u>.4140</u>	.4438	.4166	.4178	<b>.4113</b>	<u>.4122</u>	.4325	.4212	<u>.4198</u>	<b>.4150</b>	.4212
	336	.5093	.4754	.4751	<b>.4569</b>	<u>.4643</u>	.5183	.4799	.4803	<b>.4753</b>	<u>.4858</u>	.5005	.4827	.4789	<b>.4661</b>	<u>.4775</u>
	720	.6065	.5562	.5553	<b>.4773</b>	<u>.5102</u>	.5929	.5488	.5532	<b>.4774</b>	<u>.4991</u>	.5704	.5486	.5476	<b>.4718</b>	<u>.4982</u>
ETTh2	96	.1647	.1634	.1637	<b>.1627</b>	<u>.1632</u>	.1598	<u>.1584</u>	<u>.1584</u>	<b>.1583</b>	<u>.1586</u>	.1581	.1572	.1572	<b>.1568</b>	<u>.1569</u>
	192	.2209	.2183	.2173	<b>.2171</b>	<u>.2173</u>	.1930	.1913	.1913	<b>.1904</b>	<u>.1905</u>	.1923	.1909	.1908	<b>.1905</b>	<u>.1908</u>
	336	.2727	.2630	.2592	<b>.2435</b>	<u>.2535</u>	.2324	.2289	.2292	<b>.2083</b>	<u>.2242</u>	.2320	.2288	.2289	<b>.2098</b>	<u>.2211</u>
	720	.3451	.3305	.3332	<b>.2477</b>	<u>.2606</u>	.3062	.2968	.2963	<b>.2215</b>	<u>.2316</u>	.3012	.2916	.2926	<b>.2158</b>	<u>.2314</u>
Exchange Rate	96	.0882	.0876	.0885	<b>.0818</b>	<u>.0837</u>	.0913	.0885	.0878	<b>.0812</b>	<u>.0834</u>	.0828	.0799	.0803	<b>.0744</b>	<u>.0766</u>
	192	.1811	.1686	.1740	<b>.1403</b>	<u>.1479</u>	.1827	.1760	.1730	<b>.1459</b>	<u>.1519</u>	.1734	.1665	.1648	<b>.1366</b>	<u>.1499</u>
	336	.3428	.3079	.3097	<b>.2089</b>	<u>.2624</u>	.3277	.2941	.2920	<b>.2039</b>	<u>.2480</u>	.3240	.2930	.2923	<b>.2053</b>	<u>.2461</u>
	720	.8540	.8322	.8004	<b>.3421</b>	<u>.4460</u>	.8873	.8762	.8781	<b>.3494</b>	<u>.4481</u>	.8368	.8273	.8067	<b>.3352</b>	<u>.4458</u>
Weather	96	.1755	.1664	.1674	<b>.1597</b>	<u>.1617</u>	.1954	.1796	.1823	<b>.1773</b>	<u>.1793</u>	.1856	.1759	.1765	<b>.1724</b>	<u>.1737</u>
	192	.2232	.2101	.2128	<b>.2067</b>	<u>.2088</u>	.2403	.2244	.2254	<b>.2216</b>	<u>.2217</u>	.2310	.2165	.2192	<b>.2135</b>	.2189
	336	.2800	.2614	.2665	<b>.2503</b>	<u>.2515</u>	.2918	.2709	.2740	<b>.2567</b>	<u>.2626</u>	.2843	.2653	.2681	<b>.2561</b>	<u>.2587</u>
	720	.3571	.3458	.3459	<b>.2480</b>	<u>.2730</u>	.3643	.3500	.3497	<b>.2581</b>	<u>.2708</u>	.3599	.3490	.3488	<b>.2573</b>	<u>.2692</u>
	PatchTST					OLS					MICN					
	Baseline	TAFAS	PETSA	COSA-F	COSA-P	Baseline	TAFAS	PETSA	COSA-F	COSA-P	Baseline	TAFAS	PETSA	COSA-F	COSA-P	
ETTh	96	.4312	.4262	.4269	<u>.4242</u>	<b>.4238</b>	.4511	.4409	.4391	<u>.4390</u>	<b>.4372</b>	.5103	.4901	.4898	<u>.4693</u>	<b>.4684</b>
	192	.4955	.4865	.4854	<u>.4830</u>	<b>.4805</b>	.5046	.4934	.4937	<u>.4915</u>	<b>.4906</b>	.5954	.5617	.5620	<u>.5372</u>	<b>.5328</b>
	336	.5559	.5478	.5475	<u>.5438</u>	<b>.5320</b>	.5510	.5440	.5465	<u>.5385</u>	<b>.5320</b>	.6615	.6387	.6420	<u>.5950</u>	<b>.5878</b>
	720	.7117	.6860	.6822	<u>.6113</u>	<b>.5822</b>	.6997	.6630	.6431	<u>.5969</u>	<b>.5733</b>	.9233	.8142	.8375	<u>.7001</u>	<b>.6504</b>
ETTh2	96	.2362	.2351	.2362	<u>.2349</u>	<b>.2343</b>	.2306	.2285	.2288	<b>.2232</b>	<u>.2265</u>	.2582	.2551	.2552	<u>.2492</u>	<b>.2485</b>
	192	.2826	.2758	.2773	<u>.2665</u>	<b>.2608</b>	.2839	.2824	.2848	<u>.2796</u>	<b>.2791</b>	.3282	.3179	.3258	<u>.3049</u>	<b>.3017</b>
	336	.3199	.3125	.3132	<u>.2971</u>	<b>.2978</b>	.3258	.3182	.3189	<b>.3003</b>	<u>.3043</u>	.3732	.3482	.3497	<b>.3241</b>	<u>.3310</u>
	720	.4264	.4005	.4012	<b>.3233</b>	<u>.3428</u>	.4162	.3908	.3884	<b>.3177</b>	<u>.3453</u>	.4617	.4474	.4473	<b>.3650</b>	<u>.3885</u>
ETTh1	96	.4024	.3894	.3937	<b>.3625</b>	<u>.3626</u>	.3710	.3506	.3536	<b>.3454</b>	<u>.3475</u>	.4354	.3951	.3951	<u>.3837</u>	<b>.3831</b>
	192	.4512	.4372	.4413	<b>.4250</b>	<u>.4258</u>	.4439	.4160	.4184	<b>.4115</b>	<u>.4119</u>	.4855	.4566	.4574	<b>.4476</b>	.4514
	336	.5081	.4905	.4946	<b>.4568</b>	<u>.4697</u>	.5182	.4787	.4792	<b>.4748</b>	<u>.4749</u>	.5556	.5108	.5082	<b>.4832</b>	<u>.5054</u>
	720	.5629	.5427	.5462	<b>.4681</b>	<u>.4882</u>	.5922	.5478	.5522	<b>.4763</b>	<u>.5007</u>	.6212	.5756	.5778	<b>.5029</b>	<u>.5225</u>
ETTh2	96	.1584	.1581	.1583	<b>.1558</b>	<u>.1562</u>	.1602	.1590	.1589	<b>.1582</b>	<u>.1586</u>	.1710	.1711	.1730	<b>.1702</b>	<u>.1704</u>
	192	.2059	.2036	.2037	<b>.2007</b>	<u>.2022</u>	.1936	.1921	.1919	<b>.1906</b>	<u>.1907</u>	.2121	.2102	.2126	<b>.2102</b>	<u>.2120</u>
	336	.2458	.2451	.2452	<b>.2258</b>	<u>.2352</u>	.2331	.2299	.2302	<b>.2131</b>	<u>.2226</u>	.2530	.2501	.2520	<b>.2337</b>	<u>.2351</u>
	720	.3268	.3268	.3256	<b>.2446</b>	<u>.2645</u>	.3066	.2986	.2971	<b>.2171</b>	<u>.2349</u>	.3327	.3220	.3131	<b>.2477</b>	<u>.2643</u>
Exchange Rate	96	.0867	.0843	.0837	<b>.0765</b>	<u>.0788</u>	.0814	.0792	.0798	<b>.0756</b>	<u>.0773</u>	.1151	.1087	.1146	<b>.0955</b>	<u>.1008</u>
	192	.1877	.1805	.1832	<b>.1464</b>	<u>.1570</u>	.1727	.1658	.1653	<b>.1393</b>	<u>.1457</u>	.2150	.2198	.1999	<b>.1663</b>	<u>.1722</u>
	336	.3389	.3275	.3300	<b>.1983</b>	<u>.2445</u>	.3226	.2877	.2898	<b>.2020</b>	<u>.2323</u>	.3950	.3047	.3100	<b>.2119</b>	<u>.2660</u>
	720	.8648	.8659	.8643	<b>.3543</b>	<u>.4662</u>	.8366	.8138	.8149	<b>.3444</b>	<u>.4541</u>	1.0259	.7191	.7805	<b>.3871</b>	<u>.4815</u>
Weather	96	.1742	.1724	.1743	<b>.1624</b>	<u>.1634</u>	.1957	.1807	.1795	<b>.1772</b>	<u>.1803</u>	.1757	.1853	.1970	<b>.1636</b>	<u>.1651</u>
	192	.2195	.2147	.2167	<b>.2006</b>	<u>.2108</u>	.2404	.2244	.2274	<b>.2223</b>	<u>.2237</u>	.2237	.2161	.2265	<b>.2082</b>	<u>.2120</u>
	336	.2766	.2666	.2701	<b>.2451</b>	<u>.2488</u>	.2921	.2714	.2748	<b>.2551</b>	<u>.2642</u>	.2812	.2746	.2788	<b>.2729</b>	<u>.2737</u>
	720	.3544	.3383	.3442	<b>.2590</b>	<u>.2713</u>	.3644	.3466	.3493	<b>.2579</b>	<u>.2708</u>	.3508	.3573	.3681	<b>.2582</b>	<u>.2855</u>

All experiments were conducted on a machine with an Intel i7-7800X CPU and NVIDIA GeForce RTX 3080 10GB.

## 4.2 MAIN RESULTS

### 4.2.1 COMPARISON WITH SOTA TIME-SERIES TTA METHODS

The proposed COSA achieves best performance in all scenarios, as shown in Table 2<sup>2</sup>. The results reveal two key performance patterns:

1. **Architecture-agnostic benefits:** Consistent improvements across all base models demonstrate that effectiveness of COSA is not dependent on specific model architectures. The average improvement ranges from 10.48% to 13.05%.

<sup>2</sup>All reported results are averaged over 10 runs with different random seeds to ensure statistical reliability.

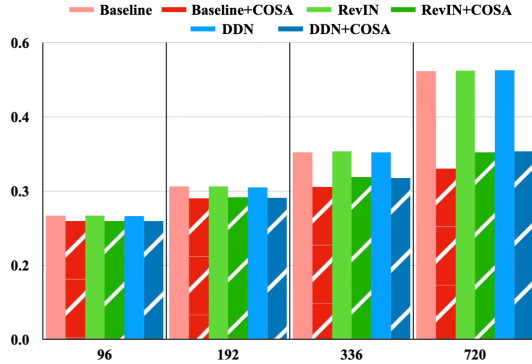


Figure 3: Prediction accuracy comparison with normalization methods.

2. **Effectiveness in long-term forecasting:** The largest performance improvements were observed at the 720 horizon, where COSA-F and COSA-P showed performance improvements of 32.24% and 26.33% compared to baseline, respectively, and 28.21% and 21.96% compared to other methods. This trend suggests that COSA becomes increasingly valuable for longer prediction horizons.

These findings demonstrate that the COSA, which performs residual correction directly in the output space, proves more effective than existing indirect dual-adapter approaches.

#### 4.2.2 COMPARISON WITH NORMALIZATION METHODS

In this section, we analyze whether COSA serves as an alternative or complementary role to existing normalizations, and demonstrate that COSA can work independently of normalizations and is compatible with normalization mechanisms. Figure 3 shows performance comparison of COSA against two representative time-series normalizers, RevIN (Kim et al., 2021) and DDN (Dai et al., 2024). The results support two claims:

1. COSA on its own generally outperforms explicit normalization, as demonstrated by the comparative analysis showing that using COSA with basic normalization achieves the lowest mean MSE across all experimental settings.
2. COSA is compatible with normalization and consistently improves accuracy within the same normalizer settings, with the addition of COSA reducing MSE by approximately 16.8~16.9%.

COSA directly optimizes MSE on revealed targets and uses the context vector to encode recent level/scale, so the linear layer learns scale and level corrections from the error signal itself. Further analysis is provided in Appendix G.7.

#### 4.3 SENSITIVITY AND ABLATIONS

We probe the following four key design choices: 1) adaptation steps  $S$ , 2) context length  $K$ , 3) batch size  $B$ , and 4) adaptive learning rate  $CALR$ .

We evaluated performance by varying each hyperparameter individually while keeping others fixed at the default settings. Figure 4 shows MSE and wall-clock time according to each hyperparameter. Figure 4a shows changes according to the number of iterative learning steps  $S$ . COSA shows a pattern where test MSE decreases as  $S$  increases. However, while wall-clock time also increases with increasing  $S$ , even at the highest  $S = 4$  setting, it showed time levels similar to PETSA.

Figure 4b examines the effect of context length  $K$ , which controls how much past information the adapter uses. Accuracy improves consistently with larger  $K$ , while wall-clock time remains unchanged. Since the adapter input concatenates base model’s prediction with the context vector (dimensionality  $L + K$ ), and  $L$  typically dominates, increasing  $K$  has negligible runtime impact. The context provides incremental but reliable gains by supplementing the level/scale information.

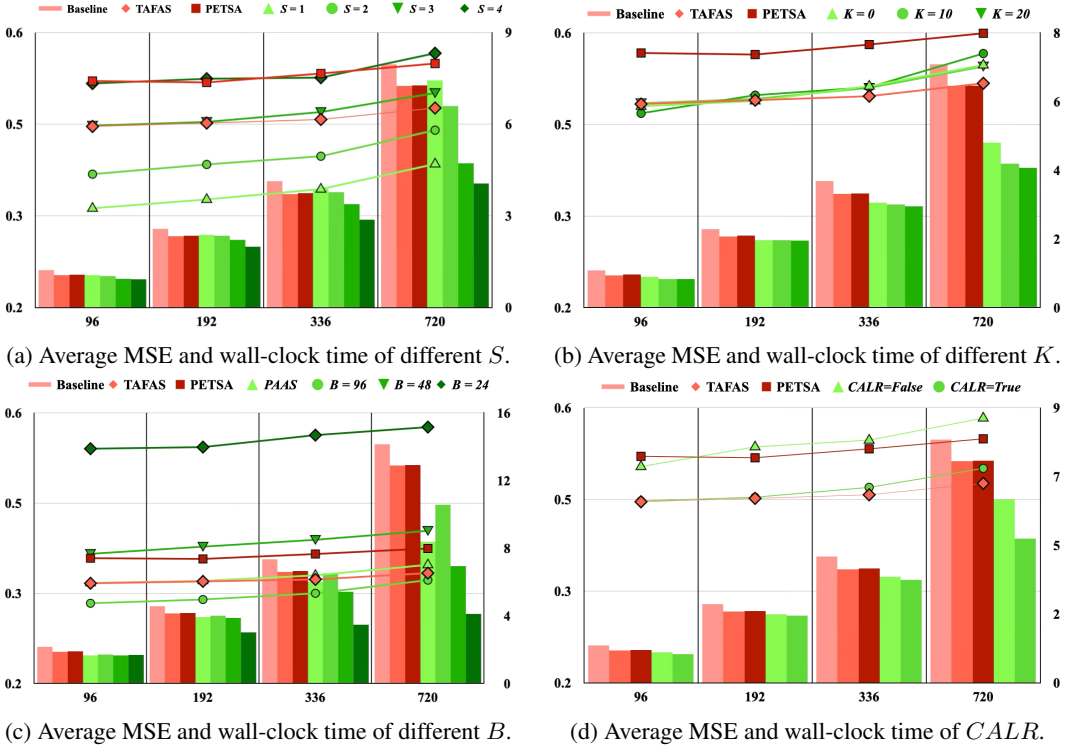


Figure 4: Hyperparameter analysis showing the trade-off between performance and efficiency. Charts display average test MSE (bars, left axis) and wall-clock time (lines, right axis, in seconds) across different parameter settings.

Figure 4c shows performance changes with batch size  $B$ . As described in Section 3.5,  $B$  determines the frequency of adaptation, collecting  $B$  {prediction, ground truth} pairs before each update. Even with  $B = 96$ , COSA outperforms the baseline, with accuracy improving as  $B$  decreases due to more frequent adaptation. This explains the superior performance of COSA-P over COSA-F on ETTh1: the average  $B$  determined by PAAS for ETTh1 is 24.55, while for other datasets the values are over 80. Detailed analysis is provided in Appendix C. However, smaller  $B$  increases wall-clock time due to both more adaptation calls and the computational cost of adaptation steps  $S$ .

Figure 4d shows performance with and without  $CALR$  (Section 3.5).  $CALR$  achieves up to 12.13% accuracy improvement as the window length increases and a 21.34% reduction in wall-clock time. This confirms that aggressive dynamic learning rate scheduling enhances performance within limited adaptation steps  $S$  while enabling early-stopping for computational efficiency.

Results for additional ablations (context aggregation methods and correction layer architecture) are provided in Appendix G.

#### 4.4 COMPUTATIONAL OVERHEAD

Table 3: Computational overhead comparison of adapter methods.

Method	# Params ↓	Peak mem (MB) ↓	Samples/sec ↑	Adaptation time/batch (ms) ↓	Inference time/batch (ms) ↓
TAFAS	1,252,958	<b>17.59</b>	<b>1,413.23 ± 92.28</b>	<b>73.23 ± 8.74</b>	10.96 ± 0.64
PETSA	<b>58,334</b>	36.09	987.91 ± 78.67	88.46 ± 7.08	12.63 ± 0.37
COSA (Ours)	1,211,287	27.07	1080.70 ± 80.66	80.12 ± 5.93	<b>1.25 ± 0.06</b>

We report 1) observed additional parameters, 2) peak memory utilization, 3) throughput (samples per second), 4) wall-clock time per batch, and 5) inference time per batch.

All hyperparameters remain at default values. Table 3 summarizes the average overhead across all datasets, base model, and horizons. COSA shows moderate overhead, falling between TAFAS and PETSA, while achieving the significantly fastest inference time. COSA performs adaptation repeat-

edly for  $S$  steps, meaning that the throughput and adaptation time are dependent on  $S$ . However, the single adapter structure and simplicity of COSA alleviate the overhead and improve the inference time, which is not affected by  $S$ . The computational complexity is  $\mathcal{O}(L \cdot (L + K))$  for the linear transformation plus  $\mathcal{O}(L)$  for the gating operation, resulting in quadratic scaling with respect to prediction horizon  $L$ . Further theoretical calculations are included in Appendix B.

## 5 CONCLUSION

We introduced COSA, an architecture-agnostic TTA module that directly corrects the prediction of base model with a single linear layer guided by short-term context and a stabilizing gate. Across six benchmarks and diverse base models, COSA improves accuracy by 13.91% to 17.03% over baselines and by 10.48% to 13.05% over prior state-of-the-art TTA methods. These gains arise from the synergy of trend-aware context, residual correction, and gated modulation.

While COSA shows strong empirical results, several areas offer room for refinement. The current adaptation relies on full ground truth, though extending to partial observations would enable real-time deployment. Performance varies with batch size  $B$ , and the fixed context length  $K$  may not be optimal for all temporal patterns. Additionally, linear corrections, while effective for many cases, could be enhanced for complex nonlinear shifts.

Future work will explore masked updates for real-time adaptation with partial targets, adaptive selection of  $K$  and  $B$  based on detected periodicity, and hybrid linear/nonlinear adapters for more complex distribution shifts. These extensions will broaden the applicability of the proposed method while maintaining its computational efficiency.

## ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00521452), and by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-15002982).

## REFERENCES

- Ammar Aamer, LuhPutu Eka Yani, and IMade Alan Priyatna. Data analytics in the supply chain management: Review of machine learning applications in demand forecasting. *Operations and Supply Chain Management: An International Journal*, 14(1):1–13, 2020.
- Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 295–305, 2022.
- Mouxiang Chen, Lefei Shen, Han Fu, Zhuo Li, Jianling Sun, and Chenghao Liu. Calibration of time-series forecasting: Detecting and adapting context-driven distribution shift. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 341–352, 2024a.
- Mouxiang Chen, Lefei Shen, Han Fu, Zhuo Li, Jianling Sun, and Chenghao Liu. Calibration of time-series forecasting: Detecting and adapting context-driven distribution shift. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 341–352, 2024b.
- Weisi Chen, Walayat Hussain, Francesco Cauteruccio, and Xu Zhang. Deep learning for financial time series prediction: A state-of-the-art review of standalone and hybrid models. *CMES-Computer Modeling in Engineering and Sciences*, 2023.
- Tao Dai, Beiliang Wu, Peiyuan Liu, Naiqi Li, Xue Yuerong, Shu-Tao Xia, and Zexuan Zhu. Ddn: Dual-domain dynamic normalization for non-stationary time series forecasting. In *Advances in Neural Information Processing Systems*, volume 37, pp. 108490–108517, 2024.
- Annalisa Di Piazza, Maria Carmela Di Piazza, Giuseppe La Tona, and Massimiliano Luna. An artificial neural network-based forecasting model of energy-related time series for electrical grid management. *Mathematics and Computers in Simulation*, 184:294–305, 2021.

- Yuntao Du, Jindong Wang, Wenjie Feng, Sinno Pan, Tao Qin, Renjun Xu, and Chongjun Wang. Adarnn: Adaptive learning and forecasting of time series. In *Proceedings of the 30th ACM international conference on information & knowledge management*, pp. 402–411, 2021.
- Federico Giannini, Giacomo Ziffer, and Emanuele Della Valle. cpnn: Continuous progressive neural networks for evolving streaming time series. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 328–340. Springer, 2023.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.
- 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.
- Peiliang Gong, Mohamed Ragab, Min Wu, Zhenghua Chen, Yongyi Su, Xiaoli Li, and Daoqiang Zhang. Augmented contrastive clustering with uncertainty-aware prototyping for time series test time adaptation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, KDD '25, pp. 390–401, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400712456. doi: 10.1145/3690624.3709239.
- Shivam Grover and Ali Etemad. Shift-aware test time adaptation and benchmarking for time-series forecasting. OpenReview, 2025. URL <https://openreview.net/forum?id=a399SmgWGl>. ICML 2025 Workshop submission (OpenReview).
- Xiaoyong Jin, Youngsuk Park, Danielle Maddix, Hao Wang, and Yuyang Wang. Domain adaptation for time series forecasting via attention sharing. In *International Conference on Machine Learning*, pp. 10280–10297. PMLR, 2022.
- HyunGi Kim, Siwon Kim, Jisoo Mok, and Sungroh Yoon. Battling the non-stationarity in time series forecasting via test-time adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 17868–17876, 2025.
- Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International conference on learning representations*, 2021.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- 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*, pp. 95–104, 2018.
- Jian Liang, Dapeng Hu, Yunbo Wang, Ran He, and Jiashi Feng. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8602–8617, 2021.
- Yong Liu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Non-stationary transformers: Exploring the stationarity in time series forecasting. In *Advances in Neural Information Processing Systems*, volume 35, pp. 9881–9893, 2022.
- 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.
- Georgios Makridis, Dimosthenis Kyriazis, and Stathis Plitsos. Predictive maintenance leveraging machine learning for time-series forecasting in the maritime industry. In *2020 IEEE 23rd international conference on intelligent transportation systems (ITSC)*, pp. 1–8. IEEE, 2020.

- Heitor Rapela Medeiros, Hossein Sharifi-Noghabi, Gabriel L. Oliveira, and Saghar Irandoust. Accurate parameter-efficient test-time adaptation for time series forecasting. In *Second Workshop on Test-Time Adaptation: Putting Updates to the Test! at ICML 2025*, 2025. URL <https://openreview.net/forum?id=uFj4EL4GTB>.
- Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Jbdc0vT0col>.
- Quang Pham, Chenghao Liu, Doyen Sahoo, and Steven CH Hoi. Learning fast and slow for online time series forecasting. *arXiv preprint arXiv:2202.11672*, 2022.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- William Toner and Luke Darlow. An analysis of linear time series forecasting models. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org, 2024.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020.
- Huiqiang Wang, Jian Peng, Feihu Huang, Jince Wang, Junhui Chen, and Yifei Xiao. MICN: Multi-scale local and global context modeling for long-term series forecasting. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=zt53IDUR1U>.
- Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7201–7211, 2022.
- Yuxuan Wang, Haixu Wu, Jiayang Dong, Yong Liu, Mingsheng Long, and Jianmin Wang. Deep time series models: A comprehensive survey and benchmark. *arXiv preprint arXiv:2407.13278*, 2024.
- Garrett Wilson, Janardhan Rao Doppa, and Diane J Cook. Multi-source deep domain adaptation with weak supervision for time-series sensor data. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1768–1778, 2020.
- Kun Yi, Qi Zhang, Wei Fan, Shoujin Wang, Pengyang Wang, Hui He, Ning An, Defu Lian, Longbing Cao, and Zhendong Niu. Frequency-domain MLPs are more effective learners in time series forecasting. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 11121–11128, 2023.
- Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test time robustness via adaptation and augmentation. In *Advances in Neural Information Processing Systems*, volume 35, pp. 38629–38642, 2022.
- YiFan Zhang, Weiqi Chen, Zhaoyang Zhu, Dalin Qin, Liang Sun, Xue Wang, Qingsong Wen, Zhang Zhang, Liang Wang, and Rong Jin. Addressing concept shift in online time series forecasting: Detect-then-adapt. *arXiv preprint arXiv:2403.14949*, 2024.
- Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 11106–11115, 2021.

## A ADAPTATION ALGORITHM OF COSA

This section details the core adaptation algorithm of COSA. This algorithm extends the optimizer mentioned in Equation 2, combining short-term loss trends, cosine annealing, and adaptive gradient clipping, batch-wise learning rate reset. We apply widely used values for the coefficients of *CALR* and the threshold value of early stopping. The base model remains frozen, operating only on adapter parameters  $\varphi = \{\mathbf{W}, \mathbf{b}, \mathbf{g}\}$ .

**Purpose:** Performs rapid adaptation of linear adapters using direct loss, adaptive learning rate scheduling, and gradient clipping to improve prediction accuracy within a few adaptation step  $S$ .

---

### Algorithm 1 COSA adaptation.

---

**Require:** Stream of batches  $(\hat{\mathbf{Y}}, \mathbf{Y}^{\text{true}})$  with length  $B$ , context vector  $\mathbf{C}$ , steps  $S$ ,  $\eta_{\min}, \eta_{\max}$ , weight decay  $\lambda$ , clip base  $c$

**Ensure:** adapted predictions with improved accuracy over the base model

- 1: **for** each batch  $(\hat{\mathbf{Y}}, \mathbf{Y}^{\text{true}})$  in Data Stream **do** ▷ Loop over new batches
- 2:     **Initialize:**  $\eta \leftarrow \eta_{\max}$ , loss history  $\mathcal{H} \leftarrow []$  ▷ Reset LR for new batch
- 3:     **for**  $s = 1$  to  $S$  **do** ▷ Main adaptation loop
- 4:         **Forward pass:** form  $\mathbf{X}^{(a)} = [\mathbf{Y}^{(0)} \parallel \mathbf{C}]$ , then  $\mathbf{H} = (\mathbf{W} \mathbf{X}^{(a)\top} + \mathbf{b})^\top$
- 5:         **Gating:**  $\hat{\mathbf{Y}} \leftarrow \mathbf{Y}^{(0)} + \tanh(\mathbf{g}) \odot \mathbf{H}$
- 6:         **MSE loss:**  $\mathcal{L} \leftarrow \|\hat{\mathbf{Y}} - \mathbf{Y}^{\text{true}}\|_F^2 + \lambda \|\varphi\|_2^2$
- 7:         **Learning rate adaptation:**
- 8:         **if**  $|\mathcal{H}| \geq 2$  **then**
- 9:              $\Delta \leftarrow \mathcal{L} - \mathcal{H}[-1]$  ▷ Recent loss change
- 10:            **if**  $\Delta > 0$  **then** ▷ Loss increased - reduce LR
- 11:                 $\eta \leftarrow \max(0.5\eta, \eta_{\min})$
- 12:                **else if**  $|\Delta| < 10^{-6}$  **then** ▷ Converged - increase LR for next batch
- 13:                 $\eta \leftarrow \min(1.1\eta, \eta_{\max})$
- 14:             **end if**
- 15:         **end if**
- 16:         **Cosine annealing:**  $\eta \leftarrow \eta_{\min} + \frac{1}{2}(\eta - \eta_{\min})(1 + \cos(\frac{s\pi}{S}))$
- 17:         **Gradient computation:**  $g_\varphi \leftarrow \nabla_\varphi \mathcal{L}$
- 18:         **Adaptive clipping:**  $\|g_\varphi\| \leftarrow \min(\|g_\varphi\|, \max(c, \mathcal{L}))$
- 19:         **Parameter update:**  $\varphi \leftarrow \varphi - \eta g_\varphi$
- 20:         **Early stopping:**
- 21:         **if**  $s > 2$  **and**  $|\mathcal{H}[-1] - \mathcal{H}[-2]| < 10^{-6}$  **then**
- 22:             **break**
- 23:         **end if**
- 24:     **end for**
- 25: **end for**

---

## B COMPUTATIONAL COST

In this section, we report theoretical calculations of parameters, FLOPs, and memory footprint.

For univariate time series, the number of parameters is as follows:

**Parameter count.**

$$\underbrace{L(L+K)}_{\mathbf{W}} + \underbrace{L}_{\mathbf{b}} + \underbrace{1}_{\mathbf{g}} \Rightarrow \#\text{params} = (L(L+K) + L + 1).$$

**FLOPs per adaptation step (batch of size  $B$ ).** The dominant cost is the linear transform and residual composition:

$$\mathcal{O}(BL(L+K)) \text{ for } \mathbf{W}\mathbf{X}^{(a)}, \quad \text{plus } \mathcal{O}(BL) \text{ for gating \& residual add.}$$

**Memory footprint.** Additional activations are modest: the linear residual  $\mathbf{H} \in \mathbb{R}^{B \times L}$  and the context vector  $\mathbf{C} \in \mathbb{R}^K$ . The total adaptation cost scales linearly with the number of adaptation steps  $S$  and variables  $V$ .

Table 4: Average  $B$  of each dataset determined by PAAS.

Dataset	ETTh1	ETTh2	ETTm1	ETTm2	Exchange Rate	Weather
Average $B$	24.55	38.41	92.73	83.41	92.80	80.38

## C BATCH SIZE ANALYSIS OF PAAS

**Adaptive vs. Fixed Batch Strategy:** While COSA-F shows the best performance in most cases, COSA-P generally performs better on the ETTh1 and ETTh2 datasets, revealing important insights about temporal adaptation strategies. The reason for this trend is that the size of  $B$  determined by PAAS is often smaller than 48 for these datasets, as shown in Table 4, enabling more frequent adaptation that better captures the higher-frequency patterns characteristic of these hourly datasets.

This differential performance validates the importance of dataset-specific adaptation scheduling: datasets with more complex temporal dynamics benefit from more frequent adaptation (smaller  $B$ ), while datasets with smoother patterns can effectively use larger batch sizes for computational efficiency.

## D BEHAVIOR ANALYSIS OF COSA

### D.1 ANALYSIS BETWEEN GATING AND LINEAR RESIDUAL LAYER

In COSA, the gating is defined as  $gating = \tanh(g) \in [-1, 1]$ , where  $g$  is a learnable parameter, and this bounded scalar modulates the correction strength by multiplying the output of the linear residual layer. If we were to use  $g$  directly instead of  $\tanh(g)$ , small variations in  $g$  could induce disproportionately large and unstable changes in the correction, making the adapter overly sensitive to noisy points. The  $\tanh$  transform keeps the gating bounded, ensuring that changes in  $g$  are reflected smoothly and gradually. When the residual magnitude spikes, the gate moves toward 0, as shown in Figure 5, thereby attenuating the residual correction and stabilizing the adaptation process.

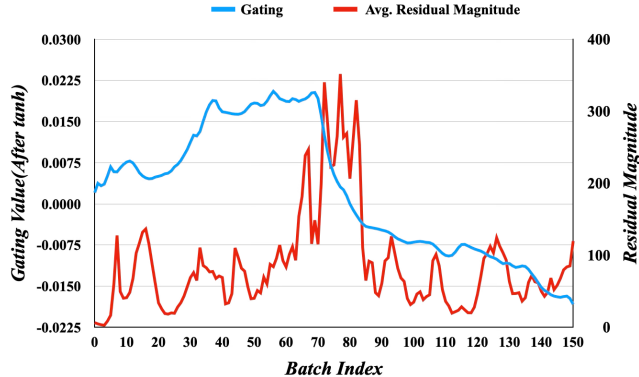


Figure 5: Trajectory of average gating and average residual magnitude of batches over time.

### D.2 ANALYSIS BETWEEN LEARNING RATE AND MSE

Figure 6 visualizes the trajectory of pre-adaptation loss and learning rate for the iTransformer-ETTm1,  $L = 96$  case. As shown in Figure 6, when a short-term loss spike occurs, CALR immediately decreases the learning rate to minimize the impact of the perturbation, and once the loss enters a stable decreasing phase, CALR increases the learning rate again to promote rapid re-adaptation. This control mechanism suppresses excessive parameter drift without requiring roll-back, enabling COSA to recover its correction performance instantly after a perturbation. The interaction between the learning rate and loss shows that, in non-stationary environments with short-term perturbations, COSA can respond and recover performance stably.

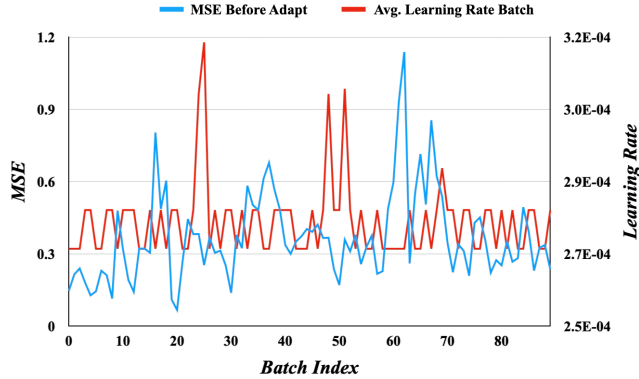


Figure 6: Trajectory of initial MSE and average learning rate of batches over time.

Table 5: Average ERV and NAR across all benchmark datasets and base models.

	TAFAS	PETSA	COSA
ERV $\uparrow$	.0100	.0160	<b>.0768</b>
NAR $\downarrow$	34.94%	36.34%	<b>20.25%</b>

## E QUALITY OF TTA

To evaluate TTA quality, we measured Explained Residual Variance (ERV) and Negative Adaptation Rate (NAR) metrics across all datasets, base models, and forecasting horizons. ERV is defined as in Equation 3, where  $\hat{R}$  represents residuals for TTA-applied predictions and  $R^{(0)}$  represents residuals for base model predictions. Specifically, ERV quantifies the extent to which TTA reduces the residual variance of the base model’s predictions. Higher ERV values indicate greater residual variance reduction and correspondingly improved prediction performance through TTA.

$$\text{ERV} = 1 - \frac{\text{Var}(\hat{R})}{\text{Var}(R^{(0)})}. \quad (3)$$

NAR is the ratio of prediction windows where the MSE worsened when TTA was applied, with smaller values indicating better performance.

$$\text{NAR} = \frac{1}{N} \sum_{t=1}^N I[MSE(\hat{\mathbf{Y}}_t) > MSE(\mathbf{Y}_t^{(0)})], \quad \text{where} \begin{cases} I[\text{con}] = 1 & \text{if con} = \text{True} \\ I[\text{con}] = 0 & \text{otherwise} \end{cases} \quad (4)$$

As shown in Table 5, COSA shows the highest ERV and the lowest NAR. This indicates that COSA provides effective TTA while decreasing the residual and improving accuracy on average.

## F FURTHER EXPERIMENTS

### F.1 COMPARISON ON A LARGER DATASET

To demonstrate that COSA can achieve stable and substantial performance gains even in larger-scale environments, we additionally conducted experiments on the Electricity dataset Lai et al. (2018); Godahewa et al. (2021). We followed the same experimental setup as in Section 4.1, and used three representative base models, iTransformer, DLinear, and FreTS. Table 6 is consistent with the findings in Section 4.2, COSA achieves either the best or second-best performance across all forecasting horizons  $L$ , and unlike other methods, which exhibit performance degradation compared to baselines in some cases, COSA improves their performances in every case. These results indicate that COSA remains robust and effective even in large-scale environments.

Table 6: Prediction accuracy comparison on Electricity dataset.

	iTransformer				DLinear				FreTS						
	No TTA	TAFAS	PETSA	COSA-F COSA-P	No TTA	TAFAS	PETSA	COSA-F COSA-P	No TTA	TAFAS	PETSA	COSA-F COSA-P			
96	.1663	<b>.1568</b>	.1596	.1571	<u>.1570</u>	.2235	<b>.2224</b>	.2242	.2232	<u>.2228</u>	.1824	<b>.1781</b>	.1802	.1801	<u>.1799</u>
192	.1794	<u>.1635</u>	.1644	<b>.1631</b>	<b>.1631</b>	.2242	.2153	.2152	.2146	<b>.2143</b>	.1794	.1894	.1781	<u>.1780</u>	<b>.1779</b>
336	.1952	.1743	.1741	<u>.1730</u>	<b>.1727</b>	.2383	.2247	<u>.2233</u>	<b>.2226</b>	<u>.2223</u>	.1905	.1885	.1885	<u>.1881</u>	<b>.1879</b>
720	.2567	.2316	.2301	<u>.2251</u>	<b>.2239</b>	.2792	.2629	.2615	<u>.2557</u>	<b>.2541</b>	.2304	.2299	.2287	<u>.2236</u>	<b>.2221</b>

## F.2 COMPARISON ON A REAL DATASET

To validate that COSA performs well not only on benchmark datasets but also in real-world deployment environments exhibiting more extreme distribution shifts, we conducted a performance evaluation on a real-world dataset. The dataset consists of power generation data collected from an actual solar power plant spanning from Jan. 24, 2024, 00:00 to Dec. 31, 2024, 23:59. It comprises hourly active power sensor readings from 30 devices with no missing values. The data was split into train, test, and validation sets at a ratio of 7:1:2. The dataset contains 8,232 rows, approximately 47.3% the size of ETTh1. We conducted experiments solely under the setting  $L = 96$ ,  $W = 96$ , with all other configurations kept identical.

The experimental results demonstrate that, consistent with the benchmark findings, COSA achieves improved performance over existing methods regardless of  $B$ . This indicates that COSA can operate effectively even under more extreme distributional changes and can reliably calibrate forecasting performance on very small-scale datasets.

Table 7: Prediction accuracy comparison.

Metric	iTransformer				DLinear				FreTS			
	No TTA	TAFAS	COSA-F	COSA-P	No TTA	TAFAS	COSA-F	COSA-P	No TTA	TAFAS	COSA-F	COSA-P
Avg. MSE	.2901	<u>.2879</u>	<b>.2730</b>	.2897	.2832	.2813	<b>.2801</b>	<u>.2803</u>	.3013	.2983	<b>.2944</b>	<u>.2980</u>

## F.3 COMPARISON WITH VARYING INPUT/PREDICTION SEQUENCE LENGTH

To further verify the performance of COSA across diverse scenarios, we conducted experiments by varying both the input window  $W$  and the prediction horizon  $L$ . To examine short-term forecasting rather than long-term forecasting, we added the setting  $W = 96$ ,  $L \in \{24, 48\}$ . To evaluate performance under longer input windows, we additionally tested  $W = 192$  with  $L \in \{192, 336, 720\}$  and  $W = 336$  with  $L \in \{336, 720\}$ . Table 8 summarizes the results across these different combinations of  $W$  and  $L$ . Consistent with Section 4.2, COSA achieves the best or second-best performance in most cases, demonstrating its ability to maintain high predictive accuracy across a wide range of settings. Unlike other methods, which show performance degradation relative to baselines in several cases, COSA improves prediction accuracy over baselines in every case. In contrast, TAFAS and PETSA, which adopt dual-adaptor architectures that modify both the input and output of base model incur substantial additional complexity as  $W$  increases, which likely contributes to their degraded performance. Since COSA operates solely in the output space of base model, it delivers consistent performance gains regardless of  $W$ .

## F.4 COMPARISON WITH SOLID

SOLID (Chen et al., 2024b) is a fine-tuning method of the prediction layer of the base model by detecting context drift. Its reconstructor estimates context drift based on the mutual information between the model’s residual and the input context. When drift is detected, SOLID selectively chooses samples and fine-tunes the model’s prediction layer at the sample level. Like COSA, SOLID operates in the output space of base model. However, unlike COSA, which keeps base model frozen and performs corrections through a lightweight adapter, SOLID directly updates the prediction layer of base model.

Table 8: Prediction accuracy comparison with different input/prediction sequence length.

		iTransformer						DLinear						FreTS					
Input		Pred	No TTA	TAFAS	PETSA	COSA-F	COSA-P	No TTA	TAFAS	PETSA	COSA-F	COSA-P	No TTA	TAFAS	PETSA	COSA-F	COSA-P		
ETTm1	96 24	.3269	.3254	.3299	<u>.3105</u>	<b>.3098</b>	.3437	.3751	.3766	<u>.3431</u>	<b>.3429</b>	.2953	.3243	.3239	<u>.2946</u>	<b>.2943</b>			
	96 48	.3732	.3768	.3776	<u>.3539</u>	<b>.3527</b>	.3838	.4177	.4166	<u>.3826</u>	<b>.3820</b>	.3424	.3770	.3754	<u>.3416</u>	<b>.3413</b>			
	192 192	.4459	.5015	.4893	<u>.4175</u>	<b>.4120</b>	.4172	.4967	.4834	<u>.4157</u>	<b>.4141</b>	.4235	.4991	.4835	<u>.4207</u>	<b>.4186</b>			
	192 336	.4729	.5709	.5422	<u>.4486</u>	<b>.4456</b>	.4537	.5558	.5362	<u>.4528</u>	<b>.4489</b>	.4596	.5677	.5398	<u>.4577</u>	<b>.4535</b>			
	192 720	.5885	.6529	.6687	<u>.5619</u>	<b>.5477</b>	.5828	.6504	.6526	<u>.5734</u>	<b>.5660</b>	.6484	.6627	.6728	<u>.6351</u>	<b>.6190</b>			
	336 336	.5354	.6139	.5686	<u>.5209</u>	<b>.5059</b>	.4392	.5652	.5174	<u>.4368</u>	<b>.4289</b>	.4837	.6114	.5444	<u>.4783</u>	<b>.4674</b>			
336 720	.6512	.6873	.6635	<u>.6115</u>	<b>.5780</b>	.5748	.6344	.6443	<u>.5589</u>	<b>.5198</b>	.7265	.6905	.6738	<u>.6960</u>	<b>.6295</b>				
ETTm2	96 24	<u>.0837</u>	.1277	.1266	<b>.0812</b>	<u>.0812</u>	.0862	.1217	.1218	<b>.0861</b>	<u>.0862</u>	.0812	.1180	.1178	<b>.0806</b>	<u>.0807</u>			
	96 48	.1047	.1518	.1520	<b>.1007</b>	<u>.1008</u>	.1031	.1444	.1452	<b>.1028</b>	<u>.1030</u>	.1004	.1452	.1447	<b>.0996</b>	<u>.0996</u>			
	192 192	.1633	.2272	.2280	<b>.1546</b>	<u>.1563</u>	.1344	.2155	.2096	<b>.1334</b>	<u>.1338</u>	.1437	.2218	.2184	<b>.1412</b>	<u>.1420</u>			
	192 336	.1658	.2541	.2487	<b>.1551</b>	<u>.1573</u>	.1474	.2445	.2394	<b>.1460</b>	<u>.1468</u>	.1556	.2494	.2460	<b>.1519</b>	<u>.1536</u>			
	192 720	.2197	.3453	.3192	<b>.2110</b>	<u>.2158</u>	.1782	.3152	.2962	<b>.1762</b>	<u>.1778</u>	.1943	.3261	.3086	<b>.1892</b>	<u>.1925</u>			
	336 336	<u>.2070</u>	.2986	.3005	<b>.2021</b>	<u>.2108</u>	.1469	.2366	.2375	<b>.1447</b>	<u>.1468</u>	.1722	.2579	.2604	<b>.1645</b>	<u>.1695</u>			
336 720	<u>.3334</u>	.4253	.4355	<b>.3255</b>	<u>.3334</u>	.1807	.3101	.3004	<b>.1776</b>	<u>.1798</u>	.1969	.3130	.3124	<b>.1920</b>	<u>.1949</u>				
ETTm1	96 24	.2543	<b>.2359</b>	.2429	<u>.2405</u>	.2450	.2578	.2616	.2599	<b>.2523</b>	<u>.2569</u>	.2494	.2503	<u>.2471</u>	<b>.2449</b>	.2492			
	96 48	.3305	<b>.3099</b>	.3135	<u>.3115</u>	.3133	.3105	.3167	.3151	<b>.3035</b>	<u>.3064</u>	.3191	.3269	.3222	<b>.3171</b>	<u>.3177</u>			
	192 192	.3800	.4006	.3961	<b>.3591</b>	<u>.3603</u>	.3694	.4029	.3942	<b>.3643</b>	<u>.3648</u>	.3551	.3863	.3815	<b>.3524</b>	<u>.3528</u>			
	192 336	.4336	.4499	.4401	<b>.4149</b>	<u>.4150</u>	.4239	.4620	.4436	<b>.4162</b>	<u>.4169</u>	.4036	.4371	.4290	<b>.3985</b>	<u>.3987</u>			
	192 720	.4992	.5385	.5090	<b>.4716</b>	<u>.4718</u>	.4797	.5376	.5101	<b>.4671</b>	<u>.4681</u>	.4631	.5152	.4941	<b>.4550</b>	<u>.4560</u>			
	336 336	.4820	.4913	<b>.4441</b>	<u>.4490</u>	.4592	.3981	.4444	.4244	<b>.3968</b>	<u>.3977</u>	.4011	.4380	.4275	<b>.3975</b>	<u>.4006</u>			
336 720	.5202	.5432	.5169	<b>.4678</b>	<u>.4904</u>	.4437	.5053	.4852	<b>.4256</b>	<u>.4423</u>	.4407	.5046	.4877	<b>.4261</b>	<u>.4397</u>				
ETTm2	96 24	.0542	.0800	.0744	<b>.0530</b>	<u>.0531</u>	.0601	.0782	.0777	<b>.0598</b>	<u>.0599</u>	.0562	.0748	.0741	<b>.0561</b>	<u>.0561</u>			
	96 48	.0735	.1044	.0988	<b>.0731</b>	<u>.0732</u>	.0767	.1027	.1015	<b>.0761</b>	<u>.0762</u>	.0736	.0992	.0987	<b>.0734</b>	<u>.0733</u>			
	192 192	.1035	.1647	.1601	<b>.1023</b>	<u>.1022</u>	.0990	.1451	.1429	<b>.0984</b>	<u>.0985</u>	.0998	.1480	.1456	<b>.0982</b>	<u>.0984</u>			
	192 336	.1343	.2191	.1970	<b>.1339</b>	<u>.1335</u>	.1201	.1791	.1741	<b>.1190</b>	<u>.1195</u>	.1210	.1863	.1781	<b>.1200</b>	<u>.1192</u>			
	192 720	.1628	.2511	.2377	<b>.1564</b>	<u>.1589</u>	.1518	.2291	.2247	<b>.1473</b>	<u>.1498</u>	.1476	.2337	.2232	<b>.1436</b>	<u>.1446</u>			
	336 336	.1354	.2280	.1987	<b>.1328</b>	<u>.1347</u>	.1176	.1782	.1704	<b>.1167</b>	<u>.1174</u>	.1198	.1815	.1760	<b>.1172</b>	<u>.1188</u>			
336 720	.1645	.2785	.2517	<b>.1455</b>	<u>.1604</u>	.1456	.2338	.2185	<b>.1317</b>	<u>.1446</u>	.1490	.2309	.2255	<b>.1289</b>	<u>.1453</u>				
Exchange Rate	96 24	.0318	<u>.0292</u>	<b>.0273</b>	.0306	.0306	.0434	<u>.0397</u>	<b>.0393</b>	.0434	.0434	.0283	<u>.0254</u>	<b>.0241</b>	.0283	.0283			
	96 48	.0526	.0542	<b>.0479</b>	<u>.0516</u>	.0516	.0606	<u>.0590</u>	<b>.0579</b>	.0606	.0606	.0481	<u>.0438</u>	<b>.0418</b>	.0480	.0481			
	192 192	.2045	<b>.1999</b>	.2189	<u>.2040</u>	.2040	.1715	.1819	.2024	<b>.1714</b>	<u>.1715</u>	<b>.1655</b>	<u>.1728</u>	.1809	<b>.1655</b>	<u>.1655</u>			
	192 336	.2963	.3221	.3511	<b>.2919</b>	<u>.2933</u>	.2878	.3316	.3484	<b>.2850</b>	<u>.2875</u>	.2804	.3129	.3318	<b>.2780</b>	<u>.2802</u>			
	192 720	.8546	<b>.7713</b>	<u>.7828</u>	.8316	.8320	<u>.9037</u>	.9393	.9488	<b>.9016</b>	<u>.9016</u>	.4506	.8200	.8315	<b>.4501</b>	<u>.4501</u>			
	336 336	.4088	<b>.3790</b>	.4204	<u>.4036</u>	.4036	.2753	.3036	.3323	<b>.2700</b>	<u>.2751</u>	.3246	.3219	.3589	<b>.3182</b>	<u>.3245</u>			
336 720	1.7882	<b>1.0041</b>	<u>1.0775</u>	1.5829	1.5828	<u>.5274</u>	.8620	.8816	<b>.5269</b>	<u>.5269</u>	.4266	.9353	.9175	<b>.4263</b>	<u>.4263</u>				
Weather	96 24	.1095	<u>.1026</u>	<b>.1016</b>	.1077	.1077	.1207	<b>.1151</b>	<b>.1151</b>	.1200	<u>.1199</u>	.1193	<b>.1085</b>	<u>.1111</u>	.1183	.1183			
	96 48	.1380	<b>.1292</b>	<u>.1310</u>	.1362	.1363	.1581	<u>.1499</u>	<b>.1475</b>	.1559	.1560	.1530	<b>.1435</b>	<u>.1452</u>	.1508	.1510			
	192 192	.2165	.2102	<b>.2051</b>	<u>.2075</u>	.2099	.2393	<u>.2225</u>	<b>.2214</b>	.2332	.2354	.2131	<u>.2058</u>	<b>.2021</b>	.2086	.2108			
	192 336	.2747	<b>.2588</b>	<b>.2549</b>	.2611	.2653	.2904	<u>.2709</u>	<b>.2691</b>	.2791	.2829	.2694	<u>.2547</u>	<b>.2531</b>	.2604	.2637			
	192 720	.3321	.3316	.3390	<b>.3151</b>	<u>.3211</u>	.3426	.3363	.3446	<b>.3276</b>	<u>.3339</u>	.3290	.3286	.3360	<b>.3158</b>	<u>.3205</u>			
	336 336	.2988	<b>.2705</b>	<u>.2708</u>	.2778	.2978	.2722	.2633	<b>.2619</b>	<u>.2627</u>	.2707	.2526	.2520	<u>.2474</u>	<b>.2444</b>	<u>.2504</u>			
336 720	.3381	.3423	.3557	<b>.2885</b>	<u>.3340</u>	.3245	.3247	.3316	<b>.2742</b>	<u>.3207</u>	.3129	.3203	<u>.3229</u>	<b>.2648</b>	<u>.3088</u>				

Using the publicly available source codes, we evaluated SOLID on three base forecasting models: DLinear, FreTS, and iTransformer. The results are summarized in Table 9. Across all settings, COSA achieves higher predictive accuracy compared to SOLID.

Moreover, COSA is significantly more efficient: it requires only  $7.44 \pm 0.0488$  seconds of wall-clock time on average, whereas SOLID incurs a much larger overhead with  $306.55 \pm 2.6183$  seconds. The substantial cost of SOLID arises from computing context drift and repeatedly updating the prediction layer of the base model, processes that are far heavier than COSA’s lightweight updates. Because COSA freezes base model and updates only a compact output-space adapter, it offers a dramatically faster execution while maintaining higher accuracy.

## F.5 COMPARISON WITH DYNATTA

This section presents complete baseline results of prediction accuracy comparison with various existing methods. DynaTTA requires setting a total of 15 hyperparameters, and due to this complexity, we report its performance using the best settings from our trials. The used hyperparameters of DynaTTA are as follows:

- HIDDEN\_DIM=64, GATING\_INIT=0.01, BASE\_LR= 0.005, MSE\_BUFFER\_SIZE=100, RTAB\_SIZE=50, RDB\_SIZE=30, METRIC\_HISTORY\_SIZE=20, ALPHA\_MIN=0.0001, ALPHA\_MAX=0.01, KAPPA=2.0, ETA=0.1, EPS=1e-8, WARMUP\_FACTOR=2, UPDATE\_BUFFERS\_INTERVAL=5, UPDATE\_METRICS\_INTERVAL=3

Table 9: Comparison with SOLID

	iTransformer		DLinear		FreTS		
	SOLID COSA-P		SOLID COSA-P		SOLID COSA-P		
ETTh1	96	.4404	<b>.4363</b>	.4595	<b>.4574</b>	<b>.4093</b>	.4371
	192	.4935	<b>.4919</b>	.5063	<b>.5066</b>	<b>.4701</b>	.4940
	336	.5420	<b>.5300</b>	.5602	<b>.5528</b>	<b>.5254</b>	.5467
	720	<b>.6523</b>	.5638	.7107	<b>.6107</b>	.6980	<b>.6259</b>
ETTh2	96	<b>.2480</b>	.2493	.2315	<b>.2281</b>	.2354	<b>.2350</b>
	192	.3061	<b>.2947</b>	.2824	<b>.2819</b>	.2830	<b>.2972</b>
	336	<b>.3339</b>	<b>.3339</b>	.3103	<b>.3083</b>	.3248	<b>.3153</b>
	720	.3701	<b>.3591</b>	<b>.3136</b>	.3477	<b>.3167</b>	.3399
ETTh1	96	<b>.3353</b>	.3455	.3634	<b>.3456</b>	<b>.3121</b>	.3525
	192	.4266	<b>.4140</b>	.4336	<b>.4222</b>	.4299	<b>.4212</b>
	336	.5019	<b>.4643</b>	.5071	<b>.4858</b>	.4915	<b>.4775</b>
	720	.5986	<b>.5102</b>	.5821	<b>.4991</b>	.5612	<b>.4982</b>
ETTh2	96	.1641	<b>.1632</b>	.1590	<b>.1583</b>	.1574	<b>.1569</b>
	192	.2291	<b>.2173</b>	.1910	<b>.1943</b>	.1933	<b>.1934</b>
	336	<b>.2525</b>	.2535	<b>.2138</b>	.2242	.2328	<b>.2211</b>
	720	.2704	<b>.2606</b>	.2419	<b>.2316</b>	.2377	<b>.2314</b>
Exchange Rate	96	.0873	<b>.0837</b>	.0913	<b>.0834</b>	.0822	<b>.0766</b>
	192	.1783	<b>.1479</b>	.1826	<b>.1519</b>	.1723	<b>.1499</b>
	336	.3294	<b>.2624</b>	.3276	<b>.2480</b>	.3218	<b>.2461</b>
	720	.7531	<b>.4460</b>	.8872	<b>.4481</b>	.8325	<b>.4458</b>
Weather	96	.1753	<b>.1617</b>	.1954	<b>.1793</b>	.1849	<b>.1737</b>
	192	.2231	<b>.2088</b>	.2403	<b>.2217</b>	.2309	<b>.2189</b>
	336	.2801	<b>.2515</b>	.2918	<b>.2626</b>	.2843	<b>.2587</b>
	720	.3450	<b>.2730</b>	.3643	<b>.2708</b>	.3561	<b>.2692</b>

Table 10: Prediction accuracy comparison with DynaTTA Grover &amp; Etemad (2025).

	Transformer-based						Linear-based				MLP-based								
	iTransformer			PatchTST			DLinear		OLS		FreTS		MICN						
	DynaTTA	COSA-F	COSA-P	DynaTTA	COSA-F	COSA-P	DynaTTA	COSA-F	COSA-P	DynaTTA	COSA-F	COSA-P	DynaTTA	COSA-F	COSA-P				
ETTh1	96	.4523	<b>.4368</b>	<b>.4363</b>	.8371	<b>.4242</b>	<b>.4238</b>	.4708	<b>.4574</b>	<b>.4482</b>	.4486	<b>.4390</b>	<b>.4372</b>	.4508	<b>.4384</b>	<b>.4371</b>	.5063	<b>.4693</b>	<b>.4684</b>
	192	.5175	<b>.4961</b>	<b>.4919</b>	.8006	<b>.4830</b>	<b>.4805</b>	.5321	<b>.5066</b>	<b>.5050</b>	.5096	<b>.4915</b>	<b>.4906</b>	.5139	<b>.4951</b>	<b>.4940</b>	.5745	<b>.5372</b>	<b>.5328</b>
	336	.5874	<b>.5651</b>	<b>.5300</b>	.8097	<b>.5438</b>	<b>.5320</b>	.5792	<b>.5528</b>	<b>.5456</b>	.5626	<b>.5385</b>	<b>.5320</b>	.5840	<b>.5467</b>	<b>.5351</b>	.6594	<b>.5950</b>	<b>.5878</b>
	720	.7123	<b>.5958</b>	<b>.5638</b>	1.0887	<b>.6113</b>	<b>.5822</b>	.7112	<b>.6107</b>	<b>.5896</b>	.6933	<b>.5969</b>	<b>.5733</b>	.7095	<b>.6259</b>	<b>.5959</b>	.8262	<b>.7001</b>	<b>.6504</b>
ETTh2	96	.2630	<b>.2504</b>	<b>.2493</b>	.4154	<b>.2349</b>	<b>.2343</b>	.2338	<b>.2300</b>	<b>.2281</b>	.2326	<b>.2232</b>	<b>.2265</b>	.2395	<b>.2367</b>	<b>.2350</b>	.2661	<b>.2492</b>	<b>.2485</b>
	192	.3210	<b>.2983</b>	<b>.2947</b>	.4315	<b>.2665</b>	<b>.2608</b>	.2888	<b>.2827</b>	<b>.2819</b>	.2911	<b>.2796</b>	<b>.2791</b>	.2916	<b>.2816</b>	<b>.2824</b>	.3417	<b>.3049</b>	<b>.3017</b>
	336	.3677	<b>.3241</b>	<b>.3339</b>	.4386	<b>.2971</b>	<b>.2978</b>	.3380	<b>.3050</b>	<b>.3083</b>	.3430	<b>.3003</b>	<b>.3043</b>	.3474	<b>.3031</b>	<b>.3153</b>	.3777	<b>.3241</b>	<b>.3310</b>
	720	.4646	<b>.3487</b>	<b>.3591</b>	.4991	<b>.3233</b>	<b>.3428</b>	.4325	<b>.3062</b>	<b>.3477</b>	.4270	<b>.3177</b>	<b>.3453</b>	.4297	<b>.3169</b>	<b>.3399</b>	.5229	<b>.3650</b>	<b>.3885</b>
ETTh1	96	.3753	<b>.3447</b>	<b>.3455</b>	.7205	<b>.3625</b>	<b>.3626</b>	.3814	<b>.3456</b>	<b>.3475</b>	.3830	<b>.3454</b>	<b>.3475</b>	.3723	<b>.3520</b>	<b>.3525</b>	.4095	<b>.3837</b>	<b>.3831</b>
	192	.4835	<b>.4124</b>	<b>.4140</b>	.6898	<b>.4250</b>	<b>.4258</b>	.4809	<b>.4113</b>	<b>.4122</b>	.4827	<b>.4115</b>	<b>.4119</b>	.4811	<b>.4150</b>	<b>.4212</b>	.5331	<b>.4476</b>	<b>.4514</b>
	336	.5843	<b>.4569</b>	<b>.4643</b>	.8073	<b>.4568</b>	<b>.4697</b>	.5711	<b>.4753</b>	<b>.4858</b>	.5719	<b>.4748</b>	<b>.4749</b>	.5739	<b>.4661</b>	<b>.4775</b>	.6418	<b>.4832</b>	<b>.5054</b>
	720	.7243	<b>.4773</b>	<b>.5102</b>	.9292	<b>.4681</b>	<b>.4882</b>	.6569	<b>.4774</b>	<b>.4991</b>	.6884	<b>.4763</b>	<b>.5007</b>	.7152	<b>.4718</b>	<b>.4982</b>	.6932	<b>.5029</b>	<b>.5225</b>
ETTh2	96	.1832	<b>.1627</b>	<b>.1632</b>	.2508	<b>.1558</b>	<b>.1562</b>	.1640	<b>.1583</b>	<b>.1586</b>	.1684	<b>.1582</b>	<b>.1586</b>	.1661	<b>.1568</b>	<b>.1569</b>	.2033	<b>.1702</b>	<b>.1704</b>
	192	.2897	<b>.2171</b>	<b>.2173</b>	.2744	<b>.2007</b>	<b>.2022</b>	.2040	<b>.1904</b>	<b>.1905</b>	.2115	<b>.1906</b>	<b>.1907</b>	.2048	<b>.1905</b>	<b>.1908</b>	.2349	<b>.2102</b>	<b>.2120</b>
	336	.3246	<b>.2435</b>	<b>.2535</b>	.3434	<b>.2258</b>	<b>.2352</b>	.2908	<b>.2083</b>	<b>.2242</b>	.2917	<b>.2131</b>	<b>.2226</b>	.2778	<b>.2098</b>	<b>.2211</b>	.2988	<b>.2337</b>	<b>.2351</b>
	720	.5344	<b>.2477</b>	<b>.2606</b>	.4447	<b>.2446</b>	<b>.2645</b>	.3765	<b>.2215</b>	<b>.2316</b>	.4571	<b>.2171</b>	<b>.2349</b>	.4020	<b>.2158</b>	<b>.2314</b>	.5362	<b>.2477</b>	<b>.2643</b>
Exchange Rate	96	.0983	<b>.0818</b>	<b>.0837</b>	.1217	<b>.0765</b>	<b>.0788</b>	.0948	<b>.0812</b>	<b>.0834</b>	.0918	<b>.0756</b>	<b>.0773</b>	.0906	<b>.0744</b>	<b>.0766</b>	.1210	<b>.0955</b>	<b>.1008</b>
	192	.1970	<b>.1403</b>	<b>.1479</b>	.2517	<b>.1464</b>	<b>.1570</b>	.1975	<b>.1459</b>	<b>.1519</b>	.1749	<b>.1393</b>	<b>.1457</b>	.1871	<b>.1366</b>	<b>.1499</b>	.2227	<b>.1663</b>	<b>.1722</b>
	336	.3251	<b>.2089</b>	<b>.2624</b>	.3728	<b>.1983</b>	<b>.2445</b>	.3001	<b>.2039</b>	<b>.2480</b>	.3024	<b>.2020</b>	<b>.2323</b>	.3111	<b>.2053</b>	<b>.2461</b>	.3536	<b>.2119</b>	<b>.2660</b>
	720	.8790	<b>.3421</b>	<b>.4460</b>	.9999	<b>.3543</b>	<b>.4662</b>	.8812	<b>.3494</b>	<b>.4481</b>	.8300	<b>.3444</b>	<b>.4541</b>	.8348	<b>.3352</b>	<b>.4458</b>	.8570	<b>.3871</b>	<b>.4815</b>
Weather	96	.1823	<b>.1597</b>	<b>.1617</b>	.2317	<b>.1624</b>	<b>.1634</b>	.1950	<b>.1773</b>	<b>.1793</b>	.1985	<b>.1772</b>	.1803	.1937	<b>.1724</b>	<b>.1737</b>	.2303	<b>.1636</b>	<b>.1651</b>
	192	.2678	<b>.2067</b>	<b>.2088</b>	.2611	<b>.2006</b>	<b>.2108</b>	.3311	<b>.2216</b>	<b>.2217</b>	.3074	<b>.2223</b>	<b>.2237</b>	.2850	<b>.2135</b>	<b>.2189</b>	.3839	<b>.2082</b>	<b>.2120</b>
	336	.3894	<b>.2503</b>	<b>.2515</b>	.3329	<b>.2451</b>	<b>.2488</b>	.4103	<b>.2567</b>	<b>.2626</b>	.9593	<b>.2851</b>	<b>.2642</b>	.3970	<b>.2561</b>	<b>.2587</b>	.5355	<b>.2729</b>	<b>.2737</b>
	720	.4996	<b>.2480</b>	<b>.2730</b>	.4067	<b>.2590</b>	<b>.2713</b>	.4915	<b>.2581</b>	<b>.2708</b>	.4974	<b>.2579</b>	<b>.2708</b>	.5251	<b>.2573</b>	<b>.2692</b>	.5773	<b>.2582</b>	<b>.2855</b>

## F.6 COMPARISON WITH VARIOUS BATCH SIZES

Table 11a and Table 11b summarize the prediction accuracy and efficiency overhead under different batch sizes  $B$ . Thanks to CALR’s stability-induced design, COSA adapts reliably even with very small batches, and the resulting increase in update frequency often leads to improved forecasting accuracy. Notably, even in extremely small settings such as  $B = 8$ , COSA maintains higher accuracy than existing TTA methods, demonstrating resilience against over-correction and short-term perturbations. On the other hand, smaller  $B$  inevitably increases the number of adaptation steps, leading

Table 11: Performance comparison with different batch sizes  $B$ .

(a) Prediction accuracy.									
	No TTA	TAFAS	PETSA	PAAS	8	16	24	48	96
<b>96</b>	0.2545	0.2471	0.2480	0.2409	<b>0.1908</b>	<u>0.2202</u>	0.2330	0.2398	0.2411
<b>192</b>	0.3144	0.3038	0.3046	0.2887	<b>0.1900</b>	<u>0.2295</u>	0.2524	0.2848	0.2946
<b>336</b>	0.3839	0.3653	0.3664	0.3280	<b>0.1837</b>	<u>0.2292</u>	0.2526	0.3013	0.3483
<b>720</b>	0.5539	0.5226	0.5232	0.3804	<b>0.1811</b>	<u>0.2384</u>	0.2672	0.3286	0.4158

(b) Wall-clock time (Seconds).									
	TAFAS	PETSA	PAAS	8	16	24	48	96	
<b>96</b>	<u>5.9271</u>	7.4129	7.0642	47.3922	25.6158	17.2297	9.6333	<b>5.1258</b>	
<b>192</b>	<u>6.0383</u>	7.3675	7.1419	49.4422	25.1222	18.0969	9.4344	<b>5.5133</b>	
<b>336</b>	<u>6.1554</u>	7.6588	7.4131	49.2211	25.1119	17.2906	9.5644	<b>5.8081</b>	
<b>720</b>	<b>6.5300</b>	7.9867	8.1578	46.3711	24.4683	17.6253	10.3036	<u>6.5397</u>	

to higher adaptation time and a clear computation–accuracy trade-off. Considering this trade-off, we adopt  $B = 48$  for COSA-F in our main experiments, which provides a balanced choice between accuracy gains and computational efficiency.

## G ADDITIONAL ABLATIONS

This section provides additional ablation studies on various design choices of COSA. Each experiment was conducted to understand the impact of specific components and determine optimal hyperparameters. The default setting uses mean-based context aggregation with  $K = 10$  and  $S = 3$ .

### G.1 CONTEXT BUILDING METHODS COMPARISON

#### G.1.1 STATISTICAL METHODS

This experiment was conducted to evaluate the impact of different context construction methods on adapter performance. We compared three methods, i.e., Mean, Median, and Weighted Average (WA), to find the optimal context construction strategy. The weight of WA was designed to assign greater weight to recent values using exponential decay weighting. Table 12 presents a comprehensive performance comparison of the three context construction methods.

Mean-based context construction demonstrated superior performance compared to median and weighted average approaches across most experimental configurations. While median-based aggregation provided robustness against outliers, it resulted in lower overall accuracy. The weighted average approach showed marginal improvements relative to its implementation complexity. These findings support the adoption of simple statistical aggregation for effective context summarization in COSA.

#### G.1.2 CONTEXT SELECTION STRATEGY

COSA aims not to model long-term time-series structure, but to perform fast and stable local residual correction for output bias observed in the current window. In non-stationary environments, the input distribution shifts continuously over time; thus, information from distant past windows may become misaligned with the current drift direction and deteriorate correction quality. For this reason, the default COSA employs a lightweight context vector constructed solely from the most recently observed batches.

To examine the effects of longer-range temporal patterns, we additionally implemented a Selective Context mechanism. This approach stores all past context values in a buffer and computes importance scores via attention between the current window and past contexts, selecting the top- $K$  values

Table 12: Prediction accuracy comparison of diverse aggregation functions.

	Transformer-based									Linear-based						MLP-based					
	iTransformer			PatchTST			DLinear			OLS			FreTS			MICN					
	Mean	Median	WA	Mean	Median	WA	Mean	Median	WA	Mean	Median	WA	Mean	Median	WA	Mean	Median	WA			
ETT1	96	<b>.4327</b>	<u>.4472</u>	<u>.4472</u>	<b>.4092</b>	<u>.4274</u>	<u>.4275</u>	<b>.4615</b>	<u>.4660</u>	<u>.4657</u>	<b>.4359</b>	<u>.4463</u>	<u>.4462</u>	<b>.4362</b>	<u>.4439</u>	<u>.4440</u>	<b>.4704</b>	<u>.4934</u>	<u>.4926</u>		
	192	<b>.4491</b>	<u>.4850</u>	<u>.4870</u>	<b>.4422</b>	<u>.4778</u>	<u>.4770</u>	<b>.4869</b>	<u>.5090</u>	<u>.5111</u>	<b>.4932</b>	<u>.4897</u>	<u>.4894</u>	<b>.4671</b>	<u>.4893</u>	<u>.4954</u>	<b>.4769</b>	<u>.5567</u>	<u>.5568</u>		
	336	<b>.4476</b>	<u>.5301</u>	<u>.5317</u>	<b>.4550</b>	<u>.5199</u>	<u>.5285</u>	<b>.4632</b>	<u>.5379</u>	<u>.5378</u>	<b>.5188</b>	<u>.5208</u>	<u>.5209</u>	<b>.4411</b>	<u>.5273</u>	<u>.5245</u>	<b>.4814</b>	<u>.5889</u>	<u>.5904</u>		
	720	<b>.4592</b>	<u>.6237</u>	<u>.6251</u>	<b>.4788</b>	<u>.6329</u>	<u>.6335</u>	<b>.4777</b>	<u>.6358</u>	<u>.6376</u>	<b>.5616</b>	<u>.6243</u>	<u>.6242</u>	<b>.4755</b>	<u>.6439</u>	<u>.6428</u>	<b>.5230</b>	<u>.7196</u>	<u>.7181</u>		
ETT2	96	<b>.2489</b>	<u>.2532</u>	<u>.2536</u>	<b>.2336</b>	<u>.2346</u>	<u>.2358</u>	<b>.2303</b>	<u>.2352</u>	<u>.2324</u>	<b>.2283</b>	<u>.2300</u>	<u>.2299</u>	<b>.2359</b>	<u>.2379</u>	<u>.2382</u>	<b>.2468</b>	<u>.2566</u>	<u>.2568</u>		
	192	<b>.3003</b>	<u>.3012</u>	<u>.3013</u>	<b>.2708</b>	<u>.2886</u>	<u>.2894</u>	<b>.2682</b>	<u>.2846</u>	<u>.2863</u>	<b>.2762</b>	<u>.2954</u>	<u>.2893</u>	<b>.2766</b>	<u>.2846</u>	<u>.2917</u>	<b>.2906</b>	<u>.3190</u>	<u>.3145</u>		
	336	<b>.3277</b>	<u>.3671</u>	<u>.3685</u>	<b>.2688</b>	<u>.3190</u>	<u>.3202</u>	<b>.2860</b>	<u>.3283</u>	<u>.3285</u>	<b>.2937</b>	<u>.3030</u>	<u>.3025</u>	<b>.2889</b>	<u>.3214</u>	<u>.3205</u>	<b>.3088</b>	<u>.3429</u>	<u>.3426</u>		
	720	<b>.3311</b>	<u>.4013</u>	<u>.4023</u>	<b>.3091</b>	<u>.3737</u>	<u>.3737</u>	<b>.3006</b>	<u>.3358</u>	<u>.3358</u>	<b>.2992</b>	<u>.3415</u>	<u>.3415</u>	<b>.2979</b>	<u>.3588</u>	<u>.3564</u>	<b>.3454</b>	<u>.4172</u>	<u>.4185</u>		
ETTm1	96	<b>.3428</b>	<u>.3687</u>	<u>.3685</u>	<b>.3604</b>	<u>.3918</u>	<u>.3921</u>	<b>.3453</b>	<u>.3575</u>	<u>.3574</u>	<b>.3440</b>	<u>.3573</u>	<u>.3572</u>	<b>.3522</b>	<u>.3633</u>	<u>.3633</u>	<b>.3804</b>	<u>.4252</u>	<u>.4248</u>		
	192	<b>.4076</b>	<u>.4307</u>	<u>.4304</u>	<b>.4171</b>	<u>.4398</u>	<u>.4398</u>	<b>.4158</b>	<u>.4266</u>	<u>.4265</u>	<b>.4125</b>	<u>.4265</u>	<u>.4263</u>	<b>.4173</b>	<u>.4257</u>	<u>.4266</u>	<b>.4402</b>	<u>.4742</u>	<u>.4740</u>		
	336	<b>.4663</b>	<u>.4930</u>	<u>.4949</u>	<b>.4662</b>	<u>.4942</u>	<u>.4944</u>	<b>.4784</b>	<u>.5009</u>	<u>.5005</u>	<b>.4700</b>	<u>.4995</u>	<u>.5006</u>	<b>.4771</b>	<u>.4898</u>	<u>.4886</u>	<b>.4937</b>	<u>.5381</u>	<u>.5381</u>		
	720	<b>.4940</b>	<u>.5614</u>	<u>.5615</u>	<b>.4776</b>	<u>.5359</u>	<u>.5316</u>	<b>.4928</b>	<u>.5640</u>	<u>.5603</u>	<b>.4693</b>	<u>.5638</u>	<u>.5653</u>	<b>.4863</b>	<u>.5414</u>	<u>.5419</u>	<b>.5083</b>	<u>.5830</u>	<u>.5838</u>		
ETTm2	96	<b>.1616</b>	<u>.1672</u>	<u>.1653</u>	<b>.1552</b>	<u>.1591</u>	<u>.1595</u>	<b>.1578</b>	<u>.1594</u>	<u>.1596</u>	<b>.1583</b>	<u>.1596</u>	<u>.1600</u>	<b>.1556</b>	<u>.1592</u>	<u>.1594</u>	<b>.1710</b>	<u>.1712</u>	<u>.1723</u>		
	192	<b>.2172</b>	<u>.2194</u>	<u>.2240</u>	<b>.1989</b>	<u>.2087</u>	<u>.2121</u>	<b>.1900</b>	<u>.1959</u>	<u>.1957</u>	<b>.1919</b>	<u>.1970</u>	<u>.1995</u>	<b>.1908</b>	<u>.1935</u>	<u>.1933</u>	<b>.2102</b>	<u>.2131</u>	<u>.2132</u>		
	336	<b>.2402</b>	<u>.2700</u>	<u>.2762</u>	<b>.2279</b>	<u>.2821</u>	<u>.2543</u>	<b>.2135</b>	<u>.2448</u>	<u>.2561</u>	<b>.2066</b>	<u>.2499</u>	<u>.2477</u>	<b>.2158</b>	<u>.2353</u>	<u>.2420</u>	<b>.2354</b>	<u>.2523</u>	<u>.2546</u>		
	720	<b>.2550</b>	<u>.3461</u>	<u>.3463</u>	<b>.2397</b>	<u>.3011</u>	<u>.3222</u>	<b>.2357</b>	<u>.2985</u>	<u>.2833</u>	<b>.2104</b>	<u>.3179</u>	<u>.2821</u>	<b>.2360</b>	<u>.2732</u>	<u>.2753</u>	<b>.2510</b>	<u>.3020</u>	<u>.3084</u>		
Exchange Rate	96	<b>.0843</b>	<u>.0852</u>	<u>.0852</u>	<b>.0803</b>	<u>.0814</u>	<u>.0815</u>	<b>.0843</b>	<u>.0884</u>	<u>.0885</u>	<b>.0728</b>	<u>.0787</u>	<u>.0787</u>	<b>.0790</b>	<u>.0785</u>	<u>.0785</u>	<b>.0980</b>	<u>.1065</u>	<u>.1068</u>		
	192	<b>.1540</b>	<u>.1606</u>	<u>.1607</u>	<b>.1480</b>	<u>.1544</u>	<u>.1546</u>	<b>.1599</b>	<u>.1653</u>	<u>.1654</u>	<b>.1335</b>	<u>.1507</u>	<u>.1509</u>	<b>.1390</b>	<u>.1500</u>	<u>.1502</u>	<b>.1827</b>	<u>.1870</u>	<u>.1872</u>		
	336	<b>.2588</b>	<u>.2754</u>	<u>.2756</u>	<b>.2611</b>	<u>.2787</u>	<u>.2789</u>	<b>.2565</b>	<u>.2918</u>	<u>.2920</u>	<b>.1833</b>	<u>.2736</u>	<u>.2738</u>	<b>.2511</b>	<u>.2683</u>	<u>.2684</u>	<b>.2660</b>	<u>.3040</u>	<u>.3043</u>		
	720	<b>.5039</b>	<u>.5113</u>	<u>.5115</u>	<b>.4937</b>	<u>.5076</u>	<u>.5078</u>	<b>.5001</b>	<u>.5267</u>	<u>.5270</u>	<b>.3349</b>	<u>.4983</u>	<u>.4986</u>	<b>.4789</b>	<u>.4977</u>	<u>.4980</u>	<b>.4815</b>	<u>.5806</u>	<u>.5809</u>		
Weather	96	<b>.1636</b>	<u>.1726</u>	<u>.1726</u>	<b>.1655</b>	<u>.1731</u>	<u>.1735</u>	<b>.1738</b>	<u>.1931</u>	<u>.1931</u>	<b>.1748</b>	<u>.1934</u>	<u>.1934</u>	<b>.1758</b>	<u>.1831</u>	<u>.1835</u>	<b>.1666</b>	<u>.1739</u>	<u>.1738</u>		
	192	<b>.2073</b>	<u>.2268</u>	<u>.2265</u>	<b>.2052</b>	<u>.2162</u>	<u>.2162</u>	<b>.2217</b>	<u>.2500</u>	<u>.2520</u>	<b>.2144</b>	<u>.2496</u>	<u>.2490</u>	<b>.2151</b>	<u>.2346</u>	<u>.2347</u>	<b>.2060</b>	<u>.2208</u>	<u>.2205</u>		
	336	<b>.2474</b>	<u>.2707</u>	<u>.2707</u>	<b>.2443</b>	<u>.2668</u>	<u>.2660</u>	<b>.2622</b>	<u>.2891</u>	<u>.2885</u>	<b>.2428</b>	<u>.2911</u>	<u>.2897</u>	<b>.2630</b>	<u>.2736</u>	<u>.2735</u>	<b>.2735</b>	<u>.2650</u>	<u>.2648</u>		
	720	<b>.2907</b>	<u>.3118</u>	<u>.3119</u>	<b>.2682</b>	<u>.3088</u>	<u>.3090</u>	<b>.2713</b>	<u>.3310</u>	<u>.3325</u>	<b>.2487</b>	<u>.3327</u>	<u>.3331</u>	<b>.2583</b>	<u>.3268</u>	<u>.3247</u>	<b>.2670</b>	<u>.3140</u>	<u>.3094</u>		
Average	<b>.3121</b>	<u>.3450</u>	<u>.3458</u>	<b>.3032</b>	<u>.3364</u>	<u>.3366</u>	<b>.3097</b>	<u>.3423</u>	<u>.3422</u>	<b>.2990</b>	<u>.3371</u>	<u>.3354</u>	<b>.3046</b>	<u>.3334</u>	<u>.3340</u>	<b>.3284</b>	<u>.3669</u>	<u>.3670</u>			

to form the context vector. Such a mechanism can leverage repeated phases or cycles in datasets with strong periodicity.

Table 13a compares COSA (the standard recent-context construction) with the dynamic context selection method. While the dynamic context selection method achieves clear improvements on datasets such as ETT, where the periodic structure is strong and easily detectable, the overall performance of the original COSA remains superior. Selective Context also introduces non-trivial overhead, since computing importance scores increases both adaptation time and inference time. Moreover, in fully non-stationary settings where distributional characteristics change continuously, older contexts may become outdated and destabilize the correction process. Nonetheless, the observed gains on datasets with pronounced periodicity show the potential of Selective Context. We provide a more detailed discussion of these observations in Section H.

### G.1.3 COMPARISON WITH ENCODER-BASED CONTEXT

To compare with encoder-based context construction, we implemented an alternative approach that replaces the original statistics-based method in COSA with a temporal encoder that directly consumes the previously observed ground-truth sequence from the past 720 steps (the longest  $L$ ). We added RNN-, LSTM-, and Attention-based encoders, each taking the past sequence in the form of  $[720, 1]$  as input and producing a  $[K, 1]$  context vector. The resulting context vector is concatenated with the base model’s prediction output, just as in the original design, and then fed into the linear correction layer. The encoder is seamlessly integrated at the front of COSA, modifying only the context-generation stage while keeping the remaining components unchanged.

Table 14 shows that, except for a few isolated cases, the original statistics-based context (0.3240) performs better than encoder-based alternatives (0.3254, 0.3260, 0.3278). Furthermore, the added architectural complexity increases both adaptation and inference overhead. In non-stationary TTA settings, where the input distribution shifts rapidly and adaptation steps are short, it is difficult for an encoder to learn stable temporal representations. Consequently, the generated embeddings may become misaligned with the current drift direction or overfit to outdated historical patterns, ultimately degrading correction performance.

Table 13: Prediction accuracy and overhead of selective context.

(a) Prediction accuracy.

	Transformer-based		Linear-based				MLP-based						
	iTransformer		DLinear		OLS		FreTS		MICN				
	Recent	Selective	Recent	Selective	Recent	Selective	Recent	Selective	Recent	Selective			
ETTh1	96	.4363	<b>.4362</b>	.4238	<b>.4234</b>	<b>.4482</b>	.4562	.4372	<b>.4359</b>	.4371	<b>.4361</b>	.4684	<b>.4673</b>
	192	<b>.4919</b>	.4927	<b>.4805</b>	.4848	<b>.5050</b>	.5088	<b>.4906</b>	.4933	<b>.4940</b>	.4965	<b>.5328</b>	.5368
	336	<b>.5300</b>	.5367	.5320	<b>.5310</b>	<b>.5456</b>	.5541	.5320	<b>.5368</b>	<b>.5351</b>	.5505	<b>.5878</b>	.5966
	720	<b>.5638</b>	.5671	<b>.5822</b>	.5881	<b>.5896</b>	.6180	<b>.5733</b>	.6093	.5959	<b>.6430</b>	<b>.6504</b>	.7137
ETTh2	96	<b>.2493</b>	.2494	<b>.2343</b>	.2349	.2281	<b>.2258</b>	.2265	<b>.2262</b>	.2350	<b>.2346</b>	<b>.2485</b>	.2486
	192	.2947	<b>.2942</b>	<b>.2608</b>	.2661	.2819	<b>.2813</b>	<b>.2791</b>	.2825	.2824	<b>.2945</b>	<b>.3017</b>	.3027
	336	<b>.3339</b>	.3367	.2978	<b>.2949</b>	.3083	<b>.3064</b>	.3043	<b>.3027</b>	<b>.3153</b>	.3175	<b>.3310</b>	.3328
	720	<b>.3591</b>	.3603	<b>.3428</b>	.3453	.3477	<b>.3462</b>	<b>.3453</b>	.3619	<b>.3399</b>	.3509	<b>.3885</b>	.3906
ETTh1	96	.3455	<b>.3440</b>	<b>.3626</b>	.3627	<b>.3475</b>	<b>.3456</b>	<b>.3475</b>	<b>.3454</b>	.3525	<b>.3522</b>	.3831	<b>.3815</b>
	192	.4140	<b>.4128</b>	<b>.4258</b>	.4278	.4122	<b>.4221</b>	<b>.4119</b>	<b>.4219</b>	<b>.4212</b>	<b>.4212</b>	<b>.4514</b>	.4533
	336	<b>.4643</b>	.4718	.4697	<b>.4693</b>	.4858	<b>.4857</b>	<b>.4749</b>	.4854	<b>.4775</b>	.4805	<b>.5054</b>	.5060
	720	<b>.5102</b>	.5246	.4882	<b>.4868</b>	<b>.4991</b>	.5065	.5007	<b>.5005</b>	.4982	<b>.4970</b>	<b>.5225</b>	.5256
ETTh2	96	.1632	<b>.1631</b>	.1562	<b>.1560</b>	.1586	<b>.1582</b>	.1586	<b>.1581</b>	<b>.1569</b>	.1570	<b>.1704</b>	.1705
	192	.2173	<b>.2165</b>	.2022	<b>.2011</b>	.1905	<b>.1930</b>	.1907	<b>.1952</b>	.1908	<b>.1929</b>	<b>.2120</b>	.2125
	336	<b>.2535</b>	.2551	.2352	<b>.2331</b>	<b>.2242</b>	.2245	.2226	<b>.2212</b>	.2211	<b>.2185</b>	<b>.2351</b>	.2555
	720	.2606	<b>.2578</b>	.2645	<b>.2633</b>	<b>.2316</b>	.2427	<b>.2349</b>	.2368	<b>.2314</b>	.2373	<b>.2643</b>	.2764
Exchange Rate	96	.0837	<b>.0835</b>	<b>.0788</b>	.0789	<b>.0834</b>	<b>.0834</b>	<b>.0773</b>	<b>.0773</b>	.0766	<b>.0763</b>	.1008	<b>.1007</b>
	192	<b>.1479</b>	.1516	.1570	<b>.1554</b>	<b>.1519</b>	.1543	<b>.1457</b>	.1462	<b>.1499</b>	.1514	<b>.1722</b>	.1770
	336	<b>.2624</b>	.2633	<b>.2445</b>	.2478	<b>.2480</b>	.2481	<b>.2323</b>	.2361	<b>.2461</b>	.2541	<b>.2660</b>	.2732
	720	<b>.4460</b>	.4749	<b>.4662</b>	.4983	<b>.4481</b>	.4984	<b>.4541</b>	.4835	<b>.4458</b>	.4914	<b>.4815</b>	.5255
Weather	96	.1617	<b>.1616</b>	.1634	<b>.1631</b>	.1793	<b>.1790</b>	.1803	<b>.1799</b>	.1737	<b>.1731</b>	<b>.1651</b>	.1654
	192	.2088	<b>.2056</b>	<b>.2108</b>	.2109	.2217	<b>.2181</b>	.2237	<b>.2222</b>	.2189	<b>.2187</b>	<b>.2120</b>	.2134
	336	<b>.2515</b>	.2524	<b>.2488</b>	.2554	<b>.2626</b>	.2665	<b>.2642</b>	.2659	<b>.2587</b>	.2603	.2737	<b>.2813</b>
	720	<b>.2730</b>	.2798	<b>.2713</b>	.2798	<b>.2708</b>	.2729	<b>.2708</b>	.2721	<b>.2692</b>	.2732	<b>.2855</b>	.2970

(b) Overhead analysis.

Method	# Params	↓ Adaptation time/batch (ms)	↓ Inference time/batch (ms)	↓ Average MSE
Recent	<b>1,211,287</b>	<b>80.12 ± 13.58</b>	<b>1.25 ± .0984</b>	<b>.3240</b>
Selective	1,212,217	83.64 ± 15.71	1.26 ± .1039	.3287

These findings confirm that the statistics-based context remains the most robust and stable choice for non-stationary adaptation, while encoder-based context generation still demonstrates potential. We discuss these observations in greater detail in Section H.

## G.2 INPUT CALIBRATION EFFECTS

COSA performs residual correction directly in the output space, and we demonstrated that output-only correction is often sufficient. However, in certain time series, we observe that input-level spikes or local noise degrade the base model’s predictions first, and this degradation subsequently propagates to the residual correction stage. To examine how such input perturbations influence the overall correction process, we conducted experiments combining COSA with the input-side GCM module from TAFAS. Table 15 reports the results.

In most cases, output-only correction achieves higher predictive accuracy. However, for datasets such as ETTh1, ETTh2, and ETTm2, where significant input noise is present, the combination with GCM produces improved results. In these cases, input GCM smooths the noisy input patterns, allowing the base model to generate more stable predictions, which in turn enhances the effectiveness of COSA.

Nevertheless, because GCM operates via distribution-shift normalization, it risks oversmoothing or removing meaningful drift signals when the input exhibits rapid or irregular changes. This behavior explains why, on average (in terms of MSE), output-only correction remains more stable across diverse non-stationary scenarios. Overall, when input disturbances are not the primary source of prediction error, output-only correction is the most robust and reliable option.

Table 14: Prediction accuracy and overhead comparison with encoder-based context.

(a) Prediction accuracy.

	Transformer-based												Linear-based						MLP-based								
	iTransformer						DLinear						FreTS														
	COSA	RNN	LSTM	Attn	COSA	RNN	LSTM	Attn	COSA	RNN	LSTM	Attn	COSA	RNN	LSTM	Attn											
ETTh1	96	.4363	.4363	.4238	<b>.4234</b>	.4482	.4562	.4372	<b>.4359</b>	<b>.4371</b>	.4361	.4684	.4674	96	.4363	.4363	.4238	<b>.4234</b>	.4482	.4562	.4372	<b>.4359</b>	<b>.4371</b>	.4361	.4684	.4674	
	192	.4919	.4927	.4805	<b>.4849</b>	.5050	.5088	.4906	<b>.4933</b>	.4940	.4965	.5328	<b>.5368</b>	192	.4919	.4927	.4805	<b>.4849</b>	.5050	.5088	.4906	<b>.4933</b>	.4940	.4965	.5328	<b>.5368</b>	
	336	<b>.5300</b>	.5367	.5320	.5310	<b>.5456</b>	.5541	.5320	.5368	.5351	.5506	.5878	<b>.5967</b>	336	<b>.5300</b>	.5367	.5320	.5310	<b>.5456</b>	.5541	.5320	.5368	.5351	.5506	.5878	<b>.5967</b>	
	720	.5638	.5672	.5822	<b>.5881</b>	<b>.5896</b>	.6180	.5733	<b>.6093</b>	.5959	<b>.6430</b>	<b>.6504</b>	.7137	720	.5638	.5672	.5822	<b>.5881</b>	<b>.5896</b>	.6180	.5733	<b>.6093</b>	.5959	<b>.6430</b>	<b>.6504</b>	.7137	
ETTh2	96	<b>.2493</b>	.2494	.2343	.2349	.2281	.2258	.2265	<b>.2263</b>	.2350	<b>.2346</b>	.2485	.2486	96	<b>.2493</b>	.2494	.2343	.2349	.2281	.2258	.2265	<b>.2263</b>	.2350	<b>.2346</b>	.2485	.2486	
	192	.2947	.2942	<b>.2608</b>	<b>.2661</b>	.2819	.2813	.2791	<b>.2825</b>	<b>.2824</b>	.2946	.3017	.3027	192	.2947	.2942	<b>.2608</b>	<b>.2661</b>	.2819	.2813	.2791	<b>.2825</b>	<b>.2824</b>	.2946	.3017	.3027	
	336	.3339	<b>.3367</b>	.2978	.2950	.3083	.3064	<b>.3043</b>	<b>.3028</b>	.3153	<b>.3175</b>	.3310	.3328	336	.3339	<b>.3367</b>	.2978	.2950	.3083	.3064	<b>.3043</b>	<b>.3028</b>	.3153	<b>.3175</b>	.3310	.3328	
	720	.3591	.3603	<b>.3428</b>	.3453	<b>.3477</b>	.3462	.3453	<b>.3619</b>	.3399	<b>.3509</b>	.3885	<b>.3906</b>	720	.3591	.3603	<b>.3428</b>	.3453	<b>.3477</b>	.3462	.3453	<b>.3619</b>	.3399	<b>.3509</b>	.3885	<b>.3906</b>	
ETTh1	96	.3455	.3440	.3626	<b>.3627</b>	.3475	.3456	.3475	<b>.3454</b>	.3525	.3522	.3831	<b>.3815</b>	96	.3455	.3440	.3626	<b>.3627</b>	.3475	.3456	.3475	<b>.3454</b>	.3525	.3522	.3831	<b>.3815</b>	
	192	.4140	.4128	.4258	<b>.4278</b>	<b>.4122</b>	.4221	.4119	.4219	<b>.4212</b>	.4212	.4514	.4533	192	.4140	.4128	.4258	<b>.4278</b>	<b>.4122</b>	.4221	.4119	.4219	<b>.4212</b>	.4212	.4514	.4533	
	336	.4643	<b>.4718</b>	.4697	<b>.4693</b>	<b>.4858</b>	.4857	.4749	.4854	<b>.4775</b>	.4806	<b>.5054</b>	.5060	336	.4643	<b>.4718</b>	.4697	<b>.4693</b>	<b>.4858</b>	.4857	.4749	.4854	<b>.4775</b>	.4806	<b>.5054</b>	.5060	
	720	<b>.5102</b>	.5246	.4882	.4868	.4991	.5066	<b>.5007</b>	<b>.5005</b>	<b>.4982</b>	.4970	.5225	<b>.5256</b>	720	<b>.5102</b>	.5246	.4882	.4868	.4991	.5066	<b>.5007</b>	<b>.5005</b>	<b>.4982</b>	.4970	.5225	<b>.5256</b>	
ETTh2	96	<b>.1632</b>	.1631	.1562	.1561	<b>.1586</b>	<b>.1582</b>	<b>.1586</b>	<b>.1581</b>	.1569	<b>.1570</b>	.1704	<b>.1705</b>	96	<b>.1632</b>	.1631	.1562	.1561	<b>.1586</b>	<b>.1582</b>	<b>.1586</b>	<b>.1581</b>	.1569	<b>.1570</b>	.1704	<b>.1705</b>	
	192	.2173	.2165	.2022	<b>.2011</b>	.1905	.1930	<b>.1907</b>	<b>.1952</b>	.1908	.1929	.2120	<b>.2125</b>	192	.2173	.2165	.2022	<b>.2011</b>	.1905	.1930	<b>.1907</b>	<b>.1952</b>	.1908	.1929	.2120	<b>.2125</b>	
	336	<b>.2535</b>	.2551	<b>.2352</b>	.2331	.2242	<b>.2245</b>	<b>.2226</b>	.2212	<b>.2211</b>	.2185	<b>.2351</b>	.2555	336	<b>.2535</b>	.2551	<b>.2352</b>	.2331	.2242	<b>.2245</b>	<b>.2226</b>	.2212	<b>.2211</b>	.2185	<b>.2351</b>	.2555	
	720	<b>.2606</b>	<b>.2578</b>	.2645	.2633	<b>.2316</b>	<b>.2427</b>	.2349	.2368	.2314	<b>.2373</b>	<b>.2643</b>	.2764	720	<b>.2606</b>	<b>.2578</b>	.2645	.2633	<b>.2316</b>	<b>.2427</b>	.2349	.2368	.2314	<b>.2373</b>	<b>.2643</b>	.2764	
Exchange Rate	96	<b>.0837</b>	<b>.0835</b>	.0788	.0789	<b>.0834</b>	<b>.0834</b>	.0773	.0773	<b>.0766</b>	<b>.0763</b>	.1008	.1007	96	<b>.0837</b>	<b>.0835</b>	.0788	.0789	<b>.0834</b>	<b>.0834</b>	.0773	.0773	<b>.0766</b>	<b>.0763</b>	.1008	.1007	
	192	<b>.1479</b>	<b>.1517</b>	.1570	.1554	<b>.1519</b>	<b>.1543</b>	.1457	.1462	<b>.1499</b>	<b>.1515</b>	.1722	.1771	192	<b>.1479</b>	<b>.1517</b>	.1570	.1554	<b>.1519</b>	<b>.1543</b>	.1457	.1462	<b>.1499</b>	<b>.1515</b>	.1722	.1771	
	336	<b>.2624</b>	<b>.2633</b>	.2445	.2478	<b>.2480</b>	.2481	<b>.2323</b>	.2361	<b>.2461</b>	<b>.2541</b>	.2660	.2732	336	<b>.2624</b>	<b>.2633</b>	.2445	.2478	<b>.2480</b>	.2481	<b>.2323</b>	.2361	<b>.2461</b>	<b>.2541</b>	.2660	.2732	
	720	<b>.4460</b>	.4749	<b>.4662</b>	.4983	<b>.4481</b>	.4984	<b>.4541</b>	.4835	<b>.4458</b>	.4914	<b>.4815</b>	.5255	720	<b>.4460</b>	.4749	<b>.4662</b>	.4983	<b>.4481</b>	.4984	<b>.4541</b>	.4835	<b>.4458</b>	.4914	<b>.4815</b>	.5255	
Weather	96	.1617	<b>.1617</b>	<b>.1634</b>	.1631	<b>.1793</b>	<b>.1790</b>	<b>.1803</b>	<b>.1799</b>	<b>.1737</b>	.1731	.1651	<b>.1654</b>	96	.1617	<b>.1617</b>	<b>.1634</b>	.1631	<b>.1793</b>	<b>.1790</b>	<b>.1803</b>	<b>.1799</b>	<b>.1737</b>	.1731	.1651	<b>.1654</b>	
	192	<b>.2088</b>	.2056	<b>.2108</b>	.2109	.2217	.2181	<b>.2237</b>	<b>.2222</b>	.2189	<b>.2187</b>	<b>.2120</b>	.2134	192	<b>.2088</b>	.2056	<b>.2108</b>	.2109	.2217	.2181	<b>.2237</b>	<b>.2222</b>	.2189	<b>.2187</b>	<b>.2120</b>	.2134	
	336	<b>.2515</b>	<b>.2524</b>	.2488	.2554	.2626	<b>.2665</b>	<b>.2642</b>	.2659	<b>.2587</b>	.2614	<b>.2737</b>	.2813	336	<b>.2515</b>	<b>.2524</b>	.2488	.2554	.2626	<b>.2665</b>	<b>.2642</b>	.2659	<b>.2587</b>	.2614	<b>.2737</b>	.2813	
	720	<b>.2730</b>	<b>.2798</b>	.2713	.2798	.2708	<b>.2729</b>	<b>.2708</b>	.2721	.2692	<b>.2660</b>	<b>.2855</b>	.2970	720	<b>.2730</b>	<b>.2798</b>	.2713	.2798	.2708	<b>.2729</b>	<b>.2708</b>	.2721	.2692	<b>.2660</b>	<b>.2855</b>	.2970	
PatchTST	COSA	RNN	LSTM	Attn	COSA	RNN	LSTM	Attn	COSA	RNN	LSTM	Attn															
	ETTh1	96	<b>.4363</b>	.4363	.4363	<b>.4357</b>	.4482	.4582	.4582	<b>.4570</b>	.4371	.4375	.4378	<b>.4362</b>	96	<b>.4363</b>	.4363	.4363	<b>.4357</b>	.4482	.4582	.4582	<b>.4570</b>	.4371	.4375	.4378	<b>.4362</b>
		192	.4919	.4813	.4825	<b>.4768</b>	.5050	.5077	.5069	<b>.5029</b>	.4940	.4951	.4953	<b>.4928</b>	192	.4919	.4813	.4825	<b>.4768</b>	.5050	.5077	.5069	<b>.5029</b>	.4940	.4951	.4953	<b>.4928</b>
336		<b>.5300</b>	.5319	.5318	<b>.5305</b>	.5456	<b>.5532</b>	.5576	.5620	<b>.5351</b>	.5445	.5487	<b>.5426</b>	336	<b>.5300</b>	.5319	.5318	<b>.5305</b>	.5456	<b>.5532</b>	.5576	.5620	<b>.5351</b>	.5445	.5487	<b>.5426</b>	
720		.5638	.5625	<b>.5626</b>	<b>.5621</b>	.5896	.6142	<b>.6154</b>	<b>.6120</b>	<b>.5959</b>	.6257	<b>.6257</b>	.6284	720	.5638	.5625	<b>.5626</b>	<b>.5621</b>	.5896	.6142	<b>.6154</b>	<b>.6120</b>	<b>.5959</b>	.6257	<b>.6257</b>	.6284	
ETTh2	96	.2493	<b>.2507</b>	.2503	.2511	.2281	<b>.2304</b>	.2280	<b>.2278</b>	.2350	<b>.2349</b>	<b>.2350</b>	.2355	96	.2493	<b>.2507</b>	.2503	.2511	.2281	<b>.2304</b>	.2280	<b>.2278</b>	.2350	<b>.2349</b>	<b>.2350</b>	.2355	
	192	.2947	<b>.2939</b>	<b>.2938</b>	.2938	.2819	.2813	<b>.2808</b>	<b>.2791</b>	<b>.2824</b>	<b>.2981</b>	.2984	.2981	192	.2947	<b>.2939</b>	<b>.2938</b>	.2938	.2819	.2813	<b>.2808</b>	<b>.2791</b>	<b>.2824</b>	<b>.2981</b>	.2984	.2981	
	336	.3339	<b>.3334</b>	<b>.3348</b>	<b>.3337</b>	.3083	.3040	<b>.3029</b>	<b>.3004</b>	.3153	<b>.3138</b>	<b>.3152</b>	.3186	336	.3339	<b>.3334</b>	<b>.3348</b>	<b>.3337</b>	.3083	.3040	<b>.3029</b>	<b>.3004</b>	.3153	<b>.3138</b>	<b>.3152</b>	.3186	
	720	<b>.3591</b>	.3551	<b>.3534</b>	.3555	.3477	<b>.3495</b>	.3536	<b>.3423</b>	<b>.3399</b>	.3314	<b>.3484</b>	.3386	720	<b>.3591</b>	.3551	<b>.3534</b>	.3555	.3477	<b>.3495</b>	.3536	<b>.3423</b>	<b>.3399</b>	.3314	<b>.3484</b>	.3386	
ETTh1	96	.3455	<b>.3459</b>	.3459	<b>.3453</b>	.3475	.3457	.3457	<b>.3455</b>	<b>.3525</b>	.3529	.3519	<b>.3519</b>	96	.3455	<b>.3459</b>	.3459	<b>.3453</b>	.3475	.3457	.3457	<b>.3455</b>	<b>.3525</b>	.3529	.3519	<b>.3519</b>	
	192	.4140	<b>.4185</b>	<b>.4183</b>	.4131	<b>.4122</b>	<b>.4224</b>	.4223	.4237	.4212	<b>.4215</b>	.4215	<b>.4216</b>	192	.4140	<b>.4185</b>	<b>.4183</b>	.4131	<b>.4122</b>	<b>.4224</b>	.4223	.4237	.4212	<b>.4215</b>	.4215	<b>.4216</b>	
	336	.4643	<b>.4588</b>	<b>.4593</b>	<b>.4550</b>	<b>.4858</b>	.4867	.4869	.4874	.4775	.4778	<b>.4761</b>	<b>.4811</b>	336	.4643	<b>.4588</b>	<b>.4593</b>	<b>.4550</b>	<b>.4858</b>	.4867	.4869	.4874	.4775	.4778	<b>.4761</b>	<b>.4811</b>	
	720	<b>.5102</b>	.5132	<b>.5138</b>	.5142	<b>.4991</b>	.4989	.4962	<b>.4913</b>	<b>.4982</b>	<b>.5043</b>	.5033	.5021	720	<b>.5102</b>	.5132	<b>.5138</b>	.5142	<b>.4991</b>	.4989	.4962	<b>.4913</b>	<b>.4982</b>	<b>.5043</b>	.5033	.5021	
ETTh2	96	.1632	<b>.1642</b>	.1641	<b>.1636</b>	<b>.1586</b>	<b>.1585</b>	<b>.1583</b>	<b>.1583</b>	.1569	.1568	.1570	<b>.1569</b>	96	.1632	<b>.1642</b>	.1641	<b>.1636</b>	<b>.1586</b>	<b>.1585</b>	<b>.1583</b>	<b>.1583</b>	.1569	.1568	.1570	<b>.1569</b>	
	192	.2173	.2174	<b>.2175</b>	<b>.2158</b>	.1905	.1951	<b>.1939</b>	<b>.1932</b>	.1908	<b>.1938</b>	<b>.1922</b>	<b>.1922</b>	192	.2173	.2174	<b>.2175</b>	<b>.2158</b>	.1905	.1951	<b>.1939</b>	<b>.1932</b>	.1908	<b>.1938</b>	<b>.1922</b>	<b>.1922</b>	
	336	<b>.2535</b>	.2536	<b>.2521</b>	.2544	.2242	<b>.2214</b>	<b>.2208</b>	<b>.2230</b>	<b>.2211</b>	.2217	.2214	<b>.2183</b>	336	<b>.2535</b>	.2536	<b>.2521</b>	.2544	.2242	<b>.2214</b>	<b>.2208</b>	<b>.2230</b>	<b>.2211</b>	.2217	.2214	<b>.2183</b>	
	720	<b>.2606</b>	<b>.2633</b>	.2641	.2653	.2316	<b>.2284</b>	<b>.2359</b>	.2383	<b>.2314</b>	.2296	.2280	<b>.2305</b>	720	<b>.2606</b>	<b>.2633</b>	.2641	.2653	.2316	<b>.2284</b>	<b>.2359</b>	.2383	<b>.2314</b>	.2296	.2280	<b>.2305</b>	
Exchange Rate	96	<b>.0837</b>	.0851	<b>.0853</b>	.0862	<b>.0834</b>	<b>.0833</b>	.0837	.0849	<b>.0766</b>	<b>.0774</b>	.0779	.0789	96	<b>.0837</b>	.0851	<b>.0853</b>	.0862	<b>.0834</b>	<b>.0833</b>	.0837	.0849	<b>.0766</b>	<b>.0774</b>	.0779	.0789	
	192	<b>.1479</b>	<b>.1499</b>	.1544	.1605	<b>.1519</b>	<b>.1550</b>	.1562	.1597	<b>.1499</b>	<b>.1548</b>	.1564	.1605	192	<b>.1479</b>	<b>.1499&lt;/</b>											

Table 15: COSA with input GCM.

		Transformer-based				Linear-based				MLP-based			
		iTransformer		PatchTST		DLinear		OLS		FreTS		MICN	
		COSA	w. Input	COSA	w. Input	COSA	w. Input	COSA	w. Input	COSA	w. Input	COSA	w. Input
ETh1	96	.4363	<b>.4362</b>	.4238	<b>.4209</b>	.4482	<b>.4460</b>	.4372	<b>.4370</b>	.4371	<b>.4343</b>	<b>.4684</b>	.4852
	192	.4919	<b>.4821</b>	.4805	<b>.4790</b>	.5050	<b>.4995</b>	<b>.4906</b>	.4907	.4940	<b>.4919</b>	<b>.5328</b>	.5530
	336	<b>.5300</b>	.5386	.5320	<b>.5216</b>	.5456	<b>.5359</b>	.5320	<b>.5193</b>	.5351	<b>.5286</b>	<b>.5878</b>	.5953
	720	<b>.5638</b>	.6287	<b>.5822</b>	.6386	<b>.5896</b>	.6513	<b>.5733</b>	.6500	<b>.5959</b>	.6726	<b>.6504</b>	.7509
ETh2	96	.2493	<b>.2001</b>	.2343	<b>.1845</b>	.2281	<b>.1819</b>	.2265	<b>.1803</b>	.2350	<b>.1849</b>	.2485	<b>.1995</b>
	192	.2947	<b>.2384</b>	.2608	<b>.2287</b>	.2819	<b>.2217</b>	.2791	<b>.2292</b>	.2824	<b>.2241</b>	.3017	<b>.2457</b>
	336	.3339	<b>.2528</b>	.2978	<b>.2405</b>	.3083	<b>.2506</b>	.3043	<b>.2474</b>	.3153	<b>.2539</b>	.3310	<b>.2755</b>
	720	.3591	<b>.3048</b>	.3428	<b>.2913</b>	.3477	<b>.2932</b>	.3453	<b>.2940</b>	.3399	<b>.2954</b>	.3885	<b>.3411</b>
ETm1	96	<b>.3455</b>	.3676	<b>.3626</b>	.3903	<b>.3475</b>	.3518	<b>.3475</b>	.3516	<b>.3525</b>	.3596	<b>.3831</b>	.4216
	192	<b>.4140</b>	.4273	<b>.4258</b>	.4396	.4122	<b>.4194</b>	.4119	<b>.4192</b>	.4212	<b>.4200</b>	<b>.4514</b>	.4742
	336	<b>.4643</b>	.4963	<b>.4697</b>	.4881	<b>.4858</b>	.4924	<b>.4749</b>	.4926	<b>.4775</b>	.4885	<b>.5054</b>	.5365
	720	<b>.5102</b>	.5869	<b>.4882</b>	.5317	<b>.4991</b>	.5678	<b>.5007</b>	.5598	<b>.4982</b>	.5452	<b>.5225</b>	.5920
ETm2	96	.1632	<b>.1242</b>	.1562	<b>.1200</b>	.1586	<b>.1217</b>	.1586	<b>.1218</b>	.1569	<b>.1203</b>	.1704	<b>.1288</b>
	192	.2173	<b>.1673</b>	.2022	<b>.1551</b>	.1905	<b>.1502</b>	.1907	<b>.1501</b>	.1908	<b>.1491</b>	.2120	<b>.1618</b>
	336	.2535	<b>.2017</b>	.2352	<b>.1900</b>	.2242	<b>.1813</b>	.2226	<b>.1831</b>	.2211	<b>.1823</b>	.2351	<b>.1975</b>
	720	.2606	<b>.2554</b>	.2645	<b>.2468</b>	<b>.2316</b>	.2465	<b>.2349</b>	.2531	<b>.2314</b>	.2364	.2643	<b>.2635</b>
Exchange Rate	96	<b>.0837</b>	.0859	<b>.0788</b>	.0835	<b>.0834</b>	.0883	<b>.0773</b>	.0789	<b>.0766</b>	.0798	<b>.1008</b>	.1105
	192	<b>.1479</b>	.1706	<b>.1570</b>	.1754	<b>.1519</b>	.1807	<b>.1457</b>	.1644	<b>.1499</b>	.1660	<b>.1722</b>	.2083
	336	<b>.2624</b>	.2974	<b>.2445</b>	.3017	<b>.2480</b>	.2986	<b>.2323</b>	.2937	<b>.2461</b>	.2951	<b>.2660</b>	.3358
	720	<b>.4460</b>	.5884	<b>.4662</b>	.5969	<b>.4481</b>	.6107	<b>.4541</b>	.5837	<b>.4458</b>	.5823	<b>.4815</b>	.6885
Weather	96	<b>.1617</b>	.1723	<b>.1634</b>	.1728	<b>.1793</b>	.1924	<b>.1803</b>	.1927	<b>.1737</b>	.1822	<b>.1651</b>	.1741
	192	<b>.2088</b>	.2214	<b>.2108</b>	.2181	<b>.2217</b>	.2363	<b>.2237</b>	.2367	<b>.2189</b>	.2269	<b>.2120</b>	.2289
	336	<b>.2515</b>	.2752	<b>.2488</b>	.2809	<b>.2626</b>	.2951	<b>.2642</b>	.2946	<b>.2587</b>	.2845	<b>.2737</b>	.2728
	720	<b>.2730</b>	.3274	<b>.2713</b>	.3316	<b>.2708</b>	.3356	<b>.2708</b>	.3361	<b>.2692</b>	.3298	<b>.2855</b>	.3364

Table 16: Prediction accuracy comparison of single linear adapter and 2-layer MLP adapter.

		Transformer-based				Linear-based				MLP-based				Transformer-based				Linear-based				MLP-based			
		iTransformer		PatchTST		DLinear		OLS		FreTS		MICN		iTransformer		PatchTST		DLinear		OLS		FreTS		MICN	
		Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP	Linear	MLP
ETh1	96	<b>.4363</b>	.4376	<b>.4238</b>	.4267	<b>.4482</b>	.4644	<b>.4372</b>	.4460	<b>.4371</b>	.4412	<b>.4684</b>	.4841	<b>2.35</b>	3.79	<b>2.35</b>	3.83	<b>2.24</b>	3.45	<b>2.15</b>	3.10	<b>2.14</b>	3.49	<b>2.38</b>	3.84
	192	<b>.4919</b>	.4968	<b>.4805</b>	.4857	<b>.5050</b>	.5135	<b>.4906</b>	.4980	<b>.4940</b>	.4967	<b>.5328</b>	.5496	<b>2.39</b>	3.73	<b>2.44</b>	3.83	<b>2.23</b>	3.49	<b>2.21</b>	3.42	<b>2.16</b>	3.47	<b>2.59</b>	3.83
	336	<b>.5300</b>	.5471	.5320	<b>.5294</b>	<b>.5456</b>	.5501	.5320	<b>.5282</b>	.5351	<b>.5346</b>	<b>.5878</b>	.6067	<b>2.40</b>	3.68	<b>2.43</b>	3.77	<b>2.22</b>	3.53	<b>2.18</b>	3.45	<b>2.32</b>	3.40	<b>2.51</b>	3.91
	720	<b>.5638</b>	.6167	<b>.5822</b>	.6350	<b>.5896</b>	.6498	<b>.5733</b>	.6351	<b>.5959</b>	.6558	<b>.6504</b>	.7325	<b>2.39</b>	3.46	<b>2.25</b>	3.49	<b>2.17</b>	3.10	<b>2.23</b>	3.22	<b>2.24</b>	3.11	<b>3.69</b>	4.57
ETh2	96	<b>.2493</b>	.2517	.2343	<b>.2339</b>	<b>.2281</b>	.2299	<b>.2265</b>	.2269	<b>.2350</b>	<b>.2485</b>	.2517	<b>2.38</b>	3.76	<b>2.36</b>	3.63	<b>2.19</b>	3.49	<b>2.19</b>	3.43	<b>1.54</b>	3.47	<b>2.37</b>	3.81	
	192	<b>.2947</b>	.2981	<b>.2608</b>	.2825	<b>.2819</b>	.2828	.2791	<b>.2781</b>	.2824	<b>.2767</b>	<b>.3017</b>	.3094	<b>2.40</b>	3.74	<b>2.44</b>	3.80	<b>2.28</b>	3.51	<b>2.20</b>	3.41	<b>2.24</b>	3.43	<b>2.66</b>	3.85
	336	.3339	<b>.3306</b>	<b>.2978</b>	.3127	<b>.3083</b>	.3106	<b>.3043</b>	.3113	<b>.3153</b>	.3168	<b>.3310</b>	.3484	<b>2.48</b>	3.71	<b>2.45</b>	3.77	<b>2.25</b>	3.49	<b>2.17</b>	3.38	<b>2.30</b>	3.47	<b>2.87</b>	3.95
	720	<b>.3591</b>	.3931	<b>.3428</b>	.3837	<b>.3477</b>	.3689	<b>.3453</b>	.3750	<b>.3399</b>	.3660	<b>.3885</b>	.4170	<b>2.34</b>	3.40	<b>2.35</b>	3.48	<b>2.21</b>	3.33	<b>2.18</b>	3.24	<b>2.19</b>	3.19	<b>3.72</b>	4.62
ETm1	96	<b>.3455</b>	.3522	<b>.3626</b>	.3680	<b>.3475</b>	.3505	<b>.3475</b>	.3489	<b>.3525</b>	.3577	<b>.3831</b>	<b>.3914</b>	<b>8.49</b>	14.08	<b>8.56</b>	14.24	<b>8.14</b>	11.50	<b>8.17</b>	13.16	<b>8.22</b>	13.22	<b>8.33</b>	14.38
	192	<b>.4140</b>	<b>.4103</b>	.4258	<b>.4246</b>	<b>.4122</b>	.4147	<b>.4119</b>	.4137	.4212	<b>.4168</b>	<b>.4514</b>	<b>.4470</b>	<b>8.72</b>	14.41	<b>8.67</b>	14.31	<b>8.35</b>	13.25	<b>5.68</b>	13.38	<b>8.57</b>	13.68	<b>9.27</b>	14.39
	336	.4643	<b>.4625</b>	<b>.4697</b>	.4709	.4858	<b>.4678</b>	<b>.4749</b>	.4680	.4775	<b>.4708</b>	<b>.5054</b>	<b>.4784</b>	<b>9.07</b>	14.62	<b>9.03</b>	12.04	<b>8.72</b>	13.59	<b>8.86</b>	13.31	<b>8.87</b>	14.17	<b>10.86</b>	15.26
	720	<b>.5102</b>	<b>.4927</b>	.4882	<b>.4651</b>	.4991	.4763	.5007	.4774	.4982	<b>.4814</b>	<b>.5225</b>	<b>.4888</b>	<b>9.64</b>	15.32	<b>8.72</b>	12.74	<b>9.32</b>	14.00	<b>9.36</b>	14.31	<b>9.60</b>	14.62	<b>16.05</b>	20.23
ETm2	96	<b>.1632</b>	<b>.1632</b>	<b>.1562</b>	.1569	<b>.1586</b>	.1587	<b>.1586</b>	.1587	<b>.1569</b>	.1571	<b>.1704</b>	<b>.1703</b>	<b>8.51</b>	11.73	<b>8.72</b>	14.41	<b>8.34</b>	13.23	<b>8.06</b>	13.04	<b>8.45</b>	9.13	<b>8.25</b>	14.35
	192	.2173	<b>.2140</b>	.2022	<b>.1978</b>	.1905	<b>.1880</b>	.1907	<b>.1884</b>	.1908	<b>.1850</b>	.2120	<b>.2039</b>	<b>8.75</b>	14.35	<b>8.77</b>	14.38	<b>8.32</b>	13.31	<b>5.27</b>	13.28	<b>8.51</b>	13.65	<b>9.32</b>	14.42
	336	.2535	<b>.2412</b>	.2352	<b>.2283</b>	.2242	<b>.2103</b>	.2226	<b>.2142</b>	.2211	<b>.2092</b>	.2351	<b>.2331</b>	<b>8.10</b>	14.64	<b>9.10</b>	14.56	<b>8.84</b>	13.50	<b>8.89</b>	13.57	<b>8.86</b>	13.64	<b>10.80</b>	15.24
	720	<b>.2606</b>	.2759	<b>.2645</b>	.2711	<b>.2316</b>	.2532	<b>.2349</b>	.2551	<b>.2314</b>	.2477	<b>.2643</b>	.2709	<b>7.74</b>	15.14	<b>9.51</b>	15.04	<b>9.29</b>	13.86	<b>9.32</b>	14.26	<b>9.50</b>	14.48	<b>16.04</b>	20.21
Exchange Rate	96	<b>.0837</b>	.0869	<b>.0788</b>	.0838	<b>.0834</b>	.0890	<b>.0773</b>	.0789	<b>.0766</b>	.0803	<b>.1008</b>	.1106	<b>1.38</b>	2.01	<b>1.38</b>	1.99	<b>1.22</b>	1.79	<b>1.20</b>	1.77	<b>1.26</b>	1.81	<b>1.39</b>	2.00
	192	<b>.1479</b>	.1730	<b>.1570</b>	.1696	<b>.1519</b>	.1761	<b>.1457</b>	.1570	<b>.1499</b>	.1603	<b>.1722</b>	.1956	<b>1.31</b>	1.89	<b>1.36</b>	1.88	<b>1.17</b>	1.68	<b>1.15</b>	1.69	<b>1.17</b>	1.70	<b>1.42</b>	1.96
	336	<b>.2624</b>	.3103	<b>.2445</b>	.2941	<b>.2480</b>	.2872	<b>.2323</b>	.2856	<b>.2461</b>	.2910	<b>.2660</b>	.3399	<b>1.26</b>	1.70	<b>1.32</b>	1.81	<b>1.09</b>	1.57	<b>1.10</b>	1.58	<b>1.10</b>	1.62	<b>1.48</b>	1.60
	720	<b>.4460</b>	.8099	<b>.4662</b>	.8268	<b>.4481</b>	.8373	<b>.4541</b>	.7980	<b>.4458</b>	.7983	<b>.4815</b>	.9665	<b>1.09</b>	1.47	<b>1.09</b>	1.43	<b>.91</b>	1.24	<b>.90</b>	1.26	<b>.93</b>	1.27	<b>1.46</b>	1.76
Weather	96	<b>.1617</b>	.1691	<b>.1634</b>	.1691	<b>.1793</b>	.1904	<b>.1803</b>	.1904	<b>.1737</b>	.1767	<b>.1651</b>	.1703	<b>13.05</b>	20.86	<b>15.19</b>	23.35	<b>12.81</b>	20.34	<b>12.76</b>	18.41	<b>12.84</b>	20.90	<b>12.88</b>	20.56
	192	.2088	<b>.2045</b>	.2108	<b>.2026</b>	<b>.2217</b>	.2285	<b>.2237</b>	.2189	<b>.2120</b>	.2120	<b>.2070</b>	<b>13.33</b>	20.79	<b>15.33</b>	23.72	<b>13.10</b>	20.43	<b>12.89</b>	20.55	<b>13.12</b>	20.72	<b>13.92</b>	20.44	
	336	.2515	<b>.2390</b>	.2488	<b>.2360</b>	.2626	<b>.2469</b>	.2642	<b>.2499</b>	.2587	<b>.2440</b>	<b>.2737</b>	<b>.2386</b>	<b>13.69</b>	21.12	<b>15.93</b>	24.21	<b>13.49</b>	20.83	<b>13.43</b>	20.98	<b>13.53</b>	20.96	<b>14.31</b>	22.14
	720	<b>.2730</b>	.2752	.2713	<b>.2660</b>	<b>.2708</b>	.2764	<b>.2708</b>	.2785	<b>.2692</b>	.2737	<b>.2855</b>	<b>.2582</b>	<b>14.68</b>	17.06	<b>16.71</b>	25.44	<b>14.05</b>	21.26	<b>14.36</b>	21.61	<b>14.41</b>	21.81	<b>18.09</b>	26.57
Average	<b>.3218</b>	.3438	<b>.3166</b>	.3383	<b>.3196</b>	.3425	<b>.3158</b>	.3368	<b>.3176</b>	.3369	<b>.3421&lt;/</b>														

#### G.4 COMPARISON WITH VARIOUS LEARNING-RATE SCHEDULERS

To validate the effectiveness of CALR, we compared it against several official PyTorch learning-rate schedulers: CosineAnnealingLR, ExponentialLR, ReduceLRonPlateau, StepLR, and a fixed learning rate. All schedulers were configured with the same base learning rate of 0.005 for a fair comparison. One-Cycle, although widely used, was excluded because it requires a predefined learning-rate schedule; in a streaming TTA scenario where samples arrive continuously and the batch size changes dynamically under PAAS, such predefinition is not feasible.

As summarized in Table 17, each scheduler achieves improvements in some individual cases. However, when results are averaged across all datasets and prediction lengths, the proposed CALR achieves the best or second-best accuracy in the vast majority of settings. These findings support CALR’s stability-induced design and its suitability for non-stationary TTA environments.

#### G.5 EXTENTION TO MULTIVARIATE TIME-SERIES FORECASTING

COSA is originally introduced as an output-space residual correction module, treating each variable as an independent univariate forecasting task, for fair comparison with existing SOTA methods that assume univariate forecasting. However, in real multivariate time-series settings, correlations among variables may influence drift patterns, suggesting that modeling cross-variable interactions could potentially benefit COSA. Basically, COSA can be extended to multivariate forecasting; to examine this possibility, we implemented it.

We first incorporate *Cross-Variable Context Attention*, allowing the context of each variable to reference information from other variables and thereby capture correlation-driven contextual interactions. Additionally, we introduce a mixed structure composed of a low-rank shared component and variable-specific components: the shared component captures drift patterns common across variables in an efficient low-dimensional form, while the specific components model idiosyncratic behavior unique to each variable. These two components are combined through a learnable mixing coefficient that automatically balances global and variable-specific contributions.

Table 18a presents the comparison results. For datasets where meaningful cross-variable dependencies exist, the multivariate structure achieves higher predictive accuracy than the univariate version of COSA. However, for datasets with weak inter-variable correlations, such as the Exchange Rate, the univariate structure remains more stable. In such cases, the shared component struggles to learn useful common patterns, which can lead to performance degradation. Moreover, as shown in Table 18b, the additional complexity leads to increased adaptation time and inference latency.

These results indicate that multivariate-based correlation modeling can indeed provide accuracy gains, but further design improvements are required for effective deployment under TTA constraints. We discuss these limitations and potential future directions in Section H.

#### G.6 EXTENSION TO VECTOR GATING

To evaluate whether finer-grained control over correction strength could provide additional benefits, we implemented an element-wise gating vector as an extension of the scalar gating mechanism in COSA. This vector shares the same dimensionality as the prediction length, allowing each time step within the prediction window to modulate its correction intensity independently. Such a design is intended to handle scenarios where drift occurs in only a specific portion of the horizon.

However, as shown in Table 19a, vector gating yields degraded overall accuracy compared to the original scalar gating, and also exhibits reduced stability. We attribute this performance degradation to the propagation of local noise: a noise spike at a particular horizon position can influence the entire gating vector over successive adaptation steps, causing cumulative negative impact throughout the correction process. In contrast, scalar gating provides consistent batch-level modulation that effectively bounds the influence of noise and maintains stable behavior across adaptation windows.

#### G.7 VISUALIZATION OF NORMALIZATION

Figure 7 visualizes the learned weights of the linear layer in COSA with and without representative time-series normalization modules (RevIN (Kim et al., 2021) and DDN (Dai et al., 2024)).

Table 17: Prediction accuracy comparison with various learning rate schedulers.

	Transformer-based												Linear-based				MLP-based			
	iTransformer						DLinear						FreTS							
	CALR	Cosine	Exp	Fixed	Plateau	Step	CALR	Cosine	Exp	Fixed	Plateau	Step	CALR	Cosine	Exp	Fixed	Plateau	Step		
ETT <sub>h1</sub>	96	<b>.4363</b>	.4858	.4961	.4848	.4848	.4885	<b>.4482</b>	.4498	.4556	.4513	.4513	.4537	<b>.4371</b>	.4400	.4441	.4398	.4398	.4405	
	192	<b>.4919</b>	.5405	.5613	.5413	.5413	.5426	<b>.5050</b>	<b>.5069</b>	.5183	.5090	.5090	.5110	.4940	<b>.4914</b>	.5002	<b>.4915</b>	<b>.4915</b>	.4931	
	336	<b>.5300</b>	.5741	.5952	.5751	.5751	.5826	<b>.5456</b>	<b>.5560</b>	.5714	.5586	.5586	.5642	<b>.5351</b>	<b>.5429</b>	.5570	.5442	.5442	.5494	
	720	<b>.5638</b>	.6216	.6835	<b>.6163</b>	<b>.6163</b>	.6336	<b>.5896</b>	.6033	.6511	<b>.5953</b>	<b>.5953</b>	.6166	<b>.5959</b>	.6116	.6593	<b>.6041</b>	<b>.6041</b>	.6200	
ETT <sub>h2</sub>	96	<b>.2493</b>	.2984	.3021	.2984	.2984	.2987	<b>.2281</b>	.2308	<b>.2287</b>	.2321	.2321	.2333	<b>.2350</b>	.2378	<b>.2359</b>	.2376	.2376	.2372	
	192	<b>.2947</b>	.3386	.3520	.3399	.3400	.3394	<b>.2819</b>	.2823	<b>.2811</b>	.2836	.2836	.2845	<b>.2824</b>	.2837	.2829	.2834	.2834	.2831	
	336	<b>.3339</b>	.3823	.3992	.3856	.3856	<b>.3802</b>	.3083	<b>.3078</b>	.3093	.3104	.3105	<b>.3081</b>	.3153	<b>.3115</b>	.3151	.3131	.3131	<b>.3087</b>	
	720	<b>.3591</b>	.4094	.4278	.4113	.4113	<b>.4070</b>	.3477	.3490	.3562	<b>.3462</b>	<b>.3463</b>	.3488	.3399	.3421	.3524	.3391	<b>.3389</b>	<b>.3387</b>	
ETT <sub>m1</sub>	96	<b>.3455</b>	.4006	.4087	<b>.3978</b>	<b>.3978</b>	.4048	<b>.3475</b>	.3501	.3610	<b>.3497</b>	<b>.3497</b>	.3569	<b>.3525</b>	.3548	.3567	<b>.3539</b>	<b>.3539</b>	.3555	
	192	<b>.4170</b>	.4650	.4772	<b>.4630</b>	<b>.4630</b>	.4683	<b>.4122</b>	<b>.4140</b>	.4283	.4146	.4146	.4191	<b>.4212</b>	.4236	.4279	.4231	<b>.4230</b>	.4246	
	336	<b>.4643</b>	.5110	.5286	<b>.5100</b>	<b>.5100</b>	.5160	.4858	<b>.4828</b>	.5028	.4838	.4838	.4875	.4775	<b>.4740</b>	.4821	<b>.4740</b>	<b>.4740</b>	.4750	
	720	<b>.5102</b>	<b>.5321</b>	.5732	.5350	.5349	.5392	.4991	<b>.4861</b>	.5153	<b>.4889</b>	<b>.4889</b>	.4915	.4982	<b>.4923</b>	.5073	.4950	.4950	.4951	
ETT <sub>m2</sub>	96	<b>.1632</b>	.2138	.2140	.2139	.2139	<b>.2136</b>	<b>.1586</b>	.1621	<b>.1591</b>	.1629	.1629	.1646	<b>.1569</b>	.1611	<b>.1572</b>	.1604	.1604	.1606	
	192	<b>.2173</b>	.2697	.2705	.2692	.2693	.2704	<b>.1905</b>	.1948	<b>.1916</b>	.1955	.1955	.1956	.1973	<b>.1908</b>	.1958	<b>.1921</b>	.1951	.1951	
	336	<b>.2535</b>	<b>.3028</b>	.3125	.3047	.3045	<b>.3065</b>	<b>.2242</b>	.2305	<b>.2268</b>	.2310	.2310	.2324	<b>.2211</b>	.2285	<b>.2247</b>	.2276	.2276	.2278	
	720	<b>.2606</b>	.3000	.3159	.2987	.2986	<b>.2966</b>	<b>.2316</b>	.2369	.2383	.2405	.2405	<b>.2348</b>	<b>.2314</b>	.2388	.2397	.2405	.2403	<b>.2334</b>	
Exchange Rate	96	<b>.0837</b>	.1335	.1344	.1335	.1335	<b>.1334</b>	<b>.0834</b>	.0865	<b>.0848</b>	.0873	.0873	.0886	<b>.0766</b>	.0800	<b>.0775</b>	.0794	.0794	.0792	
	192	<b>.1479</b>	.1961	.2025	.1960	.1960	<b>.1952</b>	<b>.1519</b>	.1523	<b>.1552</b>	.1533	.1533	.1542	<b>.1499</b>	.1518	.1545	.1513	.1513	<b>.1505</b>	
	336	<b>.2624</b>	.3120	.3495	.3116	.3117	<b>.3102</b>	<b>.2480</b>	<b>.2521</b>	.2710	.2528	.2527	.2535	<b>.2461</b>	.2501	.2735	.2491	.2492	<b>.2484</b>	
720	<b>.4460</b>	.4481	.6222	.4498	<b>.4460</b>	<b>.4448</b>	.4481	.4185	.5919	<b>.4164</b>	<b>.4162</b>	.4179	.4458	.4128	.5679	.4072	<b>.4063</b>	<b>.4068</b>		
Weather	96	<b>.1617</b>	.2118	.2151	<b>.2115</b>	<b>.2115</b>	.2123	<b>.1793</b>	<b>.1824</b>	.1827	.1827	.1827	.1858	<b>.1737</b>	.1770	.1767	<b>.1761</b>	<b>.1761</b>	.1776	
	192	<b>.2088</b>	.2514	.2592	.2517	.2517	<b>.2517</b>	<b>.2217</b>	<b>.2226</b>	.2270	.2232	.2232	.2259	<b>.2189</b>	.2201	.2235	.2191	.2192	.2204	
	336	<b>.2515</b>	<b>.2916</b>	.3063	.2949	.2948	.2920	.2626	<b>.2510</b>	.2622	<b>.2532</b>	<b>.2532</b>	.2541	<b>.2587</b>	<b>.2497</b>	.2606	.2505	.2505	<b>.2496</b>	
720	<b>.2730</b>	.3230	.3460	<b>.3225</b>	<b>.3225</b>	.3267	<b>.2708</b>	<b>.2774</b>	.2965	.2802	.2802	.2810	<b>.2692</b>	.2726	.2921	.2755	.2755	<b>.2725</b>		
ETT <sub>h1</sub>	PatchTST				OLS				MICN											
	CALR	Cosine	Exp	Fixed	Plateau	Step	CALR	Cosine	Exp	Fixed	Plateau	Step	CALR	Cosine	Exp	Fixed	Plateau	Step		
	96	<b>.4238</b>	.4723	.4294	.5723	.4923	<b>.4232</b>	<b>.4372</b>	<b>.4375</b>	.4448	.4401	.4401	.4413	<b>.4684</b>	.4739	.4904	<b>.4708</b>	<b>.4708</b>	.4769	
192	<b>.4805</b>	.5307	.4958	.6301	.5501	<b>.4838</b>	<b>.4906</b>	<b>.4912</b>	.5041	.4944	.4944	.4951	.5328	.5320	.5709	<b>.5297</b>	<b>.5297</b>	.5390		
336	<b>.5320</b>	.5851	.5535	.6826	.6026	<b>.5396</b>	<b>.5320</b>	<b>.5342</b>	.5515	.5384	.5384	.5424	.5878	<b>.5864</b>	.6417	<b>.5861</b>	<b>.5861</b>	.5974		
720	<b>.5822</b>	.6412	.6553	.7330	.6533	<b>.6031</b>	<b>.5733</b>	.5872	.6371	<b>.5808</b>	<b>.5808</b>	.5987	<b>.6504</b>	.6726	.8268	<b>.6506</b>	<b>.6506</b>	.6977		
ETT <sub>h2</sub>	96	<b>.2343</b>	.2840	.2352	.3839	.3039	<b>.2345</b>	<b>.2265</b>	.2275	<b>.2268</b>	.2301	.2300	.2296	<b>.2485</b>	.2499	.2499	.2521	.2521	<b>.2466</b>	
	192	<b>.2608</b>	.3114	.2641	.4092	.3292	<b>.2617</b>	<b>.2791</b>	<b>.2804</b>	<b>.2804</b>	.2828	.2828	.2823	.3017	<b>.2990</b>	.3077	.3025	.3026	<b>.2960</b>	
	336	<b>.2978</b>	.3514	.3086	.4543	.3743	<b>.2993</b>	.3043	<b>.3030</b>	.3064	.3063	.3063	<b>.3029</b>	.3310	<b>.3303</b>	.3399	.3388	.3388	<b>.3243</b>	
	720	<b>.3428</b>	.3913	.3663	.4951	.4151	<b>.3485</b>	<b>.3453</b>	.3487	.3589	.3473	<b>.3472</b>	.3476	<b>.3885</b>	.3965	.4019	.3929	.3929	<b>.3884</b>	
ETT <sub>m1</sub>	96	<b>.3626</b>	.4104	.3741	.5183	.4383	<b>.3655</b>	<b>.3475</b>	<b>.3487</b>	.3608	.3495	.3495	.3552	<b>.3831</b>	.3912	.3982	.3900	.3900	.3913	
	192	<b>.4258</b>	.4747	.4427	.5764	.4964	<b>.4315</b>	<b>.4119</b>	<b>.4125</b>	.4283	.4142	.4142	.4174	<b>.4514</b>	.4533	.4641	<b>.4524</b>	<b>.4524</b>	.4541	
	336	<b>.4697</b>	.5149	.4825	.6203	.5403	<b>.4700</b>	.4749	<b>.4706</b>	.4920	<b>.4728</b>	<b>.4728</b>	.4751	.5054	<b>.5053</b>	.5249	<b>.5048</b>	<b>.5048</b>	.5079	
	720	<b>.4882</b>	.5349	.5077	.6391	.5550	<b>.4877</b>	.5007	<b>.4831</b>	.5136	<b>.4873</b>	<b>.4873</b>	.4882	.5225	<b>.5152</b>	.5443	.5188	.5188	<b>.5185</b>	
ETT <sub>m2</sub>	96	<b>.1562</b>	.2065	.1566	.3064	.2265	<b>.1566</b>	<b>.1586</b>	.1608	<b>.1590</b>	.1627	.1627	.1630	<b>.1704</b>	.1742	.1709	.1762	.1763	<b>.1709</b>	
	192	<b>.2022</b>	.2527	.2056	.3532	.2732	<b>.2034</b>	<b>.1907</b>	.1933	.1914	.1954	.1953	.1956	<b>.2120</b>	.2154	.2138	.2180	.2180	<b>.2122</b>	
	336	<b>.2352</b>	.2863	.2367	.3863	.3063	<b>.2363</b>	<b>.2226</b>	.2278	<b>.2256</b>	.2298	.2298	.2298	<b>.2351</b>	.2383	.2365	.2408	.2408	<b>.2354</b>	
	720	<b>.2645</b>	.3145	.2692	.4114	.3314	<b>.2619</b>	<b>.2349</b>	.2390	.2417	.2445	.2445	<b>.2367</b>	.2643	<b>.2628</b>	.2672	.2647	.2647	<b>.2565</b>	
Exchange Rate	96	<b>.0788</b>	.1278	.0802	.2278	.1478	<b>.0776</b>	<b>.0773</b>	.0789	<b>.0780</b>	.0809	.0809	.0808	<b>.1008</b>	.1031	.1027	.1051	.1051	<b>.0994</b>	
	192	<b>.1570</b>	.2051	.1642	.3052	.2251	<b>.1544</b>	<b>.1457</b>	<b>.1452</b>	.1490	.1474	.1473	.1466	.1722	<b>.1715</b>	.1779	.1735	.1735	<b>.1670</b>	
	336	<b>.2445</b>	.2966	.2780	.3962	.3164	<b>.2448</b>	<b>.2323</b>	<b>.2331</b>	.2578	.2349	.2348	.2334	<b>.2660</b>	.2730	.3114	.2745	.2745	<b>.2675</b>	
720	<b>.4662</b>	.4743	.5963	.5724	.4923	<b>.4218</b>	.4541	.4224	.5806	<b>.4190</b>	<b>.4190</b>	<b>.4189</b>	.4815	.4515	.6597	<b>.4488</b>	.4511	<b>.4460</b>		
Weather	96	<b>.1634</b>	.2124	.1651	.3124	.2324	<b>.1629</b>	<b>.1803</b>	<b>.1822</b>	.1840	.1838	.1838	.1855	<b>.1651</b>	.1678	.1669	.1695	.1695	<b>.1644</b>	
	192	<b>.2108</b>	.2564	.2122	.3569	.2769	<b>.2065</b>	<b>.2237</b>	<b>.2232</b>	.2290	.2250	.2250	.2263	.2120	<b>.2096</b>	.2105	.2109	.2109	<b>.2050</b>	
	336	<b>.2488</b>	.2924	.2540	.3950	.3150	<b>.2419</b>	<b>.2642</b>	<b>.2618</b>	.2743	.2651	.2651	.2647	<b>.2737</b>	<b>.2545</b>	.2599	.2556	.2556	<b>.2499</b>	
720	<b>.2713</b>	.3230	.2913	.4275	.3475	<b>.2737</b>	<b>.2708</b>	<b>.2783</b>	.2987	.2824	.2824	.2817	<b>.2855</b>	<b>.2864</b>	.3035	.2959	.2959	.2906		

RevIN performs standard normalization on each input time series and then applies a corresponding denormalization step on the output. This procedure mitigates train-test distribution mismatch while restoring the information removed during normalization at the prediction stage, preventing degradation in forecasting performance.

DDN, in contrast, operates jointly in the time and frequency domains. It decomposes the input into low-frequency and high-frequency components and computes local statistics from each domain to remove non-stationarity. DDN then reconstructs non-stationary patterns in the predicted outputs using distribution statistics estimated from the model’s predictions, enabling dynamic tracking of distribution drift.

Table 18: Prediction accuracy and overhead of multivariate consideration.

(a) Prediction accuracy.

		Transformer-based				Linear-based				MLP-based			
		iTransformer		PatchTST		DLinear		OLS		FreTS		MICN	
		Indiv.	Corr.	Indiv.	Corr.	Indiv.	Corr.	Indiv.	Corr.	Indiv.	Corr.	Indiv.	Corr.
<b>ETTh1</b>	96	.4363	<b>.4351</b>	.4238	<b>.4157</b>	.4482	<b>.4202</b>	<b>.4372</b>	.4388	<b>.4371</b>	.4408	.4684	<b>.4533</b>
	192	.4919	<b>.4762</b>	.4805	<b>.4589</b>	.5050	<b>.4646</b>	.4906	<b>.4806</b>	.4940	<b>.4769</b>	.5328	<b>.5039</b>
	336	.5300	<b>.4759</b>	.5320	<b>.4798</b>	.5456	<b>.4823</b>	.5320	<b>.4884</b>	.5351	<b>.4987</b>	.5878	<b>.5105</b>
	720	.5638	<b>.4371</b>	.5822	<b>.5244</b>	.5896	<b>.4695</b>	.5733	<b>.4660</b>	.5959	<b>.4954</b>	.6504	<b>.5314</b>
<b>ETTh2</b>	96	.2493	<b>.2453</b>	.2343	<b>.1836</b>	<b>.2281</b>	.2342	.2265	<b>.2249</b>	.2350	<b>.2311</b>	.2485	<b>.2411</b>
	192	.2947	<b>.2871</b>	.2608	<b>.2172</b>	.2819	<b>.2578</b>	.2791	<b>.2770</b>	.2824	<b>.2922</b>	.3017	<b>.2923</b>
	336	<b>.3339</b>	.3341	.2978	<b>.2361</b>	.3083	<b>.2866</b>	.3043	<b>.2911</b>	.3153	<b>.2971</b>	.3310	<b>.3122</b>
	720	.3591	<b>.3306</b>	.3428	<b>.2638</b>	.3477	<b>.3094</b>	.3453	<b>.3083</b>	.3399	<b>.2979</b>	.3885	<b>.3441</b>
<b>ETTh1</b>	96	.3455	<b>.3186</b>	.3626	<b>.3595</b>	.3475	<b>.3335</b>	.3475	<b>.3330</b>	.3525	<b>.3381</b>	.3831	<b>.3385</b>
	192	.4140	<b>.3988</b>	.4258	<b>.4159</b>	.4122	<b>.4058</b>	.4119	<b>.4125</b>	.4212	<b>.4153</b>	.4514	<b>.4188</b>
	336	.4643	<b>.4491</b>	<b>.4697</b>	.4739	.4858	<b>.4561</b>	.4749	<b>.4648</b>	.4775	<b>.4618</b>	.5054	<b>.4628</b>
	720	.5102	<b>.4372</b>	<b>.4882</b>	.4892	.4991	<b>.4406</b>	.5007	<b>.4323</b>	.4982	<b>.4533</b>	.5225	<b>.4496</b>
<b>ETTh2</b>	96	<b>.1632</b>	.1633	.1562	<b>.1195</b>	.1586	<b>.1557</b>	.1586	<b>.1574</b>	.1569	<b>.1560</b>	.1704	<b>.1697</b>
	192	.2173	<b>.2141</b>	.2022	<b>.1526</b>	<b>.1905</b>	.2005	.1907	<b>.1935</b>	.1908	<b>.1921</b>	.2120	<b>.2097</b>
	336	.2535	<b>.2347</b>	.2352	<b>.1783</b>	.2242	<b>.2235</b>	.2226	<b>.2145</b>	.2211	<b>.2149</b>	.2351	<b>.2429</b>
	720	.2606	<b>.2110</b>	.2645	<b>.2162</b>	.2316	<b>.2295</b>	.2349	<b>.2006</b>	.2314	<b>.2022</b>	.2643	<b>.2265</b>
<b>Exchange Rate</b>	96	<b>.0837</b>	.0840	<b>.0788</b>	.0851	.0834	<b>.0791</b>	<b>.0773</b>	<b>.0773</b>	.0766	<b>.0765</b>	.1008	<b>.0995</b>
	192	<b>.1479</b>	.1493	<b>.1570</b>	.1828	<b>.1519</b>	.1609	<b>.1457</b>	<b>.1499</b>	.1516	<b>.1722</b>	.1726	<b>.1726</b>
	336	<b>.2624</b>	.2838	<b>.2445</b>	.3162	<b>.2480</b>	.2599	<b>.2323</b>	.2456	<b>.2461</b>	.2627	<b>.2660</b>	.2955
	720	<b>.4460</b>	.5221	<b>.4662</b>	.7731	<b>.4481</b>	.5280	<b>.4541</b>	.5184	<b>.4458</b>	.5079	<b>.4815</b>	.5711
<b>Weather</b>	96	.1617	<b>.1547</b>	<b>.1634</b>	.1656	.1793	<b>.1566</b>	.1803	<b>.1731</b>	.1737	<b>.1673</b>	.1651	<b>.1591</b>
	192	.2088	<b>.1938</b>	.2108	<b>.2073</b>	.2217	<b>.2003</b>	.2237	<b>.2158</b>	.2189	<b>.2095</b>	.2120	<b>.1979</b>
	336	.2515	<b>.2300</b>	<b>.2488</b>	.2539	.2626	<b>.2321</b>	.2642	<b>.2461</b>	.2587	<b>.2331</b>	.2737	<b>.2565</b>
	720	.2730	<b>.2236</b>	<b>.2713</b>	.2868	.2708	<b>.2254</b>	.2708	<b>.2172</b>	.2692	<b>.2165</b>	.2855	<b>.2420</b>

(b) Overhead analysis.

Method	# Params ↓	Adaptation time/batch (ms) ↓	Inference time/batch (ms) ↓	Average MSE ↓
Univariate	<b>1,211,287</b>	<b>80.12 ± 13.58</b>	<b>1.25 ± .0984</b>	.3240
Multivariate	1,214,851	186.28 ± 15.36	6.35 ± .2452	<b>.3071</b>

Each heatmap entry  $(i, j)$  shows the weight connecting the  $j$ -th input to the  $i$ -th output; columns 1: $L$  correspond to the original prediction of base model  $\mathbf{Y}^{(0)}$  and columns  $L+1:L+K$  to the context vector  $\mathbf{C}$ . The example is taken from a single variable of ETTh1 with look-back  $W=96$  and horizon  $L=96$ . Notably, the rightmost block (context columns) is strongly attenuated when RevIN or DDN is applied, whereas without a normalizer, the same block exhibits structured, non-negligible weights. This pattern indicates that explicit normalization reduces the marginal utility of the context (level/scale cues are already standardized), while in the non-normalized COSA leverages  $\mathbf{C}$  to perform level-shift correction directly in the output space—supporting our claim that the adapter can subsume normalization effects when needed and remain compatible with them when present.

## H DISCUSSION

Although COSA is built around a simple output-space linear adapter, the extended experiments in the Appendix examined multiple alternative design choices. While low-rank adaptation, input-side calibration, vector gating, selective/encoder-based context, and multivariate extensions each provide potential benefits in specific scenarios, our overall findings show that, considering average accuracy, stability, and latency, the architecture adopted in this paper remains the most consistent and robust choice for TSF-TTA.

The key observations are summarized below:

Table 19: Prediction accuracy and overhead of vector gating.

(a) Performance comparison.

	iTransformer		DLinear		FreTS		
	Scalar	Vector	Scalar	Vector	Scalar	Vector	
<b>ETTh1</b>	96	.4363	<b>.4336</b>	.4574	<b>.4436</b>	.4371	<b>.4336</b>
	192	.4919	<b>.4875</b>	.5066	<b>.4985</b>	.4940	<b>.4905</b>
	336	<b>.5300</b>	.5430	.5528	<b>.5466</b>	.5467	<b>.5353</b>
	720	<b>.5638</b>	.6013	<b>.6107</b>	.6178	.6259	<b>.6236</b>
<b>ETTh2</b>	96	.2493	<b>.1990</b>	.2281	<b>.1798</b>	.2350	<b>.1842</b>
	192	.2947	<b>.2318</b>	.2819	<b>.2198</b>	.2972	<b>.2202</b>
	336	.3339	<b>.2604</b>	.3083	<b>.2472</b>	.3153	<b>.2501</b>
	720	.3591	<b>.3035</b>	.3477	<b>.2893</b>	.3399	<b>.2898</b>
<b>ETTm1</b>	96	<b>.3455</b>	.3611	.3456	<b>.3425</b>	<b>.3525</b>	.3551
	192	.4140	<b>.4119</b>	.4222	<b>.4068</b>	.4212	<b>.4137</b>
	336	<b>.4643</b>	.4720	.4858	<b>.4706</b>	.4775	<b>.4767</b>
	720	<b>.5102</b>	.5484	<b>.4991</b>	.5370	<b>.4982</b>	.5405
<b>ETTm2</b>	96	.1632	<b>.1245</b>	.1583	<b>.1215</b>	.1569	<b>.1202</b>
	192	.2173	<b>.1683</b>	.1943	<b>.1487</b>	.1934	<b>.1484</b>
	336	.2535	<b>.2010</b>	.2242	<b>.1798</b>	.2211	<b>.1795</b>
	720	.2606	<b>.2488</b>	<b>.2316</b>	.2319	.2314	<b>.2280</b>
<b>Exchange Rate</b>	96	<b>.0837</b>	.0875	<b>.0834</b>	.0903	<b>.0766</b>	.0818
	192	<b>.1479</b>	.1774	<b>.1519</b>	.1790	<b>.1499</b>	.1698
	336	<b>.2624</b>	.3258	<b>.2480</b>	.3134	<b>.2461</b>	.3094
	720	<b>.4460</b>	.7649	<b>.4481</b>	.7904	<b>.4458</b>	.7500
<b>Weather</b>	96	<b>.1617</b>	.1718	<b>.1793</b>	.1908	<b>.1737</b>	.1825
	192	<b>.2088</b>	.2176	<b>.2217</b>	.2326	<b>.2189</b>	.2250
	336	<b>.2515</b>	.2706	<b>.2626</b>	.2808	<b>.2587</b>	.2745
	720	<b>.2730</b>	.3359	<b>.2708</b>	.3418	<b>.2692</b>	.3378

(b) Overhead analysis.

Method	# Params ↓	Adaptation time/batch (ms) ↓	Inference time/batch (ms) ↓	Average MSE ↓
Scalar	<b>1,211,287</b>	<b>80.12 ± 13.58</b>	<b>1.25 ± .0984</b>	<b>.3240</b>
Vector	1,212,446	96.34 ± 12.48	1.89 ± .0745	.3287

- **Low-rank factorization.** Despite its parameter-efficiency appeal, reducing representational capacity can lead to unstable adaptation. A joint adapter that integrates low-rank structure without compromising stability is a meaningful direction for future work.
- **Input-side calibration (GCM).** Combining COSA with input GCM improves performance on datasets with strong input noise (e.g., ETTh1/h2/m2) by smoothing perturbations before prediction. However, for fast-drifting or irregular series, GCM may oversmooth important variations and degrade performance, reaffirming that output-only correction is a reasonable and stable default.
- **Gating and its variants.** Although vector gating was expected to modulate drift at a finer temporal resolution, local noise propagated across the gating vector and reduced stability compared to scalar gating. Scalar gating’s batch-level modulation limits the influence of noise and achieves more reliable behavior.
- **Context construction.** Selective context (phase-aligned retrieval) and encoder-based context (RNN/LSTM/Attention) showed improvements in certain periodic datasets, but both suffered from outdated information, overfitting, or latency overhead in non-stationary settings. These results highlight that additional complexity does not guarantee better TTA performance unless paired with drift-aware representations and online update strategies. Future directions include sub-sequence vector gating, structure-aware context that explicitly

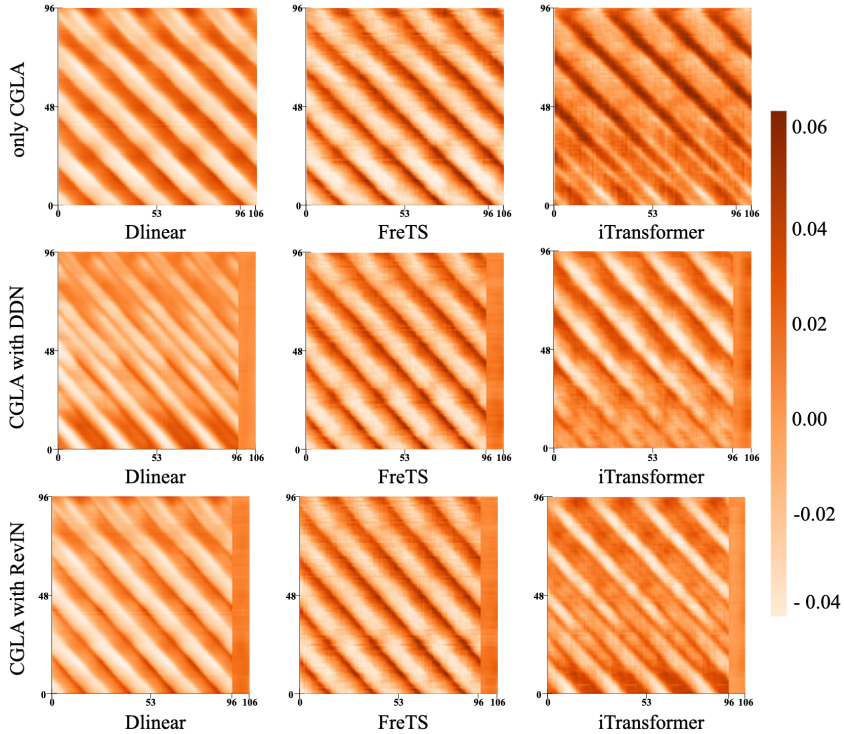


Figure 7: Weight heatmaps of the COSA linear layer for one ETTH1 variable ( $W=96, L=96$ ). Columns  $1:L$  are base-prediction inputs  $\mathbf{Y}^{(0)}$ ; columns  $L+1:L+K$  are context inputs  $\mathbf{C}$ . Each cell shows the weight from input  $j$  to output  $i$ . Color scales are kept identical across panels to allow magnitude comparison.

encodes trend and seasonality, and lightweight encoders capable of tracking drift without incurring high overhead.

- **Multivariate residual correction.** While multivariate modeling with cross-variable attention and shared components improved performance in most datasets, it degraded both accuracy and efficiency in datasets with weak inter-variable correlations (e.g., Exchange Rate). This suggests the need for selective correlation modeling or structural sparsity to suppress unnecessary cross-variable interactions.

Taken together, the extended Appendix experiments reinforce that the proposed simple architecture is particularly well-suited for TSF-TTA. They also indicate substantial room for generalizing COSA through carefully integrated low-rank structures, input calibration modules, selective or encoder-based context modeling, and vector gating, while preserving the efficiency and stability crucial for non-stationary test-time adaptation.

## I CONFIDENCE INTERVAL OF MAIN RESULTS

Table 20 reports the 95% confidence intervals of the main accuracy comparison results over 10 independent runs for each method–dataset–horizon combination. Overall, the intervals are narrow, indicating that the run-to-run variability of all methods is small, and COSA-F/P consistently retain their advantage over baselines even when accounting for this uncertainty.

## THE USE OF LARGE LANGUAGE MODELS

**Tool & Version:** Claude Sonnet 4 (Anthropic, 2025-09)

**Research Stage:** Not used.

Table 20: 95% confidence interval of main accuracy comparison results

	Transformer-based												Linear-based												MLP-based											
	TTransfomer				PatchTST				DLinear				OLS				FreTS				MICN															
	TAFAS	PETSa	COSA-F	COSA-P	TAFAS	PETSa	COSA-F	COSA-P	TAFAS	PETSa	COSA-F	COSA-P	TAFAS	PETSa	COSA-F	COSA-P	TAFAS	PETSa	COSA-F	COSA-P	TAFAS	PETSa	COSA-F	COSA-P												
ETH1	96	~4409	~4390	~4366	<b>~4361</b>	~4243	~4247	~4226	<b>~4222</b>	~4571	~4405	~4385	~4377	~4397	~4381	~4377	~4890	~4886	~4683	<b>~4682</b>	~4377	~4381	~4377	~4377												
	192	~4916	~4937	~4948	<b>~4908</b>	~4825	~4805	~4825	<b>~4775</b>	~5029	~4620	~4596	~4572	~4413	~4397	~4395	~4922	~4918	~4905	<b>~4906</b>	~4922	~4905	~4922	~4906												
	336	~5380	~5386	~5604	<b>~5248</b>	~5444	~5399	~5288	<b>~5288</b>	~5508	~5587	~5601	~5508	~5511	~5444	~5444	~5508	~5505	~5511	<b>~5448</b>	~5451	~5505	~5511	~5448												
	720	~6581	~6570	~5936	<b>~5612</b>	~6805	~6762	~6062	<b>~5760</b>	~6003	~6928	~6832	~6511	~6540	~6548	~6542	~6548	~6678	~6665	~6128	<b>~6099</b>	~6678	~6665	~6128	~6099											
		~6643	~6622	~5980	<b>~5664</b>	~6915	~6882	~6164	<b>~5884</b>	~6521	~6928	~6832	~6214	~6540	~6548	~6542	~6548	~6678	~6665	~6128	<b>~6099</b>	~6678	~6665	~6128	~6099											
ETH2	96	~2546	~2548	~2502	<b>~2490</b>	~2347	~2358	~2345	<b>~2339</b>	~2298	~2300	~2295	~2275	~2270	~2273	~2248	~2363	~2360	~2360	<b>~2344</b>	~2549	~2550	~2490	<b>~2484</b>												
	192	~2952	~2990	~2970	<b>~2934</b>	~2734	~2747	~2634	<b>~2579</b>	~2828	~2828	~2813	~2807	~2827	~2827	~2760	~2827	~2827	~2827	<b>~2798</b>	~2956	~2956	~2956	<b>~2906</b>												
	336	~3320	~3311	~3215	<b>~3308</b>	~3091	~3098	~2940	<b>~2949</b>	~3116	~3124	~2985	~3021	~3170	~3179	~2990	~3032	~3207	~3008	<b>~3008</b>	~3129	~3448	~3457	<b>~3202</b>												
	720	~3995	~4024	~3460	<b>~3567</b>	~3956	~3962	~3177	<b>~3380</b>	~3824	~3842	~3824	~3030	~3448	~3841	~3830	~3811	~3811	~3817	<b>~3124</b>	~3356	~4435	~4437	<b>~3616</b>												
		~4051	~4062	~3514	<b>~3289</b>	~4054	~4062	~3289	<b>~3289</b>	~3824	~3904	~3882	~3506	~3975	~3938	~3938	~3240	~3516	~3903	~3903	<b>~3214</b>	~4435	~4437	~3616	~3926											
ETH1	96	~3556	~3568	~3445	<b>~3453</b>	~3883	~3927	~3616	<b>~3616</b>	~3496	~3523	~3455	~3455	~3505	~3505	~3453	~3583	~3583	~3583	<b>~3521</b>	~3945	~3946	~3832	<b>~3826</b>												
	192	~4139	~4136	~4118	<b>~4118</b>	~4334	~4351	~4393	<b>~4234</b>	~4241	~4153	~4164	~4103	~4159	~4159	~4116	~4210	~4210	~4210	<b>~4147</b>	~4559	~4567	~4470	<b>~4508</b>												
	336	~4715	~4717	~4522	<b>~4522</b>	~4603	~4901	~4942	<b>~4564</b>	~4693	~4792	~4794	~4745	~4773	~4780	~4735	~4838	~4791	~4754	<b>~4631</b>	~5099	~5073	~4823	<b>~5045</b>												
	720	~5460	~5426	~4633	<b>~4633</b>	~4977	~5416	~5452	<b>~4672</b>	~4821	~5415	~5459	~4692	~4923	~5424	~5464	~4944	~5442	~5442	<b>~4669</b>	~5686	~5709	~4945	<b>~5152</b>												
		~5664	~5680	~4913	<b>~4690</b>	~5438	~5472	~4690	<b>~4690</b>	~4893	~5561	~5605	~4856	~5059	~5532	~5580	~4827	~5070	~5530	~5510	<b>~4767</b>	~5826	~5847	~5113	~5298											
ETH2	96	~1630	~1633	~1623	<b>~1628</b>	~1577	~1579	~1555	<b>~1577</b>	~1578	~1578	~1577	~1577	~1588	~1587	~1580	~1580	~1571	~1571	<b>~1567</b>	~1710	~1729	~1701	<b>~1703</b>												
	192	~2189	~2179	~2177	<b>~2177</b>	~2027	~2027	~1998	<b>~1998</b>	~2027	~1998	~1998	~1998	~2027	~2027	~1998	~2027	~2027	~2027	<b>~1998</b>	~2103	~2127	~2103	<b>~2127</b>												
	336	~2616	~2576	~2420	<b>~2420</b>	~2447	~2447	~2255	<b>~2420</b>	~2447	~2447	~2447	~2447	~2576	~2576	~2447	~2447	~2447	~2447	<b>~2087</b>	~2492	~2511	~2338	<b>~2543</b>												
	720	~3285	~3312	~2453	<b>~2453</b>	~3250	~3250	~2439	<b>~2439</b>	~2639	~2914	~2919	~2161	~2260	~2844	~2669	~2349	~2882	~2900	~2129	~2283	~3166	~3071	~2417												
		~3325	~3352	~2501	<b>~2627</b>	~3275	~3262	~2453	<b>~2453</b>	~2651	~3022	~3007	~2269	~2372	~2844	~2673	~2459	~2950	~2952	~2187	~2345	~3274	~3191	~2537												
Exchange Rate	96	~0873	~0883	~0816	<b>~0816</b>	~0835	~0828	~0756	<b>~0756</b>	~0729	~0881	~0874	~0809	~0795	~0801	~0753	~0770	~0793	~0797	<b>~0738</b>	~1071	~1130	~0936	<b>~0991</b>												
	192	~1645	~1706	~1363	<b>~1440</b>	~1711	~1742	~1394	<b>~1487</b>	~1487	~1708	~1677	~1408	~1604	~1601	~1344	~1403	~1602	~1594	~1431	~1436	~2094	~1925	~1586												
	336	~2995	~3029	~2006	<b>~2346</b>	~3186	~3225	~1886	<b>~2333</b>	~2904	~2891	~2911	~2448	~2787	~2821	~1935	~2341	~2813	~2833	~1936	~2331	~2880	~2959	~1980												
	720	~3857	~3816	~3170	<b>~3170</b>	~4152	~4152	~3280	<b>~3170</b>	~4152	~4152	~4152	~4152	~4152	~4152	~3170	~4152	~4152	~4152	~3170	~4152	~4152	~4152	~4152												
		~3857	~3816	~3170	<b>~3170</b>	~4152	~4152	~3280	<b>~3170</b>	~4152	~4152	~4152	~4152	~4152	~4152	~3170	~4152	~4152	~4152	~3170	~4152	~4152	~4152	~4152												
Weather	96	~1663	~1673	~1596	<b>~1616</b>	~1721	~1740	~1621	<b>~1621</b>	~1790	~1803	~1792	~1771	~1800	~1798	~1768	~1800	~1755	~1761	<b>~1719</b>	~1832	~1969	~1635	<b>~1650</b>												
	192	~2078	~2107	~2045	<b>~2067</b>	~2140	~2159	~1998	<b>~2109</b>	~2100	~2190	~2157	~2157	~2162	~2160	~2106	~2190	~2132	~2160	~2106	~2156	~2126	~2231	~1637												
	336	~2595	~2641	~2482	<b>~2482</b>	~2994	~2994	~2647	<b>~2482</b>	~2428	~2682	~2719	~2544	~2601	~2703	~2325	~2384	~2198	~2224	~2164	~2222	~2196	~2299	~2114												
	720	~3633	~3689	~2854	<b>~2854</b>	~3548	~3548	~2520	<b>~2854</b>	~3548	~3548	~3548	~3548	~3548	~3548	~2854	~3548	~3548	~3548	~2854	~3548	~3548	~3548	~3548												
		~3404	~3405	~2415	<b>~2415</b>	~3673	~3307	~3369	<b>~2496</b>	~2626	~3477	~3471	~2560	~2684	~3471	~3462	~2540	~2674	~3479	~3478	~2563	~2680	~3508	~3602	~2492											

**Writing Stage:** Language editing of author-drafted text for clarity and conciseness.

**Human Oversight:** All outputs reviewed/edited by the authors; authors accept full responsibility for the content.