# DYNAGUARD: REDUCING NEGATIVE TRANSFER THROUGH CONTROLLED CONTEXT INJECTION IN MULTIMODAL TIME SERIES FORECASTING

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

In time series forecasting, exogenous text such as news and reports can provide predictive signals not yet reflected in numerical histories, but unreliable, stale, or temporally misaligned context can increase forecasting error relative to a matched numerical-only reference, causing negative transfer. To reduce negative transfer, we address three linked questions: what content to retain, how to assess its temporal relevance, and where and how strongly to apply corrections. We propose DynaGuard, a backbone-compatible framework for controlled context injection: it models target dynamics from numerical histories and uses text to provide corrections governed by these three decisions. Time-Slice Context Compression (TCC) selects and compresses forecasting-relevant content within each historical interval. Context-Adaptive Temporal Bias (CATB) uses these compressed representations to predict salience and decay, guiding retrieval according to historical lag and numerical-token position. Content-Aware Gated Injection (CAGI) then uses the numerical state and retrieved context to control correction strength at each numerical token and feature channel. These stages learn jointly while preserving a direct numerical path to prediction. Our main comparisons cover 16 tasks from three data collections and six numerical backbones, with five random seeds per configuration. DynaGuard reduces negative transfer and improves forecasting accuracy. Ablations support each component's contribution, and the components also benefit other text-integration methods. Code is available at https://anonymous.4open.science/r/DynaGuard2027/.

## 1 INTRODUCTION

Multimodal time series forecasting augments numerical histories with auxiliary modalities. We focus on exogenous text, which typically comes from external sources such as news, announcements, and reports (Wang et al., 2026). Such text may convey relevant information that numerical histories do not readily capture or have yet to reflect (Liu et al., 2024a; Jia et al., 2024; Dong et al., 2024). The two modalities play asymmetric roles: numerical histories record the target's temporal dynamics and anchor the forecast, while text provides conditional context whose value depends on its content, timing, and relevance to the target. Irrelevant text can increase error over the same numerical backbone trained without text, causing negative transfer (Wang et al., 2019; Lee et al., 2026). In our evaluation, four text-integration baselines exhibit negative transfer in 37.5–71.9% of 96 task–backbone settings (Section 4). The central problem is how to **retain useful textual signals and apply them to forecasting while limiting interference**.

We break down the central problem of reducing negative transfer into **three linked subproblems**: what content to retain, how to assess its temporal relevance, and where and how strongly to apply corrections. First, text may mix predictive information with repeated background and unrelated details, so document-level relevance does not determine which content to retain. Second, the relevance of retained text changes with age and varies across events. Figure 1 illustrates this difference: a feed-supply outlook may remain useful for weeks, whereas a maintenance notice concerns a brief disruption. Both reports may be relevant but call for different temporal weights. Third, retrieved context need not affect every numerical token or feature equally; the required correction also depends on the current numerical state.



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

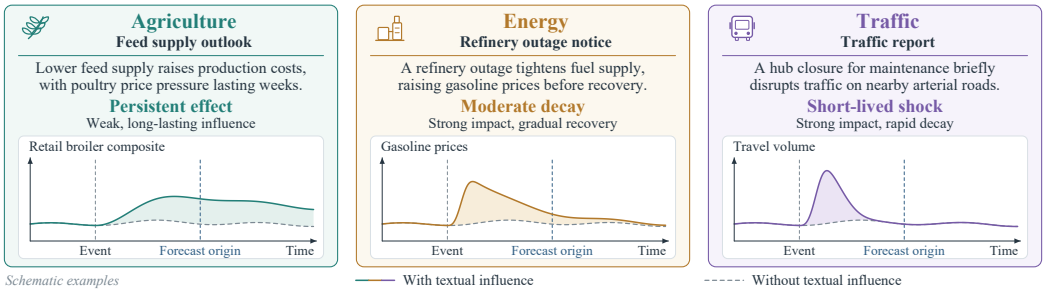


Figure 1: Illustrative event effects in Agriculture, Energy, and Traffic. Solid and dashed trajectories show series with and without textual influence; the examples vary in initial strength and persistence.

Existing methods regulate text at different stages. MM-TSFlib combines modality-specific forecasts (Liu et al., 2024a); TaTS represents paired text as auxiliary temporal variables (Li et al., 2026b); and Cross-MoE combines expert-based transformations with text-token selection (Huang et al., 2025). Concurrent work, Controlled Fusion Adapter (CFA), constrains the capacity of textual residuals through a low-rank bottleneck (Lee et al., 2026). Content selection and capacity constraints leave open how to weight retained text as it ages and apply it to the current numerical state.

We therefore formulate this problem as controlled context injection and propose DynaGuard for numerical backbones that expose representations before the prediction head. DynaGuard retains numerical representations through an additive residual path and controls textual influence through a three-stage context adapter. **Time-Slice Context Compression (TCC)** uses a shared learnable query to select and compress each interval's text into one context token, preserving the interval structure. **Context-Adaptive Temporal Bias (CATB)** uses these compressed tokens to predict content-dependent salience and decay, guiding retrieval according to historical lag and numerical-token position. Because attention normalizes over available contexts, relative relevance alone does not determine correction strength. **Content-Aware Gated Injection (CAGI)** therefore uses a gate conditioned on the numerical state and retrieved context to control correction strength at each numerical token and feature channel. The three stages jointly learn what to retain, how to retrieve it over time, and how to apply the resulting correction.

Our experiments cover nine Time-MMD domains, financial forecasting in FNSPID, and news-count forecasting in GPT4MTS: 16 tasks across six numerical backbones. Each main-comparison configuration uses five different random seeds to assess the consistency of the results.

Our contributions are threefold:

- We formulate numerical–text forecasting with unreliable, timestamped context as a problem of **controlled context injection**, involving content selection, temporal relevance, and injection location and strength. We evaluate negative transfer against a matched numerical-only reference.
- We propose DynaGuard, **a backbone-compatible framework** that jointly learns what text contributes, how its relevance changes over time, and where and how strongly it modifies numerical features, while retaining a direct numerical path to prediction.
- DynaGuard achieves a **14.6% negative-transfer rate** versus 37.5–71.9% for four text-integration baselines, and **lower MSE** than each baseline in 72–89 of 96 settings. Ablations support each control, and component transfer shows that individual controls benefit other text-integration methods.

## 2 Related Work

Text in multimodal time series forecasting comes from two sources: **endogenous text** derived from the numerical series and **exogenous text** obtained from outside it (Wang et al., 2026). Endogenous text describes properties of the observed series. For example, Time-LLM includes summary statistics and trend information in its prompts (Jin et al., 2024). Exogenous text provides external context,

such as news, announcements, and reports. GPT4MTS jointly uses numerical and textual inputs (Jia et al., 2024); Time-MMD provides timestamp-aligned data across nine domains and the MM-TSFlib framework (Liu et al., 2024a); and FNSPID pairs financial news with market observations (Dong et al., 2024). DynaGuard focuses on using such exogenous context to improve numerical forecasting.

Methods that use exogenous text differ in how they combine it with numerical histories. MM-TSFlib combines modality-specific forecasts through learned weighting (Liu et al., 2024a). TaTS represents paired text as auxiliary temporal variables (Li et al., 2026b), while Cross-MoE transforms numerical and textual features with experts before selecting text tokens by cross-modal similarity (Huang et al., 2025). Spectral Text Fusion (SpecTF) uses cross-attention to reweight numerical frequency components with textual context (Nguyen et al., 2026). DynaGuard applies controlled context injection to the **numerical representation before the prediction head**, coordinating content selection, temporal weighting, and the location and strength of the correction.

Negative transfer provides a criterion for evaluating this integration: auxiliary information can reduce performance on the target task relative to learning without it (Wang et al., 2019). In numerical–text forecasting, a matched numerical-only model provides this reference. Controlled Fusion Adapter (CFA) preserves numerical representations and adds textual residuals through a low-rank bottleneck (Lee et al., 2026). This bottleneck filters textual information by restricting residual capacity, but does not explicitly model how relevance varies with context age. Timestamp alignment selects **which text enters the historical window**; modeling relevance determines **how that history is weighted**. DynaGuard preserves the numerical foundation and integrates **content-dependent temporal relevance** into a joint design: selection within each time slice determines which information is available for retrieval, temporal decay adjusts its weight across historical slices, and gating controls the correction at each numerical token and feature. Thus, DynaGuard coordinates the construction and application of textual corrections while preserving the numerical prediction path.

Appendix B discusses numerical backbones and broader text-based forecasting systems.

# 3 DynaGuard

## 3.1 Problem Formulation

Given a numerical history $X = [x_1, \ldots, x_T] \in \mathbb{R}^{T \times C}$ and associated historical text, we predict $Y = [y_{T+1}, \ldots, y_{T+H}] \in \mathbb{R}^{H \times C_y}$. The vector $x_t \in \mathbb{R}^C$ contains the observed variables at time $t$, and $y_t \in \mathbb{R}^{C_y}$ contains the target variables. Here, $T$ and $H$ denote the look-back and forecast horizons. A historical time slice spans one sampling interval, and $S_j$ collects the text associated with slice $j$.

We reuse pretrained text features and learn their forecasting-specific contribution in the context adapter. A frozen pretrained input embedding table maps each slice to

$$E_j = \mathrm{Embed}(S_j) \in \mathbb{R}^{L_j \times d_{\mathrm{lm}}}. \tag{1}$$

Here, $L_j$ is the token count and $d_{\mathrm{lm}}$ the embedding dimension. Text inputs use rows dated before the prediction interval under the benchmark alignment (Appendix D). Missing entries use the valid textual placeholder "No information available.", while padding tokens are masked. Appendix D.3 gives the preprocessing details. The historical lag of slice $j$ is $r_j = T - j$, measured in sampling intervals from the last observed slice; the latest slice has lag zero.

The numerical history provides the forecasting foundation, and text supplies a conditional correction. Let $\mathcal{E}_{\mathrm{ts},\theta}$ denote the numerical backbone with parameters $\theta$. At the adapter's insertion point, it produces $N_q$ numerical tokens, each with $d$ features:

$$H_{\mathrm{ts}} = \mathcal{E}_{\mathrm{ts},\theta}(X) \in \mathbb{R}^{N_q \times d}. \tag{2}$$

The backbone-specific prediction mapping $\mathcal{P}_\phi$ produces the forecast, with $\phi$ denoting its parameters. To evaluate negative transfer, we independently train a numerical-only (Uni-modal) reference with the same numerical architecture, data, forecast horizon, and validation protocol:

$$\widehat{Y}_{\mathrm{uni}} = \mathcal{P}_{\phi_{\mathrm{uni}}}\left(\mathcal{E}_{\mathrm{ts},\theta_{\mathrm{uni}}}(X)\right). \tag{3}$$

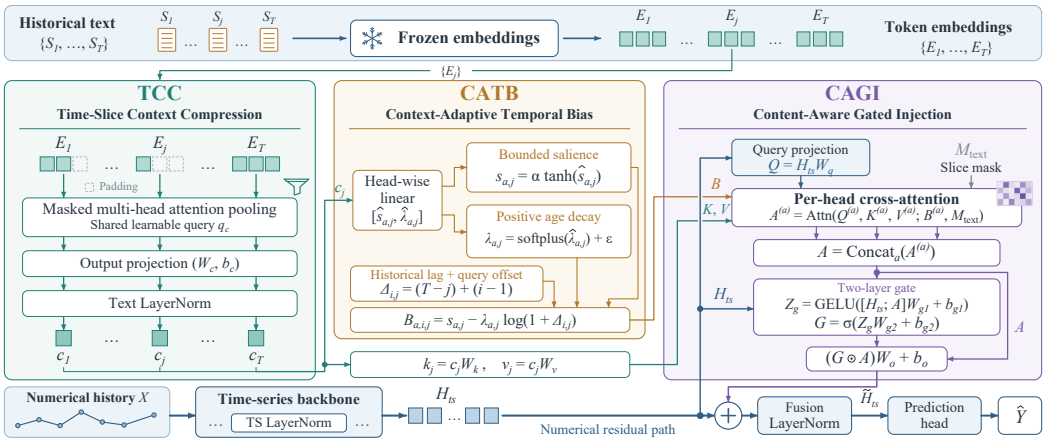


Figure 2: DynaGuard architecture. TCC produces one context token per historical time slice; CATB supplies a content-dependent temporal bias for multi-head retrieval; CAGI gates the retrieved features before residual fusion. The numerical branch carries backbone representations directly to fusion and the prediction head.

For a sample with at least one valid text slice, DynaGuard adds a context-dependent correction $\mathcal{I}$ at this representation:

$$\widetilde{H}_{\text{ts}} = \text{LN}\Big(H_{\text{ts}} + \mathcal{I}(H_{\text{ts}}, \{E_j, r_j\}_{j=1}^{T})\Big), \qquad \widehat{Y} = \mathcal{P}_{\phi}(\widetilde{H}_{\text{ts}}). \tag{4}$$

Here, LN denotes layer normalization. The additive path **carries the numerical representation directly to fusion**, while $\mathcal{I}$ supplies the text-conditioned correction.

### 3.2 Framework Overview

DynaGuard inserts a context adapter at a backbone-specific representation (Figure 2), preserving the numerical model's input and output interface. The numerical tokens may represent variables, temporal patches, variable–patch pairs, or latent states. TCC supplies the retrieval content and CATB inputs; CATB shapes the retrieved context; CAGI controls its application. Component ablations support each control's contribution within DynaGuard (Tables 2 and 48), while temporal-bias comparisons evaluate the CATB parameterization against the tested alternatives (Table 55).

### 3.3 Time-Slice Context Compression

Text within a historical slice may mix predictive information with repeated background and irrelevant details. TCC uses learned-query attention (Jaegle et al., 2021) to **select and compress** each slice before numerical queries retrieve historical context. One token per slice provides a compact key–value history and preserves the indices required for temporal weighting. Selection is learned from forecasting error without token-, document-, or event-level labels (Equations 5–6). A shared learnable query $q_c \in \mathbb{R}^{d_{\text{lm}}}$ reads each slice separately through $H_c$ attention heads of width $d_c = d/H_c$. Let $e_{j,\ell}$ be row $\ell$ of $E_j$. We use row vectors, with $\|$ and $[\cdot;\cdot]$ denoting feature concatenation. For clarity, the equations omit query, key, and value projection biases and training-time dropout. For compression head $a$,

$$\beta_{a,j,\ell} = \text{softmax}_{\ell}\left(\frac{(q_c W_a^Q)(e_{j,\ell} W_a^K)^{\top}}{\sqrt{d_c}} + m_{j,\ell}\right), \tag{5}$$

$$c_j = \text{LN}_{\text{text}}\left(\left[\Big\|_{a=1}^{H_c} \sum_{\ell=1}^{L_j} \beta_{a,j,\ell}(e_{j,\ell} W_a^V)\right] W_c + b_c\right), \qquad k_j = c_j W_k, \quad v_j = c_j W_v. \tag{6}$$

The mask $m_{j,\ell}$ is zero for valid tokens and $-\infty$ for padding. Each $W_a^{Q,K,V}$ maps $d_{\text{lm}}$ features to $d_c$; $W_c$, $W_k$, and $W_v$ have shape $d \times d$.

Combining the head summaries and applying text-side layer normalization produces the Time-Slice Context Token (TCT), $c_j \in \mathbb{R}^d$. If every token in a slice is explicitly masked, we set its compressed output to zero and exclude it from cross-attention.

### 3.4 Context-Adaptive Temporal Bias

Content selection does not determine how relevance changes with age. A shared decay profile may overweight short-lived reports and underweight persistent ones. Drawing on logarithmic relative-position biases (Chi et al., 2022) and context-adaptive positional encoding (Zheng et al., 2024), CATB **models salience and persistence separately**, predicting a bounded salience offset and a positive decay coefficient from each compressed slice. Its distance combines historical lag with numerical-token position. For fixed content and query, the bias decreases with lag (Equations 7–9). For numerical token $i \in \{1, \dots, N_q\}$ and historical slice $j \in \{1, \dots, T\}$, it uses the distance

$$\Delta_{ij} = r_j + u_i, \qquad r_j = T - j, \qquad u_i = i - 1. \tag{7}$$

The offset $u_i$ follows the numerical-token order of the backbone. At each query, older text has a larger distance; a single numerical query gives $\Delta_{1j} = T - j$. Appendix D.3 specifies the token sequences for each backbone. For each retrieval head, CATB computes:

$$\begin{gathered}[\widehat{s}_{a,j}, \widehat{\lambda}_{a,j}] = c_j W_b^{(a)} + b_b^{(a)}, \\ s_{a,j} = \alpha \tanh(\widehat{s}_{a,j}), \qquad \lambda_{a,j} = \operatorname{softplus}(\widehat{\lambda}_{a,j}) + \epsilon, \\ B_{a,i,j} = s_{a,j} - \lambda_{a,j} \log(1 + \Delta_{ij}).\end{gathered} \tag{8}$$

Here, $a$ indexes $H_r$ retrieval heads of width $d_r = d/H_r$, and $W_b^{(a)} \in \mathbb{R}^{d \times 2}$. The constants $\alpha > 0$ and $\epsilon > 0$ bound salience and set the decay floor. Salience and decay depend on the text slice, while $\Delta_{ij}$ also depends on numerical-token position. Let $q_{a,i}$ and $k_{a,j}$ denote the numerical query and context key for head $a$. At a valid context position, the unnormalized retrieval weight factorizes as

$$\widetilde{w}_{a,i,j} = \exp\left(\frac{q_{a,i} k_{a,j}^\top}{\sqrt{d_r}}\right) \exp(s_{a,j}) (1 + \Delta_{ij})^{-\lambda_{a,j}}. \tag{9}$$

The three factors capture query–context compatibility, content salience, and temporal attenuation. The slice-specific salience and decay allow **equally old reports to receive different temporal weights**. For fixed content and query, $\partial B_{a,i,j}/\partial r_j = -\lambda_{a,j}/(1+\Delta_{ij}) < 0$: increasing a slice's historical lag lowers its bias, with a larger decay coefficient giving a stronger penalty. The normalized attention weights also depend on query–context compatibility and competing slices. Appendix C gives the decay factor's mixture interpretation and the effect of attention normalization.

### 3.5 Content-Aware Gated Injection

Retrieval determines relative weights, but softmax assigns total weight one across valid slices even when all matches are weak (Vaswani et al., 2017). Because correction needs vary across numerical tokens and feature channels, CAGI uses input-dependent gating (Arevalo et al., 2017) conditioned on both the numerical state and retrieved context. The gate **scales only the retrieved features** at each numerical token before the output projection, while an additive residual connection retains a direct numerical path to prediction (Equations 10–11).

We form $Q = H_{\text{ts}} W_q$ with $W_q \in \mathbb{R}^{d \times d}$ and stack the context keys and values as the rows of $K, V \in \mathbb{R}^{T \times d}$. Splitting their features into $H_r$ heads gives $Q^{(a)} \in \mathbb{R}^{N_q \times d_r}$ and $K^{(a)}, V^{(a)} \in \mathbb{R}^{T \times d_r}$. Numerical queries retrieve context through multi-head cross-attention (Vaswani et al., 2017) under the CATB bias:

$$A^{(a)} = \operatorname{softmax}_j\left(\frac{Q^{(a)} K^{(a)\top}}{\sqrt{d_r}} + B^{(a)} + M_{\text{text}}\right) V^{(a)}. \tag{10}$$

The mask $M_{\text{text}}$ is zero at valid slices and $-\infty$ at excluded slices, broadcast over queries and heads. We concatenate head outputs as $A = \|_{a=1}^{H_r} A^{(a)} \in \mathbb{R}^{N_q \times d}$ and gate them before projection:

$$\begin{aligned} Z_g &= \text{GELU}([H_{\text{ts}}; A]W_{g1} + b_{g1}), \\ G &= \sigma(Z_g W_{g2} + b_{g2}), \\ \widetilde{H}_{\text{ts}} &= \text{LN}(H_{\text{ts}} + (G \odot A)W_o + b_o)\,. \end{aligned} \tag{11}$$

The two-layer gate has hidden width $d$, with $W_{g1} \in \mathbb{R}^{2d \times d}$, $W_{g2}, W_o \in \mathbb{R}^{d \times d}$, and $Z_g \in \mathbb{R}^{N_q \times d}$. Here, $\sigma$ is the sigmoid function and $\odot$ denotes elementwise multiplication. The gate $G \in (0,1)^{N_q \times d}$ **separately scales each retrieved feature at each numerical token** before $W_o$ mixes features. If all context positions in a sample are explicitly masked, we return $H_{\text{ts}}$ directly, bypassing fusion and its layer normalization.

### 3.6 Training and Computational Cost

All three controls learn jointly from forecasting error. For the nonempty set $\Omega$ of valid target indices in a minibatch, we minimize

$$\mathcal{L}_{\text{pred}} = \frac{1}{|\Omega|} \sum_{(n,h,c)\in\Omega} (\widehat{y}_{n,h,c} - y_{n,h,c})^2. \tag{12}$$

Here, $n$, $h$, and $c$ index the sample, forecast step, and target channel. We use staged optimization to establish the numerical forecasting path before fitting the context adapter. We first train the numerical backbone and prediction head without text fusion, then train the backbone-specific head and adapter with the remaining numerical parameters frozen. Finally, we jointly refine the full model at relatively low learning rates. Each stage selects the checkpoint with the lowest validation loss. The pretrained text embeddings remain frozen throughout. This schedule controls optimization order; it does not define textual relevance or provide a reliability guarantee. Appendix D.3 gives the stage budgets and learning-rate candidates.

We analyze attention-score computation and weighted sums at fixed head counts, excluding projections, embedding lookup, and the numerical backbone and prediction head. With at most $L_{\text{text}}$ tokens per slice, direct attention from $N_q$ numerical queries to all historical text tokens costs $O(N_q T L_{\text{text}} d)$ time. TCC compresses each slice before retrieval, giving $O(T L_{\text{text}} d + N_q T d)$ time. If attention scores are materialized, their storage changes from $O(N_q T L_{\text{text}})$ to $O(T L_{\text{text}} + N_q T)$. Compression thus removes the within-slice token-length factor from numerical-to-text attention.

For large-scale experiments, we cache precomputed frozen text embeddings on the GPU to avoid repeated lookup and host-to-device transfers. The cache contributes to measured peak GPU memory. In a paired WMT/iTransformer check with one epoch per stage, host caching reduces peak allocated GPU memory by approximately 19%, with identical final model states and test errors. Appendix I reports resource measurements and this check.

## 4 Experiments and Analysis

### 4.1 Experimental Setup

**Data and comparisons.** We evaluate 16 tasks: nine Time-MMD domains, four FNSPID stock-price targets, and three GPT4MTS news-count targets (Liu et al., 2024a; Dong et al., 2024; Jia et al., 2024). Time-MMD and FNSPID use four forecast horizons per task; GPT4MTS uses seven-day forecasts. Each task uses six numerical backbones: iTransformer, PatchTST, Crossformer, SOFTS, TimeXer, and Transformer (Liu et al., 2024b; Nie et al., 2023; Zhang & Yan, 2023; Han et al., 2024; Wang et al., 2024; Vaswani et al., 2017). This yields 330 task–backbone–horizon configurations and 96 task–backbone settings after horizon averaging.

We compare DynaGuard with Uni-modal, MM-TSFlib, Cross-MoE, TaTS, and CFA (Liu et al., 2024a; Huang et al., 2025; Li et al., 2026b; Lee et al., 2026). All methods share forecast origins, numerical windows, chronological 70%/10%/20% splits, and evaluation scales; standardization uses training data only. Text-based methods share historical text and frozen GPT-2-XL input embeddings (Radford et al., 2019). Text windows follow benchmark row dates (Appendix D.1). Each

Table 1: Relative MSE and MAE against Uni-modal across six backbones (Section 4.1). Red bold and blue underlined central values mark the best and second-best results, including ties.

| Dataset | Uni-modal (Numerical only) | | MM-TSFlib (NeurIPS 2024) | | Cross-MoE (EMNLP 2025) | | TaTS (ICLR 2026) | | CFA (KDD Workshop 2026) | | DynaGuard (Ours) | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | rMSE | rMAE | rMSE | rMAE | rMSE | rMAE | rMSE | rMAE | rMSE | rMAE | rMSE | rMAE |
| *Time-MMD (9 tasks; 6 backbones; 4 horizons; 216 configurations; 1,080 runs/method)* | | | | | | | | | | | | |
| Agriculture | 1.000±0.000 | 1.000±0.000 | 1.151±0.035 | 1.096±0.027 | 1.124±0.032 | 1.092±0.019 | 0.948±0.018 | 0.970±0.011 | 0.928±0.032 | 0.964±0.017 | **0.812**±0.038 | **0.895**±0.027 |
| Climate | 1.000±0.000 | 1.000±0.000 | 1.078±0.007 | 1.043±0.003 | 1.003±0.006 | 1.002±0.003 | 0.978±0.005 | 0.988±0.002 | 1.036±0.003 | 1.017±0.003 | **0.931**±0.009 | **0.952**±0.003 |
| Economy | 1.000±0.000 | 1.000±0.000 | 1.147±0.031 | 1.077±0.015 | 1.103±0.018 | 1.054±0.009 | 0.903±0.029 | 0.946±0.017 | 1.000±0.025 | 1.000±0.014 | **0.879**±0.015 | **0.929**±0.008 |
| Energy | 1.000±0.000 | 1.000±0.000 | 0.817±0.013 | 0.871±0.008 | 0.878±0.010 | 0.903±0.006 | 0.975±0.005 | 0.986±0.004 | 1.061±0.020 | 1.029±0.011 | **0.772**±0.012 | **0.837**±0.006 |
| Environment | 1.000±0.000 | 1.000±0.000 | 1.207±0.009 | 1.112±0.007 | 1.192±0.002 | 1.117±0.003 | 1.001±0.002 | 1.000±0.001 | 1.019±0.008 | 1.022±0.008 | **0.993**±0.003 | **0.998**±0.004 |
| Health | 1.000±0.000 | 1.000±0.000 | 1.214±0.019 | 1.089±0.013 | 1.012±0.012 | 1.022±0.006 | **0.975**±0.011 | **0.994**±0.010 | 1.046±0.019 | 1.038±0.015 | 0.982±0.007 | 0.995±0.005 |
| Security | 1.000±0.000 | 1.000±0.000 | 1.044±0.007 | 1.070±0.010 | 1.095±0.008 | 1.125±0.014 | 0.989±0.006 | 0.989±0.009 | 1.006±0.008 | 1.008±0.008 | **0.959**±0.006 | **0.930**±0.009 |
| Social Good | **1.000**±0.000 | 1.000±0.000 | 1.103±0.006 | 1.106±0.019 | 1.140±0.013 | 1.135±0.022 | 1.006±0.006 | **0.978**±0.016 | 1.119±0.008 | 1.092±0.009 | 1.004±0.004 | 0.982±0.008 |
| Traffic | 1.000±0.000 | 1.000±0.000 | 1.147±0.012 | 1.156±0.013 | 1.090±0.018 | 1.099±0.022 | 0.994±0.004 | 0.990±0.005 | 0.977±0.012 | **0.972**±0.006 | **0.941**±0.003 | 0.980±0.006 |
| Geo. mean | 1.000±0.000 | 1.000±0.000 | 1.094±0.006 | 1.066±0.004 | 1.067±0.002 | 1.059±0.001 | 0.974±0.002 | 0.982±0.003 | 1.020±0.004 | 1.015±0.003 | **0.916**±0.007 | **0.943**±0.004 |
| NTR (MSE, %) | – | | 79.6 | | 66.7 | | 25.9 | | 68.5 | | **18.5** | |
| Mean Rank (MSE) | 3.22 | | 4.98 | | 4.19 | | 2.88 | | 4.03 | | **1.70** | |
| *FNSPID (4 tasks; 6 backbones; 4 horizons; 96 configurations; 480 runs/method)* | | | | | | | | | | | | |
| WMT | 1.000±0.000 | 1.000±0.000 | 0.983±0.029 | 0.994±0.020 | **0.915**±0.043 | **0.956**±0.027 | 1.001±0.048 | 1.001±0.030 | 1.106±0.027 | 1.055±0.017 | 0.974±0.057 | 0.986±0.033 |
| GILD | 1.000±0.000 | 1.000±0.000 | 1.015±0.025 | 1.008±0.012 | 1.034±0.030 | 1.014±0.015 | 1.029±0.029 | 1.017±0.018 | 1.049±0.020 | 1.029±0.011 | **0.985**±0.015 | **0.991**±0.010 |
| WFC | 1.000±0.000 | 1.000±0.000 | 1.007±0.003 | 1.003±0.002 | 0.988±0.010 | 0.995±0.006 | **0.920**±0.026 | **0.957**±0.014 | 1.069±0.037 | 1.045±0.018 | 0.947±0.027 | 0.977±0.017 |
| SBUX | 1.000±0.000 | 1.000±0.000 | 0.995±0.010 | 1.000±0.005 | 0.951±0.021 | 0.969±0.010 | 0.929±0.010 | 0.961±0.006 | 1.080±0.015 | 1.021±0.004 | **0.912**±0.019 | **0.958**±0.011 |
| Geo. mean | 1.000±0.000 | 1.000±0.000 | 1.000±0.008 | 1.001±0.005 | 0.971±0.020 | 0.983±0.012 | 0.969±0.013 | 0.984±0.007 | 1.076±0.009 | 1.037±0.005 | **0.954**±0.020 | **0.978**±0.011 |
| NTR (MSE, %) | – | | 50.0 | | 29.2 | | 50.0 | | 87.5 | | **4.2** | |
| Mean Rank (MSE) | 3.52 | | 3.69 | | 2.62 | | 3.75 | | 5.44 | | **1.98** | |
| *GPT4MTS (3 tasks; 6 backbones; 1 horizon; 18 configurations; 90 runs/method)* | | | | | | | | | | | | |
| NumArticles | 1.000±0.000 | 1.000±0.000 | 1.181±0.005 | 1.016±0.002 | 1.119±0.020 | 0.993±0.012 | 0.977±0.008 | 0.992±0.003 | 0.990±0.012 | 0.992±0.008 | **0.937**±0.020 | **0.954**±0.014 |
| NumMentions | 1.000±0.000 | 1.000±0.000 | 1.174±0.005 | 1.012±0.002 | 1.120±0.019 | 0.991±0.012 | 0.980±0.008 | 0.993±0.004 | 0.988±0.010 | 0.991±0.008 | **0.915**±0.012 | **0.946**±0.003 |
| NumSources | 1.000±0.000 | 1.000±0.000 | 1.165±0.010 | 1.033±0.004 | 1.110±0.029 | 1.007±0.013 | 0.979±0.014 | 0.989±0.008 | 0.980±0.019 | 0.988±0.010 | **0.942**±0.012 | **0.968**±0.009 |
| Geo. mean | 1.000±0.000 | 1.000±0.000 | 1.173±0.006 | 1.020±0.002 | 1.116±0.021 | 0.997±0.012 | 0.979±0.009 | 0.992±0.004 | 0.986±0.013 | 0.990±0.008 | **0.931**±0.013 | **0.956**±0.008 |
| NTR (MSE, %) | – | | 77.8 | | 83.3 | | 55.6 | | 38.9 | | **16.7** | |
| Mean Rank (MSE) | 3.19 | | 5.17 | | 4.86 | | 3.28 | | 2.92 | | **1.58** | |

This table aggregates 9,900 runs across 330 task–backbone–horizon configurations, 6 methods, and five seeds (Appendix E.1–E.2). NTR and ranks use task–backbone settings after horizon averaging. Collection summaries weight tasks equally. The ± term measures variation across five matched seeds.

main-comparison experiment is repeated with five different random seeds to assess consistency. Appendix D.1 gives the full settings, and Appendix D.3 specifies the standard hyperparameter search spaces for all six methods.

**Metrics.** We average MSE and MAE over horizons and runs. Negative transfer compares error with a reference trained without the auxiliary source (Wang et al., 2019). In our evaluation, it occurs when aggregate MSE, rounded to three decimals, exceeds the matched Uni-modal error. For settings $\mathcal{S}$, the negative-transfer rate is

$$\mathrm{NTR}_{\mathcal{S}}(m) = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \mathbf{1}[\widetilde{E}_{m,s} > \widetilde{E}_{\mathrm{uni},s}], \tag{13}$$

where $\widetilde{E}$ is rounded aggregate MSE. Ranks and win/tie/loss counts use the same rule. Relative errors follow geometric aggregation of baseline-normalized errors (Davydenko & Fildes, 2013; Das et al., 2024):

$$\mathrm{rMSE}_{d,m} = \left( \prod_{b=1}^{B} \frac{E^{(\mathrm{MSE})}_{d,b,m}}{E^{(\mathrm{MSE})}_{d,b,\mathrm{uni}}} \right)^{1/B}, \qquad \mathrm{rMAE}_{d,m} = \left( \prod_{b=1}^{B} \frac{E^{(\mathrm{MAE})}_{d,b,m}}{E^{(\mathrm{MAE})}_{d,b,\mathrm{uni}}} \right)^{1/B}, \tag{14}$$

where $E^{(q)}_{d,b,m}$ averages metric $q$ over horizons and runs for task $d$, backbone $b$, and method $m$, and $B = 6$. Values below one indicate improvement. The ± terms are sample standard deviations of five seed-specific relative errors (Appendix D.2).

### 4.2 Reducing Negative Transfer across Domains and Backbones

Table 1 shows that DynaGuard achieves the **lowest negative-transfer rate** among text-based methods and the **best mean MSE rank** among all methods in each collection. Its first-place results span all three collections (Figure 3). On FNSPID, the best method varies by stock, while DynaGuard achieves the lowest collection-level rMSE and rMAE. Its advantage combines accuracy across targets with fewer harmful textual updates.

**Pairwise comparisons.** Figure 4 shows that DynaGuard reduces MSE against every reference in most settings. For statistical testing, the 16 tasks form 16 blocks; each block averages unrounded log-MSE ratios across backbones. Two-sided exact Wilcoxon tests with Holm correction support lower MSE against all five references (Demšar, 2006; Wilcoxon, 1945; Holm, 1979). Appendix D.1 reports the tests and a sensitivity analysis grouping targets that share a source.

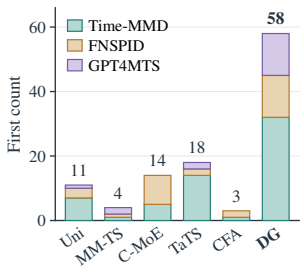


Figure 3: MSE first-place counts across 96 settings, including ties.

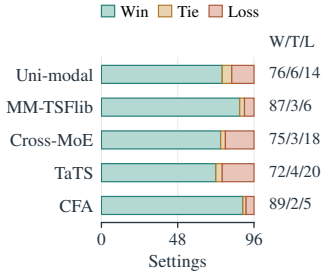


Figure 4: DynaGuard's MSE wins, ties, and losses across 96 settings.

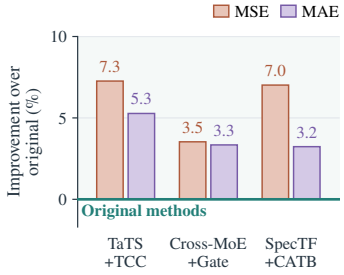


Figure 5: Mean transfer gains across three domains.

### 4.3 Ablation Studies and Component Transfer

We independently train ablations on nine Time-MMD domains with iTransformer and PatchTST, four horizons, and five runs. Replacing TCC with masked pooling, removing CATB, or removing CAGI while retaining cross-attention increases average error (Figure 6(a)). Each removal increases mean MSE in at least eight domains (Table 2), **supporting all three controls**.

Under the same protocol, CATB yields lower mean MSE in every domain than no bias, fixed per-head linear slopes (*ALiBi-style*; Press et al., 2022), or learned slopes shared across inputs (*learned linear*; Table 2). Figure 6(b) supports jointly conditioning temporal relevance on content and applying logarithmic distance weighting.

Table 2: Domain-level MSE increase (%) over full DynaGuard, averaged across iTransformer and PatchTST. Groups remove components or replace CATB. Details: Appendices F.1 and F.2.

| Variant | Agr. | Clim. | Econ. | Ener. | Env. | Hlth. | Sec. | Soc. | Traf. | Mean |
|---|---|---|---|---|---|---|---|---|---|---|
| **Component removal** | | | | | | | | | | |
| w/o TCC | +10.3 | +5.0 | +2.4 | +29.5 | +4.8 | +5.2 | +2.2 | +19.2 | +6.3 | +9.4 |
| w/o CATB | +9.5 | +5.4 | +4.0 | +27.3 | +3.8 | +5.0 | +3.1 | +19.8 | +1.8 | +8.9 |
| w/o CAGI | +31.6 | +3.2 | +6.4 | +29.7 | +5.1 | +6.3 | −0.1 | +13.5 | +2.6 | +10.9 |
| **Temporal bias** | | | | | | | | | | |
| ALiBi-style | +8.8 | +5.1 | +4.4 | +27.6 | +4.8 | +5.3 | +3.1 | +18.3 | +2.9 | +8.9 |
| Learned linear | +8.5 | +5.1 | +4.4 | +30.0 | +4.6 | +5.7 | +3.1 | +18.2 | +3.0 | +9.2 |
| None | +9.5 | +5.4 | +4.0 | +27.3 | +3.8 | +5.0 | +3.1 | +19.8 | +1.8 | +8.9 |

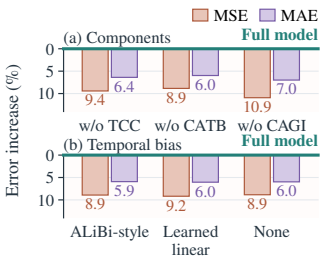


Figure 6: Mean error increases in the two ablations.

**Transfer to other methods.** We add TCC to TaTS, a content gate to Cross-MoE, and CATB to SpecTF on Economy, Energy, and Traffic (Nguyen et al., 2026). TaTS and Cross-MoE use iTransformer; SpecTF retains its spectral architecture. Across four horizons and five matched runs, **all**

**nine method–domain pairs improve both metrics** (Figure 5; Appendix F.3). The controls thus benefit variable-based, expert-based, and spectral fusion.

### 4.4 Contribution of Historical Text

**Observed text versus learned memory.** With both backbones, nine domains, and four horizons, we compare DynaGuard with TSOnly, which substitutes learned memory for observed text and TCC while retaining the remaining adapter. Agriculture and Energy improve both metrics on both backbones (Figure 7; Appendix E.3), demonstrating the value of observed context. Appendix H.1 examines text availability and numerical persistence in Social Good and Traffic.

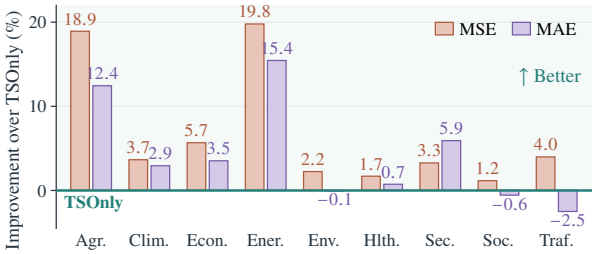


Figure 7: Error reduction over TSOnly, averaged over two backbones within each domain.

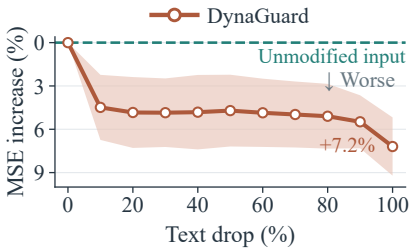


Figure 8: Text drop: mean MSE increase with 95% confidence intervals.

**Missing context.** We evaluate fixed iTransformer models on Agriculture, Energy, and Health, replacing a random fraction of historical texts with *No information available.* Each nonzero drop rate uses 20 draws per model across four horizons and five training seeds. Figure 8 reports MSE increases over each model's unmodified input. The mean increase stays below 6% with up to 90% text removal and reaches 7.2% at complete removal. Appendix G gives the protocol and domain-level results.

### 4.5 Case Study

Figure 9 shows a 24-week gasoline-price forecast from an archived iTransformer checkpoint. For this forecast, disabling CATB, masking all text, and jointly masking the two highlighted reports increase RMSE by 2.7%, 36.4%, and 22.0%, respectively. The checkpoint and numerical input remain fixed. Both text-masking interventions increase RMSE in all five runs (Table 71), supporting the contribution of historical reports; the displayed forecast illustrates the role of temporal weighting. Appendix H.2 defines the interventions and CATB's distance for the single numerical query.

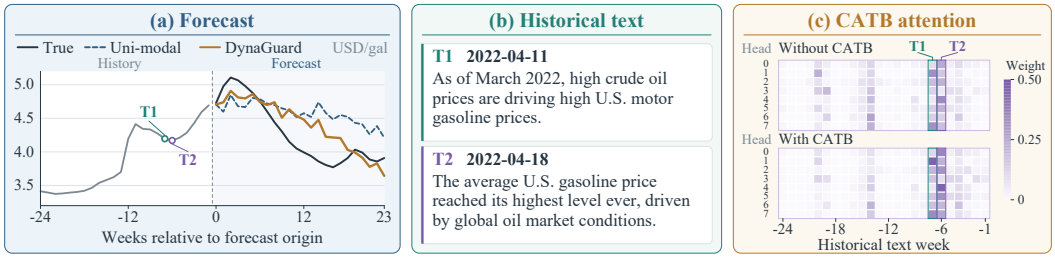


Figure 9: Energy case study: a 24-week gasoline-price forecast with iTransformer and eight-head attention over historical text. Dates denote dataset positions.

## 5 Conclusion

DynaGuard controls text selection, temporal relevance, and injection while preserving the numerical path. It achieves the **lowest negative-transfer rate among text-based methods** and best mean MSE rank in all three collections. Finer event representations and richer text are future directions. Ablations support the design, and component transfer benefits other methods.

### AI USE STATEMENT

Generative AI tools were used to assist with generating portions of the code, debugging and refining code, improving the manuscript's structure and language, and monitoring automated experiment runs. They also assisted with LaTeX formatting, figure styling, reference cross-checking, scripts for assembling and validating tables and figures, consistency checks among mathematical expressions, experimental descriptions, and the implementation, and the summarization and analysis of recorded experimental results. All AI-assisted outputs were reviewed and verified by the authors. The authors developed the scientific ideas, designed the methodology and experiments, made all final implementation and analysis decisions, and take full responsibility for the content and results of this work.

### REPRODUCIBILITY STATEMENT

Code and reproduction instructions are available at https://anonymous.4open.science/r/DynaGuard2027/. Section 3 defines DynaGuard, and Appendix C derives CATB's temporal decay. Appendix D documents the data and evaluation protocols and FNSPID baseline adaptations; Appendix D.3 specifies preprocessing, the standard Optuna search spaces for all six compared methods, validation-based configuration selection, staged training, randomness control, and the hardware and software environment. The remaining appendices report aggregate and per-horizon results, seed variation, text-drop sensitivity, and case-study interventions. Appendix I describes the independent resource measurements and cache-placement check.

## REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2623–2631, 2019. doi: 10.1145/3292500.3330701.

John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A. González. Gated multimodal units for information fusion. In *International Conference on Learning Representations Workshop*, 2017.

James S. Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, volume 24, 2011. URL https://proceedings.neurips.cc/paper/2011/hash/86e8f7ab32cfd12577bc2619bc635690-Abstract.html.

Ta-Chung Chi, Ting-Han Fan, Peter J. Ramadge, and Alexander Rudnicky. KERPLE: Kernelized relative positional embedding for length extrapolation. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8386–8399, 2022. doi: 10.52202/068431-0610.

Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10148–10167, 2024.

Andrey Davydenko and Robert Fildes. Measuring forecasting accuracy: The case of judgmental adjustments to SKU-level demand forecasts. *International Journal of Forecasting*, 29(3):510–522, 2013. doi: 10.1016/j.ijforecast.2012.09.002.

Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7:1–30, 2006.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019. doi: 10.18653/v1/N19-1423. URL https://aclanthology.org/N19-1423/.

Zihan Dong, Xinyu Fan, and Zhiyuan Peng. FNSPID: A comprehensive financial news dataset in time series. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4918–4927, 2024. doi: 10.1145/3637528.3671629.

Lu Han, Xu-Yang Chen, Han-Jia Ye, and De-Chuan Zhan. SOFTS: Efficient multivariate time series forecasting with series-core fusion. In *Advances in Neural Information Processing Systems*, volume 37, pp. 64145–64175, 2024. doi: 10.52202/079017-2046.

Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979.

Ruizheng Huang, Zhicheng Zhang, and Yong Wang. Cross-MoE: An efficient temporal prediction framework integrating textual modality. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 29927–29938, 2025. doi: 10.18653/v1/2025.emnlp-main.1520.

Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and João Carreira. Perceiver: General perception with iterative attention. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4651–4664. PMLR, 2021. URL https://proceedings.mlr.press/v139/jaegle21a/jaegle21a.pdf.

Furong Jia, Kevin Wang, Yixiang Zheng, Defu Cao, and Yan Liu. GPT4MTS: Prompt-based large language model for multimodal time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 23343–23351, 2024. doi: 10.1609/aaai.v38i21.30383.

Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y. Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen. Time-LLM: Time series forecasting by reprogramming large language models. In *International Conference on Learning Representations*, 2024.

Seunghan Lee, Jun Seo, Jaehoon Lee, Sungdong Yoo, Minjae Kim, Tae Yoon Lim, Dongwan Kang, Hwanil Choi, SoonYoung Lee, and Wonbin Ahn. Rethinking multimodal fusion for time series forecasting: Text modalities need constrained fusion. In *12th KDD Workshop on Mining and Learning from Time Series*, 2026. URL https://kdd-milets.github.io/milets2026/accepted%20papers/6paper_6_Rethinking%20Multimodal%20Fusion%20for%20Time%20Series%20Forecasting;%20Text%20Modalities%20Need%20Constrained%20Fusion.pdf.

Lehui Li, Yuyao Wang, Jisheng Yan, Wei Zhang, Jinliang Deng, Haoliang Sun, Zhongyi Han, and Yongshun Gong. From text to forecasts: Bridging modality gap with temporal evolution semantic space. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026a. URL https://icml.cc/virtual/2026/poster/63982.

Zihao Li, Xiao Lin, Zhining Liu, Jiaru Zou, Ziwei Wu, Lecheng Zheng, Dongqi Fu, Yada Zhu, Hendrik Hamann, Hanghang Tong, and Jingrui He. Language in the flow of time: Time-series-paired texts weaved into a unified temporal narrative. In *International Conference on Learning Representations*, 2026b.

Haoxin Liu, Shangqing Xu, Zhiyuan Zhao, Lingkai Kong, Harshavardhan Kamarthi, Aditya B. Sasanur, Megha Sharma, Jiaming Cui, Qingsong Wen, Chao Zhang, and B. Aditya Prakash. Time-MMD: Multi-domain multimodal dataset for time series analysis. In *Advances in Neural Information Processing Systems*, volume 37, pp. 77888–77933, 2024a. doi: 10.52202/079017-2476. Datasets and Benchmarks Track.

Haoxin Liu, Yichen Zhou, Rajat Sen, B. Aditya Prakash, and Abhimanyu Das. Rethinking post-training recipes for multimodal time-series forecasting. *arXiv preprint arXiv:2605.29401*, 2026. doi: 10.48550/arXiv.2605.29401.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. iTransformer: Inverted transformers are effective for time series forecasting. In *International Conference on Learning Representations*, 2024b.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Huu Hiep Nguyen, Minh Hoang Nguyen, Dung Nguyen, and Hung Le. Spectral text fusion: A frequency-aware approach to multimodal time-series forecasting. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 4213–4221. PMLR, 2026.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *International Conference on Learning Representations*, 2023.

Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=R8sQPpGCv0.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026. URL https://qwen.ai/blog?id=qwen3.5.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019. URL https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.

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

Siyuan Wang, Peng Chen, Yihang Wang, Wanghui Qiu, Chenjuan Guo, Bin Yang, and Yang Shu. Unlocking the value of text: Event-driven reasoning and multi-level alignment for time series forecasting. In *International Conference on Learning Representations*, 2026.

Yuxuan Wang, Haixu Wu, Jiaxiang Dong, Guo Qin, Haoran Zhang, Yong Liu, Yunzhong Qiu, Jianmin Wang, and Mingsheng Long. TimeXer: Empowering transformers for time series forecasting with exogenous variables. In *Advances in Neural Information Processing Systems*, volume 37, pp. 469–498, 2024. doi: 10.52202/079017-0015.

Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. Characterizing and avoiding negative transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11293–11302, 2019.

Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945. doi: 10.2307/3001968.

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

Chuanyang Zheng, Yihang Gao, Han Shi, Minbin Huang, Jingyao Li, Jing Xiong, Xiaozhe Ren, Michael Ng, Xin Jiang, Zhenguo Li, and Yu Li. DAPE: Data-adaptive positional encoding for length extrapolation. In *Advances in Neural Information Processing Systems*, volume 37, pp. 26659–26700, 2024. doi: 10.52202/079017-0838.

Xin Zhou, Weiqing Wang, Shilin Qu, Zhiqiang Zhang, and Christoph Bergmeir. Unveiling the potential of text in high-dimensional time series forecasting. In *NeurIPS Workshop on Time Series in the Age of Large Models*, 2024.

# APPENDIX CONTENTS

**A Limitations** **14**

**B Extended Related Work** **14**
B.1 Numerical Forecasting as the Foundation . . . . . 14
B.2 External Text Resources and Forecasting Systems . . . . . 14
B.3 Controlled Fusion and Negative Transfer . . . . . 15

**C Interpretation of the CATB Temporal Decay** **15**

**D Experimental Protocol** **16**
D.1 Datasets and Evaluation Protocol . . . . . 16
D.2 Relative Error Metrics and Aggregation . . . . . 20
D.3 Reproducibility Details . . . . . 21

**E Matched Framework and TSOnly Comparisons** **23**
E.1 Matched Framework Comparison . . . . . 23
E.2 Per-Horizon Matched Framework Comparison . . . . . 27
E.3 Comparison with TSOnly . . . . . 42

**F Component and Temporal-Bias Studies** **44**
F.1 Component Ablations . . . . . 44
F.2 Temporal-Bias Ablations . . . . . 49
F.3 Component Transfer . . . . . 54

**G Sensitivity to Missing Context** **56**
G.1 Text-Drop Results by Forecast Horizon . . . . . 57

**H Data Characteristics and Case Study** **59**
H.1 Text Availability and Numerical Persistence . . . . . 59
H.2 Case Study . . . . . 59

**I Resource Cost and Cache Placement** **59**

**J Complete-System Comparison** **61**
J.1 Per-Horizon Complete-System Comparison . . . . . 61

**K Pretrained Text-Embedding Ablation** **63**
K.1 Per-Horizon Pretrained Text-Embedding Ablation Results . . . . . 63
K.2 Additional FNSPID and GPT4MTS Results . . . . . 66

# A LIMITATIONS

TCC compresses the text in each time slice into a single token, while CATB imposes a temporal prior that decays with age when content and query are fixed. Finer-grained event representations and more flexible temporal priors could help distinguish concurrent events and capture delayed effects.

The forecasting benchmarks used in this study also have limitations in text quality, including incomplete coverage in some domains. Richer text and more precise event timestamps would support a more detailed evaluation of historical context.

# B EXTENDED RELATED WORK

This appendix relates DynaGuard's three controls—within-slice content selection, temporal weighting, and feature-wise correction—to numerical forecasting and external-text forecasting systems.

## B.1 NUMERICAL FORECASTING AS THE FOUNDATION

Numerical forecasting backbones organize the observed series in different ways. PatchTST represents a series as temporal patches and uses channel-independent processing (Nie et al., 2023); Crossformer models dependencies across both time segments and variables (Zhang & Yan, 2023); and iTransformer treats variables as tokens to model multivariate correlations (Liu et al., 2024b). SOFTS uses a centralized series representation for efficient cross-variable interaction (Han et al., 2024), whereas TimeXer is designed to incorporate generic exogenous variables into a forecasting Transformer (Wang et al., 2024). DynaGuard uses the representations learned by these backbones for text-conditioned correction. Its residual interface accommodates their different token organizations while preserving the direct numerical path to the prediction head.

An independently trained Uni-modal model with the same numerical architecture, data, and evaluation protocol provides the reference for negative transfer. Comparisons across backbones test whether context injection benefits different numerical representations.

## B.2 EXTERNAL TEXT RESOURCES AND FORECASTING SYSTEMS

Forecasting methods use endogenous text derived from the observed series and exogenous text obtained from external sources (Wang et al., 2026). Time-LLM includes series statistics and trend information in its prompts (Jin et al., 2024). The setting studied here uses exogenous documents, such as news and reports, to supply information beyond numerical history. GPT4MTS jointly uses numerical and textual inputs (Jia et al., 2024), while Time-MMD provides timestamp-aligned numerical–text data across multiple domains together with the MM-TSFlib evaluation framework (Liu et al., 2024a). FNSPID pairs financial news with market observations (Dong et al., 2024).

**Representations and interaction.** MM-TSFlib learns weights for modality-specific forecasts. TaTS represents paired texts as auxiliary temporal variables, allowing numerical forecasting architectures to process them together with the series (Li et al., 2026b). Cross-MoE uses a Cross-Ranker to select text tokens after expert transformations of numerical and textual features (Huang et al., 2025). Zhou et al. (2024) use a dual-tower structure to combine numerical and textual representations for high-dimensional forecasting. Together, these approaches show how text can enter as an additional input, an interacting representation, or a separate forecasting branch. DynaGuard places its controls at the representation stage, before the prediction head.

**Event reasoning and structured context.** VoT uses historical examples to guide event-driven LLM reasoning and aligns textual and numerical information at representation and prediction levels (Wang et al., 2026). TESS maps text into temporal primitives describing shifts, volatility, shape, lag and decay, then applies confidence-aware filtering (Li et al., 2026a). These methods connect language to forecasting through explicit intermediate descriptions. DynaGuard learns its compressed context and injection controls directly from forecasting error, using the supplied text embeddings.

**Frequency fusion and forecast revision.** SpecTF projects text embeddings into the frequency domain and uses cross-attention to reweight numerical frequency components (Nguyen et al., 2026).

PostTime trains an LLM to revise or preserve a numerical foundation model's forecast using external context (Liu et al., 2026). These approaches place control at different stages: frequency components in SpecTF and forecast outputs in PostTime. DynaGuard controls a residual update to the numerical representation, with an explicit age-dependent bias for historical text slices.

**Scope of the main comparison.** These methods use exogenous text to improve numerical forecasts, but change different parts of the prediction process. Our main comparison evaluates fusion mechanisms across matched numerical backbones, with shared historical text, frozen pretrained input embeddings, and evaluation protocols. The matched numerical-only models provide the reference for negative transfer. TESS introduces a structured text representation and additional gating supervision, while PostTime combines a numerical foundation model's forecast with an LLM trained through supervised fine-tuning and reinforcement learning to revise it (Li et al., 2026a; Liu et al., 2026). Retaining their original designs would vary the text representation or numerical prior alongside the fusion mechanism. TESS and PostTime therefore fall outside this controlled fusion comparison. Appendix J provides complementary comparisons with selected complete systems.

### B.3 Controlled Fusion and Negative Transfer

The negative-transfer perspective evaluates fusion by comparing prediction error with a reference trained without the auxiliary source (Wang et al., 2019).

CFA applies this perspective to forecasting through asymmetric, low-rank fusion: numerical features form the anchor, and constrained textual residuals supply additional information (Lee et al., 2026). DynaGuard also treats numerical features as the foundation and controls the content, temporal weighting, and strength of each update. The matched framework comparison evaluates these designs on common numerical backbones.

## C Interpretation of the CATB Temporal Decay

**Content-dependent persistence.** CATB assigns each text slice a bounded salience offset and a positive decay coefficient. For one-based query and slice indices $i$ and $j$, the distance is $\Delta_{ij} = r_j + u_i$, where $r_j = T - j$ is historical lag and $u_i = i - 1$ is the numerical-token offset. The temporal factor is $f_\lambda(\Delta_{ij}) = (1 + \Delta_{ij})^{-\lambda}$ for $\lambda > 0$. Holding content and query fixed gives $\partial \log f_\lambda(\Delta_{ij})/\partial r_j = -\lambda/(1 + \Delta_{ij}) < 0$. A larger $\lambda$ therefore imposes a stronger penalty as historical lag increases. The query offset follows the backbone's token sequence, while query–context compatibility also depends on the numerical representation.

**Mixture of decay rates.** The power-law factor has the exact representation

$$(1 + \Delta_{ij})^{-\lambda} = \frac{1}{\Gamma(\lambda)} \int_0^\infty r^{\lambda-1} e^{-r} e^{-r\Delta_{ij}}\, dr. \tag{15}$$

To obtain this identity, substitute $u = (1 + \Delta_{ij})r$ in the Gamma integral. The remaining integral is $\Gamma(\lambda)$, leaving the factor $(1 + \Delta_{ij})^{-\lambda}$. Equivalently, if $R \sim \mathrm{Gamma}(\lambda, \mathrm{rate} = 1)$, then $f_\lambda(\Delta_{ij}) = \mathbb{E}[e^{-R\Delta_{ij}}]$. CATB therefore admits a mixture interpretation over exponential decay rates. The mixing distribution describes an attention factor whose parameters learn from forecasting error; it does not estimate physical event lifetimes.

**Effect of attention normalization.** For a fixed head, suppress its index and let $\widetilde{w}_{ij} > 0$ denote the unnormalized weight at a valid slice. The normalized weight is $w_{ij} = \widetilde{w}_{ij} / \sum_{k \in \mathcal{V}} \widetilde{w}_{ik}$, where $\mathcal{V}$ is the set of valid slices. Holding the query, content scores, and all other distances fixed yields

$$\frac{\partial w_{ij}}{\partial \Delta_{ij}} = -\frac{\lambda_j}{1 + \Delta_{ij}}\, w_{ij}(1 - w_{ij}) \le 0. \tag{16}$$

Thus, a slice's normalized weight is nonincreasing in its distance when all other scores remain fixed. With one valid slice, the weight is always one. Changes in queries, content, or competing slices can change this relationship. CATB controls relative weights across slices, while CAGI scales the retrieved features before output projection. The monotonicity result concerns temporal weighting under the stated conditions, not the resulting forecast.

# D EXPERIMENTAL PROTOCOL

## D.1 DATASETS AND EVALUATION PROTOCOL

**Task selection.** We use all nine Time-MMD domains and selected tasks from the larger FNSPID and GPT4MTS collections. The additions follow two criteria: *text quality*, assessed through readability, coverage after cleaning, and stock- or event-related content; and *temporal diversity*, favoring trends and fluctuations that complement Time-MMD. The selected tasks add daily financial-price dynamics and event-related news counts to the domains in Time-MMD.

**Time-MMD.** We use the processed Time-MMD version adopted by the official TaTS implementation (Liu et al., 2024a; Li et al., 2026b). Environment retains the official target sequence from January 1982 through September 2023 (15,248 daily observations). This extends the earlier MM-TSFlib archive, which ends in May 2012 with 11,102 observations. All Time-MMD tasks use a 24-step history and label length 12. The four forecast horizons are $\{6, 8, 10, 12\}$ for monthly domains, $\{12, 24, 36, 48\}$ for Energy and Health, and $\{48, 96, 192, 336\}$ for Environment. Table 3 gives the dataset statistics; Figure 10 shows the target trajectories.

Table 3: Time-MMD dataset statistics.

| Domain | Modeled variables | Points | Frequency |
|---|---|---|---|
| Agriculture | 1 | 496 | Monthly |
| Climate | 1 | 496 | Monthly |
| Economy | 1 | 423 | Monthly |
| Energy | 1 | 1479 | Weekly |
| Environment | 1 | 15,248 | Daily |
| Health | 1 | 1389 | Weekly |
| Security | 1 | 297 | Monthly |
| Social Good | 1 | 900 | Monthly |
| Traffic | 1 | 531 | Monthly |

**FNSPID.** FNSPID pairs historical stock prices with dated financial news (Dong et al., 2024). We select four stock windows spanning retail (WMT), biopharmaceuticals (GILD), banking (WFC), and food service (SBUX). Each task forecasts the adjusted closing price (`avg_close`) from a 24-trading-day history at horizons $\{6, 8, 10, 12\}$. Daily text combines up to two distinct articles, prioritizing usable summaries and falling back to headlines; each article is capped at 80 whitespace-delimited words, including its headline. Text retains the source calendar-date alignment. Table 4 gives the selected periods and text coverage.

Figure 11 contrasts WMT's overall rise with GILD's decline and range fluctuations, WFC's turn downward, and SBUX's rise and abrupt final drop. These windows combine gradual level changes with sharp trading-day movements.

Table 4: Selected FNSPID and GPT4MTS data. Text coverage is the percentage of observations with nonempty text after preprocessing. All tasks are univariate.

| Task | Start date | End date | Observations | Text (%) |
|---|---|---|---|---|
| **FNSPID** (trading day) | | | | |
| WMT | 2019-11-13 | 2023-11-02 | 1,000 | 62.9 |
| GILD | 2014-09-12 | 2020-08-26 | 1,500 | 62.9 |
| WFC | 2015-12-11 | 2020-11-27 | 1,250 | 52.1 |
| SBUX | 2016-04-05 | 2020-03-24 | 1,000 | 51.5 |
| **GPT4MTS** (daily) | | | | |
| NumArticles | 2022-08-18 | 2023-07-31 | 348 | 97.7 |
| NumMentions | 2022-08-18 | 2023-07-31 | 348 | 97.7 |
| NumSources | 2022-08-18 | 2023-07-31 | 348 | 97.7 |

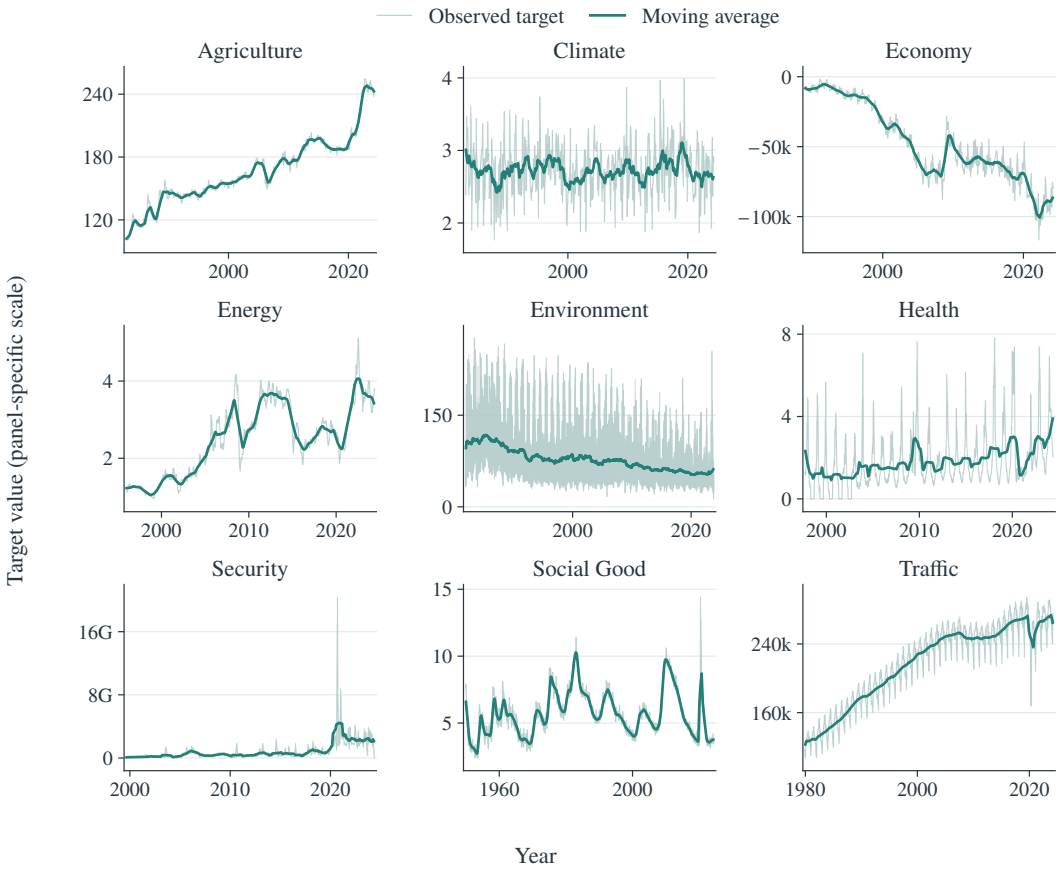


Figure 10: Time-MMD target series. Thin lines show observed targets; dark lines show centered moving averages over 12 months, 52 weeks, or 365 days, according to sampling frequency. Smoothing is used only for visualization. Eight Social Good rows with unavailable targets are omitted. Each panel has its own vertical scale.

**GPT4MTS.** The GDELT-based dataset introduced with GPT4MTS pairs event-related news counts with daily news summaries across regions and event types (Jia et al., 2024). We use event 07 (*Provide Aid*) from the official public US demo. NumArticles, NumMentions, and NumSources measure related articles, mentions, and sources, respectively. Each is a separate univariate forecasting task with a 15-day history and a 7-day horizon. The three targets share the same dates and news summaries. Their comparison tests three measures of news coverage for one event category. The source-group sensitivity analysis treats them as one group (Section D.1).

We retain the released summaries after removing boilerplate, uninformative fragments, and within-day duplicates. Cleaning preserves the retained wording and dates. Figure 12 shows daily spikes and weekly fluctuations around a changing level. The count series add daily event-related variation to the selected numerical forecasting tasks.

**Matched comparison and leakage control.** All three collections use chronological 70%/10%/20% train, validation, and test splits, with standardization fitted on training data only. The framework comparison fixes the split, forecast origin, historical numerical window, text source, and five random seeds. At each forecast origin, text inputs use dataset rows dated before the prediction interval. Unavailable slices are replaced by the fixed placeholder "No information available." Each task uses six backbones and six methods. The collection summaries cover 54 Time-MMD, 24 FNSPID, and 18 GPT4MTS task–backbone settings separately.

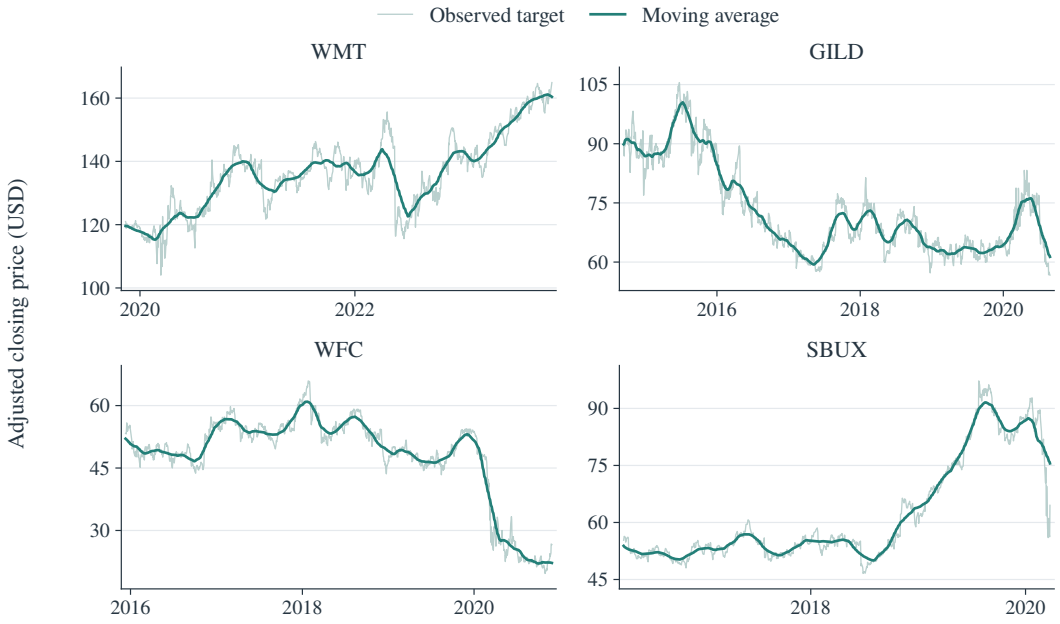


Figure 11: Selected FNSPID targets over the complete experimental windows. Thin lines show adjusted closing prices; dark lines show centered 60-trading-day moving averages for visualization only. Dates and vertical scales differ across stocks.

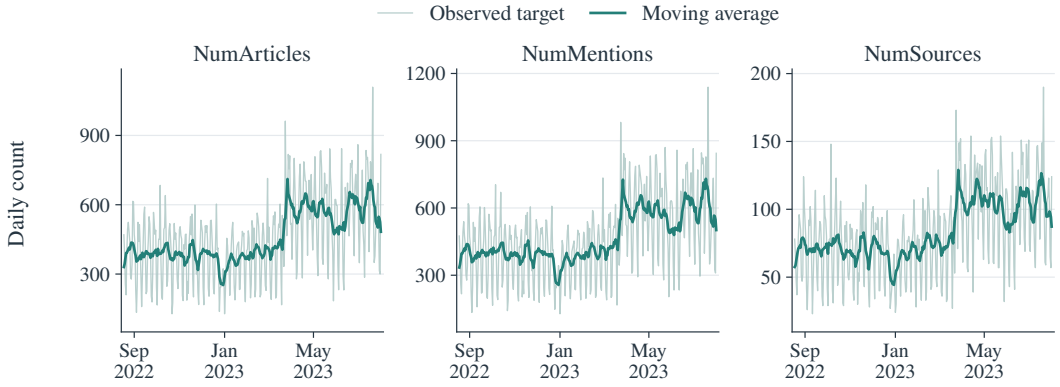


Figure 12: Selected GPT4MTS targets for the US Provide Aid event category. Thin lines show daily counts; dark lines show centered 7-day moving averages for visualization only. The three targets share dates and text; each panel uses its own vertical scale.

**Scope of temporal alignment.** Benchmark dates index the historical input slices; they do not by themselves certify the original publication time of every source document. The comparisons evaluate text integration under the provided benchmark alignment. Deployment-time availability requires separately verified source timestamps.

**FNSPID baseline adaptations.** FNSPID provides no external numerical prior. MM-TSFlib adds the unnormalized text-head output, scaled by its existing fusion weight, to the backbone forecast. Cross-MoE uses its multimodal forecast directly and bypasses the final prior-mixing step; its text experts and auxiliary losses remain active. These adaptations use the same historical inputs and evaluation scale as the other methods.

The FNSPID and GPT4MTS embedding ablations use a common fixed training configuration, separate from the main comparison (Section K.2). Other ablations and diagnostics concern Time-MMD.

**Diagnostic and ablation protocol.** The TSOnly, component, and temporal-bias comparisons use iTransformer and PatchTST on nine domains, four domain-specific horizons, and five seeds. Each variant is independently trained, and full DynaGuard uses the corresponding main-comparison runs. The component and temporal-bias summary averages domain-level percentage changes within each backbone, then averages the two backbones equally. The TSOnly summary averages the two backbones' percentage reductions within each domain. Sections E.3, F.1, and F.2 define the variants.

**Component transfer.** We add TCC to TaTS, a content gate to Cross-MoE, and CATB to SpecTF. TaTS and Cross-MoE use iTransformer; SpecTF retains its spectral architecture. Each comparison uses a 24-step history, five seeds, and four horizons: $\{6, 8, 10, 12\}$ for Economy and Traffic and $\{12, 24, 36, 48\}$ for Energy. Each method–domain pair contains 20 original and 20 adapted runs, matched by data version, horizon, seed, and metric scale. For either metric, the reduction is $100(1 - \bar{E}_{\text{added}}/\bar{E}_{\text{original}})$, with each mean weighting the 20 runs equally. The main-text figure averages these reductions equally over the three domains.

The nine method–domain combinations contain 180 matched pairs (360 runs). Each method–domain–horizon group uses one configuration across all five seeds. The records include run identifiers, source hashes, configurations, and per-horizon errors. This experiment is separate from the matched framework summaries and statistical tests. Section F.3 describes the SpecTF adaptation and reports the results.

**Aggregation.** For Time-MMD and FNSPID, an absolute-error aggregate reports the arithmetic mean of 20 runs and the mean of four horizon-specific sample standard deviations, each computed over five seeds. A per-horizon cell reports the mean and sample standard deviation of its five runs. GPT4MTS uses one horizon, $H = 7$, so its aggregate is also its per-horizon result. Within each main-comparison task–backbone–horizon group, MSE and MAE use the same configuration and five runs.

**Reporting and comparison rules.** MSE determines negative transfer, ranks, and statistical tests; MAE complements it. Negative-transfer decisions, MSE ranks, and win/tie/loss counts use aggregate MSE rounded to three decimals. Best and second-best highlights use the same precision for the corresponding metric. Paired tests use unrounded means. Complete systems with different encoders, objectives, or forecasting structures are compared separately in Section J.

**Statistical testing.** The primary analysis uses one block for each of the 16 tasks. For reference method $m$, task $d$, and $B = 6$ backbones, define the task effect

$$\delta_{d,m} = \frac{1}{B} \sum_{b=1}^{B} \log \frac{E_{\text{DG},d,b}}{E_{m,d,b}}. \tag{17}$$

We use unrounded horizon- and run-averaged MSE and apply a two-sided exact Wilcoxon signed-rank test to the 16 equally weighted task effects (Wilcoxon, 1945; Demšar, 2006). Holm's procedure corrects the five comparisons (Holm, 1979). Table 5 reports these tests alongside descriptive NTR, ranks, and win/tie/loss counts.

The sensitivity analysis averages these task effects within 11 source groups: one per Time-MMD domain, one for FNSPID's four stocks, and one for GPT4MTS's three shared-event targets. We repeat the tests and Holm correction on the 11 equally weighted group effects. This checks whether the conclusion persists when targets sharing a source form one test unit. Both analyses support lower MSE for DynaGuard against every reference.

Table 5: MSE statistics across 16 tasks and six backbones (96 settings). W/T/L counts DynaGuard's wins, ties, and losses after rounding MSE to three decimals. Holm-adjusted $p$-values use unrounded log-MSE ratios averaged over backbones, with 16 tasks as the primary comparison units and 11 source groups as a sensitivity analysis.

| Method | NTR (%) | Mean Rank (MSE) | W/T/L | $p_{\text{Holm}}$ | |
|---|---|---|---|---|---|
| | | | | 16 tasks | 11 source groups |
| Uni-modal | – | 3.29 | 76/6/14 | $< 0.001$ | 0.005 |
| MM-TSFlib | 71.9 | 4.69 | 87/3/6 | $< 0.001$ | 0.005 |
| Cross-MoE | 60.4 | 3.92 | 75/3/18 | $< 0.001$ | 0.005 |
| TaTS | 37.5 | 3.17 | 72/4/20 | 0.001 | 0.005 |
| CFA | 67.7 | 4.17 | 89/2/5 | $< 0.001$ | 0.005 |
| **DynaGuard** | **14.6** | **1.75** | – | – | – |

### D.2 Relative Error Metrics and Aggregation

**Matched errors before normalization.** Let $e^{(q)}_{d,b,m,h,s}$ denote the test error for task $d$, backbone $b$, method $m$, horizon $h$, seed $s$, and metric $q \in \{\text{MSE}, \text{MAE}\}$. For horizon set $\mathcal{H}_d$ and seed set $\mathcal{S}_d$, first compute

$$E^{(q)}_{d,b,m} = \frac{1}{|\mathcal{H}_d|\,|\mathcal{S}_d|} \sum_{h\in\mathcal{H}_d} \sum_{s\in\mathcal{S}_d} e^{(q)}_{d,b,m,h,s}. \tag{18}$$

Each method and its Uni-modal reference share the task, backbone, horizons, seeds, splits, and evaluation scale. Time-MMD and FNSPID use four horizons and five seeds; GPT4MTS uses one horizon and five seeds. We retain the unrounded means when computing ratios.

**Backbone aggregation.** Normalize each mean by the independently trained Uni-modal reference on the same backbone, then take an equally weighted geometric mean over the $B = 6$ backbones:

$$r^{(q)}_{d,b,m} = \frac{E^{(q)}_{d,b,m}}{E^{(q)}_{d,b,\text{uni}}}, \qquad R^{(q)}_{d,m} = \exp\left(\frac{1}{B}\sum_{b=1}^{B} \log r^{(q)}_{d,b,m}\right). \tag{19}$$

We report $R^{(\text{MSE})}_{d,m}$ as rMSE and $R^{(\text{MAE})}_{d,m}$ as rMAE. Geometric aggregation follows the relative-error approach of Davydenko & Fildes (2013) and its use for baseline-normalized forecasting errors in TimesFM (Das et al., 2024); here, the aggregation is over matched backbones.

**Interpretation and weighting.** Uni-modal equals one. Values below or above one indicate a geometric relative error reduction of $100(1 - R)\%$ or increase of $100(R - 1)\%$, respectively. Each backbone contributes one log ratio, independent of its absolute error scale. The result summarizes relative performance across backbones; individual backbones may still exhibit negative transfer when the task summary is below one. Averaging errors before normalization also distinguishes this measure from an average of run-level percentage changes.

**Collection aggregation.** For the tasks $\mathcal{D}$ in a collection, the summary row is

$$R^{(q)}_{\mathcal{D},m} = \exp\left(\frac{1}{|\mathcal{D}|}\sum_{d\in\mathcal{D}} \log R^{(q)}_{d,m}\right). \tag{20}$$

Each task receives equal weight within its collection. All reported tasks have complete matched results, and all errors in the log ratios are finite and positive.

**Variation across seeds.** To obtain the $\pm$ term, repeat the horizon and backbone aggregation separately for each matched seed:

$$R^{(q)}_{d,m,s} = \exp\left[\frac{1}{B}\sum_{b=1}^{B} \log \frac{\sum_{h\in\mathcal{H}_d} e^{(q)}_{d,b,m,h,s}}{\sum_{h\in\mathcal{H}_d} e^{(q)}_{d,b,\text{uni},h,s}}\right], \qquad \sigma^{(q)}_{d,m} = \sqrt{\frac{\sum_{s\in\mathcal{S}_d}(R^{(q)}_{d,m,s} - \overline{R}^{(q)}_{d,m})^2}{|\mathcal{S}_d| - 1}}, \tag{21}$$

where $\overline{R}^{(q)}_{d,m}$ is the arithmetic mean of the seed-specific ratios. We report $R^{(q)}_{d,m} \pm \sigma^{(q)}_{d,m}$. The central estimate pools errors before normalization (Equation 19); the sample standard deviation describes variation among seed-specific ratios. It is not a standard error or confidence interval. For a collection summary, we first geometrically average its task ratios for each seed, then compute the sample standard deviation over five seeds. Uni-modal is $1.000 \pm 0.000$ because each run is normalized by itself.

**Precision and negative transfer.** Ratios and seed standard deviations are displayed to three decimals; ties receive the same best or second-best highlighting. The relative-error standard deviations are computed from seed-specific ratios, separately from the absolute-error spreads in Tables 7–12. NTR and mean rank use task–backbone MSE under the three-decimal comparison rule. Paired tests use unrounded means and task-level log effects, with each task contributing one effect after averaging across its six backbones.

### D.3 Reproducibility Details

**Baseline implementations.** We implement the main-comparison baselines with reference to their official repositories where available: MM-TSFlib[1], Cross-MoE[2], TaTS[3], and CFA[4].

**Hyperparameter search space.** Table 6 specifies the standard Optuna (Akiba et al., 2019) search spaces for Uni-modal, MM-TSFlib, Cross-MoE, TaTS, CFA, and DynaGuard. Each method searches its applicable learning-rate groups, weight decay, and dropout. The backbone architecture, epoch budgets, batch size, and early-stopping settings remain fixed within each search; DynaGuard also fixes its TCC attention-head count.

Table 6: Standard Optuna search spaces for all six compared methods. Learning rates and weight decay use continuous logarithmic domains; dropout is categorical. The base learning rate applies to stage 1 in DynaGuard; its later stages use the additional rates listed below.

| Methods | Hyperparameter | Search domain |
|---|---|---|
| All six methods | Base learning rate | $[10^{-5}, 10^{-3}]$ |
| | Weight decay | $[10^{-5}, 10^{-1}]$ |
| | Dropout | $\{0, 0.1, 0.2, 0.3\}$ |
| MM-TSFlib, TaTS | Text MLP learning rate | $[10^{-4}, 2\times10^{-2}]$ |
| CFA | Text projection learning rate | $[10^{-4}, 2\times10^{-2}]$ |
| | Adapter learning rate | $[10^{-5}, 10^{-3}]$ |
| DynaGuard | Stage 2 TCC learning rate | $[10^{-4}, 2\times10^{-2}]$ |
| | Stage 2 fusion learning rate | $[10^{-4}, 10^{-2}]$ |
| | Stage 3 numerical-model learning rate | $[10^{-7}, 10^{-4}]$ |
| | Stage 3 TCC learning rate | $[10^{-6}, 10^{-3}]$ |
| | Stage 3 fusion learning rate | $[10^{-5}, 10^{-2}]$ |

**Hyperparameter selection.** For all six methods, the standard script runs an independent Optuna study for each task–backbone–horizon–method group, with a fixed trial budget set before the search. It includes the initial configuration and uses a tree-structured Parzen estimator (TPE) sampler (Bergstra et al., 2011) with random startup trials. Candidates are trained with one fixed seed and ranked by validation MSE; ties are resolved by validation MAE and then candidate order. The selected configuration is locked before validation runs with the remaining seeds. Once these runs complete, the script retrains the locked configuration with five seeds for final testing. Both metrics use the same configuration and runs. Test results do not enter the search or change the selected configuration.

[1] https://github.com/AdityaLab/MM-TSFlib
[2] https://github.com/Kilosigh/Cross-MoE
[3] https://github.com/iDEA-iSAIL-Lab-UIUC/TaTS
[4] https://github.com/seunghan96/cfa

**Numerical preprocessing.** We split each processed series chronologically into 70% training, 10% validation, and 20% test data. Validation and test inputs may include the immediately preceding history, while prediction targets stay within their assigned split. Each task predicts one target: `OT` in Time-MMD, `avg_close` in FNSPID, or `value` in GPT4MTS. A standard scaler fitted to the training target transforms all splits. MSE and MAE are evaluated on this standardized scale.

**Text preprocessing.** Historical text comes from the `fact` field in Time-MMD and GPT4MTS and `Lexrank_summary` in FNSPID. Missing entries use "No information available." We tokenize with `openai-community/gpt2-xl`, truncate each slice to 256 tokens, and pad with the end-of-sequence token. Padding is masked; the missing-text placeholder remains a valid input. GPT-2-XL supplies frozen input embeddings, without running its contextual Transformer layers. TCC learns the forecasting-specific compression in Equation 6.

**Adapter configuration and token order.** CATB uses $\alpha = 3$ and $\epsilon = 10^{-4}$. For one-based indices, $\Delta_{ij} = (T - j) + (i - 1)$ combines the text slice's historical lag with its numerical query's position in the backbone output. In our single-target settings, iTransformer and SOFTS each supply one target-variable token, so $\Delta_{1j} = T - j$. PatchTST supplies historical patches in order, and TimeXer supplies historical patches followed by a global token. Transformer supplies historical encoder states before its decoder. Crossformer supplies future segments from its final decoder layer, before that layer's prediction projection. CAGI uses a two-layer MLP with hidden width $d$, GELU, and a final sigmoid. It gates the retrieved features before the output projection. The final gate layer starts with zero weights and biases, giving an initial gate value of $0.5$.

**Prediction mapping and trainable head.** The prediction mapping $\mathcal{P}_\phi$ includes any remaining decoding and output aggregation; backbone-specific decoder inputs and numerical branches are implicit in the notation. This abstraction does not define the parameter groups used in staged optimization. Transformer retains its decoder after injection. Crossformer retains the sum of predictions from all decoder layers, with context injected only before the final layer's prediction projection. The numerical parameters trained in stage 2 are backbone-specific: the output projection in iTransformer and SOFTS, the forecasting head in PatchTST and TimeXer, `decoder.projection` in Transformer, and the full decoder and its positional embeddings in Crossformer. All other numerical parameters remain frozen.

**Randomness control.** Five different random seeds initialize Python, NumPy, PyTorch, CUDA, the data-loader generator, and worker processes. We set `CUBLAS_WORKSPACE_CONFIG=:4096:8`, enable deterministic cuDNN behavior, disable cuDNN benchmarking and TF32, and request deterministic PyTorch algorithms with warning-only fallback. Logs record warnings for operations without deterministic implementations.

**Staged optimization.** The default epoch budgets are 20, 10, and 10 for numerical pretraining, head-and-adapter fitting, and joint refinement, respectively. The standard DynaGuard configuration uses AdamW (Loshchilov & Hutter, 2019), cosine learning-rate decay with step-wise updates and a warmup fraction of 0.1, a batch size of 64, and eight TCC attention heads. The stage-2 prediction-head learning rate is fixed at $10^{-5}$. Each stage uses its own early-stopping state and restores the checkpoint with the lowest validation loss before proceeding. During the second stage, the frozen backbone modules run in evaluation mode, while the prediction head and adapter remain trainable. The pretrained token embeddings remain frozen in all three stages.

**Hardware and software.** Our A100 environment uses an NVIDIA A100 80GB PCIe GPU on a server with two Intel Xeon Silver 4210R CPUs at 2.40 GHz (20 physical cores and 40 hardware threads in total) and approximately 503.5 GiB of system memory. It runs Ubuntu 22.04.5 LTS, Python 3.12.7, PyTorch 2.6.0+cu124 with CUDA runtime 12.4, and cuDNN 9.1.0. Appendix I specifies the environments and measurement boundaries for the resource experiments and the separate cache-placement check.

# E MATCHED FRAMEWORK AND TSONLY COMPARISONS

The result sections follow the main paper: framework and TSOnly comparisons, component and temporal-bias ablations, component transfer, text drop, and the Energy case study. Each experiment's aggregate results precede its horizon-specific tables. Resource measurements, complete-system comparisons, and text-embedding comparisons follow as additional experiments.

## E.1 MATCHED FRAMEWORK COMPARISON

Figure 13 shows Time-MMD MSE changes relative to Uni-modal. Tables 7–12 report the absolute errors underlying the three collections' relative-error summaries in the main paper.




Figure 13: Time-MMD MSE change relative to Uni-modal, computed from unrounded means. Blue indicates lower error; warm colors indicate higher error. Colors saturate at $\pm 20\%$; crosses identify negative transfer under the three-decimal comparison rule. MM-TS, C-MoE, and DG denote MM-TSFlib, Cross-MoE, and DynaGuard.

Table 7: Time-MMD absolute errors (1/2): iTransformer, PatchTST, Crossformer. Means average four horizons and five seeds; uncertainty averages the four horizon-specific seed standard deviations. Bold and underlining mark the lowest and second-lowest displayed means.

| Dataset | Method | iTransformer | | PatchTST | | Crossformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| Agriculture | Uni-modal | 0.312±0.006 | 0.371±0.003 | 0.382±0.010 | 0.411±0.007 | **1.150±0.085** | **0.748±0.028** |
| | MM-TSFlib | 0.342±0.045 | 0.411±0.038 | 0.383±0.019 | 0.428±0.014 | 1.931±0.154 | 0.977±0.041 |
| | Cross-MoE | 0.262±0.010 | 0.342±0.008 | 0.326±0.017 | 0.389±0.012 | 2.144±0.173 | 1.030±0.065 |
| | TaTS | 0.300±0.005 | 0.365±0.003 | 0.381±0.008 | 0.411±0.006 | 1.175±0.067 | 0.749±0.028 |
| | CFA | 0.290±0.011 | 0.358±0.011 | 0.366±0.009 | 0.401±0.007 | 1.210±0.095 | 0.767±0.033 |
| | DynaGuard | **0.254±0.014** | **0.332±0.013** | **0.293±0.012** | **0.350±0.012** | 1.257±0.081 | 0.781±0.036 |
| Climate | Uni-modal | 1.001±0.012 | 0.813±0.005 | 0.963±0.009 | 0.796±0.004 | 0.903±0.006 | 0.753±0.004 |
| | MM-TSFlib | 1.034±0.015 | 0.823±0.008 | 1.058±0.021 | 0.838±0.008 | 0.977±0.016 | 0.790±0.006 |
| | Cross-MoE | 1.014±0.016 | 0.819±0.007 | 0.987±0.017 | 0.807±0.009 | 0.900±0.005 | 0.751±0.003 |
| | TaTS | 0.996±0.013 | 0.807±0.009 | 0.946±0.004 | 0.786±0.003 | **0.896±0.004** | 0.748±0.002 |
| | CFA | 0.996±0.010 | 0.808±0.004 | 0.950±0.013 | 0.787±0.007 | 0.975±0.021 | 0.781±0.008 |
| | DynaGuard | **0.894±0.004** | **0.751±0.002** | **0.896±0.004** | **0.753±0.002** | 0.897±0.003 | **0.746±0.002** |
| Economy | Uni-modal | **0.179±0.004** | **0.330±0.004** | 0.213±0.007 | 0.368±0.008 | 0.388±0.014 | 0.499±0.011 |
| | MM-TSFlib | 0.204±0.024 | 0.351±0.015 | 0.207±0.009 | 0.360±0.008 | 0.581±0.053 | 0.627±0.034 |
| | Cross-MoE | 0.189±0.005 | 0.340±0.005 | **0.186±0.004** | **0.338±0.005** | 0.595±0.039 | 0.639±0.023 |
| | TaTS | 0.184±0.009 | 0.334±0.010 | 0.216±0.009 | 0.370±0.009 | 0.384±0.012 | 0.496±0.010 |
| | CFA | 0.185±0.006 | 0.335±0.007 | 0.220±0.006 | 0.375±0.007 | 0.382±0.022 | 0.494±0.017 |
| | DynaGuard | 0.180±0.005 | **0.330±0.006** | 0.199±0.010 | 0.351±0.012 | **0.373±0.022** | **0.487±0.016** |
| Energy | Uni-modal | 0.298±0.002 | 0.418±0.001 | 0.312±0.002 | 0.427±0.001 | 0.351±0.027 | 0.462±0.015 |
| | MM-TSFlib | 0.248±0.006 | 0.368±0.006 | 0.254±0.004 | 0.372±0.004 | 0.296±0.027 | 0.402±0.019 |
| | Cross-MoE | 0.243±0.005 | 0.362±0.005 | **0.247±0.004** | 0.364±0.004 | 0.312±0.026 | 0.409±0.016 |
| | TaTS | 0.299±0.002 | 0.419±0.002 | 0.312±0.001 | 0.427±0.001 | 0.352±0.027 | 0.461±0.018 |
| | CFA | 0.314±0.010 | 0.433±0.007 | 0.327±0.011 | 0.439±0.007 | 0.380±0.037 | 0.470±0.018 |
| | DynaGuard | **0.241±0.012** | **0.355±0.008** | 0.252±0.004 | **0.363±0.004** | **0.272±0.034** | **0.387±0.024** |
| Environment | Uni-modal | 0.271±0.001 | 0.370±0.001 | **0.272±0.000** | **0.371±0.000** | 0.289±0.006 | 0.415±0.009 |
| | MM-TSFlib | 0.292±0.003 | 0.394±0.003 | 0.323±0.001 | 0.396±0.002 | 0.391±0.021 | 0.506±0.022 |
| | Cross-MoE | 0.289±0.003 | 0.394±0.002 | 0.319±0.001 | 0.394±0.002 | 0.399±0.016 | 0.512±0.019 |
| | TaTS | 0.272±0.000 | **0.369±0.001** | **0.272±0.000** | **0.371±0.001** | 0.289±0.005 | 0.415±0.008 |
| | CFA | 0.271±0.003 | 0.379±0.004 | 0.274±0.001 | 0.375±0.002 | 0.305±0.009 | 0.432±0.013 |
| | DynaGuard | **0.259±0.000** | 0.370±0.001 | **0.272±0.000** | **0.371±0.001** | **0.288±0.009** | **0.406±0.013** |
| Health | Uni-modal | 1.370±0.021 | 0.770±0.006 | 1.297±0.010 | 0.765±0.007 | 1.242±0.051 | 0.739±0.018 |
| | MM-TSFlib | 1.720±0.070 | 0.822±0.011 | 1.657±0.058 | 0.865±0.014 | 1.438±0.081 | 0.803±0.022 |
| | Cross-MoE | 1.456±0.052 | 0.814±0.015 | 1.308±0.015 | 0.785±0.006 | 1.222±0.031 | 0.732±0.015 |
| | TaTS | **1.256±0.063** | **0.752±0.019** | **1.285±0.010** | **0.762±0.005** | **1.216±0.018** | **0.728±0.008** |
| | CFA | 1.428±0.087 | 0.785±0.024 | 1.317±0.064 | 0.785±0.020 | 1.361±0.066 | 0.797±0.027 |
| | DynaGuard | 1.273±0.010 | 0.753±0.005 | 1.304±0.023 | 0.770±0.006 | 1.232±0.017 | **0.728±0.006** |
| Security | Uni-modal | 81.377±1.176 | 4.764±0.077 | 72.506±0.535 | 4.111±0.046 | 86.710±0.950 | 5.340±0.088 |
| | MM-TSFlib | 90.708±4.947 | 5.406±0.297 | 76.020±1.218 | 4.511±0.115 | 88.201±1.067 | 5.484±0.090 |
| | Cross-MoE | 90.728±2.499 | 5.461±0.125 | 84.917±1.986 | 5.091±0.172 | 87.013±0.827 | 5.364±0.077 |
| | TaTS | 76.016±1.215 | 4.442±0.083 | 75.019±1.336 | 4.313±0.117 | 86.887±0.990 | 5.354±0.094 |
| | CFA | 83.247±7.276 | 4.940±0.501 | 74.016±1.833 | 4.198±0.133 | 87.250±1.489 | 5.376±0.125 |
| | DynaGuard | **71.734±0.644** | **3.963±0.097** | **71.322±0.786** | **3.925±0.120** | **85.207±1.051** | **5.194±0.095** |
| Social Good | Uni-modal | 0.957±0.005 | **0.425±0.006** | **0.848±0.009** | 0.434±0.007 | 0.750±0.012 | 0.385±0.022 |
| | MM-TSFlib | 1.116±0.028 | 0.482±0.019 | 0.899±0.014 | 0.459±0.011 | 0.766±0.029 | 0.427±0.050 |
| | Cross-MoE | 1.238±0.052 | 0.489±0.023 | 0.968±0.044 | 0.478±0.032 | 0.778±0.035 | 0.429±0.040 |
| | TaTS | 0.961±0.005 | 0.426±0.003 | **0.848±0.009** | 0.434±0.007 | 0.755±0.014 | 0.387±0.022 |
| | CFA | 1.229±0.060 | 0.521±0.026 | 0.940±0.042 | 0.486±0.038 | 0.757±0.029 | 0.393±0.039 |
| | DynaGuard | **0.928±0.016** | 0.428±0.007 | 0.863±0.013 | **0.431±0.011** | **0.731±0.006** | **0.347±0.012** |
| Traffic | Uni-modal | 0.166±0.002 | 0.218±0.003 | 0.171±0.003 | 0.224±0.002 | 0.170±0.002 | 0.302±0.005 |
| | MM-TSFlib | 0.178±0.003 | 0.240±0.005 | 0.186±0.005 | 0.261±0.004 | 0.214±0.008 | 0.364±0.012 |
| | Cross-MoE | 0.184±0.004 | 0.235±0.006 | 0.163±0.003 | 0.227±0.004 | 0.216±0.007 | 0.369±0.009 |
| | TaTS | 0.169±0.002 | 0.221±0.003 | 0.171±0.002 | 0.225±0.002 | 0.170±0.002 | 0.302±0.005 |
| | CFA | 0.169±0.006 | **0.211±0.005** | 0.166±0.005 | **0.223±0.009** | **0.165±0.003** | **0.290±0.007** |
| | DynaGuard | **0.162±0.004** | 0.218±0.005 | **0.155±0.003** | 0.230±0.007 | 0.167±0.004 | 0.296±0.009 |

Table 8: Time-MMD absolute errors (2/2): SOFTS, TimeXer, Transformer. Means average four horizons and five seeds; uncertainty averages the four horizon-specific seed standard deviations. Bold and underlining mark the lowest and second-lowest displayed means.

| Dataset | Method | SOFTS | | TimeXer | | Transformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| Agriculture | Uni-modal | 0.369±0.011 | 0.401±0.009 | 0.387±0.005 | 0.416±0.004 | 1.320±0.162 | 0.801±0.070 |
| | MM-TSFlib | 0.312±0.010 | 0.377±0.007 | 0.367±0.014 | 0.413±0.011 | 2.079±0.353 | 0.987±0.094 |
| | Cross-MoE | 0.539±0.018 | 0.555±0.019 | 0.327±0.013 | 0.387±0.009 | 1.615±0.465 | 0.878±0.113 |
| | TaTS | 0.369±0.010 | 0.401±0.008 | 0.387±0.005 | 0.416±0.004 | **0.973±0.094** | **0.678±0.024** |
| | CFA | 0.308±0.006 | 0.374±0.005 | 0.340±0.007 | 0.385±0.005 | 1.228±0.153 | 0.771±0.053 |
| | DynaGuard | **0.266±0.019** | **0.338±0.013** | **0.287±0.018** | **0.352±0.011** | 1.037±0.391 | 0.730±0.152 |
| Climate | Uni-modal | 1.001±0.009 | 0.808±0.004 | 0.950±0.011 | 0.788±0.006 | 0.954±0.050 | 0.772±0.020 |
| | MM-TSFlib | 1.090±0.017 | 0.846±0.008 | 1.033±0.018 | 0.826±0.008 | 1.031±0.032 | 0.809±0.012 |
| | Cross-MoE | 1.018±0.011 | 0.816±0.006 | 0.935±0.008 | 0.781±0.005 | 0.938±0.032 | 0.765±0.012 |
| | TaTS | 0.991±0.009 | 0.804±0.004 | 0.938±0.008 | 0.782±0.005 | **0.880±0.007** | **0.744±0.003** |
| | CFA | 1.007±0.016 | 0.809±0.006 | 0.965±0.015 | 0.793±0.008 | 1.092±0.058 | 0.831±0.024 |
| | DynaGuard | **0.895±0.005** | **0.752±0.002** | **0.895±0.005** | **0.752±0.002** | 0.897±0.019 | 0.747±0.006 |
| Economy | Uni-modal | 0.195±0.004 | 0.346±0.004 | 0.222±0.005 | 0.377±0.005 | 0.686±0.142 | 0.676±0.075 |
| | MM-TSFlib | 0.194±0.005 | 0.346±0.004 | 0.218±0.020 | 0.371±0.019 | 0.964±0.140 | 0.818±0.066 |
| | Cross-MoE | 0.227±0.009 | 0.372±0.005 | **0.185±0.003** | **0.338±0.003** | 0.905±0.093 | 0.795±0.045 |
| | TaTS | 0.195±0.004 | 0.346±0.004 | 0.222±0.006 | 0.376±0.005 | **0.361±0.030** | **0.478±0.023** |
| | CFA | 0.184±0.016 | 0.335±0.017 | 0.226±0.006 | 0.381±0.006 | 0.681±0.116 | 0.672±0.066 |
| | DynaGuard | **0.177±0.006** | **0.328±0.007** | 0.199±0.011 | 0.352±0.011 | 0.432±0.072 | 0.525±0.053 |
| Energy | Uni-modal | 0.318±0.002 | 0.431±0.001 | 0.318±0.001 | 0.430±0.001 | 0.409±0.044 | 0.500±0.023 |
| | MM-TSFlib | **0.256±0.005** | **0.375±0.004** | 0.253±0.002 | 0.371±0.002 | 0.332±0.041 | 0.437±0.031 |
| | Cross-MoE | 0.407±0.012 | 0.491±0.006 | **0.246±0.004** | 0.364±0.004 | 0.330±0.047 | 0.432±0.037 |
| | TaTS | 0.318±0.002 | 0.431±0.001 | 0.318±0.001 | 0.430±0.001 | 0.348±0.018 | 0.460±0.010 |
| | CFA | 0.326±0.013 | 0.439±0.010 | 0.348±0.026 | 0.449±0.015 | 0.433±0.063 | 0.516±0.039 |
| | DynaGuard | 0.260±0.011 | **0.375±0.012** | 0.250±0.008 | **0.360±0.008** | **0.267±0.029** | **0.389±0.023** |
| Environment | Uni-modal | 0.272±0.000 | **0.370±0.000** | **0.272±0.000** | **0.371±0.001** | **0.277±0.005** | **0.395±0.011** |
| | MM-TSFlib | 0.322±0.001 | 0.394±0.001 | 0.323±0.001 | 0.396±0.001 | 0.353±0.026 | 0.472±0.024 |
| | Cross-MoE | 0.328±0.001 | 0.431±0.001 | 0.320±0.001 | 0.394±0.001 | 0.324±0.007 | 0.443±0.009 |
| | TaTS | 0.272±0.000 | **0.370±0.000** | **0.272±0.000** | **0.371±0.001** | **0.277±0.004** | **0.395±0.008** |
| | CFA | 0.275±0.001 | 0.375±0.003 | 0.274±0.001 | 0.375±0.002 | 0.287±0.007 | 0.409±0.012 |
| | DynaGuard | **0.270±0.002** | 0.374±0.004 | **0.272±0.000** | **0.371±0.001** | 0.281±0.003 | 0.396±0.003 |
| Health | Uni-modal | 1.367±0.023 | 0.774±0.006 | 1.282±0.008 | **0.757±0.005** | 1.218±0.038 | 0.753±0.015 |
| | MM-TSFlib | 1.690±0.087 | 0.837±0.018 | 1.654±0.069 | 0.850±0.014 | 1.315±0.086 | 0.790±0.028 |
| | Cross-MoE | 1.412±0.073 | 0.808±0.012 | **1.266±0.012** | 0.770±0.006 | 1.217±0.058 | 0.752±0.018 |
| | TaTS | 1.307±0.041 | 0.772±0.010 | 1.313±0.012 | 0.770±0.006 | **1.202±0.062** | **0.749±0.025** |
| | CFA | 1.415±0.087 | 0.792±0.025 | 1.334±0.057 | 0.782±0.021 | 1.280±0.071 | 0.789±0.027 |
| | DynaGuard | **1.281±0.017** | **0.753±0.007** | 1.308±0.057 | 0.774±0.023 | 1.234±0.071 | 0.756±0.025 |
| Security | Uni-modal | 78.557±2.410 | 4.604±0.163 | 72.641±0.882 | 4.139±0.083 | 85.814±2.992 | 5.236±0.268 |
| | MM-TSFlib | 79.447±1.152 | 4.689±0.080 | 76.720±2.033 | 4.578±0.175 | 87.726±4.897 | 5.443±0.407 |
| | Cross-MoE | 87.723±2.458 | 5.262±0.151 | 85.098±1.691 | 5.128±0.099 | 86.228±4.098 | 5.266±0.362 |
| | TaTS | 77.309±1.934 | 4.490±0.120 | 73.006±0.801 | 4.152±0.073 | **83.890±1.924** | **5.097±0.181** |
| | CFA | 75.846±2.009 | 4.390±0.172 | 72.488±0.874 | 4.135±0.079 | 87.763±2.973 | 5.407±0.259 |
| | DynaGuard | **72.861±1.507** | **4.040±0.112** | **71.277±0.392** | **3.906±0.069** | 86.330±3.695 | 5.280±0.337 |
| Social Good | Uni-modal | **0.882±0.006** | **0.419±0.003** | 0.839±0.008 | 0.439±0.004 | 0.742±0.027 | 0.424±0.026 |
| | MM-TSFlib | 1.063±0.026 | 0.459±0.006 | 0.888±0.014 | 0.464±0.007 | 0.830±0.124 | 0.505±0.078 |
| | Cross-MoE | 1.033±0.030 | 0.529±0.007 | 0.947±0.026 | 0.478±0.024 | 0.803±0.041 | 0.466±0.032 |
| | TaTS | 0.913±0.005 | **0.419±0.002** | 0.838±0.008 | 0.438±0.004 | **0.736±0.020** | **0.371±0.022** |
| | CFA | 1.018±0.037 | 0.457±0.022 | 0.950±0.045 | 0.497±0.040 | 0.773±0.035 | 0.417±0.030 |
| | DynaGuard | 0.889±0.024 | 0.423±0.011 | **0.812±0.008** | **0.407±0.008** | 0.808±0.019 | 0.451±0.025 |
| Traffic | Uni-modal | 0.167±0.003 | 0.227±0.004 | 0.172±0.001 | 0.229±0.003 | 0.173±0.004 | 0.314±0.008 |
| | MM-TSFlib | 0.178±0.005 | 0.243±0.007 | 0.188±0.003 | 0.269±0.005 | 0.233±0.009 | 0.389±0.010 |
| | Cross-MoE | 0.199±0.007 | 0.269±0.009 | 0.168±0.002 | 0.238±0.002 | 0.189±0.017 | 0.340±0.030 |
| | TaTS | 0.167±0.004 | 0.227±0.004 | 0.172±0.001 | 0.230±0.003 | 0.165±0.002 | 0.292±0.004 |
| | CFA | 0.168±0.003 | **0.214±0.003** | 0.162±0.005 | 0.230±0.011 | 0.166±0.004 | 0.303±0.008 |
| | DynaGuard | **0.164±0.004** | 0.231±0.009 | **0.152±0.002** | **0.227±0.006** | **0.160±0.005** | **0.276±0.013** |

Table 9: Absolute MSE and MAE on FNSPID (1/2): iTransformer, PatchTST, Crossformer. Means average four horizons and five seeds; uncertainty averages the four horizon-specific seed standard deviations. DynaGuard uses the same configurations as the main table. Bold and underlining mark the lowest and second-lowest displayed means. DG, MM-TS, and C-MoE denote DynaGuard, MM-TSFlib, and Cross-MoE.

| Task | Method | iTransformer | | PatchTST | | Crossformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| WMT | Uni-modal | 0.092±0.004 | 0.240±0.006 | 0.088±0.001 | 0.235±0.002 | 1.059±0.143 | 0.831±0.063 |
| | MM-TS | 0.091±0.004 | 0.239±0.004 | 0.085±0.001 | 0.231±0.002 | 1.078±0.118 | 0.844±0.052 |
| | C-MoE | **0.087**±0.004 | **0.232**±0.005 | **0.083**±0.002 | **0.227**±0.003 | **0.811**±0.091 | **0.745**±0.048 |
| | TaTS | 0.089±0.002 | 0.235±0.003 | 0.088±0.001 | 0.235±0.002 | 1.090±0.151 | 0.839±0.065 |
| | CFA | 0.098±0.008 | 0.245±0.009 | 0.093±0.005 | 0.241±0.006 | 1.107±0.133 | 0.857±0.057 |
| | DG | 0.091±0.002 | 0.237±0.003 | 0.087±0.002 | 0.233±0.002 | 0.971±0.157 | 0.808±0.065 |
| GILD | Uni-modal | **0.056**±0.001 | **0.166**±0.001 | 0.057±0.000 | 0.168±0.001 | 0.063±0.010 | 0.180±0.018 |
| | MM-TS | 0.057±0.001 | 0.167±0.002 | 0.058±0.000 | 0.169±0.001 | 0.063±0.010 | 0.181±0.016 |
| | C-MoE | 0.057±0.001 | **0.166**±0.002 | 0.057±0.001 | **0.166**±0.002 | 0.066±0.005 | 0.188±0.010 |
| | TaTS | 0.057±0.001 | **0.166**±0.002 | 0.057±0.000 | 0.168±0.001 | 0.064±0.007 | 0.183±0.013 |
| | CFA | 0.060±0.002 | 0.173±0.003 | **0.056**±0.001 | 0.167±0.002 | 0.071±0.009 | 0.193±0.015 |
| | DG | 0.057±0.002 | 0.168±0.003 | **0.056**±0.001 | **0.166**±0.002 | **0.058**±0.003 | **0.171**±0.005 |
| WFC | Uni-modal | **0.362**±0.016 | **0.392**±0.009 | 0.439±0.011 | 0.417±0.005 | 13.268±0.440 | 3.191±0.052 |
| | MM-TS | 0.365±0.020 | 0.394±0.011 | 0.446±0.014 | 0.421±0.007 | 13.383±0.374 | 3.206±0.046 |
| | C-MoE | 0.363±0.020 | **0.392**±0.012 | 0.452±0.024 | 0.426±0.014 | **12.459**±0.738 | **3.094**±0.096 |
| | TaTS | 0.389±0.020 | 0.401±0.011 | 0.439±0.013 | 0.416±0.007 | 13.379±0.399 | 3.206±0.047 |
| | CFA | 0.396±0.045 | 0.411±0.024 | 0.461±0.041 | 0.435±0.020 | 12.978±0.567 | 3.159±0.070 |
| | DG | **0.362**±0.013 | 0.395±0.009 | **0.413**±0.014 | **0.407**±0.007 | 12.479±0.636 | 3.095±0.081 |
| SBUX | Uni-modal | 1.173±0.060 | 0.680±0.022 | 1.325±0.016 | 0.720±0.006 | 31.857±1.324 | 5.443±0.118 |
| | MM-TS | 1.162±0.069 | 0.679±0.023 | 1.314±0.020 | 0.721±0.006 | 32.298±1.445 | 5.482±0.129 |
| | C-MoE | 1.099±0.051 | **0.656**±0.013 | 1.363±0.068 | 0.720±0.015 | 30.113±3.119 | 5.273±0.285 |
| | TaTS | 1.187±0.062 | 0.677±0.017 | 1.338±0.023 | 0.724±0.007 | 31.947±0.975 | 5.452±0.085 |
| | CFA | 1.231±0.099 | 0.702±0.030 | 1.418±0.106 | 0.721±0.019 | 32.650±2.062 | 5.508±0.178 |
| | DG | **1.091**±0.032 | 0.667±0.010 | **1.225**±0.039 | **0.689**±0.007 | **29.822**±2.415 | **5.254**±0.223 |

Table 10: Absolute MSE and MAE on FNSPID (2/2): SOFTS, TimeXer, Transformer. Means average four horizons and five seeds; uncertainty averages the four horizon-specific seed standard deviations. DynaGuard uses the same configurations as the main table. Bold and underlining mark the lowest and second-lowest displayed means. DG, MM-TS, and C-MoE denote DynaGuard, MM-TSFlib, and Cross-MoE.

| Task | Method | SOFTS | | TimeXer | | Transformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| WMT | Uni-modal | 0.093±0.002 | 0.239±0.002 | 0.087±0.001 | 0.234±0.001 | 0.804±0.247 | 0.711±0.137 |
| | MM-TS | **0.090**±0.002 | 0.237±0.002 | 0.084±0.001 | 0.230±0.001 | 0.789±0.221 | 0.709±0.120 |
| | C-MoE | **0.090**±0.004 | **0.236**±0.006 | **0.082**±0.002 | **0.226**±0.002 | **0.757**±0.367 | **0.678**±0.217 |
| | TaTS | 0.092±0.002 | 0.239±0.002 | 0.087±0.001 | 0.234±0.002 | 0.815±0.158 | 0.724±0.077 |
| | CFA | 0.099±0.009 | 0.246±0.009 | 0.091±0.006 | 0.238±0.007 | 1.116±0.270 | 0.866±0.124 |
| | DG | 0.093±0.002 | 0.239±0.003 | 0.085±0.002 | 0.231±0.003 | 0.779±0.264 | 0.691±0.132 |
| GILD | Uni-modal | 0.057±0.001 | **0.166**±0.002 | 0.057±0.001 | 0.168±0.001 | **0.059**±0.005 | 0.178±0.009 |
| | MM-TS | 0.057±0.001 | 0.167±0.002 | 0.058±0.001 | 0.169±0.002 | 0.062±0.008 | 0.181±0.013 |
| | C-MoE | **0.056**±0.001 | **0.166**±0.003 | **0.056**±0.001 | **0.166**±0.002 | 0.071±0.019 | 0.191±0.026 |
| | TaTS | 0.057±0.001 | **0.166**±0.002 | 0.057±0.001 | 0.168±0.001 | 0.069±0.009 | 0.193±0.016 |
| | CFA | 0.061±0.003 | 0.173±0.005 | **0.056**±0.001 | 0.167±0.002 | 0.064±0.010 | 0.183±0.013 |
| | DG | 0.057±0.001 | 0.167±0.002 | **0.056**±0.001 | **0.166**±0.001 | **0.059**±0.002 | **0.177**±0.005 |
| WFC | Uni-modal | 0.383±0.005 | 0.396±0.003 | 0.440±0.010 | 0.416±0.005 | 13.635±1.288 | 3.247±0.157 |
| | MM-TS | 0.390±0.007 | 0.399±0.004 | 0.453±0.013 | 0.421±0.006 | 13.112±1.249 | 3.182±0.152 |
| | C-MoE | 0.374±0.014 | **0.393**±0.007 | 0.439±0.024 | 0.418±0.013 | 13.358±1.631 | 3.190±0.198 |
| | TaTS | 0.386±0.004 | 0.397±0.003 | 0.442±0.011 | 0.416±0.005 | **7.559**±0.875 | **2.415**±0.141 |
| | CFA | 0.437±0.050 | 0.427±0.026 | 0.479±0.055 | 0.445±0.028 | 14.642±1.144 | 3.371±0.134 |
| | DG | **0.372**±0.007 | **0.393**±0.004 | **0.414**±0.012 | **0.406**±0.006 | 12.201±2.434 | 3.052±0.310 |
| SBUX | Uni-modal | 1.140±0.014 | 0.670±0.004 | 1.324±0.029 | 0.715±0.009 | 32.157±1.646 | 5.440±0.145 |
| | MM-TS | 1.134±0.033 | 0.670±0.010 | 1.291±0.025 | 0.709±0.006 | 32.356±1.773 | 5.457±0.157 |
| | C-MoE | 1.102±0.024 | **0.660**±0.006 | 1.353±0.068 | 0.714±0.017 | 26.486±3.858 | 4.908±0.397 |
| | TaTS | 1.133±0.020 | 0.668±0.005 | 1.315±0.026 | 0.712±0.007 | **20.495**±1.237 | **4.287**±0.139 |
| | CFA | 1.396±0.182 | 0.702±0.030 | 1.396±0.105 | 0.712±0.019 | 34.368±1.994 | 5.645±0.176 |
| | DG | **1.097**±0.034 | 0.664±0.009 | **1.186**±0.041 | **0.682**±0.009 | 26.644±3.419 | 4.894±0.337 |

Table 11: Absolute MSE and MAE on GPT4MTS (1/2): iTransformer, PatchTST, Crossformer. Entries are means and sample standard deviations over five seeds at $H = 7$. DynaGuard uses the same configurations as the main table. Bold and underlining mark the lowest and second-lowest displayed means. DG, MM-TS, and C-MoE denote DynaGuard, MM-TSFlib, and Cross-MoE.

| Task | Method | iTransformer | | PatchTST | | Crossformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| NumArticles | Uni-modal | 0.647±0.009 | 0.644±0.003 | 0.806±0.016 | 0.756±0.008 | 0.695±0.005 | 0.675±0.003 |
| | MM-TS | 0.787±0.016 | 0.660±0.008 | 1.034±0.009 | 0.756±0.006 | **0.688**±0.002 | **0.663**±0.004 |
| | C-MoE | 0.734±0.014 | 0.646±0.007 | 0.946±0.015 | 0.719±0.009 | 0.717±0.010 | 0.685±0.009 |
| | TaTS | 0.673±0.003 | 0.659±0.004 | 0.808±0.018 | 0.757±0.006 | 0.707±0.006 | 0.677±0.005 |
| | CFA | 0.646±0.016 | **0.638**±0.007 | 0.808±0.023 | 0.756±0.013 | 0.699±0.023 | 0.681±0.013 |
| | DG | **0.635**±0.019 | **0.638**±0.013 | **0.706**±0.010 | **0.693**±0.013 | 0.698±0.018 | 0.674±0.015 |
| NumMentions | Uni-modal | 0.669±0.010 | 0.655±0.005 | 0.827±0.018 | 0.764±0.007 | 0.713±0.003 | 0.682±0.001 |
| | MM-TS | 0.805±0.017 | 0.667±0.007 | 1.055±0.010 | 0.762±0.006 | 0.706±0.003 | 0.672±0.004 |
| | C-MoE | 0.767±0.006 | 0.656±0.003 | 0.961±0.014 | 0.726±0.007 | 0.738±0.011 | 0.692±0.008 |
| | TaTS | 0.693±0.004 | 0.668±0.005 | 0.835±0.028 | 0.768±0.009 | 0.723±0.005 | 0.686±0.005 |
| | CFA | 0.664±0.016 | **0.648**±0.007 | 0.832±0.025 | 0.765±0.012 | 0.710±0.012 | 0.682±0.006 |
| | DG | **0.654**±0.016 | **0.648**±0.010 | **0.754**±0.024 | **0.717**±0.009 | **0.689**±0.017 | **0.664**±0.011 |
| NumSources | Uni-modal | **0.497**±0.002 | **0.556**±0.004 | 0.694±0.009 | 0.685±0.004 | 0.582±0.009 | 0.596±0.006 |
| | MM-TS | 0.624±0.011 | 0.594±0.006 | 0.868±0.006 | 0.699±0.004 | **0.561**±0.003 | **0.580**±0.005 |
| | C-MoE | 0.568±0.007 | 0.564±0.004 | 0.777±0.015 | **0.653**±0.008 | 0.604±0.011 | 0.610±0.010 |
| | TaTS | 0.545±0.019 | 0.582±0.013 | 0.695±0.010 | 0.685±0.004 | 0.588±0.007 | 0.596±0.006 |
| | CFA | 0.498±0.033 | 0.562±0.023 | 0.693±0.015 | 0.677±0.010 | 0.573±0.014 | 0.592±0.007 |
| | DG | 0.505±0.017 | 0.569±0.012 | **0.651**±0.014 | **0.653**±0.009 | 0.604±0.029 | 0.609±0.012 |

Table 12: Absolute MSE and MAE on GPT4MTS (2/2): SOFTS, TimeXer, Transformer. Entries are means and sample standard deviations over five seeds at $H = 7$. DynaGuard uses the same configurations as the main table. Bold and underlining mark the lowest and second-lowest displayed means. DG, MM-TS, and C-MoE denote DynaGuard, MM-TSFlib, and Cross-MoE.

| Task | Method | SOFTS | | TimeXer | | Transformer | |
|---|---|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE | MSE | MAE |
| NumArticles | Uni-modal | 0.694±0.006 | 0.666±0.003 | 0.778±0.013 | 0.748±0.010 | 0.882±0.027 | 0.734±0.011 |
| | MM-TS | 0.921±0.003 | 0.714±0.001 | 1.017±0.019 | 0.753±0.008 | 0.893±0.026 | 0.742±0.014 |
| | C-MoE | 0.895±0.030 | 0.706±0.023 | 0.924±0.011 | 0.716±0.010 | 0.820±0.063 | 0.714±0.022 |
| | TaTS | 0.693±0.005 | 0.665±0.002 | 0.778±0.015 | 0.747±0.011 | **0.723**±0.021 | 0.684±0.010 |
| | CFA | 0.675±0.010 | 0.654±0.006 | 0.784±0.022 | 0.749±0.015 | 0.845±0.056 | 0.713±0.027 |
| | DG | **0.647**±0.009 | **0.635**±0.005 | **0.735**±0.033 | **0.711**±0.025 | 0.783±0.053 | **0.675**±0.024 |
| NumMentions | Uni-modal | 0.714±0.005 | 0.677±0.003 | 0.794±0.015 | 0.755±0.010 | 0.904±0.023 | 0.747±0.010 |
| | MM-TS | 0.934±0.003 | 0.718±0.001 | 1.036±0.020 | 0.759±0.009 | 0.914±0.028 | 0.753±0.013 |
| | C-MoE | 0.911±0.037 | 0.712±0.025 | 0.940±0.010 | 0.723±0.008 | 0.859±0.041 | 0.728±0.013 |
| | TaTS | 0.714±0.004 | 0.676±0.002 | 0.795±0.017 | 0.754±0.011 | 0.756±0.016 | 0.699±0.011 |
| | CFA | 0.693±0.010 | **0.665**±0.007 | 0.800±0.023 | 0.755±0.014 | 0.866±0.052 | 0.727±0.027 |
| | DG | **0.684**±0.004 | **0.665**±0.004 | **0.681**±0.009 | **0.675**±0.005 | **0.752**±0.042 | **0.674**±0.020 |
| NumSources | Uni-modal | 0.560±0.004 | 0.588±0.002 | 0.666±0.007 | 0.671±0.006 | 0.750±0.034 | 0.666±0.018 |
| | MM-TS | 0.736±0.008 | 0.651±0.003 | 0.843±0.016 | 0.691±0.009 | 0.741±0.031 | 0.667±0.016 |
| | C-MoE | 0.706±0.013 | 0.633±0.005 | 0.756±0.010 | 0.645±0.008 | 0.736±0.068 | 0.675±0.023 |
| | TaTS | 0.560±0.003 | 0.588±0.001 | 0.664±0.007 | 0.669±0.007 | **0.597**±0.010 | **0.598**±0.007 |
| | CFA | 0.542±0.006 | 0.579±0.006 | 0.668±0.020 | 0.667±0.014 | 0.692±0.063 | 0.636±0.035 |
| | DG | **0.528**±0.007 | **0.577**±0.004 | **0.587**±0.007 | **0.617**±0.005 | 0.636±0.041 | 0.607±0.024 |

## E.2 Per-Horizon Matched Framework Comparison

### E.2.1 Time-MMD

The Time-MMD tables cover nine domains, six backbones, six methods (including Uni-modal), and four forecast horizons. GPT4MTS uses one horizon, $H = 7$, so Tables 11 and 12 also provide its per-horizon results.

Figure 14 summarizes DynaGuard's MSE reductions at each horizon; the following tables give the corresponding absolute errors.

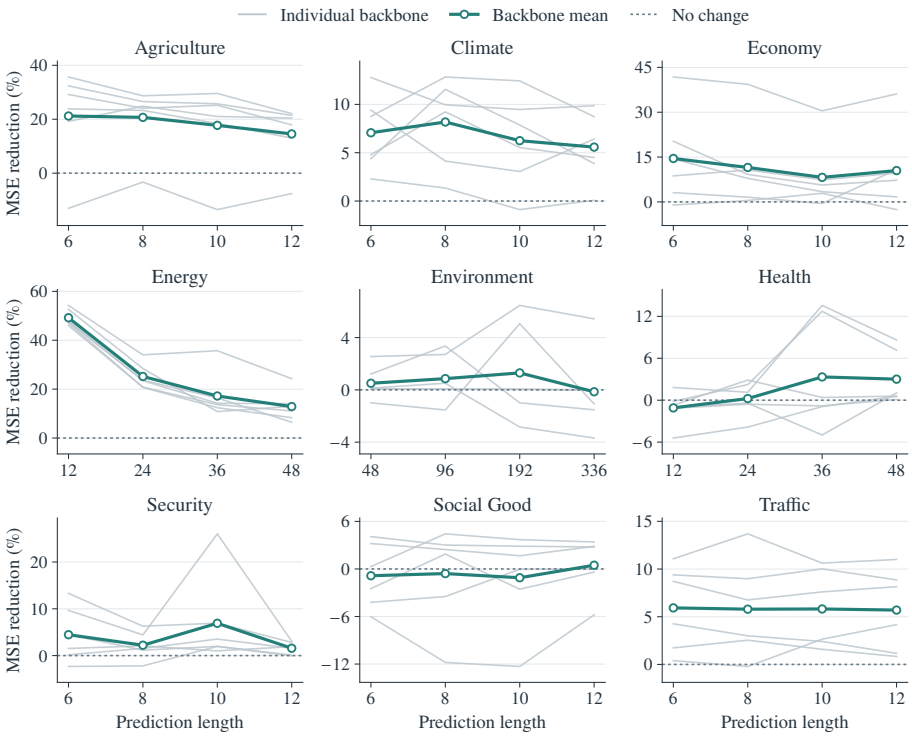


Figure 14: Time-MMD MSE reduction relative to Uni-modal, computed from unrounded five-seed means. Thin lines show individual backbones; the dark line shows their arithmetic mean. Positive values favor DynaGuard. Each panel has its own vertical scale.

Table 13: Framework results by prediction length $H$ for Agriculture, Climate, Economy with iTransformer. Entries are mean $\pm$ standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.192 \pm 0.005$ | $0.223 \pm 0.040$ | $0.153 \pm 0.007$ | $0.189 \pm 0.005$ | $0.165 \pm 0.007$ | $\mathbf{0.146 \pm 0.005}$ |
| | | MAE | $0.292 \pm 0.003$ | $0.338 \pm 0.044$ | $0.267 \pm 0.007$ | $0.292 \pm 0.004$ | $0.272 \pm 0.004$ | $\mathbf{0.255 \pm 0.007}$ |
| | 8 | MSE | $0.274 \pm 0.008$ | $0.300 \pm 0.039$ | $0.222 \pm 0.010$ | $0.266 \pm 0.004$ | $0.239 \pm 0.005$ | $\mathbf{0.210 \pm 0.018}$ |
| | | MAE | $0.349 \pm 0.005$ | $0.387 \pm 0.040$ | $0.315 \pm 0.008$ | $0.344 \pm 0.003$ | $0.325 \pm 0.005$ | $\mathbf{0.304 \pm 0.012}$ |
| | 10 | MSE | $0.353 \pm 0.006$ | $0.387 \pm 0.050$ | $0.296 \pm 0.012$ | $0.337 \pm 0.004$ | $0.323 \pm 0.011$ | $\mathbf{0.288 \pm 0.014}$ |
| | | MAE | $0.400 \pm 0.003$ | $0.440 \pm 0.038$ | $0.367 \pm 0.008$ | $0.390 \pm 0.002$ | $0.383 \pm 0.017$ | $\mathbf{0.362 \pm 0.017}$ |
| | 12 | MSE | $0.428 \pm 0.004$ | $0.459 \pm 0.050$ | $0.376 \pm 0.012$ | $0.411 \pm 0.006$ | $0.433 \pm 0.021$ | $\mathbf{0.372 \pm 0.019}$ |
| | | MAE | $0.445 \pm 0.002$ | $0.479 \pm 0.029$ | $0.420 \pm 0.010$ | $0.436 \pm 0.004$ | $0.452 \pm 0.017$ | $\mathbf{0.406 \pm 0.016}$ |
| Climate | 6 | MSE | $0.969 \pm 0.009$ | $1.027 \pm 0.021$ | $0.996 \pm 0.011$ | $0.950 \pm 0.006$ | $0.988 \pm 0.007$ | $\mathbf{0.884 \pm 0.006}$ |
| | | MAE | $0.795 \pm 0.003$ | $0.820 \pm 0.015$ | $0.812 \pm 0.008$ | $0.787 \pm 0.006$ | $0.807 \pm 0.007$ | $\mathbf{0.750 \pm 0.002}$ |
| | 8 | MSE | $1.018 \pm 0.028$ | $1.044 \pm 0.015$ | $1.001 \pm 0.015$ | $0.967 \pm 0.007$ | $0.985 \pm 0.007$ | $\mathbf{0.888 \pm 0.001}$ |
| | | MAE | $0.821 \pm 0.009$ | $0.826 \pm 0.004$ | $0.815 \pm 0.005$ | $0.795 \pm 0.007$ | $0.809 \pm 0.003$ | $\mathbf{0.749 \pm 0.000}$ |
| | 10 | MSE | $1.026 \pm 0.007$ | $1.036 \pm 0.006$ | $1.006 \pm 0.014$ | $1.039 \pm 0.019$ | $1.011 \pm 0.015$ | $\mathbf{0.899 \pm 0.007}$ |
| | | MAE | $0.828 \pm 0.006$ | $0.825 \pm 0.006$ | $0.816 \pm 0.005$ | $0.824 \pm 0.012$ | $0.807 \pm 0.003$ | $\mathbf{0.752 \pm 0.004}$ |
| | 12 | MSE | $0.991 \pm 0.005$ | $1.028 \pm 0.017$ | $1.053 \pm 0.024$ | $1.028 \pm 0.018$ | $1.002 \pm 0.013$ | $\mathbf{0.905 \pm 0.003}$ |
| | | MAE | $0.809 \pm 0.003$ | $0.823 \pm 0.008$ | $0.834 \pm 0.009$ | $0.823 \pm 0.011$ | $0.810 \pm 0.002$ | $\mathbf{0.754 \pm 0.001}$ |
| Economy | 6 | MSE | $\mathbf{0.148 \pm 0.002}$ | $0.177 \pm 0.028$ | $0.154 \pm 0.002$ | $0.150 \pm 0.002$ | $0.151 \pm 0.004$ | $0.150 \pm 0.007$ |
| | | MAE | $\mathbf{0.299 \pm 0.003}$ | $0.324 \pm 0.018$ | $0.306 \pm 0.004$ | $0.301 \pm 0.002$ | $0.301 \pm 0.005$ | $0.303 \pm 0.008$ |
| | 8 | MSE | $\mathbf{0.169 \pm 0.002}$ | $0.197 \pm 0.027$ | $0.182 \pm 0.007$ | $0.183 \pm 0.024$ | $0.175 \pm 0.004$ | $\mathbf{0.169 \pm 0.006}$ |
| | | MAE | $\mathbf{0.319 \pm 0.002}$ | $0.344 \pm 0.018$ | $0.332 \pm 0.004$ | $0.333 \pm 0.024$ | $0.324 \pm 0.005$ | $0.322 \pm 0.006$ |
| | 10 | MSE | $0.196 \pm 0.008$ | $0.218 \pm 0.025$ | $0.206 \pm 0.004$ | $0.197 \pm 0.008$ | $0.198 \pm 0.003$ | $\mathbf{0.190 \pm 0.003}$ |
| | | MAE | $0.347 \pm 0.009$ | $0.365 \pm 0.016$ | $0.357 \pm 0.004$ | $0.348 \pm 0.009$ | $0.347 \pm 0.003$ | $\mathbf{0.342 \pm 0.003}$ |
| | 12 | MSE | $\mathbf{0.204 \pm 0.002}$ | $0.225 \pm 0.017$ | $0.216 \pm 0.007$ | $0.205 \pm 0.003$ | $0.217 \pm 0.011$ | $0.209 \pm 0.005$ |
| | | MAE | $\mathbf{0.354 \pm 0.003}$ | $0.372 \pm 0.009$ | $0.364 \pm 0.006$ | $0.355 \pm 0.003$ | $0.366 \pm 0.015$ | $0.356 \pm 0.006$ |

Table 14: Framework results by prediction length $H$ for Energy, Environment, Health with iTransformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | 0.176 ± 0.001 | 0.102 ± 0.002 | 0.102 ± 0.005 | 0.176 ± 0.001 | 0.183 ± 0.004 | **0.092 ± 0.004** |
| | | MAE | 0.331 ± 0.001 | 0.238 ± 0.004 | 0.236 ± 0.007 | 0.331 ± 0.001 | 0.339 ± 0.006 | **0.212 ± 0.004** |
| | 24 | MSE | 0.257 ± 0.002 | 0.209 ± 0.008 | 0.202 ± 0.002 | 0.260 ± 0.004 | 0.271 ± 0.007 | **0.197 ± 0.016** |
| | | MAE | 0.393 ± 0.001 | 0.345 ± 0.007 | 0.336 ± 0.002 | 0.395 ± 0.002 | 0.405 ± 0.004 | **0.329 ± 0.010** |
| | 36 | MSE | 0.340 ± 0.004 | 0.299 ± 0.008 | 0.287 ± 0.006 | 0.340 ± 0.004 | 0.370 ± 0.026 | **0.284 ± 0.018** |
| | | MAE | 0.446 ± 0.003 | 0.412 ± 0.007 | 0.403 ± 0.005 | 0.447 ± 0.003 | 0.471 ± 0.013 | **0.397 ± 0.014** |
| | 48 | MSE | 0.419 ± 0.001 | 0.382 ± 0.007 | **0.379 ± 0.006** | 0.419 ± 0.001 | 0.434 ± 0.003 | 0.392 ± 0.008 |
| | | MAE | 0.502 ± 0.001 | 0.476 ± 0.005 | **0.473 ± 0.005** | 0.502 ± 0.001 | 0.517 ± 0.004 | 0.481 ± 0.005 |
| Environment | 48 | MSE | 0.269 ± 0.000 | 0.274 ± 0.001 | 0.273 ± 0.000 | 0.269 ± 0.000 | 0.273 ± 0.003 | **0.262 ± 0.000** |
| | | MAE | 0.372 ± 0.001 | 0.376 ± 0.002 | 0.376 ± 0.001 | 0.372 ± 0.001 | 0.379 ± 0.005 | **0.369 ± 0.001** |
| | 96 | MSE | 0.272 ± 0.005 | 0.289 ± 0.001 | 0.288 ± 0.003 | 0.279 ± 0.000 | 0.277 ± 0.003 | **0.265 ± 0.000** |
| | | MAE | 0.373 ± 0.001 | 0.391 ± 0.002 | 0.388 ± 0.003 | 0.374 ± 0.001 | 0.383 ± 0.005 | **0.371 ± 0.001** |
| | 192 | MSE | 0.273 ± 0.000 | 0.298 ± 0.002 | 0.293 ± 0.004 | 0.272 ± 0.000 | 0.268 ± 0.002 | **0.256 ± 0.001** |
| | | MAE | 0.370 ± 0.001 | 0.406 ± 0.001 | 0.399 ± 0.002 | **0.365 ± 0.000** | 0.379 ± 0.005 | 0.369 ± 0.002 |
| | 336 | MSE | 0.269 ± 0.000 | 0.306 ± 0.008 | 0.304 ± 0.005 | 0.269 ± 0.000 | 0.265 ± 0.005 | **0.254 ± 0.000** |
| | | MAE | **0.365 ± 0.001** | 0.401 ± 0.008 | 0.411 ± 0.004 | **0.365 ± 0.001** | 0.376 ± 0.002 | 0.370 ± 0.001 |
| Health | 12 | MSE | 0.947 ± 0.006 | 1.073 ± 0.018 | 1.012 ± 0.042 | 1.027 ± 0.082 | 1.019 ± 0.015 | **0.929 ± 0.016** |
| | | MAE | 0.657 ± 0.004 | 0.674 ± 0.009 | 0.691 ± 0.014 | 0.670 ± 0.016 | 0.656 ± 0.006 | **0.648 ± 0.006** |
| | 24 | MSE | 1.322 ± 0.013 | 1.675 ± 0.071 | 1.345 ± 0.037 | **1.204 ± 0.066** | 1.434 ± 0.094 | 1.307 ± 0.012 |
| | | MAE | 0.757 ± 0.002 | 0.813 ± 0.008 | 0.787 ± 0.013 | **0.737 ± 0.022** | 0.799 ± 0.034 | 0.760 ± 0.003 |
| | 36 | MSE | 1.624 ± 0.037 | 2.214 ± 0.155 | 1.675 ± 0.051 | **1.368 ± 0.044** | 1.719 ± 0.181 | 1.404 ± 0.004 |
| | | MAE | 0.832 ± 0.011 | 0.907 ± 0.025 | 0.868 ± 0.014 | **0.790 ± 0.022** | 0.850 ± 0.031 | **0.790 ± 0.003** |
| | 48 | MSE | 1.588 ± 0.028 | 1.919 ± 0.033 | 1.792 ± 0.077 | **1.423 ± 0.062** | 1.539 ± 0.058 | 1.451 ± 0.009 |
| | | MAE | 0.834 ± 0.008 | 0.894 ± 0.004 | 0.910 ± 0.018 | **0.809 ± 0.016** | 0.835 ± 0.026 | 0.812 ± 0.009 |

Table 15: Framework results by prediction length $H$ for Security, Social Good, Traffic with iTransformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 73.959 ± 1.175 | 78.377 ± 1.325 | 81.653 ± 1.025 | 70.672 ± 1.434 | 78.166 ± 5.954 | **66.815 ± 0.430** |
| | | MAE | 4.575 ± 0.090 | 4.913 ± 0.161 | 5.094 ± 0.038 | 4.257 ± 0.103 | 4.892 ± 0.477 | **3.847 ± 0.103** |
| | 8 | MSE | 73.111 ± 0.872 | 84.724 ± 4.962 | 86.299 ± 1.908 | 78.298 ± 1.960 | 80.683 ± 5.333 | **69.877 ± 0.567** |
| | | MAE | 4.258 ± 0.065 | 5.151 ± 0.303 | 5.294 ± 0.099 | 4.773 ± 0.130 | 4.939 ± 0.457 | **3.864 ± 0.107** |
| | 10 | MSE | 98.971 ± 2.151 | 102.548 ± 8.486 | 102.756 ± 5.506 | 75.928 ± 0.989 | 94.695 ± 16.125 | **73.225 ± 0.708** |
| | | MAE | 5.801 ± 0.111 | 5.972 ± 0.416 | 6.056 ± 0.264 | 4.319 ± 0.073 | 5.511 ± 0.938 | **3.993 ± 0.088** |
| | 12 | MSE | 79.466 ± 0.506 | 97.182 ± 5.014 | 92.204 ± 1.559 | 79.166 ± 0.478 | 79.446 ± 1.692 | **77.021 ± 0.872** |
| | | MAE | 4.420 ± 0.042 | 5.590 ± 0.308 | 5.402 ± 0.100 | 4.420 ± 0.027 | 4.418 ± 0.131 | **4.147 ± 0.091** |
| Social Good | 6 | MSE | 0.867 ± 0.004 | 0.983 ± 0.019 | 1.082 ± 0.066 | 0.871 ± 0.004 | 1.045 ± 0.027 | **0.864 ± 0.025** |
| | | MAE | 0.378 ± 0.009 | 0.424 ± 0.019 | 0.422 ± 0.024 | **0.374 ± 0.002** | 0.435 ± 0.010 | 0.394 ± 0.010 |
| | 8 | MSE | 0.942 ± 0.011 | 1.089 ± 0.029 | 1.220 ± 0.031 | 0.942 ± 0.007 | 1.185 ± 0.042 | **0.900 ± 0.014** |
| | | MAE | 0.412 ± 0.006 | 0.469 ± 0.019 | 0.479 ± 0.027 | 0.411 ± 0.002 | 0.495 ± 0.017 | **0.409 ± 0.008** |
| | 10 | MSE | 0.990 ± 0.004 | 1.156 ± 0.039 | 1.293 ± 0.058 | 0.996 ± 0.004 | 1.298 ± 0.046 | **0.954 ± 0.019** |
| | | MAE | 0.443 ± 0.004 | 0.503 ± 0.025 | 0.510 ± 0.022 | 0.445 ± 0.003 | 0.553 ± 0.031 | **0.441 ± 0.007** |
| | 12 | MSE | 1.028 ± 0.002 | 1.234 ± 0.027 | 1.356 ± 0.053 | 1.033 ± 0.007 | 1.389 ± 0.126 | **0.993 ± 0.008** |
| | | MAE | 0.469 ± 0.004 | 0.532 ± 0.015 | 0.544 ± 0.018 | 0.473 ± 0.003 | 0.600 ± 0.046 | **0.468 ± 0.004** |
| Traffic | 6 | MSE | 0.158 ± 0.003 | 0.167 ± 0.004 | 0.170 ± 0.004 | 0.161 ± 0.003 | 0.160 ± 0.005 | **0.152 ± 0.003** |
| | | MAE | 0.209 ± 0.003 | 0.225 ± 0.006 | 0.216 ± 0.006 | 0.212 ± 0.003 | **0.193 ± 0.004** | 0.202 ± 0.007 |
| | 8 | MSE | 0.164 ± 0.002 | 0.176 ± 0.004 | 0.180 ± 0.005 | 0.167 ± 0.002 | 0.165 ± 0.008 | **0.159 ± 0.004** |
| | | MAE | 0.217 ± 0.004 | 0.238 ± 0.006 | 0.232 ± 0.007 | 0.219 ± 0.003 | **0.209 ± 0.006** | 0.215 ± 0.005 |
| | 10 | MSE | 0.169 ± 0.001 | 0.181 ± 0.002 | 0.185 ± 0.006 | 0.171 ± 0.002 | 0.172 ± 0.007 | **0.165 ± 0.004** |
| | | MAE | 0.222 ± 0.003 | 0.245 ± 0.005 | 0.240 ± 0.007 | 0.224 ± 0.003 | **0.217 ± 0.006** | 0.224 ± 0.004 |
| | 12 | MSE | 0.174 ± 0.001 | 0.187 ± 0.002 | 0.199 ± 0.001 | 0.176 ± 0.001 | 0.180 ± 0.005 | **0.172 ± 0.005** |
| | | MAE | 0.227 ± 0.002 | 0.252 ± 0.003 | 0.252 ± 0.004 | 0.229 ± 0.002 | **0.225 ± 0.005** | 0.232 ± 0.006 |

Table 16: Framework results by prediction length $H$ for Agriculture, Climate, Economy with PatchTST. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.270 \pm 0.010$ | $0.277 \pm 0.026$ | $0.215 \pm 0.022$ | $0.269 \pm 0.008$ | $0.234 \pm 0.007$ | $\mathbf{0.191 \pm 0.015}$ |
| | | MAE | $0.341 \pm 0.006$ | $0.366 \pm 0.022$ | $0.315 \pm 0.018$ | $0.340 \pm 0.005$ | $0.316 \pm 0.004$ | $\mathbf{0.298 \pm 0.015}$ |
| | 8 | MSE | $0.331 \pm 0.009$ | $0.308 \pm 0.011$ | $0.276 \pm 0.011$ | $0.332 \pm 0.005$ | $0.301 \pm 0.008$ | $\mathbf{0.251 \pm 0.008}$ |
| | | MAE | $0.377 \pm 0.006$ | $0.378 \pm 0.008$ | $0.355 \pm 0.005$ | $0.377 \pm 0.003$ | $0.356 \pm 0.006$ | $\mathbf{0.321 \pm 0.005}$ |
| | 10 | MSE | $0.440 \pm 0.013$ | $0.444 \pm 0.027$ | $0.368 \pm 0.023$ | $0.435 \pm 0.010$ | $0.445 \pm 0.013$ | $\mathbf{0.329 \pm 0.010}$ |
| | | MAE | $0.452 \pm 0.010$ | $0.471 \pm 0.019$ | $0.417 \pm 0.017$ | $0.449 \pm 0.008$ | $0.456 \pm 0.009$ | $\mathbf{0.370 \pm 0.010}$ |
| | 12 | MSE | $0.487 \pm 0.006$ | $0.502 \pm 0.011$ | $0.447 \pm 0.011$ | $0.488 \pm 0.009$ | $0.486 \pm 0.009$ | $\mathbf{0.400 \pm 0.015}$ |
| | | MAE | $0.476 \pm 0.005$ | $0.499 \pm 0.007$ | $0.469 \pm 0.008$ | $0.476 \pm 0.005$ | $0.477 \pm 0.008$ | $\mathbf{0.413 \pm 0.018}$ |
| Climate | 6 | MSE | $0.927 \pm 0.008$ | $1.026 \pm 0.014$ | $0.995 \pm 0.030$ | $0.904 \pm 0.006$ | $0.961 \pm 0.015$ | $\mathbf{0.886 \pm 0.005}$ |
| | | MAE | $0.778 \pm 0.003$ | $0.823 \pm 0.005$ | $0.811 \pm 0.010$ | $0.765 \pm 0.004$ | $0.793 \pm 0.007$ | $\mathbf{0.752 \pm 0.002}$ |
| | 8 | MSE | $1.007 \pm 0.020$ | $1.099 \pm 0.030$ | $1.013 \pm 0.025$ | $0.930 \pm 0.002$ | $0.988 \pm 0.025$ | $\mathbf{0.891 \pm 0.004}$ |
| | | MAE | $0.818 \pm 0.007$ | $0.857 \pm 0.010$ | $0.817 \pm 0.016$ | $0.779 \pm 0.006$ | $0.808 \pm 0.015$ | $\mathbf{0.750 \pm 0.001}$ |
| | 10 | MSE | $0.980 \pm 0.004$ | $1.027 \pm 0.014$ | $0.965 \pm 0.004$ | $0.999 \pm 0.005$ | $0.917 \pm 0.004$ | $\mathbf{0.902 \pm 0.006}$ |
| | | MAE | $0.804 \pm 0.002$ | $0.823 \pm 0.007$ | $0.796 \pm 0.005$ | $0.811 \pm 0.002$ | $0.768 \pm 0.003$ | $\mathbf{0.755 \pm 0.003}$ |
| | 12 | MSE | $0.940 \pm 0.004$ | $1.080 \pm 0.028$ | $0.977 \pm 0.007$ | $0.951 \pm 0.003$ | $0.933 \pm 0.007$ | $\mathbf{0.904 \pm 0.001}$ |
| | | MAE | $0.782 \pm 0.003$ | $0.847 \pm 0.011$ | $0.803 \pm 0.005$ | $0.788 \pm 0.002$ | $0.779 \pm 0.005$ | $\mathbf{0.756 \pm 0.003}$ |
| Economy | 6 | MSE | $0.204 \pm 0.008$ | $0.179 \pm 0.002$ | $\mathbf{0.165 \pm 0.002}$ | $0.205 \pm 0.011$ | $0.208 \pm 0.002$ | $0.175 \pm 0.015$ |
| | | MAE | $0.358 \pm 0.008$ | $0.331 \pm 0.002$ | $\mathbf{0.318 \pm 0.002}$ | $0.359 \pm 0.011$ | $0.362 \pm 0.002$ | $0.325 \pm 0.016$ |
| | 8 | MSE | $0.217 \pm 0.010$ | $0.205 \pm 0.024$ | $\mathbf{0.176 \pm 0.004}$ | $0.216 \pm 0.010$ | $0.217 \pm 0.005$ | $0.200 \pm 0.009$ |
| | | MAE | $0.371 \pm 0.010$ | $0.359 \pm 0.020$ | $\mathbf{0.328 \pm 0.005}$ | $0.370 \pm 0.010$ | $0.372 \pm 0.006$ | $0.353 \pm 0.010$ |
| | 10 | MSE | $0.213 \pm 0.006$ | $0.219 \pm 0.004$ | $\mathbf{0.198 \pm 0.006}$ | $0.218 \pm 0.004$ | $0.222 \pm 0.006$ | $0.206 \pm 0.007$ |
| | | MAE | $0.367 \pm 0.006$ | $0.371 \pm 0.004$ | $\mathbf{0.348 \pm 0.007}$ | $0.372 \pm 0.004$ | $0.377 \pm 0.007$ | $0.358 \pm 0.009$ |
| | 12 | MSE | $0.220 \pm 0.006$ | $0.226 \pm 0.006$ | $\mathbf{0.206 \pm 0.004}$ | $0.226 \pm 0.009$ | $0.235 \pm 0.011$ | $0.216 \pm 0.009$ |
| | | MAE | $0.374 \pm 0.006$ | $0.379 \pm 0.006$ | $\mathbf{0.359 \pm 0.005}$ | $0.380 \pm 0.009$ | $0.389 \pm 0.012$ | $0.368 \pm 0.014$ |

Table 17: Framework results by prediction length $H$ for Energy, Environment, Health with PatchTST. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.196 \pm 0.002$ | $0.115 \pm 0.002$ | $0.107 \pm 0.004$ | $0.196 \pm 0.002$ | $0.209 \pm 0.013$ | $\mathbf{0.106 \pm 0.002}$ |
| | | MAE | $0.344 \pm 0.001$ | $0.252 \pm 0.004$ | $0.240 \pm 0.005$ | $0.344 \pm 0.001$ | $0.353 \pm 0.006$ | $\mathbf{0.233 \pm 0.003}$ |
| | 24 | MSE | $0.270 \pm 0.002$ | $0.212 \pm 0.002$ | $\mathbf{0.206 \pm 0.004}$ | $0.270 \pm 0.001$ | $0.280 \pm 0.007$ | $0.214 \pm 0.003$ |
| | | MAE | $0.400 \pm 0.001$ | $0.345 \pm 0.003$ | $\mathbf{0.339 \pm 0.004}$ | $0.400 \pm 0.001$ | $0.409 \pm 0.005$ | $0.340 \pm 0.002$ |
| | 36 | MSE | $0.350 \pm 0.002$ | $0.302 \pm 0.005$ | $\mathbf{0.293 \pm 0.002}$ | $0.350 \pm 0.002$ | $0.365 \pm 0.007$ | $0.302 \pm 0.007$ |
| | | MAE | $0.452 \pm 0.001$ | $0.414 \pm 0.004$ | $\mathbf{0.404 \pm 0.002}$ | $0.453 \pm 0.001$ | $0.466 \pm 0.008$ | $0.407 \pm 0.007$ |
| | 48 | MSE | $0.433 \pm 0.002$ | $0.388 \pm 0.008$ | $\mathbf{0.381 \pm 0.007}$ | $0.433 \pm 0.001$ | $0.454 \pm 0.016$ | $0.385 \pm 0.002$ |
| | | MAE | $0.511 \pm 0.001$ | $0.479 \pm 0.005$ | $\mathbf{0.472 \pm 0.005}$ | $0.511 \pm 0.001$ | $0.528 \pm 0.011$ | $0.473 \pm 0.004$ |
| Environment | 48 | MSE | $\mathbf{0.269 \pm 0.000}$ | $0.304 \pm 0.001$ | $0.300 \pm 0.000$ | $\mathbf{0.269 \pm 0.000}$ | $0.272 \pm 0.001$ | $\mathbf{0.269 \pm 0.000}$ |
| | | MAE | $\mathbf{0.373 \pm 0.001}$ | $0.390 \pm 0.003$ | $0.387 \pm 0.002$ | $\mathbf{0.373 \pm 0.001}$ | $0.378 \pm 0.002$ | $0.374 \pm 0.000$ |
| | 96 | MSE | $\mathbf{0.279 \pm 0.000}$ | $0.331 \pm 0.000$ | $0.329 \pm 0.001$ | $\mathbf{0.279 \pm 0.000}$ | $0.282 \pm 0.001$ | $\mathbf{0.279 \pm 0.000}$ |
| | | MAE | $\mathbf{0.375 \pm 0.001}$ | $0.401 \pm 0.001$ | $0.401 \pm 0.003$ | $0.376 \pm 0.001$ | $0.380 \pm 0.003$ | $0.376 \pm 0.000$ |
| | 192 | MSE | $\mathbf{0.272 \pm 0.000}$ | $0.329 \pm 0.001$ | $0.325 \pm 0.000$ | $\mathbf{0.272 \pm 0.000}$ | $0.273 \pm 0.001$ | $\mathbf{0.272 \pm 0.000}$ |
| | | MAE | $\mathbf{0.367 \pm 0.000}$ | $0.397 \pm 0.001$ | $0.394 \pm 0.001$ | $\mathbf{0.367 \pm 0.000}$ | $0.371 \pm 0.001$ | $\mathbf{0.367 \pm 0.001}$ |
| | 336 | MSE | $\mathbf{0.268 \pm 0.000}$ | $0.327 \pm 0.001$ | $0.323 \pm 0.001$ | $0.269 \pm 0.000$ | $0.269 \pm 0.000$ | $0.269 \pm 0.000$ |
| | | MAE | $0.367 \pm 0.000$ | $0.398 \pm 0.002$ | $0.395 \pm 0.001$ | $0.367 \pm 0.000$ | $0.370 \pm 0.001$ | $\mathbf{0.366 \pm 0.001}$ |
| Health | 12 | MSE | $0.996 \pm 0.005$ | $1.181 \pm 0.013$ | $\mathbf{0.958 \pm 0.019}$ | $0.965 \pm 0.008$ | $1.014 \pm 0.013$ | $1.008 \pm 0.007$ |
| | | MAE | $0.664 \pm 0.002$ | $0.715 \pm 0.003$ | $0.676 \pm 0.010$ | $\mathbf{0.660 \pm 0.003}$ | $0.675 \pm 0.007$ | $0.676 \pm 0.005$ |
| | 24 | MSE | $1.324 \pm 0.008$ | $1.700 \pm 0.065$ | $\mathbf{1.270 \pm 0.006}$ | $1.316 \pm 0.007$ | $1.308 \pm 0.050$ | $1.332 \pm 0.029$ |
| | | MAE | $0.785 \pm 0.006$ | $0.904 \pm 0.025$ | $\mathbf{0.772 \pm 0.002}$ | $0.782 \pm 0.002$ | $0.796 \pm 0.024$ | $0.776 \pm 0.010$ |
| | 36 | MSE | $1.419 \pm 0.022$ | $1.947 \pm 0.126$ | $1.519 \pm 0.027$ | $\mathbf{1.394 \pm 0.004}$ | $1.421 \pm 0.076$ | $1.431 \pm 0.051$ |
| | | MAE | $0.800 \pm 0.015$ | $0.922 \pm 0.017$ | $0.853 \pm 0.006$ | $\mathbf{0.787 \pm 0.003}$ | $0.823 \pm 0.020$ | $0.812 \pm 0.007$ |
| | 48 | MSE | $1.447 \pm 0.007$ | $1.799 \pm 0.027$ | $1.485 \pm 0.009$ | $1.465 \pm 0.022$ | $1.524 \pm 0.115$ | $\mathbf{1.446 \pm 0.005}$ |
| | | MAE | $\mathbf{0.810 \pm 0.004}$ | $0.921 \pm 0.012$ | $0.837 \pm 0.006$ | $0.819 \pm 0.012$ | $0.847 \pm 0.030$ | $0.815 \pm 0.003$ |

Table 18: Framework results by prediction length $H$ for Security, Social Good, Traffic with PatchTST. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 67.094 ± 0.697 | 71.111 ± 1.460 | 77.529 ± 1.665 | 68.618 ± 2.042 | 69.330 ± 1.913 | **66.074 ± 0.837** |
| | | MAE | 3.876 ± 0.063 | 4.384 ± 0.130 | 4.758 ± 0.164 | 4.018 ± 0.179 | 4.095 ± 0.131 | **3.724 ± 0.086** |
| | 8 | MSE | 71.398 ± 0.700 | 72.428 ± 0.828 | 81.409 ± 1.140 | 71.079 ± 0.541 | **69.581 ± 0.570** | 69.914 ± 0.658 |
| | | MAE | 4.133 ± 0.051 | 4.304 ± 0.105 | 4.887 ± 0.134 | 4.123 ± 0.070 | **3.863 ± 0.055** | 3.954 ± 0.151 |
| | 10 | MSE | 73.972 ± 0.314 | 76.789 ± 1.093 | 89.656 ± 1.984 | 82.436 ± 2.279 | **73.142 ± 0.319** | 73.227 ± 1.152 |
| | | MAE | 4.161 ± 0.022 | 4.483 ± 0.115 | 5.345 ± 0.127 | 4.778 ± 0.142 | 4.029 ± 0.043 | **3.994 ± 0.192** |
| | 12 | MSE | 77.559 ± 0.429 | 83.752 ± 1.491 | 91.073 ± 3.154 | 77.944 ± 0.484 | 84.013 ± 4.530 | **76.075 ± 0.498** |
| | | MAE | 4.276 ± 0.048 | 4.872 ± 0.108 | 5.372 ± 0.262 | 4.332 ± 0.077 | 4.807 ± 0.303 | **4.028 ± 0.051** |
| Social Good | 6 | MSE | 0.757 ± 0.008 | 0.777 ± 0.008 | 0.807 ± 0.017 | **0.756 ± 0.008** | 0.816 ± 0.009 | 0.789 ± 0.018 |
| | | MAE | 0.388 ± 0.005 | 0.399 ± 0.003 | 0.400 ± 0.009 | 0.388 ± 0.005 | 0.408 ± 0.014 | **0.381 ± 0.010** |
| | 8 | MSE | **0.813 ± 0.007** | 0.856 ± 0.016 | 0.940 ± 0.059 | 0.814 ± 0.006 | 0.883 ± 0.020 | 0.842 ± 0.010 |
| | | MAE | **0.413 ± 0.004** | 0.439 ± 0.012 | 0.473 ± 0.041 | 0.414 ± 0.003 | 0.450 ± 0.025 | 0.418 ± 0.009 |
| | 10 | MSE | 0.889 ± 0.010 | 0.962 ± 0.008 | 1.054 ± 0.068 | **0.888 ± 0.010** | 1.004 ± 0.103 | 0.890 ± 0.011 |
| | | MAE | 0.453 ± 0.010 | 0.485 ± 0.006 | 0.514 ± 0.047 | **0.452 ± 0.009** | 0.526 ± 0.079 | 0.453 ± 0.016 |
| | 12 | MSE | **0.932 ± 0.009** | 1.002 ± 0.025 | 1.070 ± 0.032 | **0.932 ± 0.011** | 1.057 ± 0.038 | 0.933 ± 0.014 |
| | | MAE | 0.483 ± 0.010 | 0.514 ± 0.020 | 0.524 ± 0.029 | 0.483 ± 0.011 | 0.559 ± 0.033 | **0.471 ± 0.008** |
| Traffic | 6 | MSE | 0.168 ± 0.003 | 0.181 ± 0.006 | 0.156 ± 0.003 | 0.168 ± 0.003 | 0.161 ± 0.005 | **0.152 ± 0.004** |
| | | MAE | 0.220 ± 0.002 | 0.256 ± 0.005 | 0.223 ± 0.004 | 0.220 ± 0.003 | **0.219 ± 0.010** | 0.231 ± 0.009 |
| | 8 | MSE | 0.168 ± 0.002 | 0.181 ± 0.003 | 0.157 ± 0.003 | 0.168 ± 0.002 | 0.161 ± 0.005 | **0.153 ± 0.002** |
| | | MAE | 0.221 ± 0.001 | 0.255 ± 0.003 | 0.221 ± 0.003 | **0.220 ± 0.001** | 0.222 ± 0.012 | 0.221 ± 0.001 |
| | 10 | MSE | 0.175 ± 0.003 | 0.190 ± 0.004 | 0.167 ± 0.004 | 0.174 ± 0.003 | 0.169 ± 0.003 | **0.157 ± 0.003** |
| | | MAE | 0.227 ± 0.001 | 0.264 ± 0.001 | 0.232 ± 0.006 | 0.228 ± 0.001 | **0.224 ± 0.009** | 0.231 ± 0.005 |
| | 12 | MSE | 0.176 ± 0.002 | 0.192 ± 0.005 | 0.170 ± 0.002 | 0.176 ± 0.002 | 0.172 ± 0.006 | **0.160 ± 0.004** |
| | | MAE | 0.230 ± 0.003 | 0.268 ± 0.006 | 0.234 ± 0.003 | 0.230 ± 0.003 | **0.228 ± 0.007** | 0.235 ± 0.014 |

Table 19: Framework results by prediction length $H$ for Agriculture, Climate, Economy with Crossformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | **1.058 ± 0.099** | 1.752 ± 0.228 | 2.054 ± 0.132 | 1.127 ± 0.038 | 1.068 ± 0.153 | 1.196 ± 0.082 |
| | | MAE | **0.713 ± 0.034** | 0.942 ± 0.043 | 1.005 ± 0.047 | 0.726 ± 0.017 | 0.724 ± 0.027 | 0.762 ± 0.038 |
| | 8 | MSE | **1.146 ± 0.053** | 2.039 ± 0.134 | 2.147 ± 0.169 | 1.192 ± 0.100 | 1.231 ± 0.109 | 1.184 ± 0.057 |
| | | MAE | **0.747 ± 0.016** | 0.996 ± 0.039 | 1.036 ± 0.060 | 0.753 ± 0.041 | 0.781 ± 0.049 | 0.753 ± 0.031 |
| | 10 | MSE | **1.175 ± 0.100** | 2.001 ± 0.160 | 2.195 ± 0.201 | 1.180 ± 0.060 | 1.279 ± 0.074 | 1.333 ± 0.071 |
| | | MAE | **0.754 ± 0.038** | 0.980 ± 0.046 | 1.040 ± 0.079 | 0.757 ± 0.028 | 0.788 ± 0.031 | 0.811 ± 0.032 |
| | 12 | MSE | 1.220 ± 0.086 | 1.933 ± 0.094 | 2.179 ± 0.192 | **1.200 ± 0.071** | 1.262 ± 0.045 | 1.313 ± 0.116 |
| | | MAE | 0.776 ± 0.026 | 0.993 ± 0.036 | 1.039 ± 0.075 | **0.761 ± 0.024** | 0.777 ± 0.027 | 0.798 ± 0.043 |
| Climate | 6 | MSE | 0.912 ± 0.005 | 0.960 ± 0.017 | 0.904 ± 0.006 | 0.903 ± 0.006 | 0.970 ± 0.011 | **0.891 ± 0.004** |
| | | MAE | 0.758 ± 0.003 | 0.785 ± 0.006 | 0.754 ± 0.004 | 0.752 ± 0.003 | 0.785 ± 0.003 | **0.745 ± 0.002** |
| | 8 | MSE | 0.906 ± 0.006 | 0.984 ± 0.009 | 0.900 ± 0.004 | **0.891 ± 0.006** | 0.975 ± 0.022 | 0.894 ± 0.001 |
| | | MAE | 0.754 ± 0.002 | 0.790 ± 0.003 | 0.751 ± 0.002 | 0.745 ± 0.003 | 0.785 ± 0.008 | **0.744 ± 0.001** |
| | 10 | MSE | **0.890 ± 0.006** | 0.983 ± 0.025 | 0.895 ± 0.005 | 0.894 ± 0.004 | 1.000 ± 0.036 | 0.898 ± 0.004 |
| | | MAE | **0.745 ± 0.004** | 0.793 ± 0.006 | 0.746 ± 0.003 | 0.747 ± 0.002 | 0.784 ± 0.015 | 0.748 ± 0.002 |
| | 12 | MSE | 0.905 ± 0.007 | 0.979 ± 0.012 | 0.902 ± 0.005 | **0.897 ± 0.002** | 0.955 ± 0.017 | 0.904 ± 0.002 |
| | | MAE | 0.754 ± 0.005 | 0.793 ± 0.008 | 0.752 ± 0.004 | **0.748 ± 0.001** | 0.771 ± 0.007 | 0.749 ± 0.001 |
| Economy | 6 | MSE | 0.364 ± 0.010 | 0.556 ± 0.033 | 0.574 ± 0.027 | 0.359 ± 0.014 | 0.355 ± 0.024 | **0.352 ± 0.010** |
| | | MAE | 0.482 ± 0.008 | 0.614 ± 0.020 | 0.628 ± 0.017 | 0.479 ± 0.013 | 0.477 ± 0.017 | **0.475 ± 0.007** |
| | 8 | MSE | 0.382 ± 0.008 | 0.552 ± 0.054 | 0.578 ± 0.038 | 0.370 ± 0.011 | **0.365 ± 0.030** | 0.376 ± 0.018 |
| | | MAE | 0.495 ± 0.006 | 0.609 ± 0.034 | 0.629 ± 0.023 | 0.486 ± 0.009 | **0.481 ± 0.024** | 0.487 ± 0.009 |
| | 10 | MSE | **0.391 ± 0.024** | 0.592 ± 0.073 | 0.604 ± 0.048 | **0.391 ± 0.014** | 0.392 ± 0.021 | 0.393 ± 0.040 |
| | | MAE | **0.501 ± 0.020** | 0.633 ± 0.048 | 0.643 ± 0.027 | **0.501 ± 0.011** | 0.502 ± 0.016 | 0.502 ± 0.033 |
| | 12 | MSE | 0.414 ± 0.016 | 0.624 ± 0.053 | 0.626 ± 0.042 | 0.417 ± 0.009 | 0.415 ± 0.014 | **0.370 ± 0.019** |
| | | MAE | 0.517 ± 0.012 | 0.652 ± 0.032 | 0.655 ± 0.025 | 0.518 ± 0.007 | 0.518 ± 0.010 | **0.484 ± 0.014** |

Table 20: Framework results by prediction length $H$ for Energy, Environment, Health with Crossformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.234 \pm 0.017$ | $0.148 \pm 0.027$ | $0.139 \pm 0.012$ | $0.221 \pm 0.011$ | $0.216 \pm 0.014$ | $\mathbf{0.111 \pm 0.006}$ |
| | | MAE | $0.376 \pm 0.015$ | $0.273 \pm 0.028$ | $0.264 \pm 0.011$ | $0.364 \pm 0.010$ | $0.360 \pm 0.010$ | $\mathbf{0.235 \pm 0.005}$ |
| | 24 | MSE | $0.321 \pm 0.051$ | $0.260 \pm 0.025$ | $0.262 \pm 0.028$ | $0.306 \pm 0.043$ | $0.328 \pm 0.023$ | $\mathbf{0.230 \pm 0.038}$ |
| | | MAE | $0.438 \pm 0.030$ | $0.378 \pm 0.023$ | $0.375 \pm 0.020$ | $0.427 \pm 0.027$ | $0.439 \pm 0.018$ | $\mathbf{0.354 \pm 0.026}$ |
| | 36 | MSE | $0.373 \pm 0.025$ | $0.351 \pm 0.019$ | $0.369 \pm 0.023$ | $0.408 \pm 0.030$ | $0.445 \pm 0.040$ | $\mathbf{0.333 \pm 0.032}$ |
| | | MAE | $0.487 \pm 0.011$ | $0.451 \pm 0.011$ | $0.462 \pm 0.012$ | $0.502 \pm 0.019$ | $0.514 \pm 0.021$ | $\mathbf{0.447 \pm 0.022}$ |
| | 48 | MSE | $0.475 \pm 0.016$ | $0.423 \pm 0.035$ | $0.479 \pm 0.043$ | $0.475 \pm 0.025$ | $0.530 \pm 0.070$ | $\mathbf{0.413 \pm 0.060}$ |
| | | MAE | $0.549 \pm 0.005$ | $\mathbf{0.507 \pm 0.014}$ | $0.536 \pm 0.022$ | $0.550 \pm 0.014$ | $0.568 \pm 0.023$ | $0.512 \pm 0.044$ |
| Environment | 48 | MSE | $\mathbf{0.273 \pm 0.003}$ | $0.305 \pm 0.007$ | $0.302 \pm 0.003$ | $\mathbf{0.273 \pm 0.003}$ | $0.294 \pm 0.009$ | $0.275 \pm 0.008$ |
| | | MAE | $0.376 \pm 0.004$ | $0.415 \pm 0.014$ | $0.406 \pm 0.012$ | $\mathbf{0.375 \pm 0.003}$ | $0.408 \pm 0.018$ | $0.378 \pm 0.012$ |
| | 96 | MSE | $\mathbf{0.284 \pm 0.002}$ | $0.366 \pm 0.016$ | $0.361 \pm 0.012$ | $\mathbf{0.284 \pm 0.002}$ | $0.309 \pm 0.004$ | $0.288 \pm 0.012$ |
| | | MAE | $0.398 \pm 0.006$ | $0.474 \pm 0.018$ | $0.468 \pm 0.021$ | $0.399 \pm 0.006$ | $0.430 \pm 0.007$ | $\mathbf{0.395 \pm 0.017}$ |
| | 192 | MSE | $0.299 \pm 0.008$ | $0.442 \pm 0.036$ | $0.467 \pm 0.016$ | $0.298 \pm 0.008$ | $0.310 \pm 0.005$ | $\mathbf{0.284 \pm 0.009}$ |
| | | MAE | $0.439 \pm 0.012$ | $0.563 \pm 0.035$ | $0.584 \pm 0.014$ | $0.437 \pm 0.011$ | $0.446 \pm 0.007$ | $\mathbf{0.406 \pm 0.014}$ |
| | 336 | MSE | $\mathbf{0.300 \pm 0.010}$ | $0.451 \pm 0.026$ | $0.465 \pm 0.032$ | $0.303 \pm 0.007$ | $0.306 \pm 0.019$ | $0.304 \pm 0.009$ |
| | | MAE | $0.447 \pm 0.014$ | $0.573 \pm 0.022$ | $0.589 \pm 0.028$ | $0.450 \pm 0.009$ | $\mathbf{0.443 \pm 0.021}$ | $0.445 \pm 0.010$ |
| Health | 12 | MSE | $1.038 \pm 0.028$ | $1.289 \pm 0.067$ | $\mathbf{1.025 \pm 0.053}$ | $1.066 \pm 0.003$ | $1.155 \pm 0.071$ | $1.045 \pm 0.016$ |
| | | MAE | $0.672 \pm 0.007$ | $0.769 \pm 0.026$ | $\mathbf{0.670 \pm 0.023}$ | $0.677 \pm 0.006$ | $0.752 \pm 0.054$ | $0.676 \pm 0.006$ |
| | 24 | MSE | $1.264 \pm 0.091$ | $1.468 \pm 0.109$ | $1.208 \pm 0.015$ | $\mathbf{1.200 \pm 0.021}$ | $1.454 \pm 0.088$ | $1.227 \pm 0.014$ |
| | | MAE | $0.759 \pm 0.034$ | $0.820 \pm 0.025$ | $0.737 \pm 0.016$ | $0.738 \pm 0.014$ | $0.832 \pm 0.030$ | $\mathbf{0.727 \pm 0.004}$ |
| | 36 | MSE | $1.296 \pm 0.037$ | $1.497 \pm 0.097$ | $1.308 \pm 0.024$ | $\mathbf{1.248 \pm 0.031}$ | $1.447 \pm 0.063$ | $1.291 \pm 0.031$ |
| | | MAE | $0.751 \pm 0.014$ | $0.814 \pm 0.027$ | $0.745 \pm 0.010$ | $\mathbf{0.733 \pm 0.010}$ | $0.812 \pm 0.012$ | $0.742 \pm 0.008$ |
| | 48 | MSE | $1.371 \pm 0.048$ | $1.497 \pm 0.052$ | $\mathbf{1.349 \pm 0.032}$ | $1.351 \pm 0.016$ | $1.389 \pm 0.041$ | $1.363 \pm 0.010$ |
| | | MAE | $0.773 \pm 0.019$ | $0.810 \pm 0.009$ | $0.774 \pm 0.011$ | $\mathbf{0.766 \pm 0.003}$ | $0.792 \pm 0.010$ | $0.769 \pm 0.006$ |

Table 21: Framework results by prediction length $H$ for Security, Social Good, Traffic with Crossformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | $81.772 \pm 2.187$ | $84.147 \pm 1.828$ | $81.758 \pm 1.509$ | $82.743 \pm 0.741$ | $82.193 \pm 1.107$ | $\mathbf{81.640 \pm 0.288}$ |
| | | MAE | $5.122 \pm 0.206$ | $5.354 \pm 0.154$ | $5.122 \pm 0.145$ | $5.216 \pm 0.070$ | $5.144 \pm 0.087$ | $\mathbf{5.108 \pm 0.025}$ |
| | 8 | MSE | $84.993 \pm 0.445$ | $87.749 \pm 0.954$ | $85.703 \pm 0.318$ | $85.501 \pm 1.343$ | $85.944 \pm 1.703$ | $\mathbf{83.683 \pm 0.581}$ |
| | | MAE | $5.274 \pm 0.045$ | $5.527 \pm 0.082$ | $5.342 \pm 0.031$ | $5.321 \pm 0.132$ | $5.357 \pm 0.140$ | $\mathbf{5.146 \pm 0.054}$ |
| | 10 | MSE | $88.813 \pm 0.907$ | $88.237 \pm 1.247$ | $89.261 \pm 0.454$ | $89.199 \pm 0.296$ | $87.567 \pm 0.827$ | $\mathbf{85.678 \pm 1.807}$ |
| | | MAE | $5.466 \pm 0.077$ | $5.421 \pm 0.102$ | $5.496 \pm 0.041$ | $5.489 \pm 0.028$ | $5.352 \pm 0.077$ | $\mathbf{5.167 \pm 0.160}$ |
| | 12 | MSE | $91.261 \pm 0.262$ | $92.669 \pm 0.239$ | $91.328 \pm 1.027$ | $90.107 \pm 1.581$ | $93.297 \pm 2.319$ | $\mathbf{89.827 \pm 1.530}$ |
| | | MAE | $5.498 \pm 0.026$ | $5.633 \pm 0.021$ | $5.498 \pm 0.091$ | $5.392 \pm 0.146$ | $5.652 \pm 0.195$ | $\mathbf{5.357 \pm 0.143}$ |
| Social Good | 6 | MSE | $0.675 \pm 0.014$ | $0.693 \pm 0.008$ | $0.721 \pm 0.038$ | $0.688 \pm 0.019$ | $0.686 \pm 0.014$ | $\mathbf{0.653 \pm 0.006}$ |
| | | MAE | $0.349 \pm 0.038$ | $0.400 \pm 0.052$ | $0.413 \pm 0.069$ | $0.347 \pm 0.032$ | $0.348 \pm 0.031$ | $\mathbf{0.302 \pm 0.008}$ |
| | 8 | MSE | $0.743 \pm 0.014$ | $0.729 \pm 0.010$ | $0.756 \pm 0.031$ | $0.739 \pm 0.011$ | $0.740 \pm 0.045$ | $\mathbf{0.725 \pm 0.006}$ |
| | | MAE | $0.345 \pm 0.004$ | $0.381 \pm 0.028$ | $0.408 \pm 0.020$ | $0.395 \pm 0.037$ | $0.386 \pm 0.062$ | $\mathbf{0.333 \pm 0.010}$ |
| | 10 | MSE | $0.772 \pm 0.006$ | $0.775 \pm 0.023$ | $0.810 \pm 0.032$ | $0.784 \pm 0.014$ | $0.779 \pm 0.019$ | $\mathbf{0.759 \pm 0.005}$ |
| | | MAE | $0.392 \pm 0.010$ | $0.418 \pm 0.033$ | $0.445 \pm 0.037$ | $0.383 \pm 0.011$ | $0.407 \pm 0.031$ | $\mathbf{0.361 \pm 0.015}$ |
| | 12 | MSE | $0.809 \pm 0.016$ | $0.869 \pm 0.076$ | $0.825 \pm 0.040$ | $0.809 \pm 0.013$ | $0.824 \pm 0.038$ | $\mathbf{0.786 \pm 0.006}$ |
| | | MAE | $0.454 \pm 0.037$ | $0.509 \pm 0.087$ | $0.449 \pm 0.034$ | $0.421 \pm 0.009$ | $0.428 \pm 0.033$ | $\mathbf{0.391 \pm 0.013}$ |
| Traffic | 6 | MSE | $\mathbf{0.158 \pm 0.002}$ | $0.178 \pm 0.009$ | $0.199 \pm 0.009$ | $\mathbf{0.158 \pm 0.002}$ | $0.159 \pm 0.003$ | $\mathbf{0.158 \pm 0.001}$ |
| | | MAE | $0.286 \pm 0.005$ | $0.323 \pm 0.017$ | $0.354 \pm 0.016$ | $0.286 \pm 0.007$ | $0.285 \pm 0.007$ | $\mathbf{0.283 \pm 0.005}$ |
| | 8 | MSE | $0.169 \pm 0.001$ | $0.215 \pm 0.011$ | $0.228 \pm 0.006$ | $0.168 \pm 0.002$ | $\mathbf{0.163 \pm 0.002}$ | $0.169 \pm 0.007$ |
| | | MAE | $0.301 \pm 0.005$ | $0.365 \pm 0.016$ | $0.383 \pm 0.008$ | $0.301 \pm 0.004$ | $\mathbf{0.287 \pm 0.005}$ | $0.303 \pm 0.011$ |
| | 10 | MSE | $0.175 \pm 0.002$ | $0.228 \pm 0.006$ | $0.223 \pm 0.005$ | $0.175 \pm 0.002$ | $\mathbf{0.170 \pm 0.004}$ | $\mathbf{0.170 \pm 0.003}$ |
| | | MAE | $0.308 \pm 0.005$ | $0.380 \pm 0.008$ | $0.377 \pm 0.005$ | $0.308 \pm 0.006$ | $\mathbf{0.297 \pm 0.009}$ | $0.298 \pm 0.007$ |
| | 12 | MSE | $0.179 \pm 0.002$ | $0.235 \pm 0.006$ | $0.212 \pm 0.007$ | $0.179 \pm 0.002$ | $\mathbf{0.170 \pm 0.002}$ | $0.172 \pm 0.004$ |
| | | MAE | $0.314 \pm 0.005$ | $0.386 \pm 0.008$ | $0.363 \pm 0.008$ | $0.312 \pm 0.005$ | $\mathbf{0.293 \pm 0.006}$ | $0.298 \pm 0.011$ |

Table 22: Framework results by prediction length $H$ for Agriculture, Climate, Economy with SOFTS. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.246 \pm 0.010$ | $0.199 \pm 0.017$ | $0.445 \pm 0.022$ | $0.246 \pm 0.010$ | $0.195 \pm 0.004$ | $\mathbf{0.158 \pm 0.006}$ |
| | | MAE | $0.327 \pm 0.009$ | $0.302 \pm 0.013$ | $0.506 \pm 0.024$ | $0.327 \pm 0.008$ | $0.299 \pm 0.007$ | $\mathbf{0.265 \pm 0.003}$ |
| | 8 | MSE | $0.329 \pm 0.009$ | $0.270 \pm 0.007$ | $0.510 \pm 0.011$ | $0.330 \pm 0.008$ | $0.269 \pm 0.005$ | $\mathbf{0.235 \pm 0.014}$ |
| | | MAE | $0.379 \pm 0.008$ | $0.353 \pm 0.005$ | $0.540 \pm 0.024$ | $0.380 \pm 0.007$ | $0.349 \pm 0.005$ | $\mathbf{0.321 \pm 0.008}$ |
| | 10 | MSE | $0.413 \pm 0.012$ | $0.347 \pm 0.010$ | $0.567 \pm 0.020$ | $0.412 \pm 0.012$ | $0.341 \pm 0.007$ | $\mathbf{0.291 \pm 0.022}$ |
| | | MAE | $0.427 \pm 0.009$ | $0.401 \pm 0.006$ | $0.573 \pm 0.017$ | $0.427 \pm 0.009$ | $0.399 \pm 0.005$ | $\mathbf{0.353 \pm 0.015}$ |
| | 12 | MSE | $0.490 \pm 0.012$ | $0.432 \pm 0.006$ | $0.634 \pm 0.018$ | $0.489 \pm 0.011$ | $0.426 \pm 0.009$ | $\mathbf{0.382 \pm 0.033}$ |
| | | MAE | $0.471 \pm 0.010$ | $0.451 \pm 0.004$ | $0.599 \pm 0.012$ | $0.471 \pm 0.009$ | $0.450 \pm 0.005$ | $\mathbf{0.413 \pm 0.025}$ |
| Climate | 6 | MSE | $1.019 \pm 0.012$ | $1.047 \pm 0.005$ | $1.016 \pm 0.008$ | $0.984 \pm 0.007$ | $1.020 \pm 0.035$ | $\mathbf{0.889 \pm 0.010}$ |
| | | MAE | $0.814 \pm 0.004$ | $0.825 \pm 0.003$ | $0.812 \pm 0.007$ | $0.799 \pm 0.002$ | $0.819 \pm 0.014$ | $\mathbf{0.751 \pm 0.005}$ |
| | 8 | MSE | $0.988 \pm 0.014$ | $1.125 \pm 0.022$ | $0.991 \pm 0.005$ | $1.002 \pm 0.019$ | $0.999 \pm 0.006$ | $\mathbf{0.890 \pm 0.005}$ |
| | | MAE | $0.804 \pm 0.007$ | $0.861 \pm 0.011$ | $0.804 \pm 0.004$ | $0.811 \pm 0.010$ | $0.803 \pm 0.002$ | $\mathbf{0.750 \pm 0.002}$ |
| | 10 | MSE | $0.991 \pm 0.003$ | $1.119 \pm 0.027$ | $1.018 \pm 0.010$ | $0.998 \pm 0.004$ | $1.021 \pm 0.015$ | $\mathbf{0.897 \pm 0.002}$ |
| | | MAE | $0.803 \pm 0.002$ | $0.855 \pm 0.010$ | $0.816 \pm 0.004$ | $0.806 \pm 0.002$ | $0.810 \pm 0.003$ | $\mathbf{0.752 \pm 0.001}$ |
| | 12 | MSE | $1.004 \pm 0.006$ | $1.070 \pm 0.012$ | $1.046 \pm 0.019$ | $0.980 \pm 0.008$ | $0.988 \pm 0.006$ | $\mathbf{0.905 \pm 0.003}$ |
| | | MAE | $0.811 \pm 0.002$ | $0.842 \pm 0.006$ | $0.831 \pm 0.009$ | $0.800 \pm 0.004$ | $0.802 \pm 0.004$ | $\mathbf{0.755 \pm 0.001}$ |
| Economy | 6 | MSE | $0.162 \pm 0.006$ | $0.166 \pm 0.006$ | $0.207 \pm 0.008$ | $0.162 \pm 0.006$ | $0.156 \pm 0.020$ | $\mathbf{0.148 \pm 0.004}$ |
| | | MAE | $0.313 \pm 0.005$ | $0.321 \pm 0.005$ | $0.348 \pm 0.006$ | $0.313 \pm 0.005$ | $0.308 \pm 0.021$ | $\mathbf{0.302 \pm 0.005}$ |
| | 8 | MSE | $0.186 \pm 0.003$ | $0.185 \pm 0.004$ | $0.223 \pm 0.008$ | $0.187 \pm 0.003$ | $0.176 \pm 0.019$ | $\mathbf{0.167 \pm 0.004}$ |
| | | MAE | $0.338 \pm 0.003$ | $0.338 \pm 0.003$ | $0.365 \pm 0.006$ | $0.338 \pm 0.003$ | $0.327 \pm 0.021$ | $\mathbf{0.318 \pm 0.003}$ |
| | 10 | MSE | $0.207 \pm 0.003$ | $0.204 \pm 0.004$ | $0.232 \pm 0.011$ | $0.207 \pm 0.003$ | $\mathbf{0.191 \pm 0.009}$ | $0.192 \pm 0.007$ |
| | | MAE | $0.359 \pm 0.002$ | $0.355 \pm 0.003$ | $0.379 \pm 0.006$ | $0.359 \pm 0.002$ | $\mathbf{0.341 \pm 0.010}$ | $0.343 \pm 0.007$ |
| | 12 | MSE | $0.224 \pm 0.005$ | $0.221 \pm 0.005$ | $0.245 \pm 0.009$ | $0.224 \pm 0.005$ | $0.214 \pm 0.016$ | $\mathbf{0.202 \pm 0.010}$ |
| | | MAE | $0.375 \pm 0.005$ | $0.372 \pm 0.005$ | $0.397 \pm 0.004$ | $0.375 \pm 0.005$ | $0.365 \pm 0.018$ | $\mathbf{0.351 \pm 0.012}$ |

Table 23: Framework results by prediction length $H$ for Energy, Environment, Health with SOFTS. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.193 \pm 0.002$ | $0.109 \pm 0.001$ | $0.273 \pm 0.008$ | $0.193 \pm 0.002$ | $0.195 \pm 0.013$ | $\mathbf{0.102 \pm 0.004}$ |
| | | MAE | $0.343 \pm 0.002$ | $0.246 \pm 0.002$ | $0.408 \pm 0.005$ | $0.343 \pm 0.002$ | $0.345 \pm 0.006$ | $\mathbf{0.232 \pm 0.005}$ |
| | 24 | MSE | $0.277 \pm 0.002$ | $\mathbf{0.215 \pm 0.003}$ | $0.375 \pm 0.010$ | $0.277 \pm 0.002$ | $0.278 \pm 0.006$ | $0.219 \pm 0.008$ |
| | | MAE | $0.405 \pm 0.001$ | $\mathbf{0.350 \pm 0.003}$ | $0.471 \pm 0.004$ | $0.405 \pm 0.001$ | $0.406 \pm 0.001$ | $0.356 \pm 0.009$ |
| | 36 | MSE | $0.358 \pm 0.002$ | $\mathbf{0.306 \pm 0.008}$ | $0.443 \pm 0.017$ | $0.358 \pm 0.002$ | $0.372 \pm 0.016$ | $0.314 \pm 0.011$ |
| | | MAE | $0.458 \pm 0.001$ | $\mathbf{0.419 \pm 0.007}$ | $0.515 \pm 0.006$ | $0.458 \pm 0.001$ | $0.473 \pm 0.016$ | $0.420 \pm 0.009$ |
| | 48 | MSE | $0.443 \pm 0.002$ | $\mathbf{0.392 \pm 0.006}$ | $0.537 \pm 0.013$ | $0.443 \pm 0.002$ | $0.458 \pm 0.020$ | $0.406 \pm 0.023$ |
| | | MAE | $0.517 \pm 0.001$ | $\mathbf{0.484 \pm 0.004}$ | $0.570 \pm 0.006$ | $0.517 \pm 0.001$ | $0.530 \pm 0.019$ | $0.493 \pm 0.026$ |
| Environment | 48 | MSE | $0.269 \pm 0.000$ | $0.301 \pm 0.001$ | $0.334 \pm 0.001$ | $0.269 \pm 0.000$ | $0.273 \pm 0.002$ | $\mathbf{0.265 \pm 0.003}$ |
| | | MAE | $0.373 \pm 0.000$ | $0.387 \pm 0.001$ | $0.437 \pm 0.001$ | $0.373 \pm 0.000$ | $0.378 \pm 0.005$ | $\mathbf{0.372 \pm 0.004}$ |
| | 96 | MSE | $0.279 \pm 0.000$ | $0.331 \pm 0.001$ | $0.331 \pm 0.001$ | $0.279 \pm 0.000$ | $0.284 \pm 0.002$ | $\mathbf{0.270 \pm 0.002}$ |
| | | MAE | $\mathbf{0.375 \pm 0.000}$ | $0.401 \pm 0.001$ | $0.430 \pm 0.001$ | $\mathbf{0.375 \pm 0.000}$ | $0.384 \pm 0.004$ | $\mathbf{0.375 \pm 0.001}$ |
| | 192 | MSE | $\mathbf{0.272 \pm 0.000}$ | $0.329 \pm 0.000$ | $0.322 \pm 0.001$ | $\mathbf{0.272 \pm 0.000}$ | $0.273 \pm 0.001$ | $0.274 \pm 0.002$ |
| | | MAE | $\mathbf{0.366 \pm 0.000}$ | $0.394 \pm 0.000$ | $0.428 \pm 0.001$ | $\mathbf{0.366 \pm 0.000}$ | $0.369 \pm 0.003$ | $0.374 \pm 0.005$ |
| | 336 | MSE | $\mathbf{0.268 \pm 0.000}$ | $0.326 \pm 0.001$ | $0.324 \pm 0.001$ | $\mathbf{0.268 \pm 0.000}$ | $0.270 \pm 0.000$ | $0.273 \pm 0.002$ |
| | | MAE | $\mathbf{0.366 \pm 0.000}$ | $0.395 \pm 0.001$ | $0.428 \pm 0.002$ | $\mathbf{0.366 \pm 0.000}$ | $0.368 \pm 0.001$ | $0.373 \pm 0.006$ |
| Health | 12 | MSE | $\mathbf{0.948 \pm 0.020}$ | $1.113 \pm 0.056$ | $0.976 \pm 0.020$ | $0.967 \pm 0.017$ | $0.983 \pm 0.033$ | $0.950 \pm 0.010$ |
| | | MAE | $0.665 \pm 0.010$ | $0.706 \pm 0.013$ | $0.680 \pm 0.002$ | $0.678 \pm 0.007$ | $0.664 \pm 0.007$ | $\mathbf{0.652 \pm 0.006}$ |
| | 24 | MSE | $1.365 \pm 0.019$ | $1.763 \pm 0.190$ | $1.440 \pm 0.194$ | $\mathbf{1.309 \pm 0.035}$ | $1.391 \pm 0.108$ | $1.334 \pm 0.016$ |
| | | MAE | $0.783 \pm 0.006$ | $0.852 \pm 0.042$ | $0.813 \pm 0.025$ | $0.776 \pm 0.006$ | $0.797 \pm 0.023$ | $\mathbf{0.767 \pm 0.010}$ |
| | 36 | MSE | $1.611 \pm 0.041$ | $2.024 \pm 0.059$ | $1.619 \pm 0.053$ | $\mathbf{1.405 \pm 0.060}$ | $1.578 \pm 0.155$ | $1.406 \pm 0.003$ |
| | | MAE | $0.832 \pm 0.007$ | $0.901 \pm 0.010$ | $0.861 \pm 0.013$ | $0.799 \pm 0.010$ | $0.869 \pm 0.063$ | $\mathbf{0.787 \pm 0.002}$ |
| | 48 | MSE | $1.546 \pm 0.012$ | $1.861 \pm 0.043$ | $1.612 \pm 0.023$ | $1.546 \pm 0.051$ | $1.708 \pm 0.051$ | $\mathbf{1.435 \pm 0.038}$ |
| | | MAE | $0.816 \pm 0.002$ | $0.889 \pm 0.005$ | $0.878 \pm 0.008$ | $0.835 \pm 0.015$ | $0.838 \pm 0.005$ | $\mathbf{0.807 \pm 0.009}$ |

Table 24: Framework results by prediction length $H$ for Security, Social Good, Traffic with SOFTS. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | $77.228 \pm 2.605$ | $73.497 \pm 0.935$ | $84.945 \pm 1.280$ | $70.842 \pm 1.433$ | $72.041 \pm 0.960$ | $\mathbf{66.965 \pm 1.230}$ |
| | | MAE | $4.677 \pm 0.172$ | $4.459 \pm 0.039$ | $5.242 \pm 0.088$ | $4.186 \pm 0.095$ | $4.325 \pm 0.074$ | $\mathbf{3.645 \pm 0.078}$ |
| | 8 | MSE | $75.244 \pm 2.089$ | $78.144 \pm 1.916$ | $83.731 \pm 1.900$ | $74.492 \pm 1.234$ | $74.837 \pm 2.854$ | $\mathbf{70.517 \pm 2.385}$ |
| | | MAE | $4.423 \pm 0.126$ | $4.661 \pm 0.095$ | $5.076 \pm 0.087$ | $4.398 \pm 0.097$ | $4.407 \pm 0.244$ | $\mathbf{3.986 \pm 0.228}$ |
| | 10 | MSE | $78.634 \pm 2.228$ | $80.299 \pm 1.222$ | $87.325 \pm 2.283$ | $78.052 \pm 1.983$ | $77.031 \pm 2.363$ | $\mathbf{73.182 \pm 0.549}$ |
| | | MAE | $4.580 \pm 0.127$ | $4.674 \pm 0.099$ | $5.184 \pm 0.156$ | $4.482 \pm 0.127$ | $4.420 \pm 0.221$ | $\mathbf{4.012 \pm 0.036}$ |
| | 12 | MSE | $83.121 \pm 2.719$ | $85.850 \pm 0.536$ | $94.891 \pm 4.369$ | $85.851 \pm 3.086$ | $\mathbf{79.474 \pm 1.858}$ | $80.781 \pm 1.862$ |
| | | MAE | $4.736 \pm 0.228$ | $4.962 \pm 0.087$ | $5.544 \pm 0.272$ | $4.896 \pm 0.161$ | $\mathbf{4.406 \pm 0.150}$ | $4.518 \pm 0.105$ |
| Social Good | 6 | MSE | $\mathbf{0.800 \pm 0.006}$ | $0.960 \pm 0.027$ | $0.968 \pm 0.036$ | $0.831 \pm 0.004$ | $0.892 \pm 0.015$ | $0.819 \pm 0.032$ |
| | | MAE | $0.372 \pm 0.003$ | $0.409 \pm 0.005$ | $0.493 \pm 0.008$ | $\mathbf{0.371 \pm 0.003}$ | $0.380 \pm 0.009$ | $0.386 \pm 0.007$ |
| | 8 | MSE | $0.863 \pm 0.008$ | $1.042 \pm 0.016$ | $1.017 \pm 0.037$ | $0.898 \pm 0.005$ | $0.968 \pm 0.031$ | $\mathbf{0.847 \pm 0.022}$ |
| | | MAE | $0.405 \pm 0.003$ | $0.443 \pm 0.006$ | $0.517 \pm 0.011$ | $0.407 \pm 0.001$ | $0.426 \pm 0.024$ | $\mathbf{0.397 \pm 0.010}$ |
| | 10 | MSE | $\mathbf{0.910 \pm 0.005}$ | $1.100 \pm 0.028$ | $1.058 \pm 0.026$ | $0.943 \pm 0.005$ | $1.094 \pm 0.045$ | $0.933 \pm 0.036$ |
| | | MAE | $\mathbf{0.434 \pm 0.003}$ | $0.479 \pm 0.006$ | $0.546 \pm 0.003$ | $0.435 \pm 0.002$ | $0.499 \pm 0.025$ | $0.444 \pm 0.021$ |
| | 12 | MSE | $\mathbf{0.954 \pm 0.006}$ | $1.151 \pm 0.034$ | $1.091 \pm 0.020$ | $0.981 \pm 0.004$ | $1.118 \pm 0.055$ | $0.958 \pm 0.007$ |
| | | MAE | $0.465 \pm 0.004$ | $0.504 \pm 0.008$ | $0.563 \pm 0.007$ | $\mathbf{0.462 \pm 0.002}$ | $0.523 \pm 0.030$ | $0.466 \pm 0.007$ |
| Traffic | 6 | MSE | $0.160 \pm 0.005$ | $0.167 \pm 0.007$ | $0.194 \pm 0.008$ | $0.160 \pm 0.005$ | $0.158 \pm 0.003$ | $\mathbf{0.157 \pm 0.004}$ |
| | | MAE | $0.217 \pm 0.006$ | $0.229 \pm 0.009$ | $0.264 \pm 0.010$ | $0.216 \pm 0.006$ | $\mathbf{0.197 \pm 0.003}$ | $0.222 \pm 0.013$ |
| | 8 | MSE | $0.165 \pm 0.003$ | $0.175 \pm 0.006$ | $0.194 \pm 0.005$ | $0.165 \pm 0.003$ | $0.163 \pm 0.003$ | $\mathbf{0.161 \pm 0.004}$ |
| | | MAE | $0.224 \pm 0.004$ | $0.240 \pm 0.008$ | $0.262 \pm 0.007$ | $0.224 \pm 0.004$ | $\mathbf{0.210 \pm 0.005}$ | $0.227 \pm 0.009$ |
| | 10 | MSE | $0.169 \pm 0.003$ | $0.181 \pm 0.005$ | $0.204 \pm 0.009$ | $0.169 \pm 0.003$ | $0.171 \pm 0.004$ | $\mathbf{0.166 \pm 0.006}$ |
| | | MAE | $0.231 \pm 0.003$ | $0.248 \pm 0.007$ | $0.274 \pm 0.013$ | $0.231 \pm 0.003$ | $\mathbf{0.220 \pm 0.002}$ | $0.232 \pm 0.007$ |
| | 12 | MSE | $0.174 \pm 0.002$ | $0.189 \pm 0.003$ | $0.204 \pm 0.006$ | $0.174 \pm 0.002$ | $0.180 \pm 0.003$ | $\mathbf{0.172 \pm 0.003}$ |
| | | MAE | $0.237 \pm 0.003$ | $0.257 \pm 0.005$ | $0.274 \pm 0.008$ | $0.237 \pm 0.003$ | $\mathbf{0.230 \pm 0.001}$ | $0.245 \pm 0.006$ |

Table 25: Framework results by prediction length $H$ for Agriculture, Climate, Economy with TimeXer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.290 \pm 0.007$ | $0.245 \pm 0.010$ | $0.213 \pm 0.013$ | $0.289 \pm 0.007$ | $0.225 \pm 0.003$ | $\mathbf{0.196 \pm 0.010}$ |
| | | MAE | $0.358 \pm 0.005$ | $0.334 \pm 0.007$ | $0.311 \pm 0.008$ | $0.357 \pm 0.005$ | $0.315 \pm 0.003$ | $\mathbf{0.293 \pm 0.007}$ |
| | 8 | MSE | $0.356 \pm 0.007$ | $0.326 \pm 0.023$ | $0.278 \pm 0.003$ | $0.355 \pm 0.007$ | $0.303 \pm 0.009$ | $\mathbf{0.261 \pm 0.014}$ |
| | | MAE | $0.397 \pm 0.005$ | $0.391 \pm 0.019$ | $0.359 \pm 0.004$ | $0.398 \pm 0.005$ | $0.363 \pm 0.005$ | $\mathbf{0.335 \pm 0.016}$ |
| | 10 | MSE | $0.422 \pm 0.004$ | $0.401 \pm 0.013$ | $0.370 \pm 0.025$ | $0.424 \pm 0.003$ | $0.379 \pm 0.008$ | $\mathbf{0.313 \pm 0.024}$ |
| | | MAE | $0.436 \pm 0.003$ | $0.435 \pm 0.010$ | $0.415 \pm 0.018$ | $0.436 \pm 0.003$ | $0.408 \pm 0.006$ | $\mathbf{0.371 \pm 0.009}$ |
| | 12 | MSE | $0.481 \pm 0.002$ | $0.494 \pm 0.011$ | $0.444 \pm 0.010$ | $0.481 \pm 0.004$ | $0.453 \pm 0.008$ | $\mathbf{0.378 \pm 0.023}$ |
| | | MAE | $0.471 \pm 0.002$ | $0.492 \pm 0.009$ | $0.462 \pm 0.008$ | $0.471 \pm 0.004$ | $0.454 \pm 0.005$ | $\mathbf{0.408 \pm 0.014}$ |
| Climate | 6 | MSE | $0.977 \pm 0.016$ | $1.050 \pm 0.029$ | $0.901 \pm 0.007$ | $0.931 \pm 0.009$ | $0.932 \pm 0.011$ | $\mathbf{0.885 \pm 0.005}$ |
| | | MAE | $0.794 \pm 0.010$ | $0.833 \pm 0.014$ | $0.764 \pm 0.002$ | $0.780 \pm 0.008$ | $0.779 \pm 0.007$ | $\mathbf{0.750 \pm 0.002}$ |
| | 8 | MSE | $0.928 \pm 0.011$ | $1.007 \pm 0.010$ | $0.910 \pm 0.008$ | $0.901 \pm 0.007$ | $0.939 \pm 0.007$ | $\mathbf{0.890 \pm 0.004}$ |
| | | MAE | $0.778 \pm 0.005$ | $0.812 \pm 0.002$ | $0.768 \pm 0.006$ | $0.761 \pm 0.003$ | $0.781 \pm 0.008$ | $\mathbf{0.750 \pm 0.002}$ |
| | 10 | MSE | $0.928 \pm 0.009$ | $1.064 \pm 0.025$ | $0.993 \pm 0.008$ | $0.950 \pm 0.009$ | $1.033 \pm 0.033$ | $\mathbf{0.900 \pm 0.006}$ |
| | | MAE | $0.780 \pm 0.006$ | $0.842 \pm 0.014$ | $0.808 \pm 0.005$ | $0.788 \pm 0.005$ | $0.819 \pm 0.011$ | $\mathbf{0.752 \pm 0.003}$ |
| | 12 | MSE | $0.968 \pm 0.009$ | $1.011 \pm 0.006$ | $0.938 \pm 0.010$ | $0.969 \pm 0.005$ | $0.956 \pm 0.010$ | $\mathbf{0.906 \pm 0.003}$ |
| | | MAE | $0.799 \pm 0.003$ | $0.818 \pm 0.003$ | $0.783 \pm 0.005$ | $0.800 \pm 0.003$ | $0.794 \pm 0.006$ | $\mathbf{0.755 \pm 0.001}$ |
| Economy | 6 | MSE | $0.204 \pm 0.010$ | $0.186 \pm 0.016$ | $\mathbf{0.158 \pm 0.002}$ | $0.202 \pm 0.006$ | $0.212 \pm 0.007$ | $0.163 \pm 0.009$ |
| | | MAE | $0.358 \pm 0.010$ | $0.339 \pm 0.017$ | $\mathbf{0.312 \pm 0.002}$ | $0.356 \pm 0.006$ | $0.366 \pm 0.007$ | $0.316 \pm 0.011$ |
| | 8 | MSE | $0.214 \pm 0.002$ | $0.211 \pm 0.024$ | $\mathbf{0.179 \pm 0.005}$ | $0.217 \pm 0.004$ | $0.220 \pm 0.006$ | $0.194 \pm 0.011$ |
| | | MAE | $0.369 \pm 0.002$ | $0.365 \pm 0.023$ | $\mathbf{0.332 \pm 0.004}$ | $0.372 \pm 0.004$ | $0.375 \pm 0.005$ | $0.346 \pm 0.013$ |
| | 10 | MSE | $0.228 \pm 0.005$ | $0.233 \pm 0.019$ | $\mathbf{0.194 \pm 0.003}$ | $0.231 \pm 0.006$ | $0.234 \pm 0.007$ | $0.216 \pm 0.007$ |
| | | MAE | $0.383 \pm 0.005$ | $0.385 \pm 0.018$ | $\mathbf{0.347 \pm 0.003}$ | $0.386 \pm 0.005$ | $0.389 \pm 0.007$ | $0.370 \pm 0.007$ |
| | 12 | MSE | $0.242 \pm 0.004$ | $0.243 \pm 0.022$ | $\mathbf{0.208 \pm 0.002}$ | $0.237 \pm 0.007$ | $0.237 \pm 0.005$ | $0.224 \pm 0.014$ |
| | | MAE | $0.396 \pm 0.003$ | $0.394 \pm 0.018$ | $\mathbf{0.361 \pm 0.003}$ | $0.392 \pm 0.006$ | $0.392 \pm 0.004$ | $0.378 \pm 0.014$ |

Table 26: Framework results by prediction length $H$ for Energy, Environment, Health with TimeXer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.197 \pm 0.001$ | $0.112 \pm 0.001$ | $0.106 \pm 0.008$ | $0.197 \pm 0.001$ | $0.208 \pm 0.023$ | $\mathbf{0.103 \pm 0.002}$ |
| | | MAE | $0.345 \pm 0.001$ | $0.248 \pm 0.002$ | $0.239 \pm 0.009$ | $0.344 \pm 0.001$ | $0.350 \pm 0.011$ | $\mathbf{0.228 \pm 0.004}$ |
| | 24 | MSE | $0.276 \pm 0.001$ | $0.213 \pm 0.002$ | $\mathbf{0.205 \pm 0.001}$ | $0.277 \pm 0.001$ | $0.303 \pm 0.030$ | $0.212 \pm 0.002$ |
| | | MAE | $0.403 \pm 0.001$ | $0.345 \pm 0.001$ | $0.338 \pm 0.001$ | $0.403 \pm 0.001$ | $0.421 \pm 0.019$ | $\mathbf{0.337 \pm 0.002}$ |
| | 36 | MSE | $0.357 \pm 0.001$ | $0.302 \pm 0.003$ | $\mathbf{0.294 \pm 0.003}$ | $0.357 \pm 0.001$ | $0.392 \pm 0.022$ | $0.306 \pm 0.014$ |
| | | MAE | $0.457 \pm 0.001$ | $0.413 \pm 0.002$ | $\mathbf{0.407 \pm 0.001}$ | $0.457 \pm 0.001$ | $0.482 \pm 0.013$ | $0.411 \pm 0.015$ |
| | 48 | MSE | $0.441 \pm 0.002$ | $0.387 \pm 0.000$ | $\mathbf{0.379 \pm 0.004}$ | $0.441 \pm 0.002$ | $0.490 \pm 0.028$ | $\mathbf{0.379 \pm 0.015}$ |
| | | MAE | $0.516 \pm 0.001$ | $0.478 \pm 0.001$ | $0.470 \pm 0.003$ | $0.516 \pm 0.001$ | $0.542 \pm 0.015$ | $\mathbf{0.464 \pm 0.010}$ |
| Environment | 48 | MSE | $\mathbf{0.269 \pm 0.000}$ | $0.303 \pm 0.001$ | $0.302 \pm 0.001$ | $\mathbf{0.269 \pm 0.000}$ | $0.272 \pm 0.000$ | $\mathbf{0.269 \pm 0.000}$ |
| | | MAE | $\mathbf{0.373 \pm 0.001}$ | $0.390 \pm 0.002$ | $0.388 \pm 0.002$ | $\mathbf{0.373 \pm 0.001}$ | $0.377 \pm 0.001$ | $0.374 \pm 0.000$ |
| | 96 | MSE | $\mathbf{0.279 \pm 0.000}$ | $0.331 \pm 0.000$ | $0.329 \pm 0.000$ | $\mathbf{0.279 \pm 0.000}$ | $0.282 \pm 0.001$ | $\mathbf{0.279 \pm 0.000}$ |
| | | MAE | $\mathbf{0.376 \pm 0.001}$ | $0.402 \pm 0.002$ | $0.399 \pm 0.001$ | $\mathbf{0.376 \pm 0.001}$ | $0.381 \pm 0.002$ | $\mathbf{0.376 \pm 0.000}$ |
| | 192 | MSE | $0.272 \pm 0.000$ | $0.329 \pm 0.001$ | $0.326 \pm 0.001$ | $0.272 \pm 0.000$ | $0.273 \pm 0.002$ | $\mathbf{0.271 \pm 0.000}$ |
| | | MAE | $\mathbf{0.367 \pm 0.000}$ | $0.397 \pm 0.001$ | $0.394 \pm 0.000$ | $\mathbf{0.367 \pm 0.000}$ | $0.372 \pm 0.004$ | $\mathbf{0.367 \pm 0.001}$ |
| | 336 | MSE | $\mathbf{0.268 \pm 0.000}$ | $0.327 \pm 0.001$ | $0.323 \pm 0.001$ | $\mathbf{0.268 \pm 0.000}$ | $0.269 \pm 0.000$ | $\mathbf{0.268 \pm 0.000}$ |
| | | MAE | $0.367 \pm 0.000$ | $0.397 \pm 0.001$ | $0.396 \pm 0.001$ | $0.367 \pm 0.000$ | $0.370 \pm 0.002$ | $\mathbf{0.366 \pm 0.001}$ |
| Health | 12 | MSE | $\mathbf{0.957 \pm 0.010}$ | $1.129 \pm 0.014$ | $0.962 \pm 0.006$ | $1.011 \pm 0.017$ | $1.015 \pm 0.020$ | $1.009 \pm 0.013$ |
| | | MAE | $\mathbf{0.650 \pm 0.002}$ | $0.689 \pm 0.008$ | $0.676 \pm 0.003$ | $0.677 \pm 0.006$ | $0.679 \pm 0.008$ | $0.677 \pm 0.009$ |
| | 24 | MSE | $1.314 \pm 0.010$ | $1.605 \pm 0.023$ | $\mathbf{1.228 \pm 0.008}$ | $1.352 \pm 0.007$ | $1.315 \pm 0.015$ | $1.365 \pm 0.009$ |
| | | MAE | $0.772 \pm 0.007$ | $0.854 \pm 0.012$ | $\mathbf{0.766 \pm 0.004}$ | $0.781 \pm 0.003$ | $0.782 \pm 0.019$ | $0.787 \pm 0.004$ |
| | 36 | MSE | $1.423 \pm 0.010$ | $2.037 \pm 0.179$ | $\mathbf{1.395 \pm 0.019}$ | $1.433 \pm 0.016$ | $1.472 \pm 0.119$ | $1.436 \pm 0.126$ |
| | | MAE | $\mathbf{0.804 \pm 0.008}$ | $0.931 \pm 0.022$ | $0.808 \pm 0.008$ | $0.809 \pm 0.009$ | $0.829 \pm 0.040$ | $0.822 \pm 0.047$ |
| | 48 | MSE | $1.432 \pm 0.003$ | $1.843 \pm 0.059$ | $1.479 \pm 0.014$ | $1.455 \pm 0.010$ | $1.535 \pm 0.074$ | $\mathbf{1.424 \pm 0.080}$ |
| | | MAE | $\mathbf{0.804 \pm 0.002}$ | $0.927 \pm 0.013$ | $0.831 \pm 0.009$ | $0.814 \pm 0.006$ | $0.839 \pm 0.018$ | $0.810 \pm 0.031$ |

Table 27: Framework results by prediction length $H$ for Security, Social Good, Traffic with TimeXer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | $69.218 \pm 1.374$ | $69.810 \pm 0.735$ | $79.301 \pm 0.918$ | $69.670 \pm 1.406$ | $68.397 \pm 0.865$ | $\mathbf{66.123 \pm 0.313}$ |
| | | MAE | $4.118 \pm 0.134$ | $4.252 \pm 0.107$ | $4.887 \pm 0.068$ | $4.193 \pm 0.139$ | $4.073 \pm 0.060$ | $\mathbf{3.655 \pm 0.113}$ |
| | 8 | MSE | $70.842 \pm 0.439$ | $74.649 \pm 1.896$ | $80.793 \pm 1.184$ | $\mathbf{69.837 \pm 0.591}$ | $71.412 \pm 1.211$ | $70.022 \pm 0.572$ |
| | | MAE | $4.086 \pm 0.056$ | $4.505 \pm 0.169$ | $4.882 \pm 0.040$ | $\mathbf{3.843 \pm 0.067}$ | $4.179 \pm 0.113$ | $3.950 \pm 0.059$ |
| | 10 | MSE | $74.195 \pm 1.335$ | $77.318 \pm 1.750$ | $87.511 \pm 2.638$ | $74.784 \pm 0.734$ | $73.149 \pm 0.664$ | $\mathbf{72.734 \pm 0.130}$ |
| | | MAE | $4.180 \pm 0.110$ | $4.530 \pm 0.158$ | $5.276 \pm 0.146$ | $4.251 \pm 0.057$ | $4.052 \pm 0.056$ | $\mathbf{3.900 \pm 0.038}$ |
| | 12 | MSE | $76.308 \pm 0.381$ | $85.101 \pm 3.752$ | $92.788 \pm 2.025$ | $77.733 \pm 0.474$ | $76.993 \pm 0.756$ | $\mathbf{76.230 \pm 0.552}$ |
| | | MAE | $4.171 \pm 0.031$ | $5.025 \pm 0.266$ | $5.466 \pm 0.142$ | $4.320 \pm 0.030$ | $4.235 \pm 0.087$ | $\mathbf{4.121 \pm 0.067}$ |
| Social Good | 6 | MSE | $0.747 \pm 0.010$ | $0.778 \pm 0.015$ | $0.820 \pm 0.018$ | $0.747 \pm 0.010$ | $0.808 \pm 0.019$ | $\mathbf{0.716 \pm 0.007}$ |
| | | MAE | $0.390 \pm 0.006$ | $0.409 \pm 0.007$ | $0.412 \pm 0.023$ | $0.390 \pm 0.005$ | $0.404 \pm 0.022$ | $\mathbf{0.364 \pm 0.006}$ |
| | 8 | MSE | $0.816 \pm 0.010$ | $0.866 \pm 0.024$ | $0.921 \pm 0.014$ | $0.815 \pm 0.010$ | $0.919 \pm 0.047$ | $\mathbf{0.792 \pm 0.009}$ |
| | | MAE | $0.423 \pm 0.004$ | $0.450 \pm 0.010$ | $0.457 \pm 0.013$ | $0.423 \pm 0.005$ | $0.480 \pm 0.050$ | $\mathbf{0.395 \pm 0.010}$ |
| | 10 | MSE | $0.870 \pm 0.007$ | $0.926 \pm 0.010$ | $1.002 \pm 0.037$ | $0.869 \pm 0.007$ | $0.998 \pm 0.089$ | $\mathbf{0.845 \pm 0.009}$ |
| | | MAE | $0.455 \pm 0.003$ | $0.483 \pm 0.005$ | $0.509 \pm 0.032$ | $0.453 \pm 0.003$ | $0.524 \pm 0.063$ | $\mathbf{0.422 \pm 0.012}$ |
| | 12 | MSE | $0.922 \pm 0.006$ | $0.980 \pm 0.008$ | $1.044 \pm 0.033$ | $0.922 \pm 0.006$ | $1.076 \pm 0.026$ | $\mathbf{0.897 \pm 0.008}$ |
| | | MAE | $0.486 \pm 0.003$ | $0.513 \pm 0.007$ | $0.533 \pm 0.028$ | $0.485 \pm 0.004$ | $0.579 \pm 0.025$ | $\mathbf{0.446 \pm 0.005}$ |
| Traffic | 6 | MSE | $0.165 \pm 0.002$ | $0.177 \pm 0.003$ | $0.158 \pm 0.002$ | $0.164 \pm 0.001$ | $0.159 \pm 0.006$ | $\mathbf{0.147 \pm 0.003}$ |
| | | MAE | $0.224 \pm 0.003$ | $0.260 \pm 0.004$ | $0.231 \pm 0.002$ | $0.225 \pm 0.004$ | $0.226 \pm 0.014$ | $\mathbf{0.219 \pm 0.003}$ |
| | 8 | MSE | $0.170 \pm 0.001$ | $0.185 \pm 0.003$ | $0.164 \pm 0.003$ | $0.170 \pm 0.001$ | $0.160 \pm 0.005$ | $\mathbf{0.147 \pm 0.004}$ |
| | | MAE | $0.228 \pm 0.002$ | $0.266 \pm 0.005$ | $0.234 \pm 0.002$ | $0.228 \pm 0.003$ | $0.226 \pm 0.009$ | $\mathbf{0.222 \pm 0.006}$ |
| | 10 | MSE | $0.175 \pm 0.001$ | $0.192 \pm 0.005$ | $0.172 \pm 0.003$ | $0.174 \pm 0.001$ | $0.164 \pm 0.006$ | $\mathbf{0.156 \pm 0.002}$ |
| | | MAE | $\mathbf{0.232 \pm 0.003}$ | $0.273 \pm 0.006$ | $0.242 \pm 0.003$ | $\mathbf{0.232 \pm 0.003}$ | $\mathbf{0.232 \pm 0.011}$ | $\mathbf{0.232 \pm 0.010}$ |
| | 12 | MSE | $0.178 \pm 0.001$ | $0.197 \pm 0.002$ | $0.177 \pm 0.002$ | $0.178 \pm 0.001$ | $0.167 \pm 0.004$ | $\mathbf{0.159 \pm 0.002}$ |
| | | MAE | $\mathbf{0.234 \pm 0.003}$ | $0.275 \pm 0.005$ | $0.244 \pm 0.002$ | $\mathbf{0.234 \pm 0.003}$ | $0.238 \pm 0.011$ | $\mathbf{0.234 \pm 0.006}$ |

Table 28: Framework results by prediction length $H$ for Agriculture, Climate, Economy with Transformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $1.173 \pm 0.172$ | $2.024 \pm 0.378$ | $1.501 \pm 0.559$ | $\mathbf{0.904 \pm 0.103}$ | $1.157 \pm 0.181$ | $0.947 \pm 0.499$ |
| | | MAE | $0.752 \pm 0.073$ | $0.971 \pm 0.082$ | $0.846 \pm 0.159$ | $\mathbf{0.658 \pm 0.027}$ | $0.758 \pm 0.061$ | $0.706 \pm 0.194$ |
| | 8 | MSE | $1.295 \pm 0.094$ | $2.232 \pm 0.494$ | $1.610 \pm 0.395$ | $\mathbf{0.944 \pm 0.047}$ | $1.178 \pm 0.127$ | $0.973 \pm 0.253$ |
| | | MAE | $0.783 \pm 0.043$ | $1.035 \pm 0.143$ | $0.897 \pm 0.073$ | $\mathbf{0.668 \pm 0.012}$ | $0.746 \pm 0.043$ | $0.716 \pm 0.116$ |
| | 10 | MSE | $1.364 \pm 0.098$ | $1.948 \pm 0.270$ | $1.724 \pm 0.488$ | $\mathbf{0.977 \pm 0.124}$ | $1.372 \pm 0.129$ | $1.077 \pm 0.494$ |
| | | MAE | $0.810 \pm 0.048$ | $0.944 \pm 0.054$ | $0.904 \pm 0.128$ | $\mathbf{0.681 \pm 0.025}$ | $0.818 \pm 0.057$ | $0.744 \pm 0.180$ |
| | 12 | MSE | $1.447 \pm 0.283$ | $2.110 \pm 0.269$ | $1.625 \pm 0.416$ | $\mathbf{1.068 \pm 0.101}$ | $1.205 \pm 0.175$ | $1.152 \pm 0.318$ |
| | | MAE | $0.859 \pm 0.115$ | $0.998 \pm 0.097$ | $0.862 \pm 0.093$ | $\mathbf{0.704 \pm 0.031}$ | $0.760 \pm 0.052$ | $0.752 \pm 0.120$ |
| Climate | 6 | MSE | $0.949 \pm 0.043$ | $1.000 \pm 0.013$ | $0.904 \pm 0.016$ | $\mathbf{0.872 \pm 0.004}$ | $1.001 \pm 0.030$ | $0.903 \pm 0.031$ |
| | | MAE | $0.767 \pm 0.016$ | $0.800 \pm 0.006$ | $0.756 \pm 0.008$ | $\mathbf{0.745 \pm 0.002}$ | $0.801 \pm 0.019$ | $0.754 \pm 0.014$ |
| | 8 | MSE | $0.980 \pm 0.052$ | $1.026 \pm 0.021$ | $0.887 \pm 0.011$ | $\mathbf{0.881 \pm 0.012}$ | $1.054 \pm 0.086$ | $0.890 \pm 0.015$ |
| | | MAE | $0.785 \pm 0.024$ | $0.807 \pm 0.008$ | $0.746 \pm 0.005$ | $\mathbf{0.744 \pm 0.006}$ | $0.813 \pm 0.039$ | $0.745 \pm 0.005$ |
| | 10 | MSE | $0.946 \pm 0.052$ | $1.017 \pm 0.036$ | $0.994 \pm 0.059$ | $\mathbf{0.884 \pm 0.010}$ | $1.104 \pm 0.031$ | $0.893 \pm 0.012$ |
| | | MAE | $0.769 \pm 0.020$ | $0.803 \pm 0.014$ | $0.781 \pm 0.022$ | $0.745 \pm 0.004$ | $0.836 \pm 0.010$ | $\mathbf{0.744 \pm 0.003}$ |
| | 12 | MSE | $0.942 \pm 0.052$ | $1.082 \pm 0.058$ | $0.967 \pm 0.041$ | $\mathbf{0.884 \pm 0.002}$ | $1.211 \pm 0.084$ | $0.899 \pm 0.015$ |
| | | MAE | $0.766 \pm 0.021$ | $0.827 \pm 0.019$ | $0.778 \pm 0.013$ | $\mathbf{0.744 \pm 0.001}$ | $0.872 \pm 0.029$ | $0.747 \pm 0.004$ |
| Economy | 6 | MSE | $0.694 \pm 0.136$ | $0.957 \pm 0.168$ | $0.895 \pm 0.093$ | $\mathbf{0.335 \pm 0.014}$ | $0.590 \pm 0.058$ | $0.404 \pm 0.062$ |
| | | MAE | $0.680 \pm 0.071$ | $0.817 \pm 0.081$ | $0.790 \pm 0.046$ | $\mathbf{0.461 \pm 0.012}$ | $0.623 \pm 0.034$ | $0.508 \pm 0.051$ |
| | 8 | MSE | $0.683 \pm 0.219$ | $0.865 \pm 0.044$ | $0.854 \pm 0.130$ | $\mathbf{0.337 \pm 0.031}$ | $0.665 \pm 0.187$ | $0.414 \pm 0.102$ |
| | | MAE | $0.669 \pm 0.119$ | $0.770 \pm 0.024$ | $0.767 \pm 0.066$ | $\mathbf{0.460 \pm 0.025}$ | $0.660 \pm 0.106$ | $0.512 \pm 0.068$ |
| | 10 | MSE | $0.639 \pm 0.110$ | $1.021 \pm 0.173$ | $0.900 \pm 0.085$ | $\mathbf{0.371 \pm 0.029}$ | $0.645 \pm 0.080$ | $0.445 \pm 0.078$ |
| | | MAE | $0.656 \pm 0.059$ | $0.844 \pm 0.076$ | $0.794 \pm 0.043$ | $\mathbf{0.485 \pm 0.022}$ | $0.654 \pm 0.047$ | $0.532 \pm 0.059$ |
| | 12 | MSE | $0.726 \pm 0.103$ | $1.013 \pm 0.173$ | $0.970 \pm 0.063$ | $\mathbf{0.403 \pm 0.046}$ | $0.824 \pm 0.142$ | $0.464 \pm 0.046$ |
| | | MAE | $0.701 \pm 0.053$ | $0.841 \pm 0.081$ | $0.829 \pm 0.026$ | $\mathbf{0.507 \pm 0.033}$ | $0.752 \pm 0.076$ | $0.550 \pm 0.033$ |

Table 29: Framework results by prediction length $H$ for Energy, Environment, Health with Transformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.281 \pm 0.030$ | $0.180 \pm 0.023$ | $0.144 \pm 0.016$ | $0.213 \pm 0.008$ | $0.285 \pm 0.063$ | $\mathbf{0.129 \pm 0.011}$ |
| | | MAE | $0.397 \pm 0.021$ | $0.306 \pm 0.024$ | $0.275 \pm 0.015$ | $0.360 \pm 0.010$ | $0.406 \pm 0.043$ | $\mathbf{0.264 \pm 0.012}$ |
| | 24 | MSE | $0.337 \pm 0.032$ | $0.313 \pm 0.055$ | $0.290 \pm 0.061$ | $0.319 \pm 0.022$ | $0.381 \pm 0.079$ | $\mathbf{0.222 \pm 0.023}$ |
| | | MAE | $0.448 \pm 0.015$ | $0.421 \pm 0.039$ | $0.402 \pm 0.059$ | $0.436 \pm 0.014$ | $0.478 \pm 0.053$ | $\mathbf{0.359 \pm 0.023}$ |
| | 36 | MSE | $0.480 \pm 0.061$ | $0.354 \pm 0.028$ | $0.404 \pm 0.052$ | $0.390 \pm 0.011$ | $0.461 \pm 0.053$ | $\mathbf{0.309 \pm 0.013}$ |
| | | MAE | $0.548 \pm 0.028$ | $0.463 \pm 0.021$ | $0.495 \pm 0.041$ | $0.495 \pm 0.008$ | $0.543 \pm 0.026$ | $\mathbf{0.431 \pm 0.015}$ |
| | 48 | MSE | $0.540 \pm 0.052$ | $0.481 \pm 0.060$ | $0.480 \pm 0.058$ | $0.469 \pm 0.031$ | $0.606 \pm 0.057$ | $\mathbf{0.409 \pm 0.068}$ |
| | | MAE | $0.608 \pm 0.028$ | $0.559 \pm 0.039$ | $0.556 \pm 0.033$ | $0.549 \pm 0.008$ | $0.638 \pm 0.035$ | $\mathbf{0.503 \pm 0.044}$ |
| Environment | 48 | MSE | $0.271 \pm 0.002$ | $0.300 \pm 0.018$ | $0.295 \pm 0.002$ | $0.272 \pm 0.002$ | $0.284 \pm 0.005$ | $\mathbf{0.270 \pm 0.002}$ |
| | | MAE | $0.374 \pm 0.005$ | $0.418 \pm 0.024$ | $0.400 \pm 0.011$ | $\mathbf{0.369 \pm 0.002}$ | $0.398 \pm 0.014$ | $0.370 \pm 0.002$ |
| | 96 | MSE | $0.281 \pm 0.003$ | $0.349 \pm 0.032$ | $0.322 \pm 0.006$ | $\mathbf{0.280 \pm 0.002}$ | $0.295 \pm 0.005$ | $\mathbf{0.280 \pm 0.002}$ |
| | | MAE | $0.392 \pm 0.010$ | $0.473 \pm 0.028$ | $0.449 \pm 0.008$ | $0.393 \pm 0.007$ | $0.417 \pm 0.008$ | $\mathbf{0.390 \pm 0.004}$ |
| | 192 | MSE | $\mathbf{0.279 \pm 0.005}$ | $0.389 \pm 0.030$ | $0.338 \pm 0.009$ | $\mathbf{0.279 \pm 0.002}$ | $0.291 \pm 0.006$ | $0.287 \pm 0.006$ |
| | | MAE | $0.406 \pm 0.008$ | $0.510 \pm 0.024$ | $0.463 \pm 0.009$ | $\mathbf{0.405 \pm 0.003}$ | $0.413 \pm 0.010$ | $0.411 \pm 0.004$ |
| | 336 | MSE | $\mathbf{0.277 \pm 0.011}$ | $0.376 \pm 0.026$ | $0.343 \pm 0.010$ | $0.279 \pm 0.010$ | $0.279 \pm 0.011$ | $0.287 \pm 0.002$ |
| | | MAE | $0.408 \pm 0.020$ | $0.487 \pm 0.021$ | $0.460 \pm 0.006$ | $0.413 \pm 0.018$ | $\mathbf{0.406 \pm 0.018}$ | $0.412 \pm 0.003$ |
| Health | 12 | MSE | $1.013 \pm 0.057$ | $1.059 \pm 0.114$ | $\mathbf{0.945 \pm 0.051}$ | $1.116 \pm 0.150$ | $1.247 \pm 0.071$ | $1.023 \pm 0.114$ |
| | | MAE | $0.693 \pm 0.020$ | $0.706 \pm 0.037$ | $\mathbf{0.658 \pm 0.017}$ | $0.734 \pm 0.057$ | $0.771 \pm 0.027$ | $0.673 \pm 0.043$ |
| | 24 | MSE | $1.265 \pm 0.051$ | $1.408 \pm 0.036$ | $1.326 \pm 0.061$ | $\mathbf{1.180 \pm 0.027}$ | $1.334 \pm 0.080$ | $1.271 \pm 0.051$ |
| | | MAE | $0.773 \pm 0.020$ | $0.818 \pm 0.021$ | $0.786 \pm 0.020$ | $\mathbf{0.726 \pm 0.004}$ | $0.827 \pm 0.016$ | $0.777 \pm 0.025$ |
| | 36 | MSE | $1.262 \pm 0.019$ | $1.400 \pm 0.118$ | $1.315 \pm 0.087$ | $\mathbf{1.213 \pm 0.014}$ | $1.301 \pm 0.097$ | $1.325 \pm 0.064$ |
| | | MAE | $0.765 \pm 0.017$ | $0.823 \pm 0.038$ | $0.782 \pm 0.025$ | $\mathbf{0.744 \pm 0.012}$ | $0.790 \pm 0.039$ | $0.796 \pm 0.015$ |
| | 48 | MSE | $1.330 \pm 0.023$ | $1.393 \pm 0.077$ | $1.281 \pm 0.033$ | $1.298 \pm 0.057$ | $\mathbf{1.239 \pm 0.037}$ | $1.317 \pm 0.055$ |
| | | MAE | $0.780 \pm 0.003$ | $0.811 \pm 0.017$ | $0.782 \pm 0.011$ | $0.790 \pm 0.028$ | $\mathbf{0.770 \pm 0.028}$ | $0.778 \pm 0.018$ |

Table 30: Framework results by prediction length $H$ for Security, Social Good, Traffic with Transformer. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 82.153 ± 6.060 | 81.319 ± 3.750 | 83.674 ± 3.673 | **79.003 ± 1.859** | 81.668 ± 2.707 | 84.062 ± 3.175 |
| | | MAE | 5.113 ± 0.530 | 5.129 ± 0.315 | 5.285 ± 0.296 | **4.875 ± 0.171** | 5.105 ± 0.238 | 5.319 ± 0.261 |
| | 8 | MSE | 83.889 ± 0.567 | 85.450 ± 3.110 | 87.241 ± 3.112 | **82.675 ± 1.859** | 85.171 ± 2.624 | 85.735 ± 4.170 |
| | | MAE | 5.179 ± 0.051 | 5.321 ± 0.255 | 5.453 ± 0.260 | **5.086 ± 0.174** | 5.267 ± 0.226 | 5.274 ± 0.361 |
| | 10 | MSE | 88.289 ± 3.974 | 90.639 ± 8.978 | **84.934 ± 6.062** | 85.538 ± 2.207 | 91.059 ± 3.900 | 86.550 ± 4.227 |
| | | MAE | 5.383 ± 0.357 | 5.609 ± 0.737 | **5.067 ± 0.563** | 5.174 ± 0.213 | 5.623 ± 0.333 | 5.252 ± 0.401 |
| | 12 | MSE | 88.924 ± 1.367 | 93.495 ± 3.749 | 89.063 ± 3.544 | **88.342 ± 1.772** | 93.156 ± 2.659 | 88.973 ± 3.207 |
| | | MAE | 5.268 ± 0.133 | 5.713 ± 0.321 | 5.259 ± 0.328 | **5.250 ± 0.164** | 5.634 ± 0.240 | 5.276 ± 0.325 |
| Social Good | 6 | MSE | 0.711 ± 0.047 | 0.740 ± 0.126 | 0.735 ± 0.046 | 0.697 ± 0.016 | **0.667 ± 0.034** | 0.754 ± 0.013 |
| | | MAE | 0.407 ± 0.057 | 0.436 ± 0.074 | 0.430 ± 0.021 | **0.340 ± 0.015** | 0.346 ± 0.032 | 0.446 ± 0.030 |
| | 8 | MSE | **0.722 ± 0.023** | 0.810 ± 0.127 | 0.795 ± 0.040 | 0.730 ± 0.027 | 0.760 ± 0.043 | 0.808 ± 0.017 |
| | | MAE | 0.397 ± 0.012 | 0.497 ± 0.102 | 0.451 ± 0.024 | **0.374 ± 0.022** | 0.395 ± 0.020 | 0.477 ± 0.020 |
| | 10 | MSE | **0.755 ± 0.024** | 0.836 ± 0.086 | 0.817 ± 0.030 | **0.755 ± 0.021** | 0.815 ± 0.029 | 0.848 ± 0.029 |
| | | MAE | 0.433 ± 0.020 | 0.504 ± 0.052 | 0.466 ± 0.025 | **0.376 ± 0.023** | 0.455 ± 0.035 | 0.442 ± 0.020 |
| | 12 | MSE | 0.779 ± 0.015 | 0.934 ± 0.157 | 0.864 ± 0.047 | **0.761 ± 0.018** | 0.849 ± 0.035 | 0.824 ± 0.018 |
| | | MAE | 0.458 ± 0.016 | 0.583 ± 0.085 | 0.519 ± 0.059 | **0.393 ± 0.025** | 0.470 ± 0.035 | 0.438 ± 0.028 |
| Traffic | 6 | MSE | 0.170 ± 0.005 | 0.220 ± 0.009 | 0.195 ± 0.020 | 0.159 ± 0.003 | 0.162 ± 0.004 | **0.155 ± 0.003** |
| | | MAE | 0.314 ± 0.012 | 0.377 ± 0.011 | 0.349 ± 0.032 | 0.285 ± 0.006 | 0.300 ± 0.008 | **0.268 ± 0.005** |
| | 8 | MSE | 0.171 ± 0.003 | 0.221 ± 0.007 | 0.190 ± 0.019 | 0.162 ± 0.001 | 0.164 ± 0.003 | **0.160 ± 0.005** |
| | | MAE | 0.312 ± 0.007 | 0.376 ± 0.009 | 0.342 ± 0.034 | 0.288 ± 0.002 | 0.300 ± 0.009 | **0.281 ± 0.015** |
| | 10 | MSE | 0.173 ± 0.003 | 0.239 ± 0.005 | 0.185 ± 0.013 | 0.169 ± 0.001 | 0.167 ± 0.004 | **0.160 ± 0.004** |
| | | MAE | 0.313 ± 0.007 | 0.397 ± 0.006 | 0.336 ± 0.028 | 0.300 ± 0.004 | 0.305 ± 0.007 | **0.278 ± 0.010** |
| | 12 | MSE | 0.178 ± 0.004 | 0.252 ± 0.013 | 0.185 ± 0.015 | 0.170 ± 0.003 | 0.170 ± 0.005 | **0.164 ± 0.007** |
| | | MAE | 0.319 ± 0.008 | 0.409 ± 0.013 | 0.331 ± 0.027 | 0.297 ± 0.006 | 0.307 ± 0.011 | **0.279 ± 0.022** |

### E.2.2 FNSPID

The following tables report the four FNSPID tasks at every horizon for the same six backbones and six methods. Their four-horizon means and mean horizon-specific standard deviations give the entries in Tables 9 and 10.

Table 31: FNSPID framework results by prediction length $H$ for WMT, GILD with iTransformer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | 0.064 ± 0.006 | 0.062 ± 0.004 | 0.061 ± 0.003 | **0.060 ± 0.002** | 0.067 ± 0.004 | 0.063 ± 0.001 |
| | | MAE | 0.204 ± 0.009 | 0.199 ± 0.005 | **0.197 ± 0.006** | **0.197 ± 0.003** | 0.207 ± 0.006 | 0.201 ± 0.002 |
| | 8 | MSE | 0.083 ± 0.002 | 0.083 ± 0.002 | **0.079 ± 0.004** | 0.080 ± 0.002 | 0.086 ± 0.006 | 0.082 ± 0.003 |
| | | MAE | 0.231 ± 0.003 | 0.231 ± 0.003 | **0.225 ± 0.006** | 0.227 ± 0.002 | 0.233 ± 0.007 | 0.229 ± 0.004 |
| | 10 | MSE | 0.101 ± 0.004 | 0.099 ± 0.003 | **0.094 ± 0.003** | 0.097 ± 0.004 | 0.108 ± 0.009 | 0.100 ± 0.001 |
| | | MAE | 0.253 ± 0.005 | 0.251 ± 0.004 | **0.243 ± 0.002** | 0.248 ± 0.004 | 0.258 ± 0.010 | 0.250 ± 0.001 |
| | 12 | MSE | 0.120 ± 0.005 | 0.121 ± 0.005 | **0.113 ± 0.006** | 0.117 ± 0.003 | 0.132 ± 0.013 | 0.119 ± 0.004 |
| | | MAE | 0.273 ± 0.006 | 0.274 ± 0.006 | **0.264 ± 0.008** | 0.269 ± 0.002 | 0.283 ± 0.013 | 0.270 ± 0.004 |
| GILD | 6 | MSE | **0.046 ± 0.001** | **0.046 ± 0.000** | 0.047 ± 0.001 | **0.046 ± 0.001** | 0.049 ± 0.002 | 0.047 ± 0.002 |
| | | MAE | **0.147 ± 0.001** | 0.148 ± 0.001 | 0.149 ± 0.001 | 0.148 ± 0.002 | 0.154 ± 0.003 | 0.151 ± 0.004 |
| | 8 | MSE | **0.052 ± 0.001** | **0.052 ± 0.001** | 0.053 ± 0.001 | **0.052 ± 0.000** | 0.054 ± 0.003 | **0.052 ± 0.001** |
| | | MAE | **0.158 ± 0.001** | 0.160 ± 0.002 | 0.159 ± 0.002 | **0.158 ± 0.001** | 0.164 ± 0.005 | 0.161 ± 0.001 |
| | 10 | MSE | **0.059 ± 0.001** | 0.060 ± 0.002 | 0.061 ± 0.001 | 0.060 ± 0.001 | 0.063 ± 0.001 | 0.062 ± 0.004 |
| | | MAE | **0.172 ± 0.002** | 0.173 ± 0.004 | 0.173 ± 0.001 | **0.172 ± 0.001** | 0.181 ± 0.002 | 0.177 ± 0.005 |
| | 12 | MSE | 0.068 ± 0.002 | 0.068 ± 0.002 | 0.067 ± 0.002 | 0.068 ± 0.002 | 0.072 ± 0.000 | **0.066 ± 0.002** |
| | | MAE | 0.186 ± 0.003 | 0.186 ± 0.003 | **0.183 ± 0.003** | 0.186 ± 0.003 | 0.194 ± 0.001 | **0.183 ± 0.003** |

Table 32: FNSPID framework results by prediction length $H$ for WFC, SBUX with iTransformer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | $0.247 \pm 0.015$ | $0.250 \pm 0.020$ | $\mathbf{0.233 \pm 0.006}$ | $0.263 \pm 0.022$ | $0.244 \pm 0.014$ | $0.242 \pm 0.004$ |
| | | MAE | $0.333 \pm 0.010$ | $0.336 \pm 0.012$ | $\mathbf{0.324 \pm 0.003}$ | $0.340 \pm 0.014$ | $0.335 \pm 0.009$ | $0.334 \pm 0.005$ |
| | 8 | MSE | $0.321 \pm 0.016$ | $0.322 \pm 0.023$ | $\mathbf{0.315 \pm 0.016}$ | $0.335 \pm 0.022$ | $0.322 \pm 0.010$ | $0.317 \pm 0.021$ |
| | | MAE | $0.377 \pm 0.008$ | $0.377 \pm 0.012$ | $\mathbf{0.372 \pm 0.010}$ | $0.380 \pm 0.012$ | $0.380 \pm 0.007$ | $0.378 \pm 0.016$ |
| | 10 | MSE | $\mathbf{0.394 \pm 0.014}$ | $0.395 \pm 0.014$ | $0.412 \pm 0.031$ | $0.440 \pm 0.012$ | $0.447 \pm 0.024$ | $0.398 \pm 0.019$ |
| | | MAE | $\mathbf{0.410 \pm 0.009}$ | $0.411 \pm 0.009$ | $0.418 \pm 0.017$ | $0.427 \pm 0.006$ | $0.440 \pm 0.012$ | $0.415 \pm 0.008$ |
| | 12 | MSE | $\mathbf{0.488 \pm 0.021}$ | $0.492 \pm 0.024$ | $0.494 \pm 0.027$ | $0.519 \pm 0.025$ | $0.569 \pm 0.131$ | $0.491 \pm 0.009$ |
| | | MAE | $\mathbf{0.448 \pm 0.009}$ | $0.451 \pm 0.010$ | $0.452 \pm 0.018$ | $0.458 \pm 0.014$ | $0.490 \pm 0.066$ | $0.454 \pm 0.006$ |
| SBUX | 6 | MSE | $0.850 \pm 0.035$ | $0.842 \pm 0.050$ | $\mathbf{0.739 \pm 0.045}$ | $0.845 \pm 0.027$ | $0.906 \pm 0.099$ | $0.787 \pm 0.031$ |
| | | MAE | $0.588 \pm 0.023$ | $0.586 \pm 0.025$ | $\mathbf{0.546 \pm 0.008}$ | $0.575 \pm 0.008$ | $0.616 \pm 0.037$ | $0.564 \pm 0.006$ |
| | 8 | MSE | $1.029 \pm 0.103$ | $1.005 \pm 0.093$ | $\mathbf{0.957 \pm 0.051}$ | $1.097 \pm 0.110$ | $1.098 \pm 0.154$ | $0.959 \pm 0.038$ |
| | | MAE | $0.640 \pm 0.028$ | $0.636 \pm 0.023$ | $\mathbf{0.621 \pm 0.012}$ | $0.650 \pm 0.031$ | $0.660 \pm 0.039$ | $0.637 \pm 0.010$ |
| | 10 | MSE | $1.282 \pm 0.049$ | $1.277 \pm 0.071$ | $1.239 \pm 0.032$ | $1.285 \pm 0.035$ | $1.408 \pm 0.054$ | $\mathbf{1.199 \pm 0.022}$ |
| | | MAE | $0.708 \pm 0.014$ | $0.713 \pm 0.019$ | $\mathbf{0.698 \pm 0.013}$ | $0.706 \pm 0.007$ | $0.748 \pm 0.016$ | $0.706 \pm 0.013$ |
| | 12 | MSE | $1.533 \pm 0.051$ | $1.524 \pm 0.060$ | $1.462 \pm 0.075$ | $1.523 \pm 0.075$ | $1.514 \pm 0.088$ | $\mathbf{1.418 \pm 0.039}$ |
| | | MAE | $0.785 \pm 0.022$ | $0.783 \pm 0.024$ | $\mathbf{0.759 \pm 0.020}$ | $0.777 \pm 0.022$ | $0.784 \pm 0.029$ | $0.763 \pm 0.011$ |

Table 33: FNSPID framework results by prediction length $H$ for WMT, GILD with PatchTST. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | $0.063 \pm 0.001$ | $0.061 \pm 0.001$ | $\mathbf{0.058 \pm 0.002}$ | $0.062 \pm 0.001$ | $0.061 \pm 0.001$ | $0.061 \pm 0.002$ |
| | | MAE | $0.200 \pm 0.002$ | $0.197 \pm 0.001$ | $\mathbf{0.192 \pm 0.003}$ | $0.200 \pm 0.001$ | $0.198 \pm 0.001$ | $0.195 \pm 0.002$ |
| | 8 | MSE | $0.077 \pm 0.001$ | $\mathbf{0.074 \pm 0.001}$ | $\mathbf{0.074 \pm 0.002}$ | $0.077 \pm 0.001$ | $0.080 \pm 0.006$ | $0.078 \pm 0.001$ |
| | | MAE | $0.221 \pm 0.001$ | $0.218 \pm 0.001$ | $\mathbf{0.217 \pm 0.003}$ | $0.221 \pm 0.001$ | $0.225 \pm 0.008$ | $0.223 \pm 0.001$ |
| | 10 | MSE | $0.098 \pm 0.001$ | $0.094 \pm 0.001$ | $\mathbf{0.092 \pm 0.002}$ | $0.098 \pm 0.001$ | $0.101 \pm 0.004$ | $0.097 \pm 0.001$ |
| | | MAE | $0.250 \pm 0.002$ | $0.245 \pm 0.002$ | $\mathbf{0.241 \pm 0.003}$ | $0.250 \pm 0.002$ | $0.255 \pm 0.006$ | $0.248 \pm 0.002$ |
| | 12 | MSE | $0.113 \pm 0.001$ | $0.109 \pm 0.001$ | $\mathbf{0.108 \pm 0.003}$ | $0.114 \pm 0.002$ | $0.130 \pm 0.010$ | $0.114 \pm 0.003$ |
| | | MAE | $0.267 \pm 0.002$ | $0.263 \pm 0.002$ | $\mathbf{0.260 \pm 0.003}$ | $0.268 \pm 0.002$ | $0.285 \pm 0.009$ | $0.266 \pm 0.002$ |
| GILD | 6 | MSE | $0.047 \pm 0.000$ | $0.047 \pm 0.000$ | $0.047 \pm 0.001$ | $0.047 \pm 0.000$ | $0.046 \pm 0.000$ | $\mathbf{0.044 \pm 0.002}$ |
| | | MAE | $0.148 \pm 0.000$ | $0.149 \pm 0.000$ | $0.149 \pm 0.002$ | $0.148 \pm 0.000$ | $0.147 \pm 0.000$ | $\mathbf{0.143 \pm 0.003}$ |
| | 8 | MSE | $0.053 \pm 0.000$ | $0.053 \pm 0.000$ | $\mathbf{0.052 \pm 0.001}$ | $0.053 \pm 0.000$ | $0.053 \pm 0.002$ | $0.053 \pm 0.000$ |
| | | MAE | $0.162 \pm 0.001$ | $0.162 \pm 0.002$ | $\mathbf{0.159 \pm 0.002}$ | $0.162 \pm 0.002$ | $0.161 \pm 0.004$ | $0.160 \pm 0.002$ |
| | 10 | MSE | $0.061 \pm 0.000$ | $0.061 \pm 0.000$ | $\mathbf{0.059 \pm 0.001}$ | $0.061 \pm 0.000$ | $0.060 \pm 0.001$ | $0.061 \pm 0.001$ |
| | | MAE | $0.174 \pm 0.001$ | $0.175 \pm 0.001$ | $\mathbf{0.171 \pm 0.001}$ | $0.174 \pm 0.001$ | $0.174 \pm 0.002$ | $0.174 \pm 0.001$ |
| | 12 | MSE | $0.069 \pm 0.000$ | $0.070 \pm 0.001$ | $0.068 \pm 0.001$ | $0.069 \pm 0.001$ | $\mathbf{0.067 \pm 0.001}$ | $0.068 \pm 0.001$ |
| | | MAE | $0.187 \pm 0.001$ | $0.189 \pm 0.001$ | $\mathbf{0.185 \pm 0.002}$ | $0.187 \pm 0.001$ | $\mathbf{0.185 \pm 0.001}$ | $\mathbf{0.185 \pm 0.002}$ |

Table 34: FNSPID framework results by prediction length $H$ for WFC, SBUX with PatchTST. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | $0.316 \pm 0.007$ | $0.324 \pm 0.012$ | $0.333 \pm 0.034$ | $0.312 \pm 0.007$ | $0.292 \pm 0.013$ | $\mathbf{0.264 \pm 0.009}$ |
| | | MAE | $0.364 \pm 0.004$ | $0.368 \pm 0.006$ | $0.374 \pm 0.020$ | $0.361 \pm 0.005$ | $0.357 \pm 0.007$ | $\mathbf{0.338 \pm 0.006}$ |
| | 8 | MSE | $0.394 \pm 0.012$ | $0.406 \pm 0.017$ | $\mathbf{0.385 \pm 0.027}$ | $0.396 \pm 0.012$ | $0.398 \pm 0.018$ | $0.392 \pm 0.018$ |
| | | MAE | $0.400 \pm 0.006$ | $0.406 \pm 0.009$ | $\mathbf{0.398 \pm 0.015}$ | $0.401 \pm 0.007$ | $0.410 \pm 0.013$ | $\mathbf{0.398 \pm 0.010}$ |
| | 10 | MSE | $0.476 \pm 0.016$ | $0.487 \pm 0.012$ | $0.497 \pm 0.010$ | $0.481 \pm 0.020$ | $0.527 \pm 0.055$ | $\mathbf{0.446 \pm 0.013}$ |
| | | MAE | $0.435 \pm 0.007$ | $0.441 \pm 0.007$ | $0.449 \pm 0.006$ | $0.437 \pm 0.011$ | $0.469 \pm 0.023$ | $\mathbf{0.427 \pm 0.004}$ |
| | 12 | MSE | $0.569 \pm 0.009$ | $0.569 \pm 0.014$ | $0.593 \pm 0.024$ | $0.567 \pm 0.012$ | $0.627 \pm 0.078$ | $\mathbf{0.548 \pm 0.017}$ |
| | | MAE | $0.467 \pm 0.004$ | $0.468 \pm 0.007$ | $0.483 \pm 0.014$ | $0.466 \pm 0.005$ | $0.505 \pm 0.038$ | $\mathbf{0.465 \pm 0.007}$ |

Continued on next page

Table 34 (continued)

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| SBUX | 6 | MSE | $1.007 \pm 0.015$ | $1.007 \pm 0.027$ | $1.075 \pm 0.060$ | $1.041 \pm 0.024$ | $0.935 \pm 0.044$ | $\mathbf{0.837 \pm 0.009}$ |
| | | MAE | $0.620 \pm 0.010$ | $0.625 \pm 0.010$ | $0.628 \pm 0.013$ | $0.631 \pm 0.008$ | $0.597 \pm 0.008$ | $\mathbf{0.572 \pm 0.001}$ |
| | 8 | MSE | $1.207 \pm 0.016$ | $1.199 \pm 0.032$ | $1.226 \pm 0.085$ | $1.223 \pm 0.016$ | $1.205 \pm 0.083$ | $\mathbf{1.116 \pm 0.038}$ |
| | | MAE | $0.681 \pm 0.004$ | $0.681 \pm 0.006$ | $0.680 \pm 0.021$ | $0.686 \pm 0.005$ | $0.673 \pm 0.016$ | $\mathbf{0.659 \pm 0.003}$ |
| | 10 | MSE | $1.460 \pm 0.020$ | $1.445 \pm 0.016$ | $1.477 \pm 0.053$ | $1.464 \pm 0.033$ | $1.728 \pm 0.122$ | $\mathbf{1.431 \pm 0.047}$ |
| | | MAE | $0.763 \pm 0.004$ | $0.762 \pm 0.005$ | $0.755 \pm 0.014$ | $0.764 \pm 0.009$ | $0.791 \pm 0.025$ | $\mathbf{0.750 \pm 0.011}$ |
| | 12 | MSE | $1.625 \pm 0.014$ | $1.605 \pm 0.007$ | $1.672 \pm 0.075$ | $1.624 \pm 0.018$ | $1.803 \pm 0.173$ | $\mathbf{1.514 \pm 0.064}$ |
| | | MAE | $0.816 \pm 0.004$ | $0.814 \pm 0.001$ | $0.816 \pm 0.014$ | $0.817 \pm 0.006$ | $0.825 \pm 0.026$ | $\mathbf{0.776 \pm 0.013}$ |

Table 35: FNSPID framework results by prediction length $H$ for WMT, GILD with Crossformer. Entries are mean $\pm$ sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | $0.632 \pm 0.094$ | $0.716 \pm 0.052$ | $\mathbf{0.567 \pm 0.074}$ | $0.643 \pm 0.114$ | $0.742 \pm 0.154$ | $0.603 \pm 0.071$ |
| | | MAE | $0.640 \pm 0.057$ | $0.688 \pm 0.028$ | $\mathbf{0.618 \pm 0.043}$ | $0.641 \pm 0.060$ | $0.699 \pm 0.082$ | $0.630 \pm 0.044$ |
| | 8 | MSE | $0.938 \pm 0.104$ | $0.941 \pm 0.118$ | $\mathbf{0.740 \pm 0.094}$ | $0.966 \pm 0.099$ | $0.985 \pm 0.142$ | $0.795 \pm 0.068$ |
| | | MAE | $0.789 \pm 0.056$ | $0.779 \pm 0.056$ | $\mathbf{0.717 \pm 0.052}$ | $0.794 \pm 0.052$ | $0.813 \pm 0.064$ | $0.739 \pm 0.028$ |
| | 10 | MSE | $1.181 \pm 0.178$ | $1.211 \pm 0.183$ | $\mathbf{0.877 \pm 0.095}$ | $1.265 \pm 0.116$ | $1.236 \pm 0.021$ | $1.083 \pm 0.197$ |
| | | MAE | $0.890 \pm 0.063$ | $0.907 \pm 0.070$ | $\mathbf{0.785 \pm 0.048}$ | $0.922 \pm 0.045$ | $0.919 \pm 0.010$ | $0.871 \pm 0.081$ |
| | 12 | MSE | $1.486 \pm 0.197$ | $1.445 \pm 0.118$ | $\mathbf{1.061 \pm 0.100}$ | $1.487 \pm 0.276$ | $1.466 \pm 0.216$ | $1.402 \pm 0.294$ |
| | | MAE | $1.007 \pm 0.076$ | $1.000 \pm 0.054$ | $\mathbf{0.860 \pm 0.048}$ | $0.997 \pm 0.101$ | $0.996 \pm 0.074$ | $0.992 \pm 0.106$ |
| GILD | 6 | MSE | $0.055 \pm 0.013$ | $0.055 \pm 0.013$ | $0.055 \pm 0.007$ | $0.052 \pm 0.006$ | $0.058 \pm 0.012$ | $\mathbf{0.046 \pm 0.000}$ |
| | | MAE | $0.167 \pm 0.024$ | $0.167 \pm 0.025$ | $0.172 \pm 0.016$ | $0.164 \pm 0.012$ | $0.171 \pm 0.021$ | $\mathbf{0.149 \pm 0.000}$ |
| | 8 | MSE | $0.056 \pm 0.004$ | $\mathbf{0.055 \pm 0.002}$ | $0.067 \pm 0.007$ | $0.057 \pm 0.003$ | $0.063 \pm 0.005$ | $\mathbf{0.055 \pm 0.003}$ |
| | | MAE | $0.169 \pm 0.009$ | $0.168 \pm 0.006$ | $0.190 \pm 0.012$ | $0.171 \pm 0.007$ | $0.181 \pm 0.008$ | $\mathbf{0.164 \pm 0.003}$ |
| | 10 | MSE | $0.066 \pm 0.009$ | $\mathbf{0.062 \pm 0.003}$ | $0.066 \pm 0.003$ | $0.070 \pm 0.012$ | $0.080 \pm 0.015$ | $0.063 \pm 0.004$ |
| | | MAE | $0.185 \pm 0.016$ | $\mathbf{0.180 \pm 0.007}$ | $0.187 \pm 0.006$ | $0.193 \pm 0.020$ | $0.209 \pm 0.024$ | $0.183 \pm 0.006$ |
| | 12 | MSE | $0.074 \pm 0.013$ | $0.081 \pm 0.019$ | $0.076 \pm 0.005$ | $0.077 \pm 0.008$ | $0.081 \pm 0.003$ | $\mathbf{0.069 \pm 0.006}$ |
| | | MAE | $0.197 \pm 0.021$ | $0.207 \pm 0.027$ | $0.201 \pm 0.008$ | $0.203 \pm 0.013$ | $0.210 \pm 0.007$ | $\mathbf{0.189 \pm 0.011}$ |

Table 36: FNSPID framework results by prediction length $H$ for WFC, SBUX with Crossformer. Entries are mean $\pm$ sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | $12.353 \pm 0.248$ | $12.645 \pm 0.281$ | $\mathbf{11.503 \pm 0.636}$ | $12.605 \pm 0.552$ | $12.404 \pm 0.839$ | $12.161 \pm 0.494$ |
| | | MAE | $3.064 \pm 0.029$ | $3.099 \pm 0.037$ | $\mathbf{2.956 \pm 0.084}$ | $3.093 \pm 0.068$ | $3.067 \pm 0.103$ | $3.036 \pm 0.061$ |
| | 8 | MSE | $12.954 \pm 0.478$ | $12.993 \pm 0.267$ | $\mathbf{11.922 \pm 0.665}$ | $12.778 \pm 0.375$ | $12.467 \pm 0.537$ | $12.189 \pm 0.594$ |
| | | MAE | $3.148 \pm 0.059$ | $3.154 \pm 0.031$ | $\mathbf{3.022 \pm 0.087}$ | $3.131 \pm 0.043$ | $3.092 \pm 0.068$ | $3.061 \pm 0.087$ |
| | 10 | MSE | $13.692 \pm 0.448$ | $13.571 \pm 0.453$ | $12.558 \pm 0.621$ | $13.452 \pm 0.268$ | $13.102 \pm 0.541$ | $\mathbf{12.506 \pm 0.753}$ |
| | | MAE | $3.248 \pm 0.053$ | $3.238 \pm 0.059$ | $3.112 \pm 0.077$ | $3.220 \pm 0.031$ | $3.181 \pm 0.069$ | $\mathbf{3.099 \pm 0.091}$ |
| | 12 | MSE | $14.071 \pm 0.585$ | $14.322 \pm 0.497$ | $13.853 \pm 1.030$ | $14.681 \pm 0.399$ | $13.940 \pm 0.350$ | $\mathbf{13.061 \pm 0.703}$ |
| | | MAE | $3.306 \pm 0.066$ | $3.331 \pm 0.058$ | $3.287 \pm 0.137$ | $3.381 \pm 0.048$ | $3.294 \pm 0.039$ | $\mathbf{3.182 \pm 0.083}$ |
| SBUX | 6 | MSE | $31.104 \pm 1.383$ | $30.956 \pm 2.291$ | $\mathbf{27.544 \pm 2.296}$ | $30.968 \pm 0.141$ | $31.442 \pm 1.278$ | $29.002 \pm 1.996$ |
| | | MAE | $5.365 \pm 0.124$ | $5.350 \pm 0.210$ | $\mathbf{5.030 \pm 0.218}$ | $5.355 \pm 0.012$ | $5.391 \pm 0.115$ | $5.168 \pm 0.191$ |
| | 8 | MSE | $31.018 \pm 1.727$ | $31.698 \pm 1.374$ | $\mathbf{28.555 \pm 3.044}$ | $31.265 \pm 0.647$ | $31.544 \pm 1.990$ | $29.756 \pm 0.863$ |
| | | MAE | $5.366 \pm 0.156$ | $5.426 \pm 0.122$ | $\mathbf{5.122 \pm 0.290}$ | $5.389 \pm 0.058$ | $5.409 \pm 0.178$ | $5.249 \pm 0.085$ |
| | 10 | MSE | $31.856 \pm 0.798$ | $32.788 \pm 1.017$ | $31.084 \pm 4.012$ | $31.962 \pm 1.488$ | $33.047 \pm 2.558$ | $\mathbf{28.068 \pm 3.252}$ |
| | | MAE | $5.448 \pm 0.072$ | $5.530 \pm 0.088$ | $5.368 \pm 0.359$ | $5.457 \pm 0.130$ | $5.547 \pm 0.217$ | $\mathbf{5.097 \pm 0.310}$ |
| | 12 | MSE | $33.450 \pm 1.390$ | $33.752 \pm 1.097$ | $33.269 \pm 3.124$ | $33.592 \pm 1.623$ | $34.569 \pm 2.424$ | $\mathbf{32.462 \pm 3.551}$ |
| | | MAE | $5.595 \pm 0.122$ | $5.621 \pm 0.096$ | $5.570 \pm 0.272$ | $5.606 \pm 0.141$ | $5.684 \pm 0.203$ | $\mathbf{5.500 \pm 0.307}$ |

Table 37: FNSPID framework results by prediction length $H$ for WMT, GILD with SOFTS. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | $0.065 \pm 0.002$ | $\mathbf{0.063 \pm 0.001}$ | $\mathbf{0.063 \pm 0.006}$ | $0.065 \pm 0.002$ | $0.064 \pm 0.003$ | $0.065 \pm 0.002$ |
| | | MAE | $0.203 \pm 0.003$ | $0.201 \pm 0.002$ | $\mathbf{0.200 \pm 0.008}$ | $0.203 \pm 0.003$ | $0.203 \pm 0.005$ | $0.203 \pm 0.003$ |
| | 8 | MSE | $0.085 \pm 0.002$ | $0.082 \pm 0.002$ | $\mathbf{0.081 \pm 0.004}$ | $0.083 \pm 0.003$ | $0.086 \pm 0.004$ | $0.085 \pm 0.002$ |
| | | MAE | $0.232 \pm 0.002$ | $0.229 \pm 0.003$ | $\mathbf{0.226 \pm 0.005}$ | $0.230 \pm 0.004$ | $0.234 \pm 0.005$ | $0.232 \pm 0.003$ |
| | 10 | MSE | $0.102 \pm 0.001$ | $\mathbf{0.099 \pm 0.001}$ | $\mathbf{0.099 \pm 0.004}$ | $0.102 \pm 0.002$ | $0.106 \pm 0.009$ | $0.103 \pm 0.004$ |
| | | MAE | $0.252 \pm 0.002$ | $\mathbf{0.249 \pm 0.002}$ | $\mathbf{0.249 \pm 0.006}$ | $0.253 \pm 0.002$ | $0.257 \pm 0.009$ | $0.254 \pm 0.004$ |
| | 12 | MSE | $0.119 \pm 0.002$ | $\mathbf{0.116 \pm 0.002}$ | $\mathbf{0.116 \pm 0.004}$ | $0.120 \pm 0.002$ | $0.140 \pm 0.019$ | $0.118 \pm 0.001$ |
| | | MAE | $0.271 \pm 0.001$ | $0.268 \pm 0.002$ | $\mathbf{0.267 \pm 0.004}$ | $0.272 \pm 0.001$ | $0.290 \pm 0.019$ | $0.269 \pm 0.002$ |
| GILD | 6 | MSE | $\mathbf{0.046 \pm 0.001}$ | $\mathbf{0.046 \pm 0.001}$ | $0.047 \pm 0.002$ | $\mathbf{0.046 \pm 0.001}$ | $0.050 \pm 0.001$ | $0.049 \pm 0.001$ |
| | | MAE | $\mathbf{0.148 \pm 0.002}$ | $\mathbf{0.148 \pm 0.002}$ | $0.150 \pm 0.004$ | $\mathbf{0.148 \pm 0.002}$ | $0.155 \pm 0.002$ | $0.151 \pm 0.002$ |
| | 8 | MSE | $0.053 \pm 0.001$ | $0.053 \pm 0.001$ | $\mathbf{0.052 \pm 0.001}$ | $\mathbf{0.052 \pm 0.001}$ | $0.055 \pm 0.001$ | $\mathbf{0.052 \pm 0.001}$ |
| | | MAE | $0.160 \pm 0.003$ | $0.160 \pm 0.002$ | $\mathbf{0.158 \pm 0.001}$ | $0.159 \pm 0.002$ | $0.164 \pm 0.001$ | $0.161 \pm 0.003$ |
| | 10 | MSE | $0.060 \pm 0.001$ | $0.060 \pm 0.002$ | $\mathbf{0.059 \pm 0.002}$ | $0.060 \pm 0.002$ | $0.064 \pm 0.004$ | $0.061 \pm 0.002$ |
| | | MAE | $0.172 \pm 0.002$ | $0.173 \pm 0.003$ | $\mathbf{0.171 \pm 0.003}$ | $0.173 \pm 0.003$ | $0.180 \pm 0.008$ | $0.173 \pm 0.002$ |
| | 12 | MSE | $0.069 \pm 0.001$ | $0.069 \pm 0.001$ | $\mathbf{0.067 \pm 0.001}$ | $0.068 \pm 0.001$ | $0.073 \pm 0.007$ | $\mathbf{0.067 \pm 0.001}$ |
| | | MAE | $0.186 \pm 0.002$ | $0.186 \pm 0.002$ | $\mathbf{0.183 \pm 0.003}$ | $0.185 \pm 0.002$ | $0.195 \pm 0.011$ | $0.184 \pm 0.001$ |

Table 38: FNSPID framework results by prediction length $H$ for WFC, SBUX with SOFTS. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | $0.253 \pm 0.002$ | $0.261 \pm 0.008$ | $0.251 \pm 0.011$ | $0.259 \pm 0.005$ | $0.270 \pm 0.021$ | $\mathbf{0.246 \pm 0.009}$ |
| | | MAE | $0.334 \pm 0.002$ | $0.339 \pm 0.006$ | $\mathbf{0.333 \pm 0.007}$ | $0.337 \pm 0.004$ | $0.345 \pm 0.010$ | $\mathbf{0.333 \pm 0.005}$ |
| | 8 | MSE | $0.336 \pm 0.002$ | $0.343 \pm 0.006$ | $0.326 \pm 0.014$ | $0.337 \pm 0.003$ | $0.352 \pm 0.016$ | $\mathbf{0.324 \pm 0.004}$ |
| | | MAE | $0.378 \pm 0.003$ | $0.382 \pm 0.004$ | $\mathbf{0.375 \pm 0.005}$ | $0.379 \pm 0.002$ | $0.393 \pm 0.009$ | $0.377 \pm 0.002$ |
| | 10 | MSE | $0.426 \pm 0.010$ | $0.433 \pm 0.010$ | $0.412 \pm 0.013$ | $0.428 \pm 0.007$ | $0.499 \pm 0.045$ | $\mathbf{0.411 \pm 0.006}$ |
| | | MAE | $0.417 \pm 0.005$ | $0.421 \pm 0.006$ | $\mathbf{0.412 \pm 0.006}$ | $0.418 \pm 0.004$ | $0.460 \pm 0.027$ | $0.414 \pm 0.004$ |
| | 12 | MSE | $0.518 \pm 0.005$ | $0.522 \pm 0.004$ | $0.508 \pm 0.018$ | $0.517 \pm 0.003$ | $0.627 \pm 0.120$ | $\mathbf{0.506 \pm 0.010}$ |
| | | MAE | $0.453 \pm 0.001$ | $0.455 \pm 0.002$ | $0.450 \pm 0.008$ | $0.453 \pm 0.002$ | $0.511 \pm 0.059$ | $\mathbf{0.449 \pm 0.004}$ |
| SBUX | 6 | MSE | $0.794 \pm 0.011$ | $0.772 \pm 0.028$ | $0.743 \pm 0.035$ | $0.769 \pm 0.026$ | $0.894 \pm 0.099$ | $\mathbf{0.725 \pm 0.037}$ |
| | | MAE | $0.566 \pm 0.003$ | $0.559 \pm 0.006$ | $\mathbf{0.555 \pm 0.008}$ | $0.559 \pm 0.007$ | $0.591 \pm 0.025$ | $0.556 \pm 0.009$ |
| | 8 | MSE | $1.014 \pm 0.014$ | $1.041 \pm 0.058$ | $1.005 \pm 0.026$ | $1.012 \pm 0.016$ | $1.091 \pm 0.171$ | $\mathbf{1.000 \pm 0.021}$ |
| | | MAE | $0.634 \pm 0.003$ | $0.643 \pm 0.020$ | $\mathbf{0.632 \pm 0.005}$ | $0.634 \pm 0.004$ | $0.643 \pm 0.024$ | $\mathbf{0.632 \pm 0.010}$ |
| | 10 | MSE | $1.279 \pm 0.018$ | $1.262 \pm 0.035$ | $\mathbf{1.220 \pm 0.015}$ | $1.279 \pm 0.022$ | $1.575 \pm 0.222$ | $1.221 \pm 0.042$ |
| | | MAE | $0.710 \pm 0.005$ | $0.708 \pm 0.008$ | $\mathbf{0.695 \pm 0.006}$ | $0.710 \pm 0.006$ | $0.742 \pm 0.031$ | $0.702 \pm 0.009$ |
| | 12 | MSE | $1.474 \pm 0.015$ | $1.461 \pm 0.013$ | $\mathbf{1.440 \pm 0.022}$ | $1.472 \pm 0.017$ | $2.023 \pm 0.234$ | $1.441 \pm 0.036$ |
| | | MAE | $0.770 \pm 0.005$ | $0.769 \pm 0.003$ | $\mathbf{0.757 \pm 0.006}$ | $0.770 \pm 0.005$ | $0.834 \pm 0.041$ | $0.767 \pm 0.009$ |

Table 39: FNSPID framework results by prediction length $H$ for WMT, GILD with TimeXer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | $0.061 \pm 0.001$ | $0.059 \pm 0.001$ | $\mathbf{0.058 \pm 0.001}$ | $0.061 \pm 0.001$ | $0.061 \pm 0.001$ | $0.060 \pm 0.001$ |
| | | MAE | $0.196 \pm 0.001$ | $0.194 \pm 0.001$ | $\mathbf{0.191 \pm 0.001}$ | $0.197 \pm 0.002$ | $0.196 \pm 0.002$ | $0.194 \pm 0.001$ |
| | 8 | MSE | $0.079 \pm 0.001$ | $0.076 \pm 0.001$ | $\mathbf{0.074 \pm 0.001}$ | $0.079 \pm 0.001$ | $0.080 \pm 0.003$ | $0.077 \pm 0.001$ |
| | | MAE | $0.224 \pm 0.001$ | $0.221 \pm 0.002$ | $\mathbf{0.217 \pm 0.002}$ | $0.224 \pm 0.001$ | $0.225 \pm 0.004$ | $0.222 \pm 0.002$ |
| | 10 | MSE | $0.096 \pm 0.001$ | $0.094 \pm 0.001$ | $\mathbf{0.089 \pm 0.002}$ | $0.096 \pm 0.001$ | $0.101 \pm 0.006$ | $0.094 \pm 0.003$ |
| | | MAE | $0.247 \pm 0.002$ | $0.244 \pm 0.001$ | $\mathbf{0.238 \pm 0.002}$ | $0.247 \pm 0.001$ | $0.253 \pm 0.008$ | $0.245 \pm 0.005$ |
| | 12 | MSE | $0.113 \pm 0.001$ | $0.108 \pm 0.001$ | $\mathbf{0.107 \pm 0.003}$ | $0.113 \pm 0.001$ | $0.123 \pm 0.012$ | $0.109 \pm 0.004$ |
| | | MAE | $0.267 \pm 0.001$ | $0.262 \pm 0.002$ | $\mathbf{0.260 \pm 0.004}$ | $0.267 \pm 0.002$ | $0.278 \pm 0.013$ | $0.263 \pm 0.006$ |

Continued on next page

Table 39 (continued)

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| GILD | 6 | MSE | **0.046 ± 0.001** | **0.046 ± 0.001** | **0.046 ± 0.001** | **0.046 ± 0.001** | **0.046 ± 0.002** | **0.046 ± 0.000** |
| | | MAE | 0.148 ± 0.001 | 0.149 ± 0.001 | 0.148 ± 0.002 | 0.148 ± 0.001 | 0.148 ± 0.002 | **0.147 ± 0.001** |
| | 8 | MSE | 0.053 ± 0.001 | 0.053 ± 0.000 | 0.053 ± 0.001 | 0.053 ± 0.001 | **0.052 ± 0.001** | 0.053 ± 0.001 |
| | | MAE | 0.161 ± 0.001 | 0.162 ± 0.002 | 0.160 ± 0.003 | 0.161 ± 0.001 | **0.159 ± 0.001** | 0.161 ± 0.001 |
| | 10 | MSE | 0.061 ± 0.001 | 0.062 ± 0.001 | **0.059 ± 0.000** | 0.061 ± 0.001 | 0.060 ± 0.000 | 0.060 ± 0.001 |
| | | MAE | 0.175 ± 0.002 | 0.176 ± 0.001 | **0.171 ± 0.001** | 0.175 ± 0.001 | 0.174 ± 0.001 | 0.173 ± 0.002 |
| | 12 | MSE | 0.069 ± 0.001 | 0.071 ± 0.001 | **0.067 ± 0.001** | 0.069 ± 0.000 | 0.068 ± 0.001 | **0.067 ± 0.001** |
| | | MAE | 0.186 ± 0.001 | 0.190 ± 0.003 | **0.183 ± 0.001** | 0.188 ± 0.001 | 0.187 ± 0.002 | 0.184 ± 0.001 |

Table 40: FNSPID framework results by prediction length $H$ for WFC, SBUX with TimeXer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | 0.304 ± 0.007 | 0.315 ± 0.014 | 0.308 ± 0.027 | 0.306 ± 0.006 | **0.301 ± 0.019** | **0.301 ± 0.012** |
| | | MAE | 0.357 ± 0.003 | 0.362 ± 0.007 | 0.359 ± 0.015 | 0.358 ± 0.002 | 0.361 ± 0.010 | **0.355 ± 0.007** |
| | 8 | MSE | 0.398 ± 0.013 | 0.408 ± 0.014 | 0.390 ± 0.026 | 0.398 ± 0.014 | 0.437 ± 0.054 | **0.351 ± 0.004** |
| | | MAE | 0.402 ± 0.007 | 0.406 ± 0.007 | 0.400 ± 0.015 | 0.401 ± 0.006 | 0.432 ± 0.030 | **0.384 ± 0.001** |
| | 10 | MSE | 0.483 ± 0.018 | 0.501 ± 0.016 | 0.482 ± 0.015 | 0.490 ± 0.020 | 0.541 ± 0.077 | **0.447 ± 0.012** |
| | | MAE | 0.435 ± 0.008 | 0.443 ± 0.007 | 0.439 ± 0.008 | 0.438 ± 0.009 | 0.472 ± 0.039 | **0.420 ± 0.006** |
| | 12 | MSE | 0.573 ± 0.003 | 0.587 ± 0.009 | 0.577 ± 0.028 | 0.573 ± 0.001 | 0.635 ± 0.069 | **0.560 ± 0.019** |
| | | MAE | 0.468 ± 0.002 | 0.474 ± 0.004 | 0.476 ± 0.014 | 0.468 ± 0.002 | 0.514 ± 0.034 | **0.466 ± 0.011** |
| SBUX | 6 | MSE | 1.007 ± 0.028 | 0.950 ± 0.034 | 1.034 ± 0.086 | 0.984 ± 0.041 | 0.913 ± 0.054 | **0.806 ± 0.032** |
| | | MAE | 0.616 ± 0.010 | 0.601 ± 0.005 | 0.618 ± 0.019 | 0.608 ± 0.011 | 0.588 ± 0.015 | **0.564 ± 0.006** |
| | 8 | MSE | 1.218 ± 0.059 | 1.205 ± 0.052 | 1.218 ± 0.049 | 1.208 ± 0.022 | 1.295 ± 0.109 | **1.077 ± 0.076** |
| | | MAE | 0.682 ± 0.017 | 0.682 ± 0.015 | 0.675 ± 0.010 | 0.681 ± 0.007 | 0.687 ± 0.020 | **0.661 ± 0.017** |
| | 10 | MSE | 1.444 ± 0.011 | 1.418 ± 0.007 | 1.465 ± 0.071 | 1.457 ± 0.029 | 1.618 ± 0.159 | **1.289 ± 0.030** |
| | | MAE | 0.750 ± 0.005 | 0.746 ± 0.003 | 0.747 ± 0.022 | 0.752 ± 0.005 | 0.761 ± 0.022 | **0.708 ± 0.003** |
| | 12 | MSE | 1.628 ± 0.019 | 1.592 ± 0.006 | 1.694 ± 0.064 | 1.612 ± 0.011 | 1.759 ± 0.097 | **1.574 ± 0.027** |
| | | MAE | 0.810 ± 0.004 | 0.806 ± 0.004 | 0.816 ± 0.016 | 0.809 ± 0.004 | 0.810 ± 0.018 | **0.796 ± 0.007** |

Table 41: FNSPID framework results by prediction length $H$ for WMT, GILD with Transformer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | 0.482 ± 0.180 | 0.620 ± 0.101 | 0.527 ± 0.275 | **0.428 ± 0.079** | 0.808 ± 0.247 | 0.548 ± 0.173 |
| | | MAE | 0.535 ± 0.128 | 0.635 ± 0.060 | 0.556 ± 0.171 | **0.528 ± 0.054** | 0.728 ± 0.134 | 0.567 ± 0.107 |
| | 8 | MSE | 0.743 ± 0.221 | 0.692 ± 0.226 | 0.863 ± 0.404 | **0.647 ± 0.098** | 1.046 ± 0.183 | 0.823 ± 0.256 |
| | | MAE | 0.695 ± 0.135 | 0.665 ± 0.143 | 0.736 ± 0.242 | **0.651 ± 0.054** | 0.853 ± 0.074 | 0.716 ± 0.139 |
| | 10 | MSE | 0.952 ± 0.324 | 0.921 ± 0.258 | **0.648 ± 0.269** | 0.992 ± 0.190 | 1.230 ± 0.322 | 0.806 ± 0.342 |
| | | MAE | 0.791 ± 0.160 | 0.779 ± 0.126 | **0.637 ± 0.166** | 0.819 ± 0.093 | 0.922 ± 0.149 | 0.706 ± 0.147 |
| | 12 | MSE | 1.039 ± 0.265 | **0.920 ± 0.298** | 0.989 ± 0.520 | 1.192 ± 0.267 | 1.379 ± 0.327 | 0.940 ± 0.283 |
| | | MAE | 0.823 ± 0.124 | **0.757 ± 0.153** | 0.783 ± 0.289 | 0.900 ± 0.106 | 0.960 ± 0.137 | 0.773 ± 0.136 |
| GILD | 6 | MSE | **0.047 ± 0.003** | 0.049 ± 0.004 | 0.055 ± 0.019 | 0.055 ± 0.009 | 0.058 ± 0.018 | 0.049 ± 0.001 |
| | | MAE | **0.156 ± 0.006** | 0.158 ± 0.009 | 0.166 ± 0.029 | 0.171 ± 0.018 | 0.173 ± 0.031 | 0.159 ± 0.003 |
| | 8 | MSE | 0.055 ± 0.005 | 0.055 ± 0.004 | 0.060 ± 0.011 | 0.066 ± 0.012 | 0.055 ± 0.000 | **0.054 ± 0.003** |
| | | MAE | 0.172 ± 0.010 | 0.169 ± 0.008 | 0.178 ± 0.019 | 0.189 ± 0.021 | 0.170 ± 0.001 | **0.164 ± 0.005** |
| | 10 | MSE | 0.064 ± 0.005 | **0.062 ± 0.004** | 0.068 ± 0.010 | 0.070 ± 0.009 | 0.071 ± 0.018 | 0.063 ± 0.003 |
| | | MAE | 0.186 ± 0.009 | **0.182 ± 0.008** | 0.190 ± 0.016 | 0.196 ± 0.016 | 0.193 ± 0.021 | 0.186 ± 0.007 |
| | 12 | MSE | **0.072 ± 0.006** | 0.085 ± 0.018 | 0.099 ± 0.037 | 0.083 ± 0.008 | **0.072 ± 0.003** | **0.072 ± 0.002** |
| | | MAE | 0.198 ± 0.011 | 0.216 ± 0.026 | 0.229 ± 0.038 | 0.217 ± 0.011 | **0.195 ± 0.002** | 0.198 ± 0.005 |

Table 42: FNSPID framework results by prediction length $H$ for WFC, SBUX with Transformer. Entries are mean ± sample standard deviation over five seeds. Bold and underlining mark the lowest and second-lowest displayed means.

| Task | $H$ | Metric | Uni-modal | MM-TSFlib | Cross-MoE | TaTS | CFA | DynaGuard |
|---|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | $12.001 \pm 1.251$ | $11.438 \pm 0.944$ | $\underline{11.210 \pm 1.746}$ | $\mathbf{6.720 \pm 0.934}$ | $12.969 \pm 1.150$ | $11.436 \pm 1.963$ |
| | | MAE | $3.022 \pm 0.170$ | $2.954 \pm 0.127$ | $\underline{2.895 \pm 0.236}$ | $\mathbf{2.248 \pm 0.164}$ | $3.147 \pm 0.137$ | $2.937 \pm 0.262$ |
| | 8 | MSE | $12.922 \pm 1.585$ | $12.877 \pm 2.025$ | $13.747 \pm 0.581$ | $\mathbf{7.070 \pm 0.845}$ | $14.082 \pm 1.301$ | $\underline{11.809 \pm 2.601}$ |
| | | MAE | $3.164 \pm 0.194$ | $3.146 \pm 0.241$ | $3.243 \pm 0.068$ | $\mathbf{2.335 \pm 0.140}$ | $3.300 \pm 0.157$ | $\underline{2.998 \pm 0.336}$ |
| | 10 | MSE | $14.939 \pm 1.297$ | $13.420 \pm 1.001$ | $12.901 \pm 1.627$ | $\mathbf{7.824 \pm 0.775}$ | $15.319 \pm 1.108$ | $\underline{12.017 \pm 2.712}$ |
| | | MAE | $3.409 \pm 0.140$ | $3.231 \pm 0.118$ | $3.153 \pm 0.198$ | $\mathbf{2.471 \pm 0.124}$ | $3.467 \pm 0.126$ | $\underline{3.032 \pm 0.342}$ |
| | 12 | MSE | $14.678 \pm 1.021$ | $14.714 \pm 1.025$ | $15.573 \pm 2.571$ | $\mathbf{8.624 \pm 0.946}$ | $16.198 \pm 1.018$ | $\underline{13.542 \pm 2.459}$ |
| | | MAE | $3.392 \pm 0.123$ | $3.397 \pm 0.123$ | $3.469 \pm 0.288$ | $\mathbf{2.606 \pm 0.138}$ | $3.569 \pm 0.115$ | $\underline{3.239 \pm 0.299}$ |
| SBUX | 6 | MSE | $31.571 \pm 1.543$ | $31.841 \pm 1.661$ | $26.808 \pm 2.120$ | $\mathbf{19.774 \pm 1.005}$ | $32.738 \pm 2.447$ | $\underline{23.006 \pm 3.789}$ |
| | | MAE | $5.394 \pm 0.132$ | $5.419 \pm 0.156$ | $4.936 \pm 0.212$ | $\mathbf{4.188 \pm 0.109}$ | $5.497 \pm 0.226$ | $\underline{4.517 \pm 0.383}$ |
| | 8 | MSE | $31.077 \pm 1.740$ | $31.497 \pm 1.674$ | $\underline{26.480 \pm 3.625}$ | $\mathbf{20.388 \pm 1.519}$ | $33.760 \pm 1.775$ | $27.433 \pm 3.588$ |
| | | MAE | $5.342 \pm 0.149$ | $5.378 \pm 0.145$ | $\underline{4.914 \pm 0.372}$ | $\mathbf{4.275 \pm 0.170}$ | $5.595 \pm 0.151$ | $4.993 \pm 0.347$ |
| | 10 | MSE | $30.707 \pm 1.128$ | $31.038 \pm 1.642$ | $\underline{27.068 \pm 5.807}$ | $\mathbf{20.785 \pm 1.175}$ | $33.463 \pm 2.139$ | $32.332 \pm 2.497$ |
| | | MAE | $5.315 \pm 0.111$ | $5.339 \pm 0.143$ | $\underline{4.955 \pm 0.587}$ | $\mathbf{4.326 \pm 0.135}$ | $5.562 \pm 0.195$ | $5.449 \pm 0.221$ |
| | 12 | MSE | $35.273 \pm 2.172$ | $35.047 \pm 2.114$ | $25.589 \pm 3.881$ | $\mathbf{21.030 \pm 1.249}$ | $37.510 \pm 1.614$ | $\underline{23.804 \pm 3.802}$ |
| | | MAE | $5.711 \pm 0.186$ | $5.691 \pm 0.182$ | $4.826 \pm 0.419$ | $\mathbf{4.358 \pm 0.141}$ | $5.926 \pm 0.132$ | $\underline{4.616 \pm 0.396}$ |

### E.3 Comparison with TSOnly

TSOnly replaces observed text and TCC with a shared learnable memory sequence while retaining the remaining adapter. It tests whether observed context adds information beyond learned memory. Uni-modal is a separate baseline without the adapter. Table 44 reports absolute errors over four horizons and five seeds; Table 43 gives each backbone's percentage reductions and their mean.

For domain $d$, backbone $b$, and metric $q \in \{\mathrm{MSE}, \mathrm{MAE}\}$, let $\bar{E}^{m}_{d,b,q}$ be method $m$'s mean error over these 20 runs. The reductions are

$$r_{d,b,q} = 100\left(1 - \frac{\bar{E}^{\mathrm{DynaGuard}}_{d,b,q}}{\bar{E}^{\mathrm{TSOnly}}_{d,b,q}}\right), \qquad \bar{r}_{d,q} = \frac{r_{d,\mathrm{iTransformer},q} + r_{d,\mathrm{PatchTST},q}}{2}. \tag{22}$$

The main-paper bars show $\bar{r}_{d,q}$, with positive values indicating lower error than TSOnly. All calculations use unrounded errors. The abbreviations Agr., Clim., Econ., Ener., Env., Hlth., Sec., Soc., and Traf. denote Agriculture, Climate, Economy, Energy, Environment, Health, Security, Social Good, and Traffic, respectively.

Table 43: Aggregate error reduction (%) over TSOnly. Each backbone uses four horizons and five seeds. Mean averages the two backbone percentages before rounding; positive values favor observed text.

| Dataset | iTransformer | | PatchTST | | Mean | |
|---|---|---|---|---|---|---|
| | MSE | MAE | MSE | MAE | MSE | MAE |
| Agriculture | +15.7 | +10.7 | +22.1 | +14.2 | +18.9 | +12.4 |
| Climate | +6.2 | +4.9 | +1.1 | +1.0 | +3.7 | +2.9 |
| Economy | +2.7 | +1.3 | +8.7 | +5.7 | +5.7 | +3.5 |
| Energy | +20.5 | +16.1 | +19.1 | +14.8 | +19.8 | +15.4 |
| Environment | +4.5 | −0.1 | +0.0 | +0.0 | +2.2 | −0.1 |
| Health | +1.6 | +0.8 | +1.8 | +0.6 | +1.7 | +0.7 |
| Security | +6.0 | +11.2 | +0.5 | +0.6 | +3.3 | +5.9 |
| Social Good | +3.3 | −2.2 | −1.0 | +1.1 | +1.2 | −0.6 |
| Traffic | −0.5 | −2.2 | +8.5 | −2.8 | +4.0 | −2.5 |

Table 44: Complete TSOnly comparison with iTransformer and PatchTST. Entries are the 20-run mean ± the mean of four per-horizon sample standard deviations over five seeds.

| Dataset | Metric | iTransformer | | PatchTST | |
|---|---|---|---|---|---|
| | | TSOnly | DynaGuard | TSOnly | DynaGuard |
| Agriculture | MSE | $0.302 \pm 0.007$ | $0.254 \pm 0.014$ | $0.376 \pm 0.008$ | $0.293 \pm 0.012$ |
| | MAE | $0.372 \pm 0.005$ | $0.332 \pm 0.013$ | $0.408 \pm 0.005$ | $0.350 \pm 0.012$ |
| Climate | MSE | $0.953 \pm 0.005$ | $0.894 \pm 0.004$ | $0.905 \pm 0.008$ | $0.896 \pm 0.004$ |
| | MAE | $0.790 \pm 0.004$ | $0.751 \pm 0.002$ | $0.761 \pm 0.005$ | $0.753 \pm 0.002$ |
| Economy | MSE | $0.184 \pm 0.022$ | $0.180 \pm 0.005$ | $0.218 \pm 0.005$ | $0.199 \pm 0.010$ |
| | MAE | $0.335 \pm 0.022$ | $0.330 \pm 0.006$ | $0.372 \pm 0.005$ | $0.351 \pm 0.012$ |
| Energy | MSE | $0.303 \pm 0.008$ | $0.241 \pm 0.012$ | $0.311 \pm 0.002$ | $0.252 \pm 0.004$ |
| | MAE | $0.423 \pm 0.006$ | $0.355 \pm 0.008$ | $0.426 \pm 0.001$ | $0.363 \pm 0.004$ |
| Environment | MSE | $0.271 \pm 0.001$ | $0.259 \pm 0.000$ | $0.272 \pm 0.000$ | $0.272 \pm 0.000$ |
| | MAE | $0.369 \pm 0.001$ | $0.370 \pm 0.001$ | $0.371 \pm 0.001$ | $0.371 \pm 0.001$ |
| Health | MSE | $1.294 \pm 0.010$ | $1.273 \pm 0.010$ | $1.328 \pm 0.004$ | $1.304 \pm 0.023$ |
| | MAE | $0.759 \pm 0.004$ | $0.753 \pm 0.005$ | $0.775 \pm 0.002$ | $0.770 \pm 0.006$ |
| Security | MSE | $76.312 \pm 1.298$ | $71.734 \pm 0.644$ | $71.715 \pm 0.290$ | $71.322 \pm 0.786$ |
| | MAE | $4.463 \pm 0.100$ | $3.963 \pm 0.097$ | $3.951 \pm 0.036$ | $3.925 \pm 0.120$ |
| Social Good | MSE | $0.960 \pm 0.015$ | $0.928 \pm 0.016$ | $0.855 \pm 0.010$ | $0.863 \pm 0.013$ |
| | MAE | $0.419 \pm 0.008$ | $0.428 \pm 0.007$ | $0.435 \pm 0.009$ | $0.431 \pm 0.011$ |
| Traffic | MSE | $0.161 \pm 0.003$ | $0.162 \pm 0.004$ | $0.170 \pm 0.002$ | $0.155 \pm 0.003$ |
| | MAE | $0.214 \pm 0.004$ | $0.218 \pm 0.005$ | $0.224 \pm 0.002$ | $0.230 \pm 0.007$ |

### E.3.1 Results by Forecast Horizon

The following tables report DynaGuard and TSOnly on all nine Time-MMD domains for iTransformer and PatchTST. Each horizon contains five seeds.

Table 45: Per-horizon TSOnly comparison for Agriculture, Climate, Economy. Entries are mean ± sample standard deviation over five seeds.

| Dataset | $H$ | Metric | iTransformer | | PatchTST | |
|---|---|---|---|---|---|---|
| | | | TSOnly | DynaGuard | TSOnly | DynaGuard |
| Agriculture | 6 | MSE | $0.191 \pm 0.004$ | $0.146 \pm 0.005$ | $0.263 \pm 0.010$ | $0.191 \pm 0.015$ |
| | | MAE | $0.300 \pm 0.001$ | $0.255 \pm 0.007$ | $0.337 \pm 0.005$ | $0.298 \pm 0.015$ |
| | 8 | MSE | $0.260 \pm 0.005$ | $0.210 \pm 0.018$ | $0.327 \pm 0.003$ | $0.251 \pm 0.008$ |
| | | MAE | $0.347 \pm 0.005$ | $0.304 \pm 0.012$ | $0.374 \pm 0.002$ | $0.321 \pm 0.005$ |
| | 10 | MSE | $0.333 \pm 0.003$ | $0.288 \pm 0.014$ | $0.427 \pm 0.010$ | $0.329 \pm 0.010$ |
| | | MAE | $0.391 \pm 0.003$ | $0.362 \pm 0.017$ | $0.444 \pm 0.008$ | $0.370 \pm 0.010$ |
| | 12 | MSE | $0.422 \pm 0.017$ | $0.372 \pm 0.019$ | $0.487 \pm 0.008$ | $0.400 \pm 0.015$ |
| | | MAE | $0.447 \pm 0.011$ | $0.406 \pm 0.016$ | $0.476 \pm 0.005$ | $0.413 \pm 0.018$ |
| Climate | 6 | MSE | $0.946 \pm 0.005$ | $0.884 \pm 0.006$ | $0.900 \pm 0.014$ | $0.886 \pm 0.005$ |
| | | MAE | $0.787 \pm 0.005$ | $0.750 \pm 0.002$ | $0.761 \pm 0.006$ | $0.752 \pm 0.002$ |
| | 8 | MSE | $0.949 \pm 0.007$ | $0.888 \pm 0.001$ | $0.897 \pm 0.005$ | $0.891 \pm 0.004$ |
| | | MAE | $0.787 \pm 0.004$ | $0.749 \pm 0.000$ | $0.757 \pm 0.003$ | $0.750 \pm 0.001$ |
| | 10 | MSE | $0.959 \pm 0.006$ | $0.899 \pm 0.007$ | $0.910 \pm 0.006$ | $0.902 \pm 0.006$ |
| | | MAE | $0.793 \pm 0.004$ | $0.752 \pm 0.004$ | $0.761 \pm 0.003$ | $0.755 \pm 0.003$ |
| | 12 | MSE | $0.959 \pm 0.004$ | $0.905 \pm 0.003$ | $0.914 \pm 0.009$ | $0.904 \pm 0.001$ |
| | | MAE | $0.792 \pm 0.002$ | $0.754 \pm 0.001$ | $0.764 \pm 0.007$ | $0.756 \pm 0.003$ |
| Economy | 6 | MSE | $0.150 \pm 0.004$ | $0.150 \pm 0.007$ | $0.209 \pm 0.006$ | $0.175 \pm 0.015$ |
| | | MAE | $0.300 \pm 0.004$ | $0.303 \pm 0.008$ | $0.363 \pm 0.006$ | $0.325 \pm 0.016$ |
| | 8 | MSE | $0.183 \pm 0.032$ | $0.169 \pm 0.006$ | $0.217 \pm 0.006$ | $0.200 \pm 0.009$ |
| | | MAE | $0.334 \pm 0.032$ | $0.322 \pm 0.006$ | $0.371 \pm 0.005$ | $0.353 \pm 0.010$ |
| | 10 | MSE | $0.196 \pm 0.026$ | $0.190 \pm 0.003$ | $0.224 \pm 0.004$ | $0.206 \pm 0.007$ |
| | | MAE | $0.347 \pm 0.027$ | $0.342 \pm 0.003$ | $0.378 \pm 0.003$ | $0.358 \pm 0.009$ |
| | 12 | MSE | $0.210 \pm 0.026$ | $0.209 \pm 0.005$ | $0.222 \pm 0.006$ | $0.216 \pm 0.009$ |
| | | MAE | $0.360 \pm 0.026$ | $0.356 \pm 0.006$ | $0.377 \pm 0.006$ | $0.368 \pm 0.014$ |

Table 46: Per-horizon TSOnly comparison for Energy, Environment, Health. Entries are mean ± sample standard deviation over five seeds.

| Dataset | $H$ | Metric | iTransformer | | PatchTST | |
|---|---|---|---|---|---|---|
| | | | TSOnly | DynaGuard | TSOnly | DynaGuard |
| Energy | 12 | MSE | $0.187 \pm 0.013$ | $0.092 \pm 0.004$ | $0.195 \pm 0.002$ | $0.106 \pm 0.002$ |
| | | MAE | $0.341 \pm 0.011$ | $0.212 \pm 0.004$ | $0.343 \pm 0.002$ | $0.233 \pm 0.003$ |
| | 24 | MSE | $0.266 \pm 0.008$ | $0.197 \pm 0.016$ | $0.269 \pm 0.003$ | $0.214 \pm 0.003$ |
| | | MAE | $0.401 \pm 0.006$ | $0.329 \pm 0.010$ | $0.400 \pm 0.002$ | $0.340 \pm 0.002$ |
| | 36 | MSE | $0.343 \pm 0.008$ | $0.284 \pm 0.018$ | $0.348 \pm 0.002$ | $0.302 \pm 0.007$ |
| | | MAE | $0.449 \pm 0.006$ | $0.397 \pm 0.014$ | $0.451 \pm 0.001$ | $0.407 \pm 0.007$ |
| | 48 | MSE | $0.418 \pm 0.001$ | $0.392 \pm 0.008$ | $0.431 \pm 0.001$ | $0.385 \pm 0.002$ |
| | | MAE | $0.502 \pm 0.001$ | $0.481 \pm 0.005$ | $0.510 \pm 0.001$ | $0.473 \pm 0.004$ |
| Environment | 48 | MSE | $0.269 \pm 0.000$ | $0.262 \pm 0.000$ | $0.269 \pm 0.000$ | $0.269 \pm 0.000$ |
| | | MAE | $0.372 \pm 0.001$ | $0.369 \pm 0.001$ | $0.373 \pm 0.001$ | $0.374 \pm 0.000$ |
| | 96 | MSE | $0.277 \pm 0.004$ | $0.265 \pm 0.000$ | $0.279 \pm 0.000$ | $0.279 \pm 0.000$ |
| | | MAE | $0.374 \pm 0.002$ | $0.371 \pm 0.001$ | $0.376 \pm 0.001$ | $0.376 \pm 0.000$ |
| | 192 | MSE | $0.271 \pm 0.001$ | $0.256 \pm 0.001$ | $0.272 \pm 0.000$ | $0.272 \pm 0.000$ |
| | | MAE | $0.367 \pm 0.001$ | $0.369 \pm 0.002$ | $0.367 \pm 0.000$ | $0.367 \pm 0.001$ |
| | 336 | MSE | $0.269 \pm 0.000$ | $0.254 \pm 0.000$ | $0.268 \pm 0.000$ | $0.269 \pm 0.000$ |
| | | MAE | $0.365 \pm 0.001$ | $0.370 \pm 0.001$ | $0.367 \pm 0.000$ | $0.366 \pm 0.001$ |
| Health | 12 | MSE | $0.972 \pm 0.007$ | $0.929 \pm 0.016$ | $1.036 \pm 0.005$ | $1.008 \pm 0.007$ |
| | | MAE | $0.672 \pm 0.003$ | $0.648 \pm 0.006$ | $0.684 \pm 0.003$ | $0.676 \pm 0.005$ |
| | 24 | MSE | $1.315 \pm 0.010$ | $1.307 \pm 0.012$ | $1.381 \pm 0.006$ | $1.332 \pm 0.029$ |
| | | MAE | $0.763 \pm 0.004$ | $0.760 \pm 0.003$ | $0.793 \pm 0.003$ | $0.776 \pm 0.010$ |
| | 36 | MSE | $1.433 \pm 0.014$ | $1.404 \pm 0.004$ | $1.429 \pm 0.004$ | $1.431 \pm 0.051$ |
| | | MAE | $0.795 \pm 0.005$ | $0.790 \pm 0.003$ | $0.803 \pm 0.002$ | $0.812 \pm 0.007$ |
| | 48 | MSE | $1.455 \pm 0.010$ | $1.451 \pm 0.009$ | $1.465 \pm 0.003$ | $1.446 \pm 0.005$ |
| | | MAE | $0.806 \pm 0.003$ | $0.812 \pm 0.009$ | $0.818 \pm 0.001$ | $0.815 \pm 0.003$ |

Table 47: Per-horizon TSOnly comparison for Security, Social Good, Traffic. Entries are mean ± sample standard deviation over five seeds.

| Dataset | $H$ | Metric | iTransformer | | PatchTST | |
|---|---|---|---|---|---|---|
| | | | TSOnly | DynaGuard | TSOnly | DynaGuard |
| Security | 6 | MSE | $71.943 \pm 0.534$ | $66.815 \pm 0.430$ | $66.373 \pm 0.286$ | $66.074 \pm 0.837$ |
| | | MAE | $4.414 \pm 0.042$ | $3.847 \pm 0.103$ | $3.700 \pm 0.045$ | $3.724 \pm 0.086$ |
| | 8 | MSE | $74.656 \pm 1.632$ | $69.877 \pm 0.567$ | $69.739 \pm 0.260$ | $69.914 \pm 0.658$ |
| | | MAE | $4.427 \pm 0.130$ | $3.864 \pm 0.107$ | $3.891 \pm 0.039$ | $3.954 \pm 0.151$ |
| | 10 | MSE | $77.882 \pm 1.642$ | $73.225 \pm 0.708$ | $73.902 \pm 0.461$ | $73.227 \pm 1.152$ |
| | | MAE | $4.474 \pm 0.133$ | $3.993 \pm 0.088$ | $4.056 \pm 0.036$ | $3.994 \pm 0.192$ |
| | 12 | MSE | $80.767 \pm 1.386$ | $77.021 \pm 0.872$ | $76.847 \pm 0.151$ | $76.075 \pm 0.498$ |
| | | MAE | $4.538 \pm 0.096$ | $4.147 \pm 0.091$ | $4.156 \pm 0.025$ | $4.028 \pm 0.051$ |
| Social Good | 6 | MSE | $0.880 \pm 0.022$ | $0.864 \pm 0.025$ | $0.752 \pm 0.008$ | $0.789 \pm 0.018$ |
| | | MAE | $0.373 \pm 0.006$ | $0.394 \pm 0.010$ | $0.382 \pm 0.005$ | $0.381 \pm 0.010$ |
| | 8 | MSE | $0.945 \pm 0.016$ | $0.900 \pm 0.014$ | $0.817 \pm 0.007$ | $0.842 \pm 0.010$ |
| | | MAE | $0.411 \pm 0.013$ | $0.409 \pm 0.008$ | $0.410 \pm 0.006$ | $0.418 \pm 0.009$ |
| | 10 | MSE | $0.986 \pm 0.012$ | $0.954 \pm 0.019$ | $0.900 \pm 0.008$ | $0.890 \pm 0.011$ |
| | | MAE | $0.431 \pm 0.006$ | $0.441 \pm 0.007$ | $0.461 \pm 0.006$ | $0.453 \pm 0.016$ |
| | 12 | MSE | $1.028 \pm 0.009$ | $0.993 \pm 0.008$ | $0.950 \pm 0.016$ | $0.933 \pm 0.014$ |
| | | MAE | $0.460 \pm 0.008$ | $0.468 \pm 0.004$ | $0.489 \pm 0.021$ | $0.471 \pm 0.008$ |
| Traffic | 6 | MSE | $0.149 \pm 0.003$ | $0.152 \pm 0.003$ | $0.166 \pm 0.003$ | $0.152 \pm 0.004$ |
| | | MAE | $0.198 \pm 0.005$ | $0.202 \pm 0.007$ | $0.219 \pm 0.002$ | $0.231 \pm 0.009$ |
| | 8 | MSE | $0.159 \pm 0.006$ | $0.159 \pm 0.004$ | $0.166 \pm 0.002$ | $0.153 \pm 0.002$ |
| | | MAE | $0.211 \pm 0.007$ | $0.215 \pm 0.005$ | $0.219 \pm 0.001$ | $0.221 \pm 0.001$ |
| | 10 | MSE | $0.165 \pm 0.003$ | $0.165 \pm 0.004$ | $0.173 \pm 0.003$ | $0.157 \pm 0.003$ |
| | | MAE | $0.219 \pm 0.003$ | $0.224 \pm 0.004$ | $0.227 \pm 0.001$ | $0.231 \pm 0.005$ |
| | 12 | MSE | $0.171 \pm 0.002$ | $0.172 \pm 0.005$ | $0.174 \pm 0.002$ | $0.160 \pm 0.004$ |
| | | MAE | $0.226 \pm 0.002$ | $0.232 \pm 0.006$ | $0.229 \pm 0.003$ | $0.235 \pm 0.014$ |

# F Component and Temporal-Bias Studies

## F.1 Component Ablations

We compare four variants under the protocol in Section D.1. *w/o TCC* replaces learned attention pooling with masked pooling. *w/o CATB* removes the temporal bias. *w/o CAGI* removes the gate

while retaining cross-attention. *TS query* replaces TCC's shared learnable query with a query derived from numerical history, extending the three component ablations in the main paper.

For each domain, Table 2 averages the percentage MSE increases of iTransformer and PatchTST before rounding; Table 48 retains the separate backbone results.

Every ablation raises aggregate MSE in at least eight domains and MAE in at least seven on each backbone (Table 48). These results support combining the three controls and using a shared compression query. Figure 15 shows the iTransformer domain-level changes; Section F.1.1 gives the per-horizon results for both backbones.

Table 48: Aggregate component ablations with iTransformer and PatchTST. Each dataset lists all variants as rows; entries are the 20-run mean ± the mean of four horizon-specific sample standard deviations over five seeds. Lower is better; bold and underlined entries are the best and second-best results within each backbone–dataset–metric group.

| Dataset | Variant | iTransformer | | PatchTST | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| Agriculture | Full DynaGuard | **0.254 ± 0.014** | **0.332 ± 0.013** | **0.293 ± 0.012** | **0.350 ± 0.012** |
| | w/o TCC | 0.264 ± 0.013 | 0.344 ± 0.012 | 0.341 ± 0.014 | 0.379 ± 0.013 |
| | w/o CATB | 0.268 ± 0.012 | 0.346 ± 0.015 | 0.332 ± 0.016 | 0.374 ± 0.014 |
| | w/o CAGI | 0.340 ± 0.058 | 0.394 ± 0.032 | 0.379 ± 0.025 | 0.412 ± 0.012 |
| | TS query | 0.273 ± 0.014 | 0.347 ± 0.011 | 0.337 ± 0.024 | 0.386 ± 0.013 |
| Climate | Full DynaGuard | **0.894 ± 0.004** | **0.751 ± 0.002** | **0.896 ± 0.004** | **0.753 ± 0.002** |
| | w/o TCC | 0.972 ± 0.008 | 0.800 ± 0.005 | 0.906 ± 0.010 | 0.761 ± 0.005 |
| | w/o CATB | 0.972 ± 0.007 | 0.800 ± 0.004 | 0.913 ± 0.009 | 0.767 ± 0.006 |
| | w/o CAGI | 0.929 ± 0.034 | 0.775 ± 0.020 | 0.917 ± 0.014 | 0.770 ± 0.008 |
| | TS query | 0.972 ± 0.008 | 0.800 ± 0.004 | 0.925 ± 0.014 | 0.773 ± 0.006 |
| Economy | Full DynaGuard | **0.180 ± 0.005** | **0.330 ± 0.006** | 0.199 ± 0.010 | 0.351 ± 0.012 |
| | w/o TCC | 0.186 ± 0.005 | 0.335 ± 0.004 | 0.201 ± 0.008 | 0.354 ± 0.009 |
| | w/o CATB | 0.186 ± 0.006 | 0.336 ± 0.006 | 0.208 ± 0.012 | 0.361 ± 0.013 |
| | w/o CAGI | 0.188 ± 0.009 | 0.339 ± 0.012 | 0.216 ± 0.011 | 0.369 ± 0.011 |
| | TS query | 0.186 ± 0.007 | 0.335 ± 0.006 | **0.194 ± 0.012** | **0.343 ± 0.015** |
| Energy | Full DynaGuard | **0.241 ± 0.012** | **0.355 ± 0.008** | **0.252 ± 0.004** | **0.363 ± 0.004** |
| | w/o TCC | 0.324 ± 0.027 | 0.440 ± 0.017 | 0.313 ± 0.012 | 0.431 ± 0.010 |
| | w/o CATB | 0.316 ± 0.009 | 0.434 ± 0.007 | 0.311 ± 0.008 | 0.430 ± 0.006 |
| | w/o CAGI | 0.325 ± 0.018 | 0.443 ± 0.014 | 0.314 ± 0.017 | 0.433 ± 0.014 |
| | TS query | 0.311 ± 0.015 | 0.433 ± 0.011 | 0.321 ± 0.019 | 0.439 ± 0.014 |
| Environment | Full DynaGuard | **0.259 ± 0.000** | **0.370 ± 0.001** | **0.272 ± 0.000** | **0.371 ± 0.001** |
| | w/o TCC | 0.267 ± 0.004 | 0.378 ± 0.004 | 0.290 ± 0.005 | 0.398 ± 0.006 |
| | w/o CATB | 0.265 ± 0.005 | **0.370 ± 0.002** | 0.286 ± 0.005 | 0.392 ± 0.006 |
| | w/o CAGI | 0.268 ± 0.006 | 0.373 ± 0.006 | 0.290 ± 0.003 | 0.396 ± 0.005 |
| | TS query | 0.264 ± 0.004 | **0.370 ± 0.002** | 0.274 ± 0.002 | 0.376 ± 0.003 |
| Health | Full DynaGuard | **1.273 ± 0.010** | **0.753 ± 0.005** | **1.304 ± 0.023** | **0.770 ± 0.006** |
| | w/o TCC | 1.362 ± 0.049 | 0.766 ± 0.011 | 1.347 ± 0.046 | 0.794 ± 0.026 |
| | w/o CATB | 1.347 ± 0.016 | 0.761 ± 0.004 | 1.360 ± 0.042 | 0.795 ± 0.022 |
| | w/o CAGI | 1.370 ± 0.048 | 0.771 ± 0.012 | 1.369 ± 0.072 | 0.803 ± 0.029 |
| | TS query | 1.349 ± 0.018 | 0.762 ± 0.005 | 1.336 ± 0.014 | 0.781 ± 0.007 |
| Security | Full DynaGuard | 71.734 ± 0.644 | 3.963 ± 0.097 | **71.322 ± 0.786** | 3.925 ± 0.120 |
| | w/o TCC | 74.774 ± 2.111 | 4.294 ± 0.215 | 71.501 ± 0.403 | 3.942 ± 0.041 |
| | w/o CATB | 75.949 ± 1.358 | 4.420 ± 0.113 | 71.560 ± 0.342 | 3.938 ± 0.036 |
| | w/o CAGI | **71.548 ± 0.873** | **3.921 ± 0.135** | 71.399 ± 0.613 | **3.914 ± 0.049** |
| | TS query | 75.843 ± 1.302 | 4.415 ± 0.115 | 72.597 ± 1.335 | 4.057 ± 0.157 |
| Social Good | Full DynaGuard | **0.928 ± 0.016** | **0.428 ± 0.007** | **0.863 ± 0.013** | **0.431 ± 0.011** |
| | w/o TCC | 1.184 ± 0.064 | 0.490 ± 0.019 | 0.956 ± 0.017 | 0.455 ± 0.008 |
| | w/o CATB | 1.163 ± 0.054 | 0.484 ± 0.020 | 0.986 ± 0.047 | 0.479 ± 0.029 |
| | w/o CAGI | 1.123 ± 0.037 | 0.470 ± 0.012 | 0.914 ± 0.058 | 0.459 ± 0.048 |
| | TS query | 1.162 ± 0.074 | 0.483 ± 0.024 | 1.045 ± 0.059 | 0.497 ± 0.025 |
| Traffic | Full DynaGuard | **0.162 ± 0.004** | 0.218 ± 0.005 | **0.155 ± 0.003** | 0.230 ± 0.007 |
| | w/o TCC | 0.171 ± 0.003 | 0.214 ± 0.003 | 0.167 ± 0.002 | 0.250 ± 0.005 |
| | w/o CATB | 0.165 ± 0.004 | **0.211 ± 0.005** | 0.158 ± 0.004 | 0.229 ± 0.006 |
| | w/o CAGI | 0.167 ± 0.004 | 0.213 ± 0.005 | 0.159 ± 0.004 | 0.241 ± 0.009 |
| | TS query | 0.167 ± 0.003 | 0.212 ± 0.004 | 0.158 ± 0.004 | **0.223 ± 0.006** |

### F.1.1 Results by Forecast Horizon

These tables disaggregate Table 48 by prediction length for iTransformer and PatchTST.

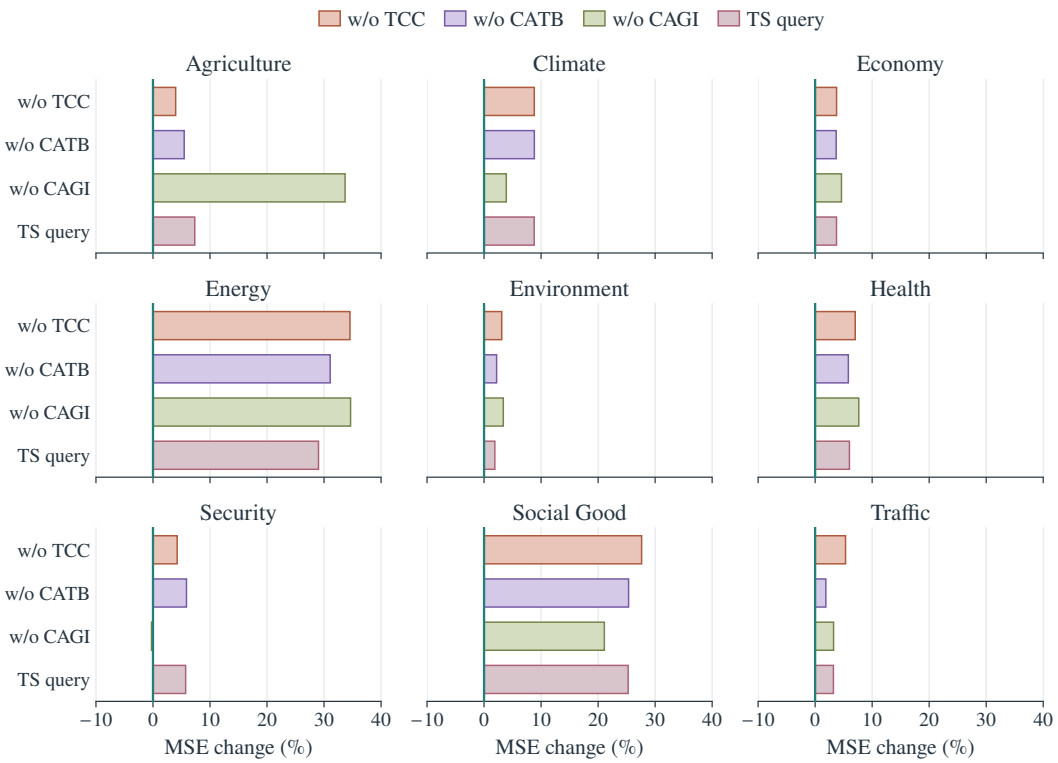


Figure 15: MSE change relative to full DynaGuard with iTransformer, computed from unrounded means over four horizons and five seeds. Positive values indicate higher error. TS query derives TCC's query from numerical history.

Table 49: Per-horizon component-ablation results with iTransformer for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | **0.146 ± 0.005** | 0.158 ± 0.005 | 0.160 ± 0.007 | 0.241 ± 0.056 | 0.162 ± 0.007 |
| | | MAE | **0.255 ± 0.007** | 0.267 ± 0.007 | 0.268 ± 0.007 | 0.325 ± 0.039 | 0.270 ± 0.006 |
| | 8 | MSE | **0.210 ± 0.018** | 0.226 ± 0.015 | 0.232 ± 0.012 | 0.289 ± 0.043 | 0.226 ± 0.019 |
| | | MAE | **0.304 ± 0.012** | 0.315 ± 0.012 | 0.321 ± 0.015 | 0.367 ± 0.031 | 0.319 ± 0.009 |
| | 10 | MSE | **0.288 ± 0.014** | 0.315 ± 0.008 | 0.297 ± 0.013 | 0.365 ± 0.054 | 0.318 ± 0.012 |
| | | MAE | **0.362 ± 0.017** | 0.385 ± 0.010 | 0.374 ± 0.013 | 0.414 ± 0.027 | 0.375 ± 0.009 |
| | 12 | MSE | 0.372 ± 0.019 | **0.358 ± 0.025** | 0.383 ± 0.017 | 0.463 ± 0.078 | 0.384 ± 0.020 |
| | | MAE | **0.406 ± 0.016** | 0.411 ± 0.016 | 0.423 ± 0.026 | 0.469 ± 0.030 | 0.424 ± 0.020 |
| Climate | 6 | MSE | **0.884 ± 0.006** | 0.964 ± 0.005 | 0.965 ± 0.004 | 0.894 ± 0.032 | 0.966 ± 0.005 |
| | | MAE | **0.750 ± 0.002** | 0.797 ± 0.005 | 0.797 ± 0.004 | 0.758 ± 0.018 | 0.798 ± 0.003 |
| | 8 | MSE | **0.888 ± 0.001** | 0.969 ± 0.004 | 0.969 ± 0.006 | 0.928 ± 0.045 | 0.967 ± 0.007 |
| | | MAE | **0.749 ± 0.000** | 0.799 ± 0.002 | 0.799 ± 0.003 | 0.773 ± 0.027 | 0.797 ± 0.004 |
| | 10 | MSE | **0.899 ± 0.007** | 0.978 ± 0.014 | 0.978 ± 0.013 | 0.953 ± 0.038 | 0.978 ± 0.012 |
| | | MAE | **0.752 ± 0.004** | 0.803 ± 0.009 | 0.802 ± 0.006 | 0.787 ± 0.024 | 0.802 ± 0.006 |
| | 12 | MSE | **0.905 ± 0.003** | 0.978 ± 0.007 | 0.978 ± 0.007 | 0.938 ± 0.020 | 0.978 ± 0.009 |
| | | MAE | **0.754 ± 0.001** | 0.802 ± 0.005 | 0.801 ± 0.004 | 0.780 ± 0.013 | 0.802 ± 0.005 |

Continued on next page

Table 49 (continued)

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Economy | 6 | MSE | **0.150 ± 0.007** | 0.154 ± 0.004 | 0.155 ± 0.004 | 0.152 ± 0.005 | 0.155 ± 0.003 |
| | | MAE | **0.303 ± 0.008** | 0.305 ± 0.004 | 0.308 ± 0.003 | 0.304 ± 0.006 | 0.307 ± 0.003 |
| | 8 | MSE | **0.169 ± 0.006** | 0.177 ± 0.003 | 0.179 ± 0.010 | 0.179 ± 0.012 | 0.175 ± 0.004 |
| | | MAE | **0.322 ± 0.006** | 0.327 ± 0.003 | 0.330 ± 0.009 | 0.329 ± 0.013 | 0.325 ± 0.003 |
| | 10 | MSE | **0.190 ± 0.003** | 0.198 ± 0.007 | 0.196 ± 0.005 | 0.199 ± 0.009 | 0.201 ± 0.015 |
| | | MAE | **0.342 ± 0.003** | 0.346 ± 0.006 | 0.344 ± 0.005 | 0.348 ± 0.013 | 0.347 ± 0.011 |
| | 12 | MSE | **0.209 ± 0.005** | 0.217 ± 0.005 | 0.215 ± 0.006 | 0.222 ± 0.011 | 0.214 ± 0.007 |
| | | MAE | **0.356 ± 0.006** | 0.363 ± 0.005 | 0.362 ± 0.005 | 0.375 ± 0.015 | 0.360 ± 0.008 |

Table 50: Per-horizon component-ablation results with iTransformer for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | **0.092 ± 0.004** | 0.182 ± 0.020 | 0.173 ± 0.009 | 0.179 ± 0.016 | 0.177 ± 0.015 |
| | | MAE | **0.212 ± 0.004** | 0.331 ± 0.015 | 0.323 ± 0.005 | 0.330 ± 0.016 | 0.329 ± 0.012 |
| | 24 | MSE | **0.197 ± 0.016** | 0.274 ± 0.018 | 0.270 ± 0.011 | 0.282 ± 0.009 | 0.265 ± 0.011 |
| | | MAE | **0.329 ± 0.010** | 0.411 ± 0.012 | 0.405 ± 0.009 | 0.417 ± 0.007 | 0.403 ± 0.009 |
| | 36 | MSE | **0.284 ± 0.018** | 0.384 ± 0.037 | 0.375 ± 0.010 | 0.382 ± 0.020 | 0.354 ± 0.017 |
| | | MAE | **0.397 ± 0.014** | 0.484 ± 0.026 | 0.479 ± 0.008 | 0.483 ± 0.014 | 0.465 ± 0.010 |
| | 48 | MSE | **0.392 ± 0.008** | 0.458 ± 0.034 | 0.447 ± 0.006 | 0.456 ± 0.026 | 0.449 ± 0.015 |
| | | MAE | **0.481 ± 0.005** | 0.536 ± 0.017 | 0.530 ± 0.004 | 0.542 ± 0.019 | 0.534 ± 0.011 |
| Environment | 48 | MSE | **0.262 ± 0.000** | 0.268 ± 0.002 | 0.263 ± 0.002 | 0.266 ± 0.004 | **0.262 ± 0.000** |
| | | MAE | **0.369 ± 0.001** | 0.377 ± 0.002 | **0.369 ± 0.002** | 0.371 ± 0.003 | 0.370 ± 0.000 |
| | 96 | MSE | **0.265 ± 0.000** | 0.270 ± 0.001 | 0.268 ± 0.006 | 0.275 ± 0.005 | 0.268 ± 0.007 |
| | | MAE | **0.371 ± 0.001** | 0.378 ± 0.002 | 0.373 ± 0.001 | 0.379 ± 0.005 | 0.372 ± 0.001 |
| | 192 | MSE | **0.256 ± 0.001** | 0.259 ± 0.002 | 0.263 ± 0.007 | 0.264 ± 0.009 | 0.263 ± 0.007 |
| | | MAE | 0.369 ± 0.002 | 0.372 ± 0.002 | 0.370 ± 0.004 | 0.372 ± 0.007 | **0.368 ± 0.003** |
| | 336 | MSE | **0.254 ± 0.000** | 0.271 ± 0.010 | 0.265 ± 0.003 | 0.267 ± 0.006 | 0.264 ± 0.003 |
| | | MAE | 0.370 ± 0.001 | 0.383 ± 0.009 | **0.369 ± 0.004** | 0.372 ± 0.007 | **0.369 ± 0.003** |
| Health | 12 | MSE | **0.929 ± 0.016** | 0.982 ± 0.041 | 0.977 ± 0.013 | 0.982 ± 0.028 | 0.977 ± 0.018 |
| | | MAE | **0.648 ± 0.006** | 0.655 ± 0.007 | 0.653 ± 0.004 | 0.657 ± 0.005 | 0.654 ± 0.006 |
| | 24 | MSE | **1.307 ± 0.012** | 1.381 ± 0.032 | 1.349 ± 0.014 | 1.337 ± 0.013 | 1.354 ± 0.020 |
| | | MAE | **0.760 ± 0.003** | 0.773 ± 0.010 | 0.761 ± 0.002 | 0.764 ± 0.005 | 0.763 ± 0.005 |
| | 36 | MSE | **1.404 ± 0.004** | 1.563 ± 0.075 | 1.523 ± 0.019 | 1.572 ± 0.083 | 1.515 ± 0.021 |
| | | MAE | **0.790 ± 0.003** | 0.816 ± 0.019 | 0.809 ± 0.006 | 0.823 ± 0.019 | 0.806 ± 0.005 |
| | 48 | MSE | **1.451 ± 0.009** | 1.523 ± 0.049 | 1.539 ± 0.020 | 1.590 ± 0.067 | 1.552 ± 0.015 |
| | | MAE | **0.812 ± 0.009** | 0.821 ± 0.010 | 0.821 ± 0.004 | 0.842 ± 0.017 | 0.825 ± 0.003 |

Table 51: Per-horizon component-ablation results with iTransformer for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 66.815 ± 0.430 | 70.904 ± 2.608 | 72.353 ± 1.051 | **66.141 ± 0.373** | 72.236 ± 0.862 |
| | | MAE | 3.847 ± 0.103 | 4.258 ± 0.291 | 4.427 ± 0.093 | **3.639 ± 0.075** | 4.421 ± 0.086 |
| | 8 | MSE | 69.877 ± 0.567 | 74.320 ± 2.939 | 75.013 ± 2.050 | **69.365 ± 0.214** | 74.865 ± 1.917 |
| | | MAE | 3.864 ± 0.107 | 4.350 ± 0.325 | 4.440 ± 0.177 | **3.770 ± 0.030** | 4.441 ± 0.188 |
| | 10 | MSE | **73.225 ± 0.708** | 77.001 ± 1.711 | 77.730 ± 1.941 | 74.029 ± 1.862 | 77.616 ± 2.043 |
| | | MAE | **3.993 ± 0.088** | 4.408 ± 0.114 | 4.463 ± 0.142 | 4.116 ± 0.287 | 4.448 ± 0.150 |
| | 12 | MSE | 77.021 ± 0.872 | 76.870 ± 1.189 | 78.701 ± 0.391 | **76.656 ± 1.042** | 78.657 ± 0.385 |
| | | MAE | **4.147 ± 0.091** | 4.161 ± 0.132 | 4.351 ± 0.041 | 4.158 ± 0.149 | 4.349 ± 0.036 |
| Social Good | 6 | MSE | **0.864 ± 0.025** | 1.044 ± 0.037 | 1.012 ± 0.024 | 1.025 ± 0.032 | 1.009 ± 0.043 |
| | | MAE | **0.394 ± 0.010** | 0.425 ± 0.014 | 0.418 ± 0.014 | 0.421 ± 0.008 | 0.417 ± 0.018 |
| | 8 | MSE | **0.900 ± 0.014** | 1.174 ± 0.048 | 1.166 ± 0.055 | 1.128 ± 0.027 | 1.137 ± 0.054 |
| | | MAE | **0.409 ± 0.008** | 0.483 ± 0.013 | 0.481 ± 0.019 | 0.465 ± 0.010 | 0.465 ± 0.017 |
| | 10 | MSE | **0.954 ± 0.019** | 1.251 ± 0.070 | 1.202 ± 0.054 | 1.140 ± 0.055 | 1.212 ± 0.060 |
| | | MAE | **0.441 ± 0.007** | 0.515 ± 0.015 | 0.503 ± 0.023 | 0.482 ± 0.015 | 0.507 ± 0.013 |
| | 12 | MSE | **0.993 ± 0.008** | 1.268 ± 0.099 | 1.272 ± 0.084 | 1.201 ± 0.035 | 1.291 ± 0.138 |
| | | MAE | **0.468 ± 0.004** | 0.537 ± 0.035 | 0.535 ± 0.025 | 0.512 ± 0.015 | 0.544 ± 0.047 |

Continued on next page

Table 51 (continued)

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Traffic | 6 | MSE | **$0.152 \pm 0.003$** | $0.158 \pm 0.002$ | $0.155 \pm 0.002$ | $0.157 \pm 0.005$ | $0.156 \pm 0.002$ |
| | | MAE | $0.202 \pm 0.007$ | **$0.193 \pm 0.003$** | **$0.193 \pm 0.003$** | $0.196 \pm 0.006$ | $0.194 \pm 0.003$ |
| | 8 | MSE | **$0.159 \pm 0.004$** | $0.167 \pm 0.003$ | $0.163 \pm 0.003$ | $0.166 \pm 0.002$ | $0.163 \pm 0.003$ |
| | | MAE | $0.215 \pm 0.005$ | $0.208 \pm 0.003$ | $0.207 \pm 0.005$ | $0.212 \pm 0.003$ | **$0.205 \pm 0.003$** |
| | 10 | MSE | **$0.165 \pm 0.004$** | $0.174 \pm 0.003$ | $0.167 \pm 0.006$ | $0.168 \pm 0.004$ | $0.171 \pm 0.005$ |
| | | MAE | $0.224 \pm 0.004$ | $0.221 \pm 0.003$ | **$0.217 \pm 0.006$** | **$0.217 \pm 0.005$** | $0.220 \pm 0.005$ |
| | 12 | MSE | **$0.172 \pm 0.005$** | $0.183 \pm 0.006$ | $0.175 \pm 0.004$ | $0.177 \pm 0.005$ | $0.179 \pm 0.003$ |
| | | MAE | $0.232 \pm 0.006$ | $0.233 \pm 0.004$ | **$0.228 \pm 0.004$** | **$0.228 \pm 0.006$** | **$0.228 \pm 0.004$** |

Table 52: Per-horizon component-ablation results with PatchTST for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | **$0.191 \pm 0.015$** | $0.250 \pm 0.023$ | $0.244 \pm 0.022$ | $0.276 \pm 0.022$ | $0.242 \pm 0.015$ |
| | | MAE | **$0.298 \pm 0.015$** | $0.328 \pm 0.022$ | $0.326 \pm 0.022$ | $0.350 \pm 0.010$ | $0.326 \pm 0.009$ |
| | 8 | MSE | **$0.251 \pm 0.008$** | $0.290 \pm 0.008$ | $0.285 \pm 0.008$ | $0.351 \pm 0.028$ | $0.295 \pm 0.023$ |
| | | MAE | **$0.321 \pm 0.005$** | $0.346 \pm 0.006$ | $0.344 \pm 0.006$ | $0.385 \pm 0.009$ | $0.356 \pm 0.015$ |
| | 10 | MSE | **$0.329 \pm 0.010$** | $0.376 \pm 0.013$ | $0.363 \pm 0.019$ | $0.415 \pm 0.034$ | $0.373 \pm 0.027$ |
| | | MAE | **$0.370 \pm 0.010$** | $0.398 \pm 0.014$ | $0.389 \pm 0.011$ | $0.441 \pm 0.019$ | $0.409 \pm 0.014$ |
| | 12 | MSE | **$0.400 \pm 0.015$** | $0.450 \pm 0.012$ | $0.437 \pm 0.016$ | $0.475 \pm 0.017$ | $0.439 \pm 0.032$ |
| | | MAE | **$0.413 \pm 0.018$** | $0.443 \pm 0.012$ | $0.437 \pm 0.016$ | $0.473 \pm 0.011$ | $0.451 \pm 0.014$ |
| Climate | 6 | MSE | **$0.886 \pm 0.005$** | $0.898 \pm 0.012$ | $0.908 \pm 0.005$ | $0.909 \pm 0.017$ | $0.918 \pm 0.016$ |
| | | MAE | **$0.752 \pm 0.002$** | $0.761 \pm 0.005$ | $0.769 \pm 0.006$ | $0.771 \pm 0.007$ | $0.772 \pm 0.006$ |
| | 8 | MSE | **$0.891 \pm 0.004$** | $0.899 \pm 0.007$ | $0.900 \pm 0.011$ | $0.910 \pm 0.020$ | $0.919 \pm 0.019$ |
| | | MAE | **$0.750 \pm 0.001$** | $0.758 \pm 0.005$ | $0.760 \pm 0.005$ | $0.765 \pm 0.009$ | $0.772 \pm 0.008$ |
| | 10 | MSE | **$0.902 \pm 0.006$** | $0.911 \pm 0.008$ | $0.912 \pm 0.010$ | $0.922 \pm 0.013$ | $0.925 \pm 0.007$ |
| | | MAE | **$0.755 \pm 0.003$** | $0.762 \pm 0.004$ | $0.763 \pm 0.006$ | $0.770 \pm 0.008$ | $0.772 \pm 0.003$ |
| | 12 | MSE | **$0.904 \pm 0.001$** | $0.915 \pm 0.011$ | $0.931 \pm 0.011$ | $0.928 \pm 0.008$ | $0.938 \pm 0.016$ |
| | | MAE | **$0.756 \pm 0.003$** | $0.765 \pm 0.007$ | $0.774 \pm 0.009$ | $0.775 \pm 0.006$ | $0.776 \pm 0.008$ |
| Economy | 6 | MSE | **$0.175 \pm 0.015$** | $0.178 \pm 0.007$ | $0.185 \pm 0.014$ | $0.191 \pm 0.015$ | $0.176 \pm 0.010$ |
| | | MAE | $0.325 \pm 0.016$ | $0.328 \pm 0.009$ | $0.336 \pm 0.015$ | $0.344 \pm 0.017$ | **$0.324 \pm 0.012$** |
| | 8 | MSE | $0.200 \pm 0.009$ | $0.193 \pm 0.013$ | $0.196 \pm 0.018$ | $0.212 \pm 0.007$ | **$0.186 \pm 0.017$** |
| | | MAE | $0.353 \pm 0.010$ | $0.346 \pm 0.014$ | $0.349 \pm 0.020$ | $0.366 \pm 0.006$ | **$0.336 \pm 0.020$** |
| | 10 | MSE | $0.206 \pm 0.007$ | $0.217 \pm 0.004$ | $0.222 \pm 0.008$ | $0.229 \pm 0.012$ | **$0.204 \pm 0.013$** |
| | | MAE | $0.358 \pm 0.009$ | $0.372 \pm 0.004$ | $0.376 \pm 0.007$ | $0.382 \pm 0.011$ | **$0.353 \pm 0.016$** |
| | 12 | MSE | $0.216 \pm 0.009$ | $0.216 \pm 0.008$ | $0.228 \pm 0.009$ | $0.230 \pm 0.008$ | **$0.207 \pm 0.011$** |
| | | MAE | $0.368 \pm 0.014$ | $0.369 \pm 0.010$ | $0.382 \pm 0.010$ | $0.384 \pm 0.008$ | **$0.357 \pm 0.012$** |

Table 53: Per-horizon component-ablation results with PatchTST for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | **$0.106 \pm 0.002$** | $0.203 \pm 0.025$ | $0.201 \pm 0.019$ | $0.190 \pm 0.003$ | $0.203 \pm 0.023$ |
| | | MAE | **$0.233 \pm 0.003$** | $0.355 \pm 0.026$ | $0.348 \pm 0.012$ | $0.343 \pm 0.002$ | $0.351 \pm 0.017$ |
| | 24 | MSE | **$0.214 \pm 0.003$** | $0.267 \pm 0.001$ | $0.269 \pm 0.004$ | $0.280 \pm 0.030$ | $0.277 \pm 0.008$ |
| | | MAE | **$0.340 \pm 0.002$** | $0.401 \pm 0.002$ | $0.404 \pm 0.005$ | $0.412 \pm 0.023$ | $0.413 \pm 0.006$ |
| | 36 | MSE | **$0.302 \pm 0.007$** | $0.347 \pm 0.002$ | $0.348 \pm 0.007$ | $0.360 \pm 0.031$ | $0.356 \pm 0.015$ |
| | | MAE | **$0.407 \pm 0.007$** | $0.454 \pm 0.002$ | $0.456 \pm 0.006$ | $0.466 \pm 0.028$ | $0.464 \pm 0.016$ |
| | 48 | MSE | **$0.385 \pm 0.002$** | $0.435 \pm 0.020$ | $0.425 \pm 0.003$ | $0.425 \pm 0.002$ | $0.449 \pm 0.031$ |
| | | MAE | **$0.473 \pm 0.004$** | $0.516 \pm 0.012$ | $0.510 \pm 0.002$ | $0.510 \pm 0.002$ | $0.527 \pm 0.018$ |
| Environment | 48 | MSE | **$0.269 \pm 0.000$** | $0.279 \pm 0.002$ | $0.275 \pm 0.002$ | $0.279 \pm 0.002$ | $0.270 \pm 0.002$ |
| | | MAE | $0.374 \pm 0.000$ | $0.389 \pm 0.002$ | $0.379 \pm 0.003$ | $0.383 \pm 0.005$ | **$0.373 \pm 0.003$** |
| | 96 | MSE | **$0.279 \pm 0.000$** | $0.299 \pm 0.004$ | $0.295 \pm 0.007$ | $0.298 \pm 0.004$ | $0.282 \pm 0.002$ |
| | | MAE | **$0.376 \pm 0.000$** | $0.404 \pm 0.005$ | $0.396 \pm 0.007$ | $0.398 \pm 0.004$ | $0.384 \pm 0.003$ |
| | 192 | MSE | **$0.272 \pm 0.000$** | $0.291 \pm 0.009$ | $0.289 \pm 0.004$ | $0.292 \pm 0.005$ | $0.274 \pm 0.001$ |
| | | MAE | **$0.367 \pm 0.001$** | $0.397 \pm 0.011$ | $0.395 \pm 0.005$ | $0.399 \pm 0.006$ | $0.375 \pm 0.003$ |
| | 336 | MSE | **$0.269 \pm 0.000$** | $0.291 \pm 0.003$ | $0.287 \pm 0.008$ | $0.293 \pm 0.003$ | $0.270 \pm 0.001$ |
| | | MAE | **$0.366 \pm 0.001$** | $0.403 \pm 0.005$ | $0.398 \pm 0.011$ | $0.405 \pm 0.005$ | $0.373 \pm 0.004$ |

Continued on next page

Table 53 (continued)

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Health | 12 | MSE | **1.008 ± 0.007** | 1.056 ± 0.033 | 1.045 ± 0.003 | 1.065 ± 0.030 | 1.036 ± 0.007 |
| | | MAE | **0.676 ± 0.005** | 0.700 ± 0.021 | 0.695 ± 0.007 | 0.713 ± 0.017 | 0.688 ± 0.005 |
| | 24 | MSE | **1.332 ± 0.029** | 1.390 ± 0.036 | 1.450 ± 0.070 | 1.420 ± 0.041 | 1.372 ± 0.018 |
| | | MAE | **0.776 ± 0.010** | 0.814 ± 0.028 | 0.823 ± 0.041 | 0.814 ± 0.032 | 0.790 ± 0.008 |
| | 36 | MSE | **1.431 ± 0.051** | 1.444 ± 0.029 | 1.443 ± 0.033 | 1.459 ± 0.067 | 1.451 ± 0.016 |
| | | MAE | **0.812 ± 0.007** | 0.823 ± 0.024 | 0.817 ± 0.013 | 0.835 ± 0.032 | 0.818 ± 0.009 |
| | 48 | MSE | **1.446 ± 0.005** | 1.498 ± 0.085 | 1.500 ± 0.061 | 1.533 ± 0.150 | 1.486 ± 0.015 |
| | | MAE | **0.815 ± 0.003** | 0.841 ± 0.029 | 0.844 ± 0.026 | 0.849 ± 0.036 | 0.827 ± 0.006 |

Table 54: Per-horizon component-ablation results with PatchTST for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | Full | w/o TCC | w/o CATB | w/o CAGI | TS query |
|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 66.074 ± 0.837 | 65.882 ± 0.577 | 66.065 ± 0.353 | **65.745 ± 0.602** | 66.749 ± 0.799 |
| | | MAE | 3.724 ± 0.086 | 3.694 ± 0.046 | 3.693 ± 0.036 | **3.645 ± 0.035** | 3.759 ± 0.132 |
| | 8 | MSE | 69.914 ± 0.658 | 69.850 ± 0.570 | 69.875 ± 0.532 | **69.397 ± 0.288** | 70.361 ± 0.572 |
| | | MAE | 3.954 ± 0.151 | 3.909 ± 0.038 | 3.893 ± 0.037 | **3.830 ± 0.039** | 3.956 ± 0.084 |
| | 10 | MSE | **73.227 ± 1.152** | 73.411 ± 0.110 | 73.444 ± 0.219 | 73.726 ± 1.084 | 74.403 ± 1.342 |
| | | MAE | **3.994 ± 0.192** | 4.022 ± 0.031 | 4.022 ± 0.032 | 4.041 ± 0.065 | 4.117 ± 0.145 |
| | 12 | MSE | **76.075 ± 0.498** | 76.861 ± 0.357 | 76.857 ± 0.265 | 76.729 ± 0.479 | 78.876 ± 2.625 |
| | | MAE | **4.028 ± 0.051** | 4.145 ± 0.048 | 4.144 ± 0.039 | 4.139 ± 0.059 | 4.394 ± 0.266 |
| Social Good | 6 | MSE | **0.789 ± 0.018** | 0.882 ± 0.022 | 0.892 ± 0.024 | 0.803 ± 0.040 | 0.967 ± 0.076 |
| | | MAE | **0.381 ± 0.010** | 0.418 ± 0.013 | 0.427 ± 0.015 | 0.405 ± 0.035 | 0.450 ± 0.034 |
| | 8 | MSE | **0.842 ± 0.010** | 0.942 ± 0.015 | 0.949 ± 0.028 | 0.878 ± 0.023 | 1.011 ± 0.072 |
| | | MAE | **0.418 ± 0.009** | 0.448 ± 0.004 | 0.458 ± 0.013 | 0.429 ± 0.021 | 0.481 ± 0.026 |
| | 10 | MSE | **0.890 ± 0.011** | 0.965 ± 0.018 | 1.017 ± 0.041 | 0.971 ± 0.088 | 1.061 ± 0.039 |
| | | MAE | **0.453 ± 0.016** | 0.459 ± 0.010 | 0.493 ± 0.021 | 0.488 ± 0.060 | 0.512 ± 0.021 |
| | 12 | MSE | **0.933 ± 0.014** | 1.034 ± 0.012 | 1.084 ± 0.094 | 1.004 ± 0.082 | 1.142 ± 0.049 |
| | | MAE | **0.471 ± 0.008** | 0.494 ± 0.006 | 0.537 ± 0.067 | 0.513 ± 0.078 | 0.546 ± 0.020 |
| Traffic | 6 | MSE | **0.152 ± 0.004** | 0.161 ± 0.003 | 0.155 ± 0.004 | 0.153 ± 0.003 | **0.152 ± 0.003** |
| | | MAE | 0.231 ± 0.009 | 0.243 ± 0.002 | 0.225 ± 0.009 | 0.236 ± 0.008 | **0.218 ± 0.008** |
| | 8 | MSE | **0.153 ± 0.002** | 0.162 ± 0.003 | 0.154 ± 0.004 | 0.155 ± 0.001 | **0.153 ± 0.005** |
| | | MAE | 0.221 ± 0.001 | 0.248 ± 0.006 | 0.226 ± 0.003 | 0.243 ± 0.005 | **0.219 ± 0.008** |
| | 10 | MSE | **0.157 ± 0.003** | 0.170 ± 0.002 | 0.160 ± 0.003 | 0.163 ± 0.007 | 0.161 ± 0.005 |
| | | MAE | 0.231 ± 0.005 | 0.253 ± 0.005 | 0.231 ± 0.004 | 0.241 ± 0.014 | **0.227 ± 0.005** |
| | 12 | MSE | **0.160 ± 0.004** | 0.174 ± 0.003 | 0.164 ± 0.004 | 0.163 ± 0.003 | 0.166 ± 0.003 |
| | | MAE | 0.235 ± 0.014 | 0.255 ± 0.007 | 0.234 ± 0.007 | 0.245 ± 0.009 | **0.229 ± 0.005** |

## F.2 Temporal-Bias Ablations

We compare four temporal-bias settings. *None* adds no temporal or positional bias. *ALiBi-style* uses fixed per-head linear decay slopes (Press et al., 2022). *Learned linear decay* starts from those slopes and learns them, sharing them across inputs. *Content-adaptive* uses CATB: each context slice determines its salience and decay, with a logarithmic distance penalty. The three decay variants use the same distance $\Delta_{ij} = (T - j) + (i - 1)$. For each backbone, CATB uses the main-comparison runs; the alternatives are independently trained on the same domains, horizons, and seeds.

On each backbone, CATB has lower aggregate MSE than every alternative in all nine domains and lower or tied MAE in eight under the three-decimal rule (Table 55). The experiment evaluates content conditioning and logarithmic decay together. Section F.2.1 reports the per-horizon results.

### F.2.1 Results by Forecast Horizon

These tables disaggregate Table 55 by prediction length under the same four temporal-bias settings for both backbones.

Table 55: Aggregate temporal-bias ablations with iTransformer and PatchTST. Each dataset lists all variants as rows; entries are the 20-run mean ± the mean of four horizon-specific sample standard deviations over five seeds. Lower is better; bold and underlined entries are the best and second-best results within each backbone–dataset–metric group.

| Dataset | Variant | iTransformer | | PatchTST | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| Agriculture | None | $0.268 \pm 0.012$ | $0.346 \pm 0.015$ | $0.332 \pm 0.016$ | $0.374 \pm 0.014$ |
| | ALiBi-style | $0.262 \pm 0.016$ | $0.344 \pm 0.014$ | $0.335 \pm 0.015$ | $0.376 \pm 0.014$ |
| | Learned linear decay | $0.261 \pm 0.016$ | $0.343 \pm 0.014$ | $0.335 \pm 0.015$ | $0.376 \pm 0.014$ |
| | Content-adaptive | $\mathbf{0.254 \pm 0.014}$ | $\mathbf{0.332 \pm 0.013}$ | $\mathbf{0.293 \pm 0.012}$ | $\mathbf{0.350 \pm 0.012}$ |
| Climate | None | $0.972 \pm 0.007$ | $0.800 \pm 0.004$ | $0.913 \pm 0.009$ | $0.767 \pm 0.006$ |
| | ALiBi-style | $0.973 \pm 0.007$ | $0.800 \pm 0.004$ | $0.908 \pm 0.010$ | $0.765 \pm 0.007$ |
| | Learned linear decay | $0.973 \pm 0.007$ | $0.800 \pm 0.004$ | $0.908 \pm 0.010$ | $0.765 \pm 0.007$ |
| | Content-adaptive | $\mathbf{0.894 \pm 0.004}$ | $\mathbf{0.751 \pm 0.002}$ | $\mathbf{0.896 \pm 0.004}$ | $\mathbf{0.753 \pm 0.002}$ |
| Economy | None | $0.186 \pm 0.006$ | $0.336 \pm 0.006$ | $0.208 \pm 0.012$ | $0.361 \pm 0.013$ |
| | ALiBi-style | $0.185 \pm 0.005$ | $0.334 \pm 0.005$ | $0.211 \pm 0.010$ | $0.364 \pm 0.011$ |
| | Learned linear decay | $0.185 \pm 0.005$ | $0.334 \pm 0.005$ | $0.211 \pm 0.010$ | $0.364 \pm 0.011$ |
| | Content-adaptive | $\mathbf{0.180 \pm 0.005}$ | $\mathbf{0.330 \pm 0.006}$ | $\mathbf{0.199 \pm 0.010}$ | $\mathbf{0.351 \pm 0.012}$ |
| Energy | None | $0.316 \pm 0.009$ | $0.434 \pm 0.007$ | $0.311 \pm 0.008$ | $0.430 \pm 0.006$ |
| | ALiBi-style | $0.314 \pm 0.010$ | $0.432 \pm 0.007$ | $0.314 \pm 0.014$ | $0.433 \pm 0.011$ |
| | Learned linear decay | $0.326 \pm 0.015$ | $0.440 \pm 0.010$ | $0.314 \pm 0.012$ | $0.433 \pm 0.011$ |
| | Content-adaptive | $\mathbf{0.241 \pm 0.012}$ | $\mathbf{0.355 \pm 0.008}$ | $\mathbf{0.252 \pm 0.004}$ | $\mathbf{0.363 \pm 0.004}$ |
| Environment | None | $0.265 \pm 0.005$ | $\mathbf{0.370 \pm 0.002}$ | $0.286 \pm 0.005$ | $0.392 \pm 0.006$ |
| | ALiBi-style | $0.269 \pm 0.003$ | $0.372 \pm 0.004$ | $0.288 \pm 0.005$ | $0.395 \pm 0.006$ |
| | Learned linear decay | $0.269 \pm 0.004$ | $0.372 \pm 0.005$ | $0.287 \pm 0.006$ | $0.393 \pm 0.007$ |
| | Content-adaptive | $\mathbf{0.259 \pm 0.000}$ | $\mathbf{0.370 \pm 0.001}$ | $\mathbf{0.272 \pm 0.000}$ | $\mathbf{0.371 \pm 0.001}$ |
| Health | None | $1.347 \pm 0.016$ | $0.761 \pm 0.004$ | $1.360 \pm 0.042$ | $0.795 \pm 0.022$ |
| | ALiBi-style | $1.341 \pm 0.036$ | $0.761 \pm 0.008$ | $1.372 \pm 0.075$ | $0.799 \pm 0.031$ |
| | Learned linear decay | $1.356 \pm 0.032$ | $0.763 \pm 0.009$ | $1.367 \pm 0.041$ | $0.799 \pm 0.025$ |
| | Content-adaptive | $\mathbf{1.273 \pm 0.010}$ | $\mathbf{0.753 \pm 0.005}$ | $\mathbf{1.304 \pm 0.023}$ | $\mathbf{0.770 \pm 0.006}$ |
| Security | None | $75.949 \pm 1.358$ | $4.420 \pm 0.113$ | $71.560 \pm 0.342$ | $3.938 \pm 0.036$ |
| | ALiBi-style | $75.885 \pm 1.335$ | $4.411 \pm 0.118$ | $71.565 \pm 0.376$ | $3.938 \pm 0.039$ |
| | Learned linear decay | $75.884 \pm 1.335$ | $4.411 \pm 0.118$ | $71.565 \pm 0.376$ | $3.938 \pm 0.039$ |
| | Content-adaptive | $\mathbf{71.734 \pm 0.644}$ | $\mathbf{3.963 \pm 0.097}$ | $\mathbf{71.322 \pm 0.786}$ | $\mathbf{3.925 \pm 0.120}$ |
| Social Good | None | $1.163 \pm 0.054$ | $0.484 \pm 0.020$ | $0.986 \pm 0.047$ | $0.479 \pm 0.029$ |
| | ALiBi-style | $1.144 \pm 0.053$ | $0.479 \pm 0.018$ | $0.979 \pm 0.037$ | $0.475 \pm 0.020$ |
| | Learned linear decay | $1.153 \pm 0.058$ | $0.482 \pm 0.023$ | $0.969 \pm 0.035$ | $0.474 \pm 0.022$ |
| | Content-adaptive | $\mathbf{0.928 \pm 0.016}$ | $\mathbf{0.428 \pm 0.007}$ | $\mathbf{0.863 \pm 0.013}$ | $\mathbf{0.431 \pm 0.011}$ |
| Traffic | None | $0.165 \pm 0.004$ | $0.211 \pm 0.005$ | $0.158 \pm 0.004$ | $\mathbf{0.229 \pm 0.006}$ |
| | ALiBi-style | $0.170 \pm 0.003$ | $\mathbf{0.210 \pm 0.004}$ | $0.157 \pm 0.003$ | $\mathbf{0.229 \pm 0.007}$ |
| | Learned linear decay | $0.170 \pm 0.003$ | $0.211 \pm 0.004$ | $0.157 \pm 0.003$ | $\mathbf{0.229 \pm 0.007}$ |
| | Content-adaptive | $\mathbf{0.162 \pm 0.004}$ | $0.218 \pm 0.005$ | $\mathbf{0.155 \pm 0.003}$ | $0.230 \pm 0.007$ |

Table 56: Per-horizon temporal-bias ablation results with iTransformer for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.160 \pm 0.007$ | $0.165 \pm 0.005$ | $0.159 \pm 0.001$ | $\mathbf{0.146 \pm 0.005}$ |
| | | MAE | $0.268 \pm 0.007$ | $0.272 \pm 0.003$ | $0.267 \pm 0.005$ | $\mathbf{0.255 \pm 0.007}$ |
| | 8 | MSE | $0.232 \pm 0.012$ | $0.228 \pm 0.017$ | $0.225 \pm 0.018$ | $\mathbf{0.210 \pm 0.018}$ |
| | | MAE | $0.321 \pm 0.015$ | $0.319 \pm 0.014$ | $0.316 \pm 0.015$ | $\mathbf{0.304 \pm 0.012}$ |
| | 10 | MSE | $0.297 \pm 0.013$ | $0.297 \pm 0.008$ | $0.299 \pm 0.013$ | $\mathbf{0.288 \pm 0.014}$ |
| | | MAE | $0.374 \pm 0.013$ | $0.371 \pm 0.014$ | $0.370 \pm 0.013$ | $\mathbf{0.362 \pm 0.017}$ |
| | 12 | MSE | $0.383 \pm 0.017$ | $\mathbf{0.359 \pm 0.033}$ | $0.362 \pm 0.033$ | $0.372 \pm 0.019$ |
| | | MAE | $0.423 \pm 0.026$ | $0.415 \pm 0.026$ | $0.418 \pm 0.022$ | $\mathbf{0.406 \pm 0.016}$ |
| Climate | 6 | MSE | $0.965 \pm 0.004$ | $0.964 \pm 0.004$ | $0.964 \pm 0.004$ | $\mathbf{0.884 \pm 0.006}$ |
| | | MAE | $0.797 \pm 0.004$ | $0.798 \pm 0.003$ | $0.798 \pm 0.003$ | $\mathbf{0.750 \pm 0.002}$ |
| | 8 | MSE | $0.969 \pm 0.006$ | $0.969 \pm 0.005$ | $0.969 \pm 0.005$ | $\mathbf{0.888 \pm 0.001}$ |
| | | MAE | $0.799 \pm 0.003$ | $0.799 \pm 0.003$ | $0.799 \pm 0.003$ | $\mathbf{0.749 \pm 0.000}$ |
| | 10 | MSE | $0.978 \pm 0.013$ | $0.978 \pm 0.012$ | $0.978 \pm 0.012$ | $\mathbf{0.899 \pm 0.007}$ |
| | | MAE | $0.802 \pm 0.006$ | $0.802 \pm 0.005$ | $0.802 \pm 0.005$ | $\mathbf{0.752 \pm 0.004}$ |
| | 12 | MSE | $0.978 \pm 0.007$ | $0.978 \pm 0.007$ | $0.978 \pm 0.007$ | $\mathbf{0.905 \pm 0.003}$ |
| | | MAE | $0.801 \pm 0.004$ | $0.803 \pm 0.006$ | $0.803 \pm 0.006$ | $\mathbf{0.754 \pm 0.001}$ |
| Economy | 6 | MSE | $0.155 \pm 0.004$ | $0.154 \pm 0.004$ | $0.154 \pm 0.004$ | $\mathbf{0.150 \pm 0.007}$ |
| | | MAE | $0.308 \pm 0.003$ | $0.306 \pm 0.004$ | $0.306 \pm 0.004$ | $\mathbf{0.303 \pm 0.008}$ |
| | 8 | MSE | $0.179 \pm 0.010$ | $0.176 \pm 0.006$ | $0.175 \pm 0.005$ | $\mathbf{0.169 \pm 0.006}$ |
| | | MAE | $0.330 \pm 0.009$ | $0.326 \pm 0.005$ | $0.325 \pm 0.004$ | $\mathbf{0.322 \pm 0.006}$ |
| | 10 | MSE | $0.196 \pm 0.005$ | $0.195 \pm 0.005$ | $0.195 \pm 0.005$ | $\mathbf{0.190 \pm 0.003}$ |
| | | MAE | $0.344 \pm 0.005$ | $0.344 \pm 0.005$ | $0.344 \pm 0.005$ | $\mathbf{0.342 \pm 0.003}$ |
| | 12 | MSE | $0.215 \pm 0.006$ | $0.215 \pm 0.006$ | $0.214 \pm 0.006$ | $\mathbf{0.209 \pm 0.005}$ |
| | | MAE | $0.362 \pm 0.005$ | $0.361 \pm 0.006$ | $0.361 \pm 0.005$ | $\mathbf{0.356 \pm 0.006}$ |

Table 57: Per-horizon temporal-bias ablation results with iTransformer for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.173 \pm 0.009$ | $0.172 \pm 0.004$ | $0.175 \pm 0.002$ | $\mathbf{0.092 \pm 0.004}$ |
| | | MAE | $0.323 \pm 0.005$ | $0.319 \pm 0.006$ | $0.321 \pm 0.003$ | $\mathbf{0.212 \pm 0.004}$ |
| | 24 | MSE | $0.270 \pm 0.011$ | $0.264 \pm 0.009$ | $0.274 \pm 0.019$ | $\mathbf{0.197 \pm 0.016}$ |
| | | MAE | $0.405 \pm 0.009$ | $0.403 \pm 0.005$ | $0.407 \pm 0.011$ | $\mathbf{0.329 \pm 0.010}$ |
| | 36 | MSE | $0.375 \pm 0.010$ | $0.375 \pm 0.012$ | $0.383 \pm 0.015$ | $\mathbf{0.284 \pm 0.018}$ |
| | | MAE | $0.479 \pm 0.008$ | $0.475 \pm 0.007$ | $0.483 \pm 0.012$ | $\mathbf{0.397 \pm 0.014}$ |
| | 48 | MSE | $0.447 \pm 0.006$ | $0.445 \pm 0.014$ | $0.470 \pm 0.025$ | $\mathbf{0.392 \pm 0.008}$ |
| | | MAE | $0.530 \pm 0.004$ | $0.532 \pm 0.010$ | $0.549 \pm 0.012$ | $\mathbf{0.481 \pm 0.005}$ |
| Environment | 48 | MSE | $0.263 \pm 0.002$ | $0.264 \pm 0.002$ | $0.263 \pm 0.002$ | $\mathbf{0.262 \pm 0.000}$ |
| | | MAE | $\mathbf{0.369 \pm 0.002}$ | $0.371 \pm 0.001$ | $0.370 \pm 0.002$ | $\mathbf{0.369 \pm 0.001}$ |
| | 96 | MSE | $0.268 \pm 0.006$ | $0.268 \pm 0.002$ | $0.266 \pm 0.002$ | $\mathbf{0.265 \pm 0.000}$ |
| | | MAE | $0.373 \pm 0.001$ | $0.374 \pm 0.003$ | $0.372 \pm 0.003$ | $\mathbf{0.371 \pm 0.001}$ |
| | 192 | MSE | $0.263 \pm 0.007$ | $0.272 \pm 0.004$ | $0.275 \pm 0.004$ | $\mathbf{0.256 \pm 0.001}$ |
| | | MAE | $0.370 \pm 0.004$ | $0.373 \pm 0.003$ | $0.374 \pm 0.004$ | $\mathbf{0.369 \pm 0.002}$ |
| | 336 | MSE | $0.265 \pm 0.003$ | $0.271 \pm 0.004$ | $0.270 \pm 0.006$ | $\mathbf{0.254 \pm 0.000}$ |
| | | MAE | $\mathbf{0.369 \pm 0.004}$ | $0.372 \pm 0.009$ | $0.370 \pm 0.010$ | $0.370 \pm 0.001$ |
| Health | 12 | MSE | $0.977 \pm 0.013$ | $0.973 \pm 0.020$ | $0.971 \pm 0.038$ | $\mathbf{0.929 \pm 0.016}$ |
| | | MAE | $0.653 \pm 0.004$ | $0.652 \pm 0.006$ | $0.651 \pm 0.007$ | $\mathbf{0.648 \pm 0.006}$ |
| | 24 | MSE | $1.349 \pm 0.014$ | $1.346 \pm 0.014$ | $1.348 \pm 0.014$ | $\mathbf{1.307 \pm 0.012}$ |
| | | MAE | $0.761 \pm 0.002$ | $0.761 \pm 0.006$ | $0.762 \pm 0.005$ | $\mathbf{0.760 \pm 0.003}$ |
| | 36 | MSE | $1.523 \pm 0.019$ | $1.532 \pm 0.038$ | $1.537 \pm 0.044$ | $\mathbf{1.404 \pm 0.004}$ |
| | | MAE | $0.809 \pm 0.006$ | $0.810 \pm 0.007$ | $0.809 \pm 0.012$ | $\mathbf{0.790 \pm 0.003}$ |
| | 48 | MSE | $1.539 \pm 0.020$ | $1.511 \pm 0.074$ | $1.566 \pm 0.034$ | $\mathbf{1.451 \pm 0.009}$ |
| | | MAE | $0.821 \pm 0.004$ | $0.820 \pm 0.014$ | $0.831 \pm 0.011$ | $\mathbf{0.812 \pm 0.009}$ |

Table 58: Per-horizon temporal-bias ablation results with iTransformer for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Security | 6 | MSE | 72.353 ± 1.051 | 72.353 ± 1.038 | 72.353 ± 1.037 | **66.815 ± 0.430** |
| | | MAE | 4.427 ± 0.093 | 4.420 ± 0.086 | 4.420 ± 0.086 | **3.847 ± 0.103** |
| | 8 | MSE | 75.013 ± 2.050 | 74.917 ± 2.041 | 74.917 ± 2.040 | **69.877 ± 0.567** |
| | | MAE | 4.440 ± 0.177 | 4.427 ± 0.198 | 4.427 ± 0.198 | **3.864 ± 0.107** |
| | 10 | MSE | 77.730 ± 1.941 | 77.711 ± 1.844 | 77.711 ± 1.844 | **73.225 ± 0.708** |
| | | MAE | 4.463 ± 0.142 | 4.462 ± 0.139 | 4.462 ± 0.139 | **3.993 ± 0.088** |
| | 12 | MSE | 78.701 ± 0.391 | 78.557 ± 0.418 | 78.557 ± 0.417 | **77.021 ± 0.872** |
| | | MAE | 4.351 ± 0.041 | 4.335 ± 0.051 | 4.335 ± 0.051 | **4.147 ± 0.091** |
| Social Good | 6 | MSE | 1.012 ± 0.024 | 0.996 ± 0.020 | 1.016 ± 0.015 | **0.864 ± 0.025** |
| | | MAE | 0.418 ± 0.014 | 0.414 ± 0.012 | 0.419 ± 0.008 | **0.394 ± 0.010** |
| | 8 | MSE | 1.166 ± 0.055 | 1.135 ± 0.060 | 1.158 ± 0.062 | **0.900 ± 0.014** |
| | | MAE | 0.481 ± 0.019 | 0.472 ± 0.021 | 0.478 ± 0.017 | **0.409 ± 0.008** |
| | 10 | MSE | 1.202 ± 0.054 | 1.187 ± 0.031 | 1.205 ± 0.065 | **0.954 ± 0.019** |
| | | MAE | 0.503 ± 0.023 | 0.493 ± 0.010 | 0.504 ± 0.031 | **0.441 ± 0.007** |
| | 12 | MSE | 1.272 ± 0.084 | 1.258 ± 0.102 | 1.233 ± 0.089 | **0.993 ± 0.008** |
| | | MAE | 0.535 ± 0.025 | 0.536 ± 0.027 | 0.528 ± 0.034 | **0.468 ± 0.004** |
| Traffic | 6 | MSE | 0.155 ± 0.002 | 0.161 ± 0.001 | 0.161 ± 0.001 | **0.152 ± 0.003** |
| | | MAE | **0.193 ± 0.003** | 0.194 ± 0.003 | **0.193 ± 0.003** | 0.202 ± 0.007 |
| | 8 | MSE | 0.163 ± 0.003 | 0.168 ± 0.003 | 0.168 ± 0.003 | **0.159 ± 0.004** |
| | | MAE | 0.207 ± 0.005 | 0.207 ± 0.005 | **0.206 ± 0.005** | 0.215 ± 0.005 |
| | 10 | MSE | 0.167 ± 0.006 | 0.170 ± 0.004 | 0.171 ± 0.005 | **0.165 ± 0.004** |
| | | MAE | 0.217 ± 0.006 | **0.215 ± 0.003** | **0.215 ± 0.003** | 0.224 ± 0.004 |
| | 12 | MSE | 0.175 ± 0.004 | 0.182 ± 0.004 | 0.182 ± 0.003 | **0.172 ± 0.005** |
| | | MAE | 0.228 ± 0.004 | **0.227 ± 0.004** | **0.227 ± 0.003** | 0.232 ± 0.006 |

Table 59: Per-horizon temporal-bias ablation results with PatchTST for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | 0.244 ± 0.022 | 0.245 ± 0.022 | 0.244 ± 0.023 | **0.191 ± 0.015** |
| | | MAE | 0.326 ± 0.022 | 0.326 ± 0.023 | 0.325 ± 0.023 | **0.298 ± 0.015** |
| | 8 | MSE | 0.285 ± 0.008 | 0.287 ± 0.006 | 0.287 ± 0.006 | **0.251 ± 0.008** |
| | | MAE | 0.344 ± 0.006 | 0.346 ± 0.006 | 0.346 ± 0.006 | **0.321 ± 0.005** |
| | 10 | MSE | 0.363 ± 0.019 | 0.370 ± 0.013 | 0.370 ± 0.014 | **0.329 ± 0.010** |
| | | MAE | 0.389 ± 0.011 | 0.394 ± 0.008 | 0.394 ± 0.008 | **0.370 ± 0.010** |
| | 12 | MSE | 0.437 ± 0.016 | 0.438 ± 0.017 | 0.437 ± 0.017 | **0.400 ± 0.015** |
| | | MAE | 0.437 ± 0.016 | 0.439 ± 0.020 | 0.438 ± 0.020 | **0.413 ± 0.018** |
| Climate | 6 | MSE | 0.908 ± 0.005 | 0.905 ± 0.009 | 0.905 ± 0.009 | **0.886 ± 0.005** |
| | | MAE | 0.769 ± 0.006 | 0.769 ± 0.005 | 0.769 ± 0.005 | **0.752 ± 0.002** |
| | 8 | MSE | 0.900 ± 0.011 | 0.900 ± 0.012 | 0.901 ± 0.013 | **0.891 ± 0.004** |
| | | MAE | 0.760 ± 0.005 | 0.762 ± 0.012 | 0.762 ± 0.012 | **0.750 ± 0.001** |
| | 10 | MSE | 0.912 ± 0.010 | 0.910 ± 0.007 | 0.910 ± 0.006 | **0.902 ± 0.006** |
| | | MAE | 0.763 ± 0.006 | 0.761 ± 0.003 | 0.761 ± 0.003 | **0.755 ± 0.003** |
| | 12 | MSE | 0.931 ± 0.011 | 0.916 ± 0.011 | 0.916 ± 0.011 | **0.904 ± 0.001** |
| | | MAE | 0.774 ± 0.009 | 0.766 ± 0.006 | 0.766 ± 0.006 | **0.756 ± 0.003** |
| Economy | 6 | MSE | 0.185 ± 0.014 | 0.192 ± 0.015 | 0.192 ± 0.015 | **0.175 ± 0.015** |
| | | MAE | 0.336 ± 0.015 | 0.344 ± 0.017 | 0.344 ± 0.017 | **0.325 ± 0.016** |
| | 8 | MSE | **0.196 ± 0.018** | 0.200 ± 0.009 | 0.202 ± 0.009 | 0.200 ± 0.009 |
| | | MAE | **0.349 ± 0.020** | 0.353 ± 0.009 | 0.354 ± 0.009 | 0.353 ± 0.010 |
| | 10 | MSE | 0.222 ± 0.008 | 0.222 ± 0.008 | 0.222 ± 0.008 | **0.206 ± 0.007** |
| | | MAE | 0.376 ± 0.007 | 0.376 ± 0.007 | 0.376 ± 0.007 | **0.358 ± 0.009** |
| | 12 | MSE | 0.228 ± 0.009 | 0.228 ± 0.009 | 0.228 ± 0.009 | **0.216 ± 0.009** |
| | | MAE | 0.382 ± 0.010 | 0.382 ± 0.010 | 0.382 ± 0.010 | **0.368 ± 0.014** |

Table 60: Per-horizon temporal-bias ablation results with PatchTST for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.201 \pm 0.019$ | $0.205 \pm 0.024$ | $0.202 \pm 0.014$ | $\mathbf{0.106 \pm 0.002}$ |
| | | MAE | $0.348 \pm 0.012$ | $0.356 \pm 0.021$ | $0.352 \pm 0.011$ | $\mathbf{0.233 \pm 0.003}$ |
| | 24 | MSE | $0.269 \pm 0.004$ | $0.267 \pm 0.002$ | $0.267 \pm 0.002$ | $\mathbf{0.214 \pm 0.003}$ |
| | | MAE | $0.404 \pm 0.005$ | $0.401 \pm 0.001$ | $0.401 \pm 0.001$ | $\mathbf{0.340 \pm 0.002}$ |
| | 36 | MSE | $0.348 \pm 0.007$ | $0.352 \pm 0.012$ | $0.350 \pm 0.009$ | $\mathbf{0.302 \pm 0.007}$ |
| | | MAE | $0.456 \pm 0.006$ | $0.457 \pm 0.008$ | $0.457 \pm 0.008$ | $\mathbf{0.407 \pm 0.007}$ |
| | 48 | MSE | $0.425 \pm 0.003$ | $0.435 \pm 0.017$ | $0.439 \pm 0.023$ | $\mathbf{0.385 \pm 0.002}$ |
| | | MAE | $0.510 \pm 0.002$ | $0.518 \pm 0.015$ | $0.524 \pm 0.024$ | $\mathbf{0.473 \pm 0.004}$ |
| Environment | 48 | MSE | $0.275 \pm 0.002$ | $0.279 \pm 0.005$ | $0.276 \pm 0.004$ | $\mathbf{0.269 \pm 0.000}$ |
| | | MAE | $0.379 \pm 0.003$ | $0.386 \pm 0.005$ | $0.381 \pm 0.005$ | $\mathbf{0.374 \pm 0.000}$ |
| | 96 | MSE | $0.295 \pm 0.007$ | $0.296 \pm 0.004$ | $0.297 \pm 0.005$ | $\mathbf{0.279 \pm 0.000}$ |
| | | MAE | $0.396 \pm 0.007$ | $0.397 \pm 0.004$ | $0.398 \pm 0.005$ | $\mathbf{0.376 \pm 0.000}$ |
| | 192 | MSE | $0.289 \pm 0.004$ | $0.289 \pm 0.003$ | $0.289 \pm 0.005$ | $\mathbf{0.272 \pm 0.000}$ |
| | | MAE | $0.395 \pm 0.005$ | $0.393 \pm 0.006$ | $0.394 \pm 0.006$ | $\mathbf{0.367 \pm 0.001}$ |
| | 336 | MSE | $0.287 \pm 0.008$ | $0.290 \pm 0.007$ | $0.288 \pm 0.009$ | $\mathbf{0.269 \pm 0.000}$ |
| | | MAE | $0.398 \pm 0.011$ | $0.402 \pm 0.006$ | $0.399 \pm 0.012$ | $\mathbf{0.366 \pm 0.001}$ |
| Health | 12 | MSE | $1.045 \pm 0.003$ | $1.026 \pm 0.014$ | $1.046 \pm 0.037$ | $\mathbf{1.008 \pm 0.007}$ |
| | | MAE | $0.695 \pm 0.007$ | $0.689 \pm 0.013$ | $0.700 \pm 0.022$ | $\mathbf{0.676 \pm 0.005}$ |
| | 24 | MSE | $1.450 \pm 0.070$ | $1.432 \pm 0.076$ | $1.417 \pm 0.032$ | $\mathbf{1.332 \pm 0.029}$ |
| | | MAE | $0.823 \pm 0.041$ | $0.813 \pm 0.026$ | $0.813 \pm 0.019$ | $\mathbf{0.776 \pm 0.010}$ |
| | 36 | MSE | $1.443 \pm 0.033$ | $1.511 \pm 0.138$ | $1.441 \pm 0.006$ | $\mathbf{1.431 \pm 0.051}$ |
| | | MAE | $0.817 \pm 0.013$ | $0.844 \pm 0.055$ | $0.818 \pm 0.019$ | $\mathbf{0.812 \pm 0.007}$ |
| | 48 | MSE | $1.500 \pm 0.061$ | $1.517 \pm 0.070$ | $1.565 \pm 0.088$ | $\mathbf{1.446 \pm 0.005}$ |
| | | MAE | $0.844 \pm 0.026$ | $0.849 \pm 0.029$ | $0.865 \pm 0.040$ | $\mathbf{0.815 \pm 0.003}$ |

Table 61: Per-horizon temporal-bias ablation results with PatchTST for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | None | ALiBi | Learned linear | Content-adaptive |
|---|---|---|---|---|---|---|
| Security | 6 | MSE | $66.065 \pm 0.353$ | $\mathbf{66.009 \pm 0.487}$ | $\mathbf{66.009 \pm 0.487}$ | $66.074 \pm 0.837$ |
| | | MAE | $3.693 \pm 0.036$ | $\mathbf{3.684 \pm 0.046}$ | $\mathbf{3.684 \pm 0.046}$ | $3.724 \pm 0.086$ |
| | 8 | MSE | $69.875 \pm 0.532$ | $\mathbf{69.871 \pm 0.501}$ | $\mathbf{69.871 \pm 0.501}$ | $69.914 \pm 0.658$ |
| | | MAE | $3.893 \pm 0.037$ | $\mathbf{3.891 \pm 0.037}$ | $\mathbf{3.891 \pm 0.037}$ | $3.954 \pm 0.151$ |
| | 10 | MSE | $73.444 \pm 0.219$ | $73.510 \pm 0.265$ | $73.510 \pm 0.266$ | $\mathbf{73.227 \pm 1.152}$ |
| | | MAE | $4.022 \pm 0.032$ | $4.030 \pm 0.032$ | $4.030 \pm 0.032$ | $\mathbf{3.994 \pm 0.192}$ |
| | 12 | MSE | $76.857 \pm 0.265$ | $76.871 \pm 0.252$ | $76.871 \pm 0.252$ | $\mathbf{76.075 \pm 0.498}$ |
| | | MAE | $4.144 \pm 0.039$ | $4.147 \pm 0.041$ | $4.147 \pm 0.041$ | $\mathbf{4.028 \pm 0.051}$ |
| Social Good | 6 | MSE | $0.892 \pm 0.024$ | $0.890 \pm 0.023$ | $0.877 \pm 0.029$ | $\mathbf{0.789 \pm 0.018}$ |
| | | MAE | $0.427 \pm 0.015$ | $0.425 \pm 0.014$ | $0.419 \pm 0.013$ | $\mathbf{0.381 \pm 0.010}$ |
| | 8 | MSE | $0.949 \pm 0.028$ | $0.930 \pm 0.048$ | $0.924 \pm 0.043$ | $\mathbf{0.842 \pm 0.010}$ |
| | | MAE | $0.458 \pm 0.013$ | $0.451 \pm 0.015$ | $0.450 \pm 0.014$ | $\mathbf{0.418 \pm 0.009}$ |
| | 10 | MSE | $1.017 \pm 0.041$ | $1.031 \pm 0.037$ | $1.013 \pm 0.034$ | $\mathbf{0.890 \pm 0.011}$ |
| | | MAE | $0.493 \pm 0.021$ | $0.504 \pm 0.026$ | $0.503 \pm 0.028$ | $\mathbf{0.453 \pm 0.016}$ |
| | 12 | MSE | $1.084 \pm 0.094$ | $1.063 \pm 0.042$ | $1.061 \pm 0.033$ | $\mathbf{0.933 \pm 0.014}$ |
| | | MAE | $0.537 \pm 0.067$ | $0.519 \pm 0.022$ | $0.523 \pm 0.031$ | $\mathbf{0.471 \pm 0.008}$ |
| Traffic | 6 | MSE | $0.155 \pm 0.004$ | $\mathbf{0.152 \pm 0.003}$ | $\mathbf{0.152 \pm 0.003}$ | $\mathbf{0.152 \pm 0.004}$ |
| | | MAE | $0.225 \pm 0.009$ | $\mathbf{0.224 \pm 0.008}$ | $\mathbf{0.224 \pm 0.008}$ | $0.231 \pm 0.009$ |
| | 8 | MSE | $0.154 \pm 0.004$ | $\mathbf{0.152 \pm 0.004}$ | $\mathbf{0.152 \pm 0.004}$ | $0.153 \pm 0.002$ |
| | | MAE | $0.226 \pm 0.003$ | $0.226 \pm 0.004$ | $0.227 \pm 0.004$ | $\mathbf{0.221 \pm 0.001}$ |
| | 10 | MSE | $0.160 \pm 0.003$ | $0.159 \pm 0.004$ | $0.159 \pm 0.004$ | $\mathbf{0.157 \pm 0.003}$ |
| | | MAE | $0.231 \pm 0.004$ | $\mathbf{0.230 \pm 0.007}$ | $\mathbf{0.230 \pm 0.007}$ | $0.231 \pm 0.005$ |
| | 12 | MSE | $0.164 \pm 0.004$ | $0.163 \pm 0.003$ | $0.163 \pm 0.003$ | $\mathbf{0.160 \pm 0.004}$ |
| | | MAE | $\mathbf{0.234 \pm 0.007}$ | $\mathbf{0.234 \pm 0.010}$ | $\mathbf{0.234 \pm 0.010}$ | $0.235 \pm 0.014$ |

### F.3 Component Transfer

Tables 62 and 63 report domain-level reductions and paired absolute errors for TaTS with TCC, Cross-MoE with a content gate, and SpecTF with CATB. Both metrics improve in all nine method–domain pairs under the protocol in Section D.1.

In SpecTF, CATB predicts salience and decay from the normalized concatenation of the real and imaginary text-key features. It adds the resulting bias before attention softmax, sharing the historical text axis across frequency queries and retaining the original query, key, and value paths. This transfer uses historical lag $r_j = T - j$ across frequency queries, following the host architecture's convention. The bias strength and decay initialization are fixed across seeds within each domain–horizon group.

Improvements use unrounded aggregate errors, $100(1 - \bar{E}_{\text{added}}/\bar{E}_{\text{original}})$. The Mean row averages the three domain percentages equally and supplies the main-text component-transfer figure. The absolute-error table reports each domain separately; its uncertainty summarizes variation across seeds within each horizon.

Table 62: Component transfer: improvement (%) over each original method. Each domain aggregates four horizons and five seeds. Mean averages the three domain percentages before rounding and gives the main-text bar values; positive values indicate lower error.

| Dataset | TaTS + TCC | | Cross-MoE + Gate | | SpecTF + CATB | |
|---|---|---|---|---|---|---|
| | MSE | MAE | MSE | MAE | MSE | MAE |
| Economy | +3.9 | +1.6 | +1.9 | +2.2 | +13.4 | +8.2 |
| Energy | +16.3 | +12.5 | +4.6 | +4.4 | +1.1 | +0.4 |
| Traffic | +1.6 | +1.6 | +4.2 | +3.4 | +6.6 | +1.1 |
| **Mean** | +7.3 | +5.3 | +3.5 | +3.3 | +7.0 | +3.2 |

Table 63: Absolute MSE and MAE for component transfer. Entries are 20-run means over four horizons and five seeds ± the mean of four horizon-specific sample standard deviations across seeds. Lower is better.

| Dataset | Adaptation | Original method | | With component | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| Economy | TaTS + TCC | $0.184 \pm 0.009$ | $0.334 \pm 0.010$ | $0.176 \pm 0.011$ | $0.329 \pm 0.011$ |
| | Cross-MoE + Gate | $0.189 \pm 0.005$ | $0.340 \pm 0.005$ | $0.186 \pm 0.006$ | $0.332 \pm 0.005$ |
| | SpecTF + CATB | $0.208 \pm 0.019$ | $0.361 \pm 0.021$ | $0.180 \pm 0.011$ | $0.331 \pm 0.013$ |
| Energy | TaTS + TCC | $0.299 \pm 0.002$ | $0.419 \pm 0.002$ | $0.250 \pm 0.010$ | $0.366 \pm 0.009$ |
| | Cross-MoE + Gate | $0.243 \pm 0.005$ | $0.362 \pm 0.005$ | $0.231 \pm 0.013$ | $0.346 \pm 0.010$ |
| | SpecTF + CATB | $0.283 \pm 0.007$ | $0.408 \pm 0.005$ | $0.280 \pm 0.015$ | $0.407 \pm 0.009$ |
| Traffic | TaTS + TCC | $0.169 \pm 0.002$ | $0.221 \pm 0.003$ | $0.166 \pm 0.003$ | $0.217 \pm 0.003$ |
| | Cross-MoE + Gate | $0.184 \pm 0.004$ | $0.235 \pm 0.006$ | $0.176 \pm 0.002$ | $0.227 \pm 0.002$ |
| | SpecTF + CATB | $0.186 \pm 0.010$ | $0.222 \pm 0.004$ | $0.174 \pm 0.010$ | $0.220 \pm 0.006$ |

#### F.3.1 Results by Forecast Horizon

These tables give paired original and adapted results at all four horizons on Economy, Energy, and Traffic. Each cell summarizes the same five seeds as the corresponding aggregate table.

Table 64: Per-horizon component-transfer results on Economy. Entries are mean ± sample standard deviation over five seeds. TaTS and Cross-MoE use iTransformer; SpecTF uses its spectral architecture.

| Adaptation | $H$ | Original method | | With component | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| TaTS + TCC | 6 | $0.150 \pm 0.002$ | $0.301 \pm 0.002$ | $0.140 \pm 0.002$ | $0.292 \pm 0.002$ |
| | 8 | $0.183 \pm 0.024$ | $0.333 \pm 0.024$ | $0.168 \pm 0.016$ | $0.324 \pm 0.016$ |
| | 10 | $0.197 \pm 0.008$ | $0.348 \pm 0.009$ | $0.197 \pm 0.009$ | $0.348 \pm 0.010$ |
| | 12 | $0.205 \pm 0.003$ | $0.355 \pm 0.003$ | $0.201 \pm 0.016$ | $0.352 \pm 0.018$ |
| Cross-MoE + Gate | 6 | $0.154 \pm 0.002$ | $0.306 \pm 0.004$ | $0.148 \pm 0.002$ | $0.296 \pm 0.002$ |
| | 8 | $0.182 \pm 0.007$ | $0.332 \pm 0.004$ | $0.168 \pm 0.005$ | $0.316 \pm 0.006$ |
| | 10 | $0.206 \pm 0.004$ | $0.357 \pm 0.004$ | $0.205 \pm 0.009$ | $0.350 \pm 0.008$ |
| | 12 | $0.216 \pm 0.007$ | $0.364 \pm 0.006$ | $0.223 \pm 0.006$ | $0.367 \pm 0.005$ |
| SpecTF + CATB | 6 | $0.163 \pm 0.032$ | $0.315 \pm 0.034$ | $0.142 \pm 0.003$ | $0.293 \pm 0.004$ |
| | 8 | $0.214 \pm 0.025$ | $0.368 \pm 0.028$ | $0.172 \pm 0.013$ | $0.323 \pm 0.015$ |
| | 10 | $0.238 \pm 0.005$ | $0.394 \pm 0.004$ | $0.192 \pm 0.015$ | $0.345 \pm 0.016$ |
| | 12 | $0.215 \pm 0.016$ | $0.366 \pm 0.020$ | $0.213 \pm 0.015$ | $0.364 \pm 0.016$ |

Table 65: Per-horizon component-transfer results on Energy. Entries are mean ± sample standard deviation over five seeds. TaTS and Cross-MoE use iTransformer; SpecTF uses its spectral architecture.

| Adaptation | $H$ | Original method | | With component | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| TaTS + TCC | 12 | $0.176 \pm 0.001$ | $0.331 \pm 0.001$ | $0.112 \pm 0.012$ | $0.248 \pm 0.014$ |
| | 24 | $0.260 \pm 0.004$ | $0.395 \pm 0.002$ | $0.207 \pm 0.005$ | $0.335 \pm 0.002$ |
| | 36 | $0.340 \pm 0.004$ | $0.447 \pm 0.003$ | $0.296 \pm 0.010$ | $0.403 \pm 0.009$ |
| | 48 | $0.419 \pm 0.001$ | $0.502 \pm 0.001$ | $0.385 \pm 0.014$ | $0.479 \pm 0.010$ |
| Cross-MoE + Gate | 12 | $0.102 \pm 0.005$ | $0.236 \pm 0.007$ | $0.086 \pm 0.002$ | $0.207 \pm 0.003$ |
| | 24 | $0.202 \pm 0.002$ | $0.336 \pm 0.002$ | $0.199 \pm 0.001$ | $0.329 \pm 0.000$ |
| | 36 | $0.287 \pm 0.006$ | $0.403 \pm 0.005$ | $0.283 \pm 0.022$ | $0.396 \pm 0.015$ |
| | 48 | $0.379 \pm 0.006$ | $0.473 \pm 0.005$ | $0.358 \pm 0.029$ | $0.453 \pm 0.021$ |
| SpecTF + CATB | 12 | $0.169 \pm 0.006$ | $0.319 \pm 0.005$ | $0.165 \pm 0.007$ | $0.317 \pm 0.005$ |
| | 24 | $0.240 \pm 0.003$ | $0.381 \pm 0.001$ | $0.251 \pm 0.026$ | $0.389 \pm 0.018$ |
| | 36 | $0.318 \pm 0.008$ | $0.435 \pm 0.006$ | $0.314 \pm 0.014$ | $0.435 \pm 0.010$ |
| | 48 | $0.404 \pm 0.011$ | $0.499 \pm 0.008$ | $0.388 \pm 0.013$ | $0.487 \pm 0.005$ |

Table 66: Per-horizon component-transfer results on Traffic. Entries are mean ± sample standard deviation over five seeds. TaTS and Cross-MoE use iTransformer; SpecTF uses its spectral architecture.

| Adaptation | $H$ | Original method | | With component | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| TaTS + TCC | 6 | $0.161 \pm 0.003$ | $0.212 \pm 0.003$ | $0.160 \pm 0.008$ | $0.202 \pm 0.004$ |
| | 8 | $0.167 \pm 0.002$ | $0.219 \pm 0.003$ | $0.163 \pm 0.002$ | $0.217 \pm 0.004$ |
| | 10 | $0.171 \pm 0.002$ | $0.224 \pm 0.003$ | $0.168 \pm 0.002$ | $0.222 \pm 0.003$ |
| | 12 | $0.176 \pm 0.001$ | $0.229 \pm 0.002$ | $0.173 \pm 0.002$ | $0.227 \pm 0.002$ |

Continued on next page

Table 66 (continued)

| Adaptation | $H$ | Original method | | With component | |
|---|---|---|---|---|---|
| | | MSE | MAE | MSE | MAE |
| Cross-MoE + Gate | 6 | $0.170 \pm 0.004$ | $0.216 \pm 0.006$ | $0.174 \pm 0.002$ | $0.215 \pm 0.002$ |
| | 8 | $0.180 \pm 0.005$ | $0.232 \pm 0.007$ | $0.181 \pm 0.002$ | $0.227 \pm 0.003$ |
| | 10 | $0.185 \pm 0.006$ | $0.240 \pm 0.007$ | $0.181 \pm 0.003$ | $0.236 \pm 0.004$ |
| | 12 | $0.199 \pm 0.001$ | $0.252 \pm 0.004$ | $0.168 \pm 0.001$ | $0.231 \pm 0.001$ |
| SpecTF + CATB | 6 | $0.175 \pm 0.007$ | $0.204 \pm 0.005$ | $0.170 \pm 0.011$ | $0.201 \pm 0.007$ |
| | 8 | $0.177 \pm 0.013$ | $0.218 \pm 0.003$ | $0.171 \pm 0.011$ | $0.218 \pm 0.007$ |
| | 10 | $0.196 \pm 0.007$ | $0.228 \pm 0.004$ | $0.169 \pm 0.004$ | $0.225 \pm 0.007$ |
| | 12 | $0.195 \pm 0.013$ | $0.238 \pm 0.005$ | $0.185 \pm 0.015$ | $0.234 \pm 0.006$ |

## G Sensitivity to Missing Context

We evaluate fixed DynaGuard checkpoints with iTransformer on Agriculture, Energy, and Health from Time-MMD, using four forecast horizons per domain and five training seeds. *Text drop* replaces the text in a random fraction $r$ of historical time slices with "No information available." We use $r \in \{0, 0.1, \ldots, 1\}$ with 20 random draws at each nonzero rate; $r = 0$ is the unmodified input. Each draw samples timestamps without replacement from the test input sequence, including the first window's history, and is shared across overlapping windows and training seeds within each setting. The numerical input and model weights remain fixed.

Let $\bar{E}_{dhs}(r)$ denote MSE averaged over the 20 draws for domain $d$, horizon $h$, and training seed $s$, and let $E_{dhs}(0)$ be the same checkpoint's unmodified-input MSE. For each seed, we compute

$$\Delta_s(r) = \frac{100}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \frac{1}{|\mathcal{H}_d|} \sum_{h \in \mathcal{H}_d} \left( \frac{\bar{E}_{dhs}(r)}{E_{dhs}(0)} - 1 \right), \tag{23}$$

where $\mathcal{D}$ contains the three domains and $\mathcal{H}_d$ contains domain $d$'s four horizons; $\bar{E}_{dhs}(0) = E_{dhs}(0)$. Thus, normalization precedes averaging across horizons and domains. The main curve averages $\Delta_s(r)$ across five training seeds. Shading gives the 95% Student-$t$ interval $\bar{\Delta}(r) \pm t_{0.975,4} s_\Delta(r)/\sqrt{5}$, where $s_\Delta$ is the sample standard deviation across seeds. Domain-level curves omit the average over domains. Zero denotes unchanged MSE, and positive increases extend downward.

The placeholder remains a valid text input at $r = 1$, unlike the all-text-masked intervention in Section H.2, which bypasses fusion. Figure 16 separates the domain means: Agriculture accounts for most of the initial increase, while Energy and Health change less under partial text loss. Table 67 reports absolute MSE and MAE, averaging draws and horizons within each seed before computing the mean and sample standard deviation across seeds.

Table 67: Text-drop results for DynaGuard with iTransformer. Entries are mean ± sample standard deviation across five training seeds. For each seed, we average 20 draws per nonzero drop rate and then four forecast horizons. At $r = 0$, the same retained models receive unmodified text.

| $r$ | Agriculture | | Energy | | Health | |
|---|---|---|---|---|---|---|
| | MSE | MAE | MSE | MAE | MSE | MAE |
| 0.0 | $0.262 \pm 0.007$ | $0.336 \pm 0.003$ | $0.245 \pm 0.002$ | $0.360 \pm 0.002$ | $1.350 \pm 0.040$ | $0.762 \pm 0.011$ |
| 0.1 | $0.293 \pm 0.008$ | $0.360 \pm 0.007$ | $0.246 \pm 0.004$ | $0.360 \pm 0.002$ | $1.363 \pm 0.053$ | $0.767 \pm 0.016$ |
| 0.2 | $0.295 \pm 0.009$ | $0.362 \pm 0.008$ | $0.246 \pm 0.005$ | $0.360 \pm 0.003$ | $1.365 \pm 0.056$ | $0.768 \pm 0.017$ |
| 0.3 | $0.295 \pm 0.009$ | $0.362 \pm 0.008$ | $0.247 \pm 0.005$ | $0.361 \pm 0.002$ | $1.365 \pm 0.057$ | $0.768 \pm 0.017$ |
| 0.4 | $0.295 \pm 0.009$ | $0.362 \pm 0.008$ | $0.246 \pm 0.005$ | $0.360 \pm 0.002$ | $1.366 \pm 0.059$ | $0.768 \pm 0.018$ |
| 0.5 | $0.294 \pm 0.008$ | $0.361 \pm 0.007$ | $0.246 \pm 0.005$ | $0.360 \pm 0.002$ | $1.366 \pm 0.059$ | $0.768 \pm 0.018$ |
| 0.6 | $0.295 \pm 0.008$ | $0.362 \pm 0.007$ | $0.247 \pm 0.005$ | $0.361 \pm 0.002$ | $1.367 \pm 0.061$ | $0.769 \pm 0.018$ |
| 0.7 | $0.294 \pm 0.008$ | $0.361 \pm 0.007$ | $0.249 \pm 0.005$ | $0.362 \pm 0.003$ | $1.366 \pm 0.060$ | $0.769 \pm 0.018$ |
| 0.8 | $0.294 \pm 0.008$ | $0.361 \pm 0.007$ | $0.250 \pm 0.006$ | $0.363 \pm 0.004$ | $1.366 \pm 0.062$ | $0.769 \pm 0.019$ |
| 0.9 | $0.293 \pm 0.008$ | $0.361 \pm 0.007$ | $0.254 \pm 0.007$ | $0.365 \pm 0.006$ | $1.370 \pm 0.067$ | $0.771 \pm 0.020$ |
| 1.0 | $0.289 \pm 0.016$ | $0.362 \pm 0.009$ | $0.264 \pm 0.013$ | $0.374 \pm 0.012$ | $1.406 \pm 0.134$ | $0.783 \pm 0.039$ |

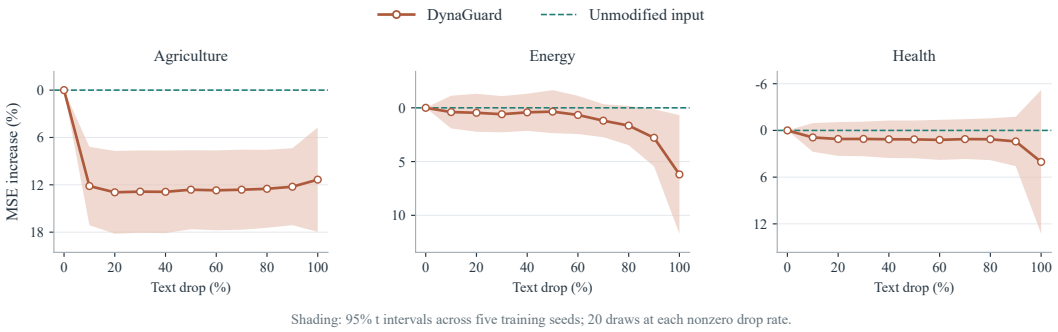


Figure 16: Text-drop sensitivity with iTransformer on three Time-MMD domains. Curves show MSE increases over each checkpoint's unmodified input, averaged over draws, horizons, and seeds. Shading gives 95% confidence intervals across five training seeds; dashed lines mark zero degradation.

### G.1 Text-Drop Results by Forecast Horizon

Tables 68–70 give the per-horizon results. Each entry averages 20 draws per nonzero rate within each seed, then reports the mean and sample standard deviation across five training seeds.

Table 68: Text-drop results by forecast horizon on Agriculture. Each entry is the mean ± sample standard deviation across five training seeds after averaging 20 draws per nonzero drop rate. The two panels cover all eleven rates for DynaGuard with iTransformer.

| $H$ | Metric | $r = 0.0$ | $r = 0.1$ | $r = 0.2$ | $r = 0.3$ | $r = 0.4$ | $r = 0.5$ |
|---|---|---|---|---|---|---|---|
| 6 | MSE | $0.152 \pm 0.005$ | $0.172 \pm 0.010$ | $0.173 \pm 0.011$ | $0.173 \pm 0.011$ | $0.173 \pm 0.011$ | $0.173 \pm 0.011$ |
| | MAE | $0.264 \pm 0.006$ | $0.277 \pm 0.007$ | $0.277 \pm 0.008$ | $0.277 \pm 0.008$ | $0.277 \pm 0.008$ | $0.277 \pm 0.008$ |
| 8 | MSE | $0.225 \pm 0.013$ | $0.242 \pm 0.014$ | $0.243 \pm 0.015$ | $0.243 \pm 0.015$ | $0.243 \pm 0.015$ | $0.243 \pm 0.015$ |
| | MAE | $0.312 \pm 0.007$ | $0.329 \pm 0.011$ | $0.330 \pm 0.012$ | $0.330 \pm 0.012$ | $0.330 \pm 0.012$ | $0.330 \pm 0.012$ |
| 10 | MSE | $0.294 \pm 0.015$ | $0.338 \pm 0.030$ | $0.340 \pm 0.032$ | $0.339 \pm 0.032$ | $0.339 \pm 0.032$ | $0.337 \pm 0.031$ |
| | MAE | $0.359 \pm 0.010$ | $0.393 \pm 0.023$ | $0.395 \pm 0.024$ | $0.395 \pm 0.024$ | $0.395 \pm 0.024$ | $0.394 \pm 0.023$ |
| 12 | MSE | $0.377 \pm 0.024$ | $0.422 \pm 0.038$ | $0.425 \pm 0.040$ | $0.425 \pm 0.040$ | $0.425 \pm 0.039$ | $0.424 \pm 0.041$ |
| | MAE | $0.409 \pm 0.016$ | $0.443 \pm 0.021$ | $0.446 \pm 0.022$ | $0.445 \pm 0.022$ | $0.446 \pm 0.022$ | $0.445 \pm 0.023$ |

| $H$ | Metric | $r = 0.6$ | $r = 0.7$ | $r = 0.8$ | $r = 0.9$ | $r = 1.0$ |
|---|---|---|---|---|---|---|
| 6 | MSE | $0.173 \pm 0.011$ | $0.173 \pm 0.011$ | $0.173 \pm 0.011$ | $0.174 \pm 0.011$ | $0.178 \pm 0.009$ |
| | MAE | $0.277 \pm 0.008$ | $0.277 \pm 0.008$ | $0.277 \pm 0.008$ | $0.278 \pm 0.007$ | $0.286 \pm 0.008$ |
| 8 | MSE | $0.244 \pm 0.016$ | $0.244 \pm 0.016$ | $0.244 \pm 0.016$ | $0.244 \pm 0.016$ | $0.244 \pm 0.023$ |
| | MAE | $0.330 \pm 0.012$ | $0.330 \pm 0.012$ | $0.331 \pm 0.012$ | $0.331 \pm 0.012$ | $0.334 \pm 0.014$ |
| 10 | MSE | $0.338 \pm 0.031$ | $0.337 \pm 0.031$ | $0.337 \pm 0.031$ | $0.335 \pm 0.030$ | $0.332 \pm 0.045$ |
| | MAE | $0.394 \pm 0.023$ | $0.393 \pm 0.023$ | $0.393 \pm 0.023$ | $0.392 \pm 0.022$ | $0.397 \pm 0.028$ |
| 12 | MSE | $0.424 \pm 0.040$ | $0.424 \pm 0.040$ | $0.422 \pm 0.042$ | $0.420 \pm 0.042$ | $0.401 \pm 0.047$ |
| | MAE | $0.445 \pm 0.022$ | $0.445 \pm 0.021$ | $0.444 \pm 0.022$ | $0.443 \pm 0.020$ | $0.432 \pm 0.016$ |

Table 69: Text-drop results by forecast horizon on Energy. Each entry is the mean ± sample standard deviation across five training seeds after averaging 20 draws per nonzero drop rate. The two panels cover all eleven rates for DynaGuard with iTransformer.

| $H$ | Metric | $r = 0.0$ | $r = 0.1$ | $r = 0.2$ | $r = 0.3$ | $r = 0.4$ | $r = 0.5$ |
|---|---|---|---|---|---|---|---|
| 12 | MSE | $0.092 \pm 0.001$ | $0.092 \pm 0.001$ | $0.092 \pm 0.001$ | $0.092 \pm 0.001$ | $0.092 \pm 0.001$ | $0.092 \pm 0.002$ |
| | MAE | $0.214 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ |
| 24 | MSE | $0.189 \pm 0.007$ | $0.189 \pm 0.005$ | $0.189 \pm 0.005$ | $0.190 \pm 0.005$ | $0.189 \pm 0.005$ | $0.190 \pm 0.005$ |
| | MAE | $0.330 \pm 0.007$ | $0.330 \pm 0.006$ | $0.330 \pm 0.006$ | $0.331 \pm 0.006$ | $0.330 \pm 0.007$ | $0.331 \pm 0.007$ |
| 36 | MSE | $0.294 \pm 0.019$ | $0.294 \pm 0.019$ | $0.294 \pm 0.020$ | $0.295 \pm 0.021$ | $0.293 \pm 0.018$ | $0.292 \pm 0.016$ |
| | MAE | $0.409 \pm 0.018$ | $0.408 \pm 0.018$ | $0.409 \pm 0.019$ | $0.409 \pm 0.019$ | $0.407 \pm 0.016$ | $0.407 \pm 0.015$ |
| 48 | MSE | $0.403 \pm 0.013$ | $0.409 \pm 0.021$ | $0.410 \pm 0.024$ | $0.410 \pm 0.024$ | $0.410 \pm 0.024$ | $0.409 \pm 0.023$ |
| | MAE | $0.487 \pm 0.007$ | $0.489 \pm 0.009$ | $0.489 \pm 0.010$ | $0.490 \pm 0.010$ | $0.490 \pm 0.010$ | $0.489 \pm 0.010$ |

| $H$ | Metric | $r = 0.6$ | $r = 0.7$ | $r = 0.8$ | $r = 0.9$ | $r = 1.0$ |
|---|---|---|---|---|---|---|
| 12 | MSE | $0.092 \pm 0.002$ | $0.092 \pm 0.002$ | $0.092 \pm 0.002$ | $0.092 \pm 0.002$ | $0.092 \pm 0.002$ |
| | MAE | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ | $0.213 \pm 0.003$ |
| 24 | MSE | $0.190 \pm 0.005$ | $0.192 \pm 0.006$ | $0.193 \pm 0.007$ | $0.195 \pm 0.008$ | $0.204 \pm 0.020$ |
| | MAE | $0.331 \pm 0.006$ | $0.332 \pm 0.007$ | $0.333 \pm 0.008$ | $0.334 \pm 0.008$ | $0.341 \pm 0.017$ |
| 36 | MSE | $0.295 \pm 0.021$ | $0.295 \pm 0.021$ | $0.297 \pm 0.025$ | $0.302 \pm 0.034$ | $0.310 \pm 0.046$ |
| | MAE | $0.408 \pm 0.018$ | $0.408 \pm 0.019$ | $0.410 \pm 0.021$ | $0.414 \pm 0.029$ | $0.421 \pm 0.042$ |
| 48 | MSE | $0.411 \pm 0.024$ | $0.416 \pm 0.026$ | $0.419 \pm 0.030$ | $0.426 \pm 0.033$ | $0.450 \pm 0.059$ |
| | MAE | $0.490 \pm 0.010$ | $0.493 \pm 0.012$ | $0.495 \pm 0.015$ | $0.500 \pm 0.021$ | $0.520 \pm 0.045$ |

Table 70: Text-drop results by forecast horizon on Health. Each entry is the mean ± sample standard deviation across five training seeds after averaging 20 draws per nonzero drop rate. The two panels cover all eleven rates for DynaGuard with iTransformer.

| $H$ | Metric | $r = 0.0$ | $r = 0.1$ | $r = 0.2$ | $r = 0.3$ | $r = 0.4$ | $r = 0.5$ |
|---|---|---|---|---|---|---|---|
| 12 | MSE | $1.094 \pm 0.102$ | $1.096 \pm 0.103$ | $1.097 \pm 0.104$ | $1.096 \pm 0.102$ | $1.096 \pm 0.101$ | $1.096 \pm 0.102$ |
| | MAE | $0.669 \pm 0.019$ | $0.670 \pm 0.020$ | $0.671 \pm 0.021$ | $0.670 \pm 0.020$ | $0.671 \pm 0.020$ | $0.671 \pm 0.020$ |
| 24 | MSE | $1.372 \pm 0.026$ | $1.404 \pm 0.084$ | $1.409 \pm 0.093$ | $1.408 \pm 0.093$ | $1.407 \pm 0.095$ | $1.406 \pm 0.096$ |
| | MAE | $0.769 \pm 0.010$ | $0.780 \pm 0.032$ | $0.781 \pm 0.035$ | $0.781 \pm 0.034$ | $0.781 \pm 0.034$ | $0.780 \pm 0.034$ |
| 36 | MSE | $1.420 \pm 0.024$ | $1.424 \pm 0.023$ | $1.426 \pm 0.024$ | $1.428 \pm 0.025$ | $1.425 \pm 0.021$ | $1.428 \pm 0.023$ |
| | MAE | $0.790 \pm 0.003$ | $0.793 \pm 0.007$ | $0.793 \pm 0.007$ | $0.794 \pm 0.006$ | $0.793 \pm 0.007$ | $0.794 \pm 0.006$ |
| 48 | MSE | $1.514 \pm 0.050$ | $1.526 \pm 0.032$ | $1.531 \pm 0.033$ | $1.530 \pm 0.040$ | $1.536 \pm 0.045$ | $1.536 \pm 0.048$ |
| | MAE | $0.819 \pm 0.016$ | $0.824 \pm 0.012$ | $0.826 \pm 0.013$ | $0.826 \pm 0.015$ | $0.828 \pm 0.017$ | $0.828 \pm 0.018$ |

| $H$ | Metric | $r = 0.6$ | $r = 0.7$ | $r = 0.8$ | $r = 0.9$ | $r = 1.0$ |
|---|---|---|---|---|---|---|
| 12 | MSE | $1.096 \pm 0.101$ | $1.095 \pm 0.101$ | $1.094 \pm 0.100$ | $1.095 \pm 0.102$ | $1.102 \pm 0.101$ |
| | MAE | $0.671 \pm 0.020$ | $0.671 \pm 0.020$ | $0.671 \pm 0.020$ | $0.671 \pm 0.021$ | $0.676 \pm 0.025$ |
| 24 | MSE | $1.406 \pm 0.095$ | $1.406 \pm 0.097$ | $1.406 \pm 0.101$ | $1.416 \pm 0.117$ | $1.533 \pm 0.367$ |
| | MAE | $0.780 \pm 0.033$ | $0.780 \pm 0.034$ | $0.780 \pm 0.034$ | $0.783 \pm 0.039$ | $0.817 \pm 0.111$ |
| 36 | MSE | $1.428 \pm 0.022$ | $1.427 \pm 0.019$ | $1.424 \pm 0.018$ | $1.427 \pm 0.020$ | $1.447 \pm 0.042$ |
| | MAE | $0.795 \pm 0.006$ | $0.795 \pm 0.006$ | $0.795 \pm 0.006$ | $0.796 \pm 0.006$ | $0.802 \pm 0.014$ |
| 48 | MSE | $1.540 \pm 0.057$ | $1.537 \pm 0.052$ | $1.540 \pm 0.054$ | $1.543 \pm 0.056$ | $1.544 \pm 0.056$ |
| | MAE | $0.830 \pm 0.022$ | $0.830 \pm 0.020$ | $0.831 \pm 0.022$ | $0.834 \pm 0.023$ | $0.837 \pm 0.025$ |

# H Data Characteristics and Case Study

## H.1 Text Availability and Numerical Persistence

We examine the raw Time-MMD `fact` field and target series to characterize Social Good and Traffic in the TSOnly comparison. A slice has available text if its trimmed text is nonempty and contains none of "NA," "no relevant information," "no direct information," "not relevant," or "no information available." Matching ignores case; "NA" must be a separate word, and the other patterns are phrases. Among nonempty slices, we also compute the fraction containing at least one such pattern. These checks measure availability and explicit missing-information indicators, not predictive content.

Social Good has a text-availability ratio of 0.552, falling to 0.374 among slices whose absolute target change is at or above the 90th percentile. Its nonempty slices average 1,244 characters. Traffic has lag-1 and lag-12 target autocorrelations of 0.951 and 0.979, while 0.141 of its nonempty slices match a missing-information pattern. Social Good therefore has less available text around large target changes, while Traffic retains strong numerical persistence. These descriptive statistics provide context for the comparison without attributing individual forecast errors to either property.

## H.2 Case Study

We examine Energy test sample 194 with iTransformer and a 24-week forecast horizon. The main figure displays seed 2025 (run 2); Table 71 reports all five runs. RMSE is measured on the original gasoline-price scale. This window-level analysis complements the preceding evaluations across all horizons.

The interventions reuse each seed's checkpoint and numerical input. *No CATB* sets the temporal-bias mode to `none`. *No text* masks every text token and bypasses the fusion layer normalization. *Joint mask* masks the two highlighted reports, assigned to April 11 and April 18, 2022 in the dataset. With iTransformer's single numerical query, CATB uses $\Delta_{1j} = T - j$, so the latest historical slice has zero lag.

Table 71: Energy case study (sample 194, iTransformer, $H = 24$): RMSE in USD/gallon. CATB and text interventions use the same DynaGuard checkpoint as the full prediction for each seed, without retraining. The two masked reports occupy 11 and 18 April 2022 in the dataset. The last row reports the five-seed mean and sample standard deviation.

| Run | Uni-modal | DynaGuard | CATB off | All text masked | Two reports masked |
|---|---|---|---|---|---|
| 1 | 0.5179 | 0.3967 | 0.3964 | 0.4863 | 0.3968 |
| 2 | 0.4742 | 0.3043 | 0.3124 | 0.4151 | 0.3712 |
| 3 | 0.5299 | 0.5291 | 0.5319 | 0.5311 | 0.5622 |
| 4 | 0.6116 | 0.5415 | 0.5412 | 0.5698 | 0.5417 |
| 5 | 0.6076 | 0.5012 | 0.5006 | 0.6003 | 0.5014 |
| Mean $\pm$ SD | $0.5483 \pm 0.0597$ | $0.4545 \pm 0.1015$ | $0.4565 \pm 0.0990$ | $0.5205 \pm 0.0728$ | $0.4747 \pm 0.0861$ |

# I Resource Cost and Cache Placement

**Measurement setup.** We profile six methods with iTransformer on Energy, WMT, and NumMentions, using history/horizon pairs of 24/24, 24/8, and 15/7, respectively. Measurements use two Intel Xeon Silver 4210R CPUs, one NVIDIA A100 80GB PCIe GPU, and PyTorch 2.6.0+cu124. This independent resource panel retains fixed configurations and the archived implementations' text-computation paths. MM-TSFlib and Cross-MoE use window-wise affine fusion in their no-prior branches; the panel is separate from the matched forecasting comparisons. DynaGuard uses training batches of 64; the other methods use 32. Each method retains its optimizer, early stopping, and training stages.

**Measurement boundaries.** Full fitting time includes process startup, text preparation, training, validation, and restoration of the final model. Peak GPU memory is the maximum memory allocated

to PyTorch tensors over the process. Inference measures synchronized forward calls with the model resident and numerical inputs prepared in host memory, including the required transfers and each method's text computation or cache access. After 20 warm-up calls, each measurement uses at least 200 requests, cycling over eight prepared batches. We first take the median latency within each run, then report its mean and sample standard deviation across seeds. Figure 17 uses the same across-seed summary for fitting time and peak memory.

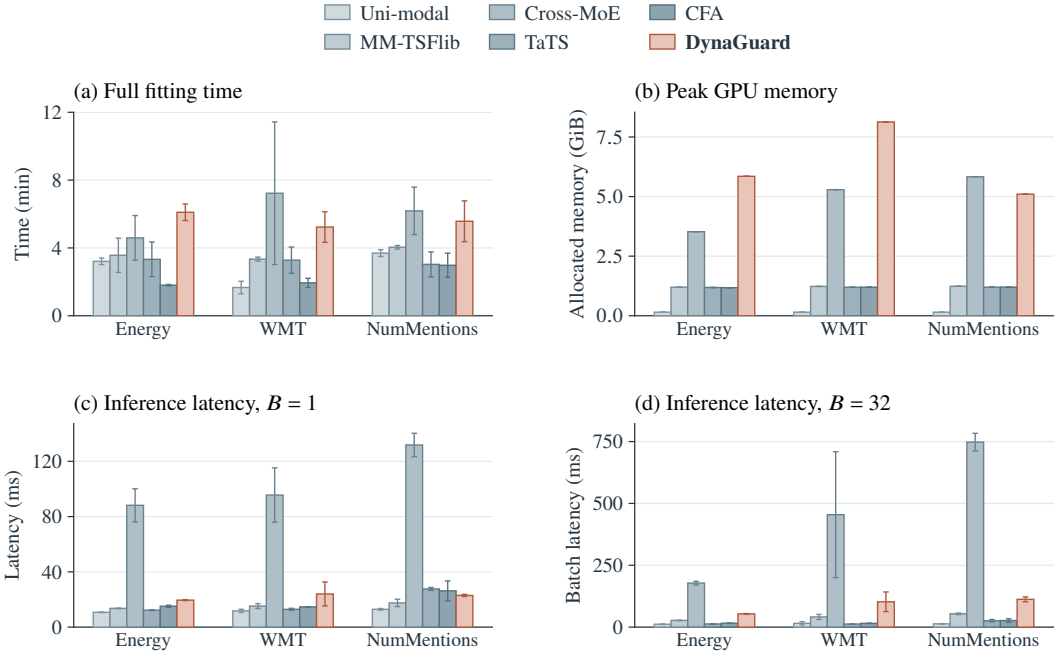


Figure 17: Resource measurements with iTransformer under the independent profiling configurations. Bars show means; error bars show sample standard deviations. Fitting time includes all training stages and data preparation. GPU memory is peak allocated tensor memory ($1\,\mathrm{GiB} = 2^{30}$ bytes). Latency is the within-run median for a batch of $B$ forecast windows. DynaGuard uses GPU-resident text embeddings.

**Inference and staged fitting.** DynaGuard has lower median inference latency than the measured Cross-MoE implementation on all three datasets, for both single-window and batched requests. The text-computation paths explain part of this advantage: Cross-MoE executes six frozen language-model Transformer blocks, whereas DynaGuard reads precomputed input embeddings and learns their compression through TCC. These measurements compare the complete execution paths. For DynaGuard's staged fitting, numerical pretraining takes the largest share of time; joint refinement accounts for approximately 18–25% of full fitting time. Total fitting time also reflects the number of epochs and evaluation calls in each method's training schedule.

**GPU memory and cache placement.** DynaGuard stores frozen token embeddings on the GPU to reuse them across epochs and forecast windows, reducing repeated embedding lookup and host-to-device transfers during large-scale experiments. The cache retains token-level inputs so that TCC remains trainable. Its size depends on the number of text slices and their padded token lengths, in addition to the memory required for model computation. Storing these embeddings in host memory and transferring each requested batch provides a lower-memory option without changing the forecasting model.

We verify this option with a paired WMT/iTransformer check: a 24-step history, an eight-step horizon, batch size 64, seed 2024, and one epoch per training stage. Both runs use an NVIDIA A100 40GB PCIe GPU and PyTorch 2.5.1+cu121, with cache location as the sole computational change. Table 72 shows the memory reduction; the final model-state hashes and test errors are identical. This check isolates cache placement under a short training schedule. Host-memory caching introduces

batch transfers, so we use GPU caching when prioritizing experiment throughput and host-memory caching when prioritizing GPU capacity.

Table 72: Paired cache-placement check on WMT with iTransformer. Each training stage runs for one epoch. Both runs produce identical final model states and test errors. GiB denotes $2^{30}$ bytes.

| Cache location | Peak allocated (GiB) | Peak reserved (GiB) | MSE | MAE |
|---|---|---|---|---|
| GPU | 8.13 | 11.44 | 0.14401 | 0.31018 |
| Host memory | 6.59 | 8.96 | 0.14401 | 0.31018 |

# J Complete-System Comparison

We compare two DynaGuard configurations with GPT4MTS, SpecTF, VoT, and TFHTS (Zhou et al., 2024) on all nine Time-MMD domains. Each complete system retains its original text encoder and forecasting design, extending the same-backbone comparisons in the main paper. Across the 18 domain–metric groups, the two DynaGuard configurations attain 10 lowest and 12 second-lowest entries. Table 73 reports aggregate MSE and MAE with seed variation.

Table 73: Model-level comparison on nine domain datasets. Complete systems retain their original text encoders; entries are the 20-run mean ± the mean of four per-horizon sample standard deviations over five seeds.

| Domain | Metric | DG+iTrans. | DG+TimeXer | GPT4MTS | SpecTF | VoT | TFHTS |
|---|---|---|---|---|---|---|---|
| Agriculture | MSE | 0.254 ± 0.014 | 0.287 ± 0.018 | 0.275 ± 0.022 | **0.245 ± 0.016** | 0.273 ± 0.002 | 0.423 ± 0.010 |
| | MAE | 0.332 ± 0.013 | 0.352 ± 0.011 | 0.354 ± 0.015 | **0.328 ± 0.012** | 0.335 ± 0.002 | 0.466 ± 0.006 |
| Climate | MSE | **0.894 ± 0.004** | 0.895 ± 0.005 | 1.113 ± 0.024 | 0.975 ± 0.008 | 1.115 ± 0.006 | 1.133 ± 0.019 |
| | MAE | **0.751 ± 0.002** | 0.752 ± 0.002 | 0.864 ± 0.011 | 0.801 ± 0.004 | 0.834 ± 0.002 | 0.858 ± 0.007 |
| Economy | MSE | **0.180 ± 0.005** | 0.199 ± 0.011 | 0.192 ± 0.004 | 0.208 ± 0.019 | 0.264 ± 0.003 | 0.271 ± 0.020 |
| | MAE | **0.330 ± 0.006** | 0.352 ± 0.011 | 0.343 ± 0.004 | 0.361 ± 0.021 | 0.408 ± 0.003 | 0.418 ± 0.017 |
| Energy | MSE | 0.241 ± 0.012 | 0.250 ± 0.008 | 0.256 ± 0.010 | 0.283 ± 0.007 | **0.217 ± 0.005** | 0.303 ± 0.006 |
| | MAE | 0.355 ± 0.008 | 0.360 ± 0.008 | 0.368 ± 0.009 | 0.408 ± 0.005 | **0.336 ± 0.004** | 0.419 ± 0.006 |
| Environment | MSE | **0.259 ± 0.000** | 0.272 ± 0.000 | 0.344 ± 0.007 | 0.265 ± 0.004 | 14.572 ± 0.333 | 0.361 ± 0.010 |
| | MAE | 0.370 ± 0.001 | 0.371 ± 0.001 | 0.417 ± 0.009 | **0.369 ± 0.002** | 2.033 ± 0.035 | 0.434 ± 0.011 |
| Health | MSE | **1.273 ± 0.010** | 1.308 ± 0.057 | 1.759 ± 0.120 | 1.290 ± 0.017 | 1.644 ± 0.014 | 2.195 ± 0.027 |
| | MAE | **0.753 ± 0.005** | 0.774 ± 0.023 | 0.878 ± 0.025 | 0.758 ± 0.006 | 0.831 ± 0.003 | 1.013 ± 0.007 |
| Security | MSE | 71.734 ± 0.644 | **71.277 ± 0.392** | 81.338 ± 7.738 | 75.849 ± 2.019 | 79.738 ± 0.217 | 75.764 ± 1.234 |
| | MAE | 3.963 ± 0.097 | **3.906 ± 0.069** | 4.908 ± 0.434 | 4.318 ± 0.193 | 4.375 ± 0.014 | 4.493 ± 0.125 |
| Social Good | MSE | 0.928 ± 0.016 | 0.812 ± 0.008 | 1.400 ± 0.329 | 1.383 ± 0.090 | **0.810 ± 0.005** | 1.135 ± 0.067 |
| | MAE | 0.428 ± 0.007 | 0.407 ± 0.008 | 0.537 ± 0.080 | 0.504 ± 0.014 | **0.386 ± 0.004** | 0.640 ± 0.038 |
| Traffic | MSE | 0.162 ± 0.004 | **0.152 ± 0.002** | 0.161 ± 0.005 | 0.186 ± 0.010 | 0.170 ± 0.001 | 0.286 ± 0.008 |
| | MAE | 0.218 ± 0.005 | 0.227 ± 0.006 | **0.215 ± 0.005** | 0.222 ± 0.004 | 0.274 ± 0.001 | 0.377 ± 0.008 |

## J.1 Per-Horizon Complete-System Comparison

The following tables give the per-horizon results for the six configurations in Table 73, using the same complete-system protocol.

Table 74: Model-level results by prediction length $H$ for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | DG+iTrans. | DG+TimeXer | GPT4MTS | SpecTF | VoT | TFHTS |
|---|---|---|---|---|---|---|---|---|
| Agriculture | 6 | MSE | 0.146 ± 0.005 | 0.196 ± 0.010 | 0.164 ± 0.017 | **0.144 ± 0.009** | 0.168 ± 0.002 | 0.326 ± 0.010 |
| | | MAE | **0.255 ± 0.007** | 0.293 ± 0.007 | 0.283 ± 0.012 | 0.259 ± 0.009 | 0.267 ± 0.002 | 0.405 ± 0.005 |
| | 8 | MSE | 0.210 ± 0.018 | 0.261 ± 0.014 | 0.246 ± 0.033 | **0.204 ± 0.009** | 0.237 ± 0.001 | 0.388 ± 0.008 |
| | | MAE | **0.304 ± 0.012** | 0.335 ± 0.016 | 0.337 ± 0.024 | 0.305 ± 0.008 | 0.315 ± 0.001 | 0.447 ± 0.006 |
| | 10 | MSE | 0.288 ± 0.014 | 0.313 ± 0.024 | 0.301 ± 0.014 | **0.270 ± 0.021** | 0.307 ± 0.004 | 0.450 ± 0.013 |
| | | MAE | 0.362 ± 0.017 | 0.371 ± 0.009 | 0.371 ± 0.010 | **0.344 ± 0.016** | 0.358 ± 0.002 | 0.483 ± 0.007 |
| | 12 | MSE | 0.372 ± 0.019 | 0.378 ± 0.023 | 0.388 ± 0.025 | **0.361 ± 0.025** | 0.380 ± 0.001 | 0.530 ± 0.010 |
| | | MAE | 0.406 ± 0.016 | 0.408 ± 0.014 | 0.424 ± 0.014 | 0.406 ± 0.015 | **0.400 ± 0.001** | 0.528 ± 0.004 |
| Climate | 6 | MSE | **0.884 ± 0.006** | 0.885 ± 0.005 | 1.123 ± 0.039 | 0.955 ± 0.012 | 1.009 ± 0.005 | 1.112 ± 0.021 |
| | | MAE | **0.750 ± 0.002** | **0.750 ± 0.002** | 0.869 ± 0.018 | 0.789 ± 0.006 | 0.798 ± 0.002 | 0.853 ± 0.007 |
| | 8 | MSE | **0.888 ± 0.001** | 0.890 ± 0.004 | 1.115 ± 0.019 | 0.973 ± 0.005 | 1.082 ± 0.006 | 1.133 ± 0.025 |
| | | MAE | **0.749 ± 0.000** | 0.750 ± 0.002 | 0.864 ± 0.007 | 0.799 ± 0.002 | 0.824 ± 0.002 | 0.858 ± 0.008 |
| | 10 | MSE | **0.899 ± 0.007** | 0.900 ± 0.006 | 1.131 ± 0.028 | 0.985 ± 0.008 | 1.154 ± 0.006 | 1.149 ± 0.019 |
| | | MAE | **0.752 ± 0.004** | **0.752 ± 0.003** | 0.871 ± 0.011 | 0.806 ± 0.005 | 0.846 ± 0.003 | 0.863 ± 0.007 |
| | 12 | MSE | **0.905 ± 0.003** | 0.906 ± 0.003 | 1.083 ± 0.011 | 0.989 ± 0.006 | 1.213 ± 0.006 | 1.138 ± 0.011 |
| | | MAE | **0.754 ± 0.001** | 0.755 ± 0.001 | 0.854 ± 0.006 | 0.809 ± 0.003 | 0.867 ± 0.002 | 0.856 ± 0.004 |
| Economy | 6 | MSE | **0.150 ± 0.007** | 0.163 ± 0.009 | 0.164 ± 0.004 | 0.163 ± 0.032 | 0.223 ± 0.003 | 0.269 ± 0.009 |
| | | MAE | **0.303 ± 0.008** | 0.316 ± 0.011 | 0.317 ± 0.004 | 0.315 ± 0.034 | 0.378 ± 0.003 | 0.416 ± 0.011 |
| | 8 | MSE | **0.169 ± 0.006** | 0.194 ± 0.011 | 0.186 ± 0.006 | 0.214 ± 0.025 | 0.249 ± 0.002 | 0.257 ± 0.018 |
| | | MAE | **0.322 ± 0.006** | 0.346 ± 0.013 | 0.336 ± 0.003 | 0.368 ± 0.028 | 0.398 ± 0.002 | 0.405 ± 0.017 |
| | 10 | MSE | **0.190 ± 0.003** | 0.216 ± 0.007 | 0.201 ± 0.003 | 0.238 ± 0.005 | 0.279 ± 0.004 | 0.277 ± 0.021 |
| | | MAE | **0.342 ± 0.003** | 0.370 ± 0.007 | 0.351 ± 0.005 | 0.394 ± 0.004 | 0.418 ± 0.004 | 0.424 ± 0.017 |
| | 12 | MSE | **0.209 ± 0.005** | 0.224 ± 0.014 | 0.218 ± 0.005 | 0.215 ± 0.016 | 0.305 ± 0.002 | 0.282 ± 0.032 |
| | | MAE | **0.356 ± 0.006** | 0.378 ± 0.014 | 0.369 ± 0.006 | 0.366 ± 0.020 | 0.440 ± 0.002 | 0.427 ± 0.025 |

Table 75: Model-level results by prediction length $H$ for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | DG+iTrans. | DG+TimeXer | GPT4MTS | SpecTF | VoT | TFHTS |
|---|---|---|---|---|---|---|---|---|
| Energy | 12 | MSE | 0.092 ± 0.004 | 0.103 ± 0.002 | 0.104 ± 0.005 | 0.169 ± 0.006 | **0.086 ± 0.002** | 0.164 ± 0.007 |
| | | MAE | 0.212 ± 0.004 | 0.228 ± 0.004 | 0.232 ± 0.008 | 0.319 ± 0.005 | **0.210 ± 0.003** | 0.310 ± 0.009 |
| | 24 | MSE | 0.197 ± 0.016 | 0.212 ± 0.002 | 0.215 ± 0.006 | 0.240 ± 0.003 | **0.174 ± 0.003** | 0.266 ± 0.007 |
| | | MAE | 0.329 ± 0.010 | 0.337 ± 0.002 | 0.343 ± 0.005 | 0.381 ± 0.001 | **0.310 ± 0.003** | 0.397 ± 0.007 |
| | 36 | MSE | 0.284 ± 0.018 | 0.306 ± 0.014 | 0.314 ± 0.024 | 0.318 ± 0.008 | **0.258 ± 0.005** | 0.350 ± 0.004 |
| | | MAE | 0.397 ± 0.014 | 0.411 ± 0.015 | 0.420 ± 0.018 | 0.435 ± 0.006 | **0.378 ± 0.004** | 0.456 ± 0.004 |
| | 48 | MSE | 0.392 ± 0.008 | 0.379 ± 0.015 | 0.390 ± 0.006 | 0.404 ± 0.011 | **0.351 ± 0.009** | 0.433 ± 0.006 |
| | | MAE | 0.481 ± 0.005 | 0.464 ± 0.010 | 0.478 ± 0.004 | 0.499 ± 0.008 | **0.445 ± 0.007** | 0.512 ± 0.004 |
| Environment | 48 | MSE | **0.262 ± 0.000** | 0.269 ± 0.000 | 0.310 ± 0.003 | 0.265 ± 0.002 | 1.323 ± 0.024 | 0.338 ± 0.013 |
| | | MAE | **0.369 ± 0.001** | 0.374 ± 0.000 | 0.393 ± 0.005 | **0.369 ± 0.002** | 0.772 ± 0.007 | 0.421 ± 0.014 |
| | 96 | MSE | **0.265 ± 0.000** | 0.279 ± 0.000 | 0.358 ± 0.009 | 0.268 ± 0.001 | 3.810 ± 0.157 | 0.371 ± 0.008 |
| | | MAE | **0.371 ± 0.001** | 0.376 ± 0.000 | 0.427 ± 0.011 | **0.371 ± 0.003** | 1.237 ± 0.025 | 0.440 ± 0.010 |
| | 192 | MSE | **0.256 ± 0.001** | 0.271 ± 0.000 | 0.356 ± 0.008 | 0.259 ± 0.005 | 12.914 ± 0.580 | 0.372 ± 0.011 |
| | | MAE | 0.369 ± 0.002 | **0.367 ± 0.001** | 0.421 ± 0.010 | 0.369 ± 0.001 | 2.233 ± 0.062 | 0.439 ± 0.010 |
| | 336 | MSE | **0.254 ± 0.000** | 0.268 ± 0.000 | 0.351 ± 0.006 | 0.267 ± 0.007 | 40.239 ± 0.571 | 0.363 ± 0.007 |
| | | MAE | 0.370 ± 0.001 | **0.366 ± 0.001** | 0.427 ± 0.009 | **0.366 ± 0.002** | 3.890 ± 0.045 | 0.437 ± 0.010 |
| Health | 12 | MSE | **0.929 ± 0.016** | 1.009 ± 0.013 | 1.172 ± 0.025 | 0.961 ± 0.019 | 1.129 ± 0.020 | 1.888 ± 0.043 |
| | | MAE | **0.648 ± 0.006** | 0.677 ± 0.009 | 0.715 ± 0.011 | 0.650 ± 0.006 | 0.684 ± 0.005 | 0.934 ± 0.012 |
| | 24 | MSE | **1.307 ± 0.012** | 1.365 ± 0.009 | 1.819 ± 0.241 | 1.327 ± 0.009 | 1.590 ± 0.010 | 2.408 ± 0.023 |
| | | MAE | **0.760 ± 0.003** | 0.787 ± 0.004 | 0.903 ± 0.047 | 0.770 ± 0.002 | 0.808 ± 0.002 | 1.057 ± 0.006 |
| | 36 | MSE | **1.404 ± 0.004** | 1.436 ± 0.126 | 2.010 ± 0.159 | 1.413 ± 0.014 | 1.826 ± 0.015 | 2.284 ± 0.016 |
| | | MAE | **0.790 ± 0.003** | 0.822 ± 0.047 | 0.947 ± 0.028 | 0.797 ± 0.008 | 0.887 ± 0.003 | 1.040 ± 0.004 |
| | 48 | MSE | 1.451 ± 0.009 | **1.424 ± 0.080** | 2.034 ± 0.056 | 1.459 ± 0.027 | 2.032 ± 0.011 | 2.197 ± 0.028 |
| | | MAE | 0.812 ± 0.009 | **0.810 ± 0.031** | 0.947 ± 0.014 | 0.816 ± 0.011 | 0.944 ± 0.002 | 1.021 ± 0.009 |

Table 76: Model-level results by prediction length $H$ for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | DG+iTrans. | DG+TimeXer | GPT4MTS | SpecTF | VoT | TFHTS |
|---|---|---|---|---|---|---|---|---|
| Security | 6 | MSE | 66.815 ± 0.430 | **66.123 ± 0.313** | 73.333 ± 6.341 | 71.611 ± 2.839 | 73.304 ± 0.141 | 69.746 ± 1.493 |
| | | MAE | 3.847 ± 0.103 | **3.655 ± 0.113** | 4.556 ± 0.393 | 4.180 ± 0.334 | 4.171 ± 0.007 | 4.195 ± 0.149 |
| | 8 | MSE | **69.877 ± 0.567** | 70.022 ± 0.572 | 75.585 ± 2.916 | 77.144 ± 1.424 | 77.637 ± 0.263 | 73.395 ± 1.131 |
| | | MAE | **3.864 ± 0.107** | 3.950 ± 0.059 | 4.632 ± 0.196 | 4.552 ± 0.120 | 4.327 ± 0.021 | 4.362 ± 0.137 |
| | 10 | MSE | 73.225 ± 0.708 | **72.734 ± 0.130** | 81.415 ± 2.863 | 78.691 ± 3.320 | 81.809 ± 0.216 | 78.490 ± 1.324 |
| | | MAE | 3.993 ± 0.088 | **3.900 ± 0.038** | 4.902 ± 0.209 | 4.490 ± 0.245 | 4.431 ± 0.018 | 4.666 ± 0.129 |
| | 12 | MSE | 77.021 ± 0.872 | 76.230 ± 0.552 | 95.018 ± 18.831 | **75.948 ± 0.492** | 86.201 ± 0.250 | 81.426 ± 0.986 |
| | | MAE | 4.147 ± 0.091 | 4.121 ± 0.067 | 5.541 ± 0.939 | **4.051 ± 0.071** | 4.573 ± 0.012 | 4.750 ± 0.083 |
| Social Good | 6 | MSE | 0.864 ± 0.025 | **0.716 ± 0.007** | 1.233 ± 0.372 | 1.273 ± 0.071 | 0.756 ± 0.007 | 0.964 ± 0.069 |
| | | MAE | 0.394 ± 0.010 | 0.364 ± 0.006 | 0.498 ± 0.161 | 0.454 ± 0.015 | **0.360 ± 0.006** | 0.542 ± 0.042 |
| | 8 | MSE | 0.900 ± 0.014 | 0.792 ± 0.009 | 1.409 ± 0.369 | 1.352 ± 0.105 | **0.769 ± 0.006** | 1.177 ± 0.083 |
| | | MAE | 0.409 ± 0.008 | 0.395 ± 0.010 | 0.523 ± 0.051 | 0.491 ± 0.014 | **0.367 ± 0.004** | 0.664 ± 0.046 |
| | 10 | MSE | 0.954 ± 0.019 | 0.845 ± 0.009 | 1.523 ± 0.401 | 1.378 ± 0.048 | **0.828 ± 0.006** | 1.169 ± 0.096 |
| | | MAE | 0.441 ± 0.007 | 0.422 ± 0.012 | 0.561 ± 0.073 | 0.513 ± 0.009 | **0.395 ± 0.002** | 0.660 ± 0.050 |
| | 12 | MSE | 0.993 ± 0.008 | 0.897 ± 0.008 | 1.436 ± 0.173 | 1.528 ± 0.136 | **0.886 ± 0.003** | 1.231 ± 0.019 |
| | | MAE | 0.468 ± 0.004 | 0.446 ± 0.005 | 0.566 ± 0.034 | 0.561 ± 0.018 | **0.422 ± 0.003** | 0.693 ± 0.013 |
| Traffic | 6 | MSE | 0.152 ± 0.003 | **0.147 ± 0.003** | 0.149 ± 0.004 | 0.175 ± 0.007 | 0.151 ± 0.001 | 0.268 ± 0.015 |
| | | MAE | **0.202 ± 0.007** | 0.219 ± 0.003 | 0.203 ± 0.007 | 0.204 ± 0.005 | 0.254 ± 0.001 | 0.362 ± 0.009 |
| | 8 | MSE | 0.159 ± 0.004 | **0.147 ± 0.004** | 0.158 ± 0.003 | 0.177 ± 0.013 | 0.163 ± 0.001 | 0.283 ± 0.005 |
| | | MAE | 0.215 ± 0.005 | 0.222 ± 0.006 | **0.211 ± 0.002** | 0.218 ± 0.003 | 0.269 ± 0.002 | 0.376 ± 0.008 |
| | 10 | MSE | 0.165 ± 0.004 | **0.156 ± 0.002** | 0.164 ± 0.005 | 0.196 ± 0.007 | 0.176 ± 0.000 | 0.290 ± 0.007 |
| | | MAE | 0.224 ± 0.004 | 0.232 ± 0.010 | **0.218 ± 0.002** | 0.228 ± 0.004 | 0.282 ± 0.001 | 0.380 ± 0.007 |
| | 12 | MSE | 0.172 ± 0.005 | **0.159 ± 0.002** | 0.174 ± 0.008 | 0.195 ± 0.013 | 0.189 ± 0.001 | 0.304 ± 0.003 |
| | | MAE | 0.232 ± 0.006 | 0.234 ± 0.006 | **0.229 ± 0.008** | 0.238 ± 0.005 | 0.292 ± 0.001 | 0.391 ± 0.009 |

# K Pretrained Text-Embedding Ablation

This ablation compares BERT (768 dimensions; Devlin et al., 2019), GPT-2 (768 dimensions), and GPT-2-XL (1600 dimensions; Radford et al., 2019) as tokenizer and input-embedding sources. It uses DynaGuard with iTransformer on nine Time-MMD domains, four domain-specific horizons, and five seeds under the staged-training protocol. Embeddings remain frozen, and no contextual language-model layers are run. BERT and GPT-2 variants are independently trained; GPT-2-XL uses the matched main-comparison runs.

GPT-2-XL has the lowest aggregate MSE in all nine domains and the lowest aggregate MAE in seven. This comparison evaluates each tokenizer and pretrained embedding space together, including its dimensionality. Table 77 reports the aggregate results, and Section K.1 gives the complete horizon-specific results.

## K.1 Per-Horizon Pretrained Text-Embedding Ablation Results

These tables give the per-horizon results underlying Table 77, including the same GPT-2-XL reference used in the main comparison.

Table 77: Aggregate text-embedding-source ablation with iTransformer across nine domains. BERT and GPT-2 are independently trained variants; GPT-2-XL directly reuses the matched main DynaGuard runs. Entries are the 20-run mean ± the mean of four horizon-specific sample standard deviations over five seeds. All settings use only tokenizer outputs and input embedding tables, not full contextual language models.

| Dataset | Variant | iTransformer | |
|---|---|---|---|
| | | MSE | MAE |
| Agriculture | BERT | $0.285 \pm 0.010$ | $0.363 \pm 0.010$ |
| | GPT-2 | $0.285 \pm 0.010$ | $0.363 \pm 0.014$ |
| | GPT-2-XL | $\mathbf{0.254 \pm 0.014}$ | $\mathbf{0.332 \pm 0.013}$ |
| Climate | BERT | $0.972 \pm 0.007$ | $0.800 \pm 0.004$ |
| | GPT-2 | $0.972 \pm 0.007$ | $0.800 \pm 0.005$ |
| | GPT-2-XL | $\mathbf{0.894 \pm 0.004}$ | $\mathbf{0.751 \pm 0.002}$ |
| Economy | BERT | $0.185 \pm 0.005$ | $0.335 \pm 0.004$ |
| | GPT-2 | $0.185 \pm 0.004$ | $0.334 \pm 0.004$ |
| | GPT-2-XL | $\mathbf{0.180 \pm 0.005}$ | $\mathbf{0.330 \pm 0.006}$ |
| Energy | BERT | $0.317 \pm 0.006$ | $0.433 \pm 0.005$ |
| | GPT-2 | $0.314 \pm 0.008$ | $0.431 \pm 0.005$ |
| | GPT-2-XL | $\mathbf{0.241 \pm 0.012}$ | $\mathbf{0.355 \pm 0.008}$ |
| Environment | BERT | $0.266 \pm 0.001$ | $\mathbf{0.368 \pm 0.000}$ |
| | GPT-2 | $0.268 \pm 0.001$ | $0.370 \pm 0.003$ |
| | GPT-2-XL | $\mathbf{0.259 \pm 0.000}$ | $0.370 \pm 0.001$ |
| Health | BERT | $1.349 \pm 0.016$ | $0.760 \pm 0.004$ |
| | GPT-2 | $1.350 \pm 0.019$ | $0.761 \pm 0.004$ |
| | GPT-2-XL | $\mathbf{1.273 \pm 0.010}$ | $\mathbf{0.753 \pm 0.005}$ |
| Security | BERT | $75.735 \pm 1.463$ | $4.395 \pm 0.128$ |
| | GPT-2 | $74.944 \pm 2.369$ | $4.322 \pm 0.217$ |
| | GPT-2-XL | $\mathbf{71.734 \pm 0.644}$ | $\mathbf{3.963 \pm 0.097}$ |
| Social Good | BERT | $1.155 \pm 0.057$ | $0.484 \pm 0.016$ |
| | GPT-2 | $1.140 \pm 0.038$ | $0.479 \pm 0.016$ |
| | GPT-2-XL | $\mathbf{0.928 \pm 0.016}$ | $\mathbf{0.428 \pm 0.007}$ |
| Traffic | BERT | $0.171 \pm 0.003$ | $\mathbf{0.211 \pm 0.004}$ |
| | GPT-2 | $0.171 \pm 0.003$ | $\mathbf{0.211 \pm 0.004}$ |
| | GPT-2-XL | $\mathbf{0.162 \pm 0.004}$ | $0.218 \pm 0.005$ |

Table 78: Per-horizon text-embedding-source ablation results with iTransformer for Agriculture, Climate, Economy. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | BERT | GPT-2 | GPT-2-XL |
|---|---|---|---|---|---|
| Agriculture | 6 | MSE | $0.172 \pm 0.003$ | $0.172 \pm 0.005$ | $\mathbf{0.146 \pm 0.005}$ |
| | | MAE | $0.281 \pm 0.010$ | $0.281 \pm 0.006$ | $\mathbf{0.255 \pm 0.007}$ |
| | 8 | MSE | $0.242 \pm 0.009$ | $0.243 \pm 0.005$ | $\mathbf{0.210 \pm 0.018}$ |
| | | MAE | $0.338 \pm 0.012$ | $0.346 \pm 0.021$ | $\mathbf{0.304 \pm 0.012}$ |
| | 10 | MSE | $0.322 \pm 0.016$ | $0.326 \pm 0.011$ | $\mathbf{0.288 \pm 0.014}$ |
| | | MAE | $0.392 \pm 0.011$ | $0.390 \pm 0.013$ | $\mathbf{0.362 \pm 0.017}$ |
| | 12 | MSE | $0.406 \pm 0.014$ | $0.399 \pm 0.018$ | $\mathbf{0.372 \pm 0.019}$ |
| | | MAE | $0.443 \pm 0.007$ | $0.434 \pm 0.015$ | $\mathbf{0.406 \pm 0.016}$ |
| Climate | 6 | MSE | $0.964 \pm 0.004$ | $0.964 \pm 0.004$ | $\mathbf{0.884 \pm 0.006}$ |
| | | MAE | $0.798 \pm 0.003$ | $0.797 \pm 0.003$ | $\mathbf{0.750 \pm 0.002}$ |
| | 8 | MSE | $0.969 \pm 0.005$ | $0.968 \pm 0.005$ | $\mathbf{0.888 \pm 0.001}$ |
| | | MAE | $0.799 \pm 0.003$ | $0.797 \pm 0.004$ | $\mathbf{0.749 \pm 0.000}$ |
| | 10 | MSE | $0.978 \pm 0.011$ | $0.979 \pm 0.012$ | $\mathbf{0.899 \pm 0.007}$ |
| | | MAE | $0.802 \pm 0.004$ | $0.803 \pm 0.007$ | $\mathbf{0.752 \pm 0.004}$ |
| | 12 | MSE | $0.978 \pm 0.009$ | $0.978 \pm 0.008$ | $\mathbf{0.905 \pm 0.003}$ |
| | | MAE | $0.803 \pm 0.006$ | $0.802 \pm 0.005$ | $\mathbf{0.754 \pm 0.001}$ |
| Economy | 6 | MSE | $0.155 \pm 0.004$ | $0.155 \pm 0.004$ | $\mathbf{0.150 \pm 0.007}$ |
| | | MAE | $0.307 \pm 0.003$ | $0.307 \pm 0.004$ | $\mathbf{0.303 \pm 0.008}$ |
| | 8 | MSE | $0.174 \pm 0.005$ | $0.175 \pm 0.003$ | $\mathbf{0.169 \pm 0.006}$ |
| | | MAE | $0.325 \pm 0.004$ | $0.325 \pm 0.003$ | $\mathbf{0.322 \pm 0.006}$ |
| | 10 | MSE | $0.196 \pm 0.005$ | $0.195 \pm 0.005$ | $\mathbf{0.190 \pm 0.003}$ |
| | | MAE | $0.345 \pm 0.004$ | $0.344 \pm 0.005$ | $\mathbf{0.342 \pm 0.003}$ |
| | 12 | MSE | $0.215 \pm 0.005$ | $0.215 \pm 0.005$ | $\mathbf{0.209 \pm 0.005}$ |
| | | MAE | $0.362 \pm 0.005$ | $0.361 \pm 0.005$ | $\mathbf{0.356 \pm 0.006}$ |

Table 79: Per-horizon text-embedding-source ablation results with iTransformer for Energy, Environment, Health. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | BERT | GPT-2 | GPT-2-XL |
|---|---|---|---|---|---|
| Energy | 12 | MSE | $0.173 \pm 0.002$ | $0.173 \pm 0.003$ | $\mathbf{0.092 \pm 0.004}$ |
| | | MAE | $0.320 \pm 0.004$ | $0.319 \pm 0.002$ | $\mathbf{0.212 \pm 0.004}$ |
| | 24 | MSE | $0.269 \pm 0.006$ | $0.269 \pm 0.009$ | $\mathbf{0.197 \pm 0.016}$ |
| | | MAE | $0.405 \pm 0.005$ | $0.406 \pm 0.006$ | $\mathbf{0.329 \pm 0.010}$ |
| | 36 | MSE | $0.378 \pm 0.007$ | $0.367 \pm 0.010$ | $\mathbf{0.284 \pm 0.018}$ |
| | | MAE | $0.475 \pm 0.006$ | $0.470 \pm 0.005$ | $\mathbf{0.397 \pm 0.014}$ |
| | 48 | MSE | $0.447 \pm 0.009$ | $0.447 \pm 0.009$ | $\mathbf{0.392 \pm 0.008}$ |
| | | MAE | $0.531 \pm 0.006$ | $0.532 \pm 0.007$ | $\mathbf{0.481 \pm 0.005}$ |
| Environment | 48 | MSE | $\mathbf{0.262 \pm 0.000}$ | $0.263 \pm 0.000$ | $\mathbf{0.262 \pm 0.000}$ |
| | | MAE | $\mathbf{0.369 \pm 0.000}$ | $\mathbf{0.369 \pm 0.001}$ | $\mathbf{0.369 \pm 0.001}$ |
| | 96 | MSE | $\mathbf{0.265 \pm 0.000}$ | $0.266 \pm 0.001$ | $\mathbf{0.265 \pm 0.000}$ |
| | | MAE | $\mathbf{0.370 \pm 0.000}$ | $0.373 \pm 0.002$ | $0.371 \pm 0.001$ |
| | 192 | MSE | $0.272 \pm 0.000$ | $0.273 \pm 0.001$ | $\mathbf{0.256 \pm 0.001}$ |
| | | MAE | $\mathbf{0.367 \pm 0.001}$ | $0.371 \pm 0.005$ | $0.369 \pm 0.002$ |
| | 336 | MSE | $0.266 \pm 0.004$ | $0.269 \pm 0.001$ | $\mathbf{0.254 \pm 0.000}$ |
| | | MAE | $\mathbf{0.367 \pm 0.001}$ | $\mathbf{0.367 \pm 0.005}$ | $0.370 \pm 0.001$ |
| Health | 12 | MSE | $0.971 \pm 0.013$ | $0.976 \pm 0.017$ | $\mathbf{0.929 \pm 0.016}$ |
| | | MAE | $0.651 \pm 0.004$ | $0.653 \pm 0.005$ | $\mathbf{0.648 \pm 0.006}$ |
| | 24 | MSE | $1.346 \pm 0.015$ | $1.338 \pm 0.019$ | $\mathbf{1.307 \pm 0.012}$ |
| | | MAE | $0.760 \pm 0.002$ | $\mathbf{0.758 \pm 0.003}$ | $0.760 \pm 0.003$ |
| | 36 | MSE | $1.527 \pm 0.022$ | $1.530 \pm 0.023$ | $\mathbf{1.404 \pm 0.004}$ |
| | | MAE | $0.807 \pm 0.007$ | $0.808 \pm 0.007$ | $\mathbf{0.790 \pm 0.003}$ |
| | 48 | MSE | $1.550 \pm 0.017$ | $1.555 \pm 0.017$ | $\mathbf{1.451 \pm 0.009}$ |
| | | MAE | $0.824 \pm 0.003$ | $0.827 \pm 0.002$ | $\mathbf{0.812 \pm 0.009}$ |

Table 80: Per-horizon text-embedding-source ablation results with iTransformer for Security, Social Good, Traffic. Entries are mean ± standard deviation over five seeds.

| Domain | $H$ | Metric | BERT | GPT-2 | GPT-2-XL |
|---|---|---|---|---|---|
| Security | 6 | MSE | 72.029 ± 1.296 | 71.169 ± 2.406 | **66.815 ± 0.430** |
| | | MAE | 4.385 ± 0.112 | 4.296 ± 0.196 | **3.847 ± 0.103** |
| | 8 | MSE | 74.696 ± 2.052 | 74.285 ± 2.745 | **69.877 ± 0.567** |
| | | MAE | 4.410 ± 0.189 | 4.383 ± 0.248 | **3.864 ± 0.107** |
| | 10 | MSE | 77.562 ± 2.056 | 76.265 ± 2.942 | **73.225 ± 0.708** |
| | | MAE | 4.440 ± 0.157 | 4.316 ± 0.259 | **3.993 ± 0.088** |
| | 12 | MSE | 78.651 ± 0.445 | 78.058 ± 1.382 | **77.021 ± 0.872** |
| | | MAE | 4.346 ± 0.055 | 4.293 ± 0.166 | **4.147 ± 0.091** |
| Social Good | 6 | MSE | 1.023 ± 0.039 | 0.999 ± 0.010 | **0.864 ± 0.025** |
| | | MAE | 0.424 ± 0.016 | 0.415 ± 0.010 | **0.394 ± 0.010** |
| | 8 | MSE | 1.133 ± 0.037 | 1.150 ± 0.056 | **0.900 ± 0.014** |
| | | MAE | 0.473 ± 0.009 | 0.477 ± 0.016 | **0.409 ± 0.008** |
| | 10 | MSE | 1.211 ± 0.065 | 1.173 ± 0.039 | **0.954 ± 0.019** |
| | | MAE | 0.501 ± 0.015 | 0.495 ± 0.022 | **0.441 ± 0.007** |
| | 12 | MSE | 1.254 ± 0.088 | 1.237 ± 0.045 | **0.993 ± 0.008** |
| | | MAE | 0.538 ± 0.024 | 0.527 ± 0.014 | **0.468 ± 0.004** |
| Traffic | 6 | MSE | 0.161 ± 0.001 | 0.161 ± 0.001 | **0.152 ± 0.003** |
| | | MAE | 0.194 ± 0.003 | **0.193 ± 0.003** | 0.202 ± 0.007 |
| | 8 | MSE | 0.169 ± 0.004 | 0.168 ± 0.003 | **0.159 ± 0.004** |
| | | MAE | **0.207 ± 0.005** | **0.207 ± 0.003** | 0.215 ± 0.005 |
| | 10 | MSE | 0.174 ± 0.004 | 0.174 ± 0.004 | **0.165 ± 0.004** |
| | | MAE | **0.217 ± 0.004** | **0.217 ± 0.004** | 0.224 ± 0.004 |
| | 12 | MSE | 0.182 ± 0.004 | 0.182 ± 0.004 | **0.172 ± 0.005** |
| | | MAE | **0.227 ± 0.006** | 0.228 ± 0.005 | 0.232 ± 0.006 |

### K.2 Additional FNSPID and GPT4MTS Results

For FNSPID and GPT4MTS, we add Qwen3.5-4B (2560 dimensions) and Qwen3.5-9B (4096 dimensions) (Qwen Team, 2026). Each of the five sources supplies its tokenizer and frozen input embedding table at the model's native width, without contextual language-model layers. Within each task, horizon, and seed, all sources use the same numerical backbone, data version, DynaGuard components, and fixed training configuration. That configuration is separate from the selected main-comparison configuration, so the GPT-2-XL reference differs from the main results.

Tables 81 and 82 report all four FNSPID stocks and three GPT4MTS targets over five seeds. FNSPID averages four horizons; GPT4MTS uses $H = 7$. These fixed-configuration ablations are separate from the main-comparison summaries and statistical tests.

Tables 83 and 84 give the FNSPID per-horizon results for all five sources. Table 82 also serves as the per-horizon table for GPT4MTS.

Table 81: Text-embedding-source ablation with DynaGuard and iTransformer on FNSPID. Entries are the 20-run mean ± the mean of four horizon-specific sample standard deviations over five seeds ($H \in \{6, 8, 10, 12\}$). GPT-2-XL reuses the formal main-comparison runs.

| Dataset | Embedding source | MSE | MAE |
|---|---|---|---|
| WMT | BERT | $\mathbf{0.091 \pm 0.002}$ | $0.238 \pm 0.002$ |
| | GPT-2 | $\mathbf{0.091 \pm 0.002}$ | $\mathbf{0.237 \pm 0.002}$ |
| | GPT-2-XL | $\mathbf{0.091 \pm 0.002}$ | $\mathbf{0.237 \pm 0.003}$ |
| | Qwen3.5-4B | $0.093 \pm 0.004$ | $0.240 \pm 0.004$ |
| | Qwen3.5-9B | $0.092 \pm 0.002$ | $\mathbf{0.237 \pm 0.003}$ |
| GILD | BERT | $0.060 \pm 0.002$ | $0.172 \pm 0.003$ |
| | GPT-2 | $0.061 \pm 0.002$ | $0.174 \pm 0.004$ |
| | GPT-2-XL | $\mathbf{0.057 \pm 0.002}$ | $\mathbf{0.168 \pm 0.003}$ |
| | Qwen3.5-4B | $0.062 \pm 0.005$ | $0.176 \pm 0.008$ |
| | Qwen3.5-9B | $0.061 \pm 0.002$ | $0.174 \pm 0.003$ |
| WFC | BERT | $0.369 \pm 0.008$ | $0.398 \pm 0.005$ |
| | GPT-2 | $0.370 \pm 0.011$ | $0.398 \pm 0.006$ |
| | GPT-2-XL | $\mathbf{0.362 \pm 0.013}$ | $\mathbf{0.395 \pm 0.009}$ |
| | Qwen3.5-4B | $0.372 \pm 0.011$ | $0.400 \pm 0.007$ |
| | Qwen3.5-9B | $0.376 \pm 0.014$ | $0.400 \pm 0.007$ |
| SBUX | BERT | $1.208 \pm 0.086$ | $0.697 \pm 0.025$ |
| | GPT-2 | $1.218 \pm 0.071$ | $0.698 \pm 0.020$ |
| | GPT-2-XL | $\mathbf{1.091 \pm 0.032}$ | $\mathbf{0.667 \pm 0.010}$ |
| | Qwen3.5-4B | $1.188 \pm 0.061$ | $0.690 \pm 0.025$ |
| | Qwen3.5-9B | $1.201 \pm 0.085$ | $0.694 \pm 0.026$ |

Table 82: Text-embedding-source ablation with DynaGuard and iTransformer on GPT4MTS at $H = 7$. Entries are mean ± sample standard deviation over five seeds; these are also the complete per-horizon results. GPT-2-XL reuses the formal main-comparison runs.

| Dataset | Embedding source | MSE | MAE |
|---|---|---|---|
| NumArticles | BERT | $0.653 \pm 0.006$ | $0.651 \pm 0.005$ |
| | GPT-2 | $0.651 \pm 0.007$ | $0.648 \pm 0.004$ |
| | GPT-2-XL | $\mathbf{0.635 \pm 0.019}$ | $\mathbf{0.638 \pm 0.013}$ |
| | Qwen3.5-4B | $0.650 \pm 0.006$ | $0.647 \pm 0.004$ |
| | Qwen3.5-9B | $0.637 \pm 0.017$ | $0.639 \pm 0.013$ |
| NumMentions | BERT | $0.667 \pm 0.005$ | $0.658 \pm 0.005$ |
| | GPT-2 | $0.667 \pm 0.005$ | $0.658 \pm 0.004$ |
| | GPT-2-XL | $0.654 \pm 0.016$ | $0.648 \pm 0.010$ |
| | Qwen3.5-4B | $0.663 \pm 0.006$ | $0.655 \pm 0.005$ |
| | Qwen3.5-9B | $\mathbf{0.644 \pm 0.011}$ | $\mathbf{0.642 \pm 0.009}$ |
| NumSources | BERT | $0.527 \pm 0.010$ | $0.579 \pm 0.007$ |
| | GPT-2 | $0.524 \pm 0.008$ | $0.577 \pm 0.005$ |
| | GPT-2-XL | $\mathbf{0.505 \pm 0.017}$ | $\mathbf{0.569 \pm 0.012}$ |
| | Qwen3.5-4B | $0.527 \pm 0.015$ | $0.580 \pm 0.008$ |
| | Qwen3.5-9B | $0.523 \pm 0.018$ | $0.582 \pm 0.009$ |

Table 83: Per-horizon text-embedding-source ablation with DynaGuard and iTransformer on FNSPID (WMT, GILD). Entries are mean ± sample standard deviation over five seeds. GPT-2-XL reuses the formal main-comparison runs.

| Dataset | $H$ | Metric | BERT | GPT-2 | GPT-2-XL | Qwen3.5-4B | Qwen3.5-9B |
|---|---|---|---|---|---|---|---|
| WMT | 6 | MSE | **0.063 ± 0.001** | **0.063 ± 0.001** | **0.063 ± 0.001** | 0.064 ± 0.001 | **0.063 ± 0.001** |
| | | MAE | **0.201 ± 0.001** | **0.201 ± 0.001** | **0.201 ± 0.002** | 0.203 ± 0.001 | **0.201 ± 0.002** |
| | 8 | MSE | **0.082 ± 0.001** | **0.082 ± 0.002** | **0.082 ± 0.003** | **0.082 ± 0.003** | **0.082 ± 0.002** |
| | | MAE | 0.229 ± 0.002 | **0.228 ± 0.002** | 0.229 ± 0.004 | 0.229 ± 0.004 | **0.228 ± 0.003** |
| | 10 | MSE | 0.102 ± 0.002 | 0.102 ± 0.002 | **0.100 ± 0.001** | 0.106 ± 0.006 | 0.103 ± 0.003 |
| | | MAE | 0.252 ± 0.002 | 0.252 ± 0.002 | **0.250 ± 0.001** | 0.256 ± 0.007 | 0.253 ± 0.002 |
| | 12 | MSE | **0.118 ± 0.003** | **0.118 ± 0.003** | 0.119 ± 0.004 | 0.122 ± 0.005 | **0.118 ± 0.004** |
| | | MAE | 0.269 ± 0.002 | 0.269 ± 0.003 | 0.270 ± 0.004 | 0.273 ± 0.006 | **0.268 ± 0.004** |
| GILD | 6 | MSE | 0.048 ± 0.000 | 0.048 ± 0.000 | **0.047 ± 0.002** | 0.051 ± 0.006 | 0.048 ± 0.001 |
| | | MAE | 0.152 ± 0.001 | 0.152 ± 0.001 | **0.151 ± 0.004** | 0.157 ± 0.010 | 0.153 ± 0.001 |
| | 8 | MSE | 0.055 ± 0.002 | 0.056 ± 0.003 | **0.052 ± 0.001** | 0.055 ± 0.002 | 0.055 ± 0.002 |
| | | MAE | 0.165 ± 0.003 | 0.167 ± 0.006 | **0.161 ± 0.001** | 0.166 ± 0.003 | 0.165 ± 0.003 |
| | 10 | MSE | 0.063 ± 0.002 | 0.064 ± 0.001 | **0.062 ± 0.004** | 0.068 ± 0.010 | 0.066 ± 0.002 |
| | | MAE | **0.177 ± 0.003** | 0.180 ± 0.001 | **0.177 ± 0.005** | 0.186 ± 0.015 | 0.184 ± 0.003 |
| | 12 | MSE | 0.074 ± 0.003 | 0.075 ± 0.005 | **0.066 ± 0.002** | 0.074 ± 0.003 | 0.074 ± 0.004 |
| | | MAE | 0.195 ± 0.004 | 0.195 ± 0.006 | **0.183 ± 0.003** | 0.195 ± 0.004 | 0.195 ± 0.006 |

Table 84: Per-horizon text-embedding-source ablation with DynaGuard and iTransformer on FNSPID (WFC, SBUX). Entries are mean ± sample standard deviation over five seeds. GPT-2-XL reuses the formal main-comparison runs.

| Dataset | $H$ | Metric | BERT | GPT-2 | GPT-2-XL | Qwen3.5-4B | Qwen3.5-9B |
|---|---|---|---|---|---|---|---|
| WFC | 6 | MSE | 0.246 ± 0.004 | 0.247 ± 0.007 | **0.242 ± 0.004** | 0.249 ± 0.006 | 0.247 ± 0.007 |
| | | MAE | **0.334 ± 0.002** | 0.335 ± 0.004 | **0.334 ± 0.005** | 0.336 ± 0.005 | 0.335 ± 0.005 |
| | 8 | MSE | 0.328 ± 0.009 | 0.326 ± 0.011 | **0.317 ± 0.021** | 0.328 ± 0.014 | 0.330 ± 0.014 |
| | | MAE | 0.382 ± 0.006 | 0.381 ± 0.007 | **0.378 ± 0.016** | 0.382 ± 0.007 | 0.381 ± 0.009 |
| | 10 | MSE | 0.406 ± 0.011 | 0.406 ± 0.009 | **0.398 ± 0.019** | 0.409 ± 0.012 | 0.416 ± 0.014 |
| | | MAE | 0.418 ± 0.005 | 0.419 ± 0.003 | **0.415 ± 0.008** | 0.422 ± 0.006 | 0.423 ± 0.007 |
| | 12 | MSE | 0.495 ± 0.010 | 0.500 ± 0.016 | **0.491 ± 0.009** | 0.500 ± 0.013 | 0.509 ± 0.019 |
| | | MAE | 0.456 ± 0.007 | 0.458 ± 0.010 | **0.454 ± 0.006** | 0.459 ± 0.010 | 0.460 ± 0.009 |
| SBUX | 6 | MSE | 0.856 ± 0.045 | 0.859 ± 0.029 | **0.787 ± 0.031** | 0.890 ± 0.107 | 0.845 ± 0.058 |
| | | MAE | 0.601 ± 0.023 | 0.596 ± 0.014 | **0.564 ± 0.006** | 0.610 ± 0.043 | 0.596 ± 0.023 |
| | 8 | MSE | 1.127 ± 0.058 | 1.094 ± 0.066 | **0.959 ± 0.038** | 1.099 ± 0.047 | 1.126 ± 0.077 |
| | | MAE | 0.666 ± 0.017 | 0.660 ± 0.019 | **0.637 ± 0.010** | 0.662 ± 0.016 | 0.668 ± 0.024 |
| | 10 | MSE | 1.320 ± 0.067 | 1.347 ± 0.048 | **1.199 ± 0.022** | 1.293 ± 0.041 | 1.308 ± 0.051 |
| | | MAE | 0.727 ± 0.016 | 0.733 ± 0.016 | **0.706 ± 0.013** | 0.713 ± 0.025 | 0.726 ± 0.014 |
| | 12 | MSE | 1.530 ± 0.176 | 1.571 ± 0.141 | **1.418 ± 0.039** | 1.469 ± 0.051 | 1.526 ± 0.154 |
| | | MAE | 0.793 ± 0.043 | 0.802 ± 0.030 | **0.763 ± 0.011** | 0.777 ± 0.017 | 0.786 ± 0.042 |