

XLinear: A Lightweight and Accurate MLP-Based Model for Long-Term Time Series Forecasting with Exogenous Inputs

Xinyang Chen^{1*}, Huidong Jin^{2*}, Yu Huang^{1,3,4†}, Zaiwen Feng^{1,3,4†}

¹College of Informatics, Huazhong Agricultural University, Wuhan, Hubei, China

²Statistical Machine Learning, Data61, CSIRO, Canberra Australia

³Hubei Key Laboratory of Agricultural Bioinformatics, Wuhan, Hubei, China

⁴Engineering Research Center of Agricultural Intelligent Technology, Ministry of Education, China
Warren.Jin@csiro.au, {YHuang, Zaiwen.Feng}@mail.hzau.edu.cn

Abstract

Despite the prevalent assumption of uniform variable importance in long-term time series forecasting models, real-world applications often exhibit asymmetric causal relationships and varying data acquisition costs. Specifically, cost-effective exogenous data (e.g., local weather) can unilaterally influence dynamics of endogenous variables, such as lake surface temperature. Exploiting these links enables more effective forecasts when exogenous inputs are readily available. Transformer-based models capture long-range dependencies but incur high computation and suffer from permutation invariance. Patch-based variants improve efficiency yet can miss local temporal patterns. To efficiently exploit informative signals across both the temporal dimension and relevant exogenous variables, this study proposes XLinear, a lightweight time series forecasting model built upon Multi-Layer Perceptrons (MLPs). XLinear uses a global token derived from an endogenous variable as a pivotal hub for interacting with exogenous variables, and employs MLPs with sigmoid activation to extract both temporal patterns and variate-wise dependencies. Its prediction head then integrates these signals to forecast the endogenous series. We evaluate XLinear on seven standard benchmarks and five real-world datasets with exogenous inputs. Compared with state-of-the-art models, XLinear delivers superior accuracy and efficiency for both multivariate forecasts and univariate forecasts influenced by exogenous inputs.

Code — <https://github.com/Zaiwen/XLinear.git>

Introduction

Time series forecasting has shown strong value across domains such as financial analysis (Liu et al. 2023a), irrigation scheduling (Shao et al. 2019, 2025), population flows (Bakar and Jin 2018), yield forecast (Jin et al. 2022), and environmental science (Genova et al. 2025; Bakar, Kokic, and Jin 2015, 2016; Kokic, Jin, and Crimp 2013; Jin and Henderson 2011). The predictability of time series data lies in two main aspects: persistent temporal patterns and cross-variate information (Chen et al. 2023). In practice, actual

*These authors contributed equally.

†Corresponding authors.

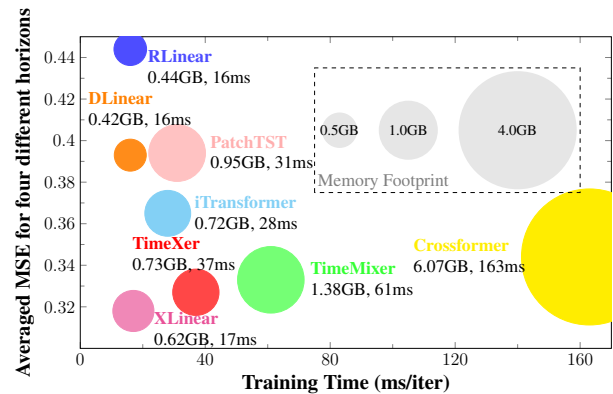


Figure 1: Model performance on the Electricity dataset for univariate forecasting with 320 exogenous variables. TimesNet and TiDE are excluded for clarity due to slow training.

forecasts require modeling both the endogenous time series and informative exogenous inputs, which often drive the target dynamics rather than vice versa. For instance, crop growth for a given genotype is influenced by environmental and management factors over time (Jin et al. 2022), whereas lake surface or crop canopy temperature is primarily driven by local weather (Shao et al. 2019, 2025). With the advent of deep learning, the field has witnessed transformative advancements, particularly with Transformer (Vaswani et al. 2017) emerging as the prevailing modeling paradigm. Leveraging its multi-head self-attention mechanism, Transformer excels at capturing long-range dependencies in sequential data, establishing itself as a benchmark for long-term time series forecasting. Early Transformer-based models (Li et al. 2019; Zhou et al. 2021; Wu et al. 2021; Zhou et al. 2022) focused on optimizing the self-attention mechanism while paying little attention to cross-variable dependencies. By assuming channel-independence, PatchTST (Nie et al. 2023) leveraged patch-based attention to gain high efficiency. Crossformer (Zhang and Yan 2023) and iTransformer (Liu et al. 2023b) focused on cross-channel modeling but both have limitations in temporal modeling: Crossformer’s unaligned patch interactions may introduce noise, while iTransformer weakly captures

temporal dependencies. The recent TimeXer (Wang et al. 2024b) model established a new state-of-the-art (SOTA) benchmark for both multivariate and univariate time series forecasting with exogenous drivers, by effectively extracting information from endogenous time series via self-attention and from exogenous drives via cross-attention. However, patch-based Transformer models excessively rely on the patching mechanism to achieve desirable performance, which limits their applicability in forecasting tasks unsuitable for patching (Luo and Wang 2024). Additionally, their permutation-invariant self-attention mechanism suffers from temporal information loss (Tang and Zhang 2025). Inspired by its global tokens for endogenous variables, we will propose a more efficient and effective model in this study.

Since the emergence of lightweight models for long-term time series forecasting (Zeng et al. 2023), doubts have been raised about the actual effectiveness of Transformer, which also inspired our research and model design. Among them, TSMixer (Chen et al. 2023) models both temporal and cross-variable dependencies simultaneously, but its forecast accuracy does not show significant improvement over PatchTST, indicating its cross-variable modeling limitation. DLinear (Zeng et al. 2023), TimeMixer (Wang et al. 2024a), and xPatch (Stitsyuk and Choi 2025) decompose the seasonal-trend components in the temporal dimension of the sequence and capture complex temporal patterns through specific processing methods. Although these models can capture temporal dependencies, for real-world forecast applications, crucial information from exogenous drives cannot be ignored.

Pre-trained foundation time series models have been made impressive progress in the last few years. Almost all of them are Transformer-based (Kottapalli et al. 2025). Their time series forecast accuracy remains limited on new and unseen data sets or some in-domain data sets (Shi et al. 2025), comparing with full-shot SOTA models like TimeXer and TimeMixer. There are various other deep models in the literature, such as those based on convolution mechanisms (Wu et al. 2022), and graph neural networks (GNN) (Yi et al. 2023; Cai et al. 2024). Their forecast accuracy is not as good as TimeXer (Wang et al. 2024b) and our proposed model.

Inspired by the advances of TimeXer and lightweight models, this study proposes XLinear, a lightweight model that integrates a Multi-Layer Perceptron (MLP) with a sigmoid activation function as a gating module. Unlike TimeXer, which employs patch-level self-attention and variate-level cross-attention to capture temporal and inter-variable dependencies, XLinear utilizes unified gating modules that sequentially filter features along the temporal and variable dimensions, achieving more precise and efficient dependency modeling. The prediction head employs a cross-channel fusion mechanism to synergistically integrate bidimensional information, effectively mitigating noise arising from interactions between features of different dimensions. Although XLinear comprises only two sets of MLP-based feature extractors, it strikes a strong balance between accuracy and efficiency. As shown in Fig. 1, in the benchmark test on the Electricity dataset (with a unified batch size of 4), XLinear matches the training speed of the simple yet ef-

fective DLinear and RLinear, consumes less memory than SOTA models, and delivers the highest forecast accuracy among all 10 comparative models.

The core contributions of this study are as follows:

Bridging the Efficiency-Accuracy Gap: We address the critical challenge in time series forecasting, especially with exogenous inputs, of balancing computational efficiency (a strength of MLP-based models) with high forecasting accuracy (a hallmark of patch-based Transformers). Our work directly tackles this by explicitly incorporating exogenous inputs and effectively reconciling complex temporal and cross-variable dependencies.

Introducing XLinear, a Novel MLP-Based Model: Its core innovation lies in a novel gating module that uses an MLP with a sigmoid activation function that enables XLinear to efficiently boost predictive accuracy by fully leveraging informative exogenous signals through cross-variable dependencies, facilitated by global tokens derived from endogenous sequences.

Demonstrating Superior Performance: Through extensive tests on 12 diverse datasets, XLinear consistently outperforms SOTA models in both accuracy and efficiency. Notably, XLinear achieves at least 30% faster training speeds than these efficient Transformers.

Method

Previously, MLP-based models have demonstrated strong capabilities in capturing internal time series dependencies. Building on this foundation, we design an MLP-based gating module to efficiently capture dependencies across both temporal and variable dimensions, enabling selective feature filtering. To effectively integrate crucial information across these dimensions, inspired by TimeXer (Wang et al. 2024b), we introduce learnable global tokens for each endogenous variable, serving as hubs for interactions with exogenous variables and thereby facilitating causal information transfer. The separation of global and temporal tokens effectively mitigates cross-variable noise, while their integration enhances feature representation and forecasting performance. The architecture of XLinear is illustrated in Fig. 2.

Forward Process

Given the historical observation sequence of the endogenous variables $X_{1:T} = \{X_{1:T}^{(1)}, X_{1:T}^{(2)}, \dots, X_{1:T}^{(M)}\}$, where $X_{1:T}^{(j)}$ denotes the sequence of the j -th endogenous variable, and the sequence of exogenous variables up to time T , denoted as $E_{1:T} = \{E_{1:T}^{(1)}, E_{1:T}^{(2)}, \dots, E_{1:T}^{(C)}\}$, where $E_{1:T}^{(i)}$ represents the sequence of the i -th exogenous variable, the model is designed to predict the values of the endogenous variable sequence for the subsequent S time steps. The entire process can be described as:

$$\hat{X}_{(T+1):(T+S)} = \mathbf{F}(X_{1:T}, E_{1:T}) \in \mathbb{R}^{M \times S}. \quad (1)$$

$\mathbf{F}(\cdot)$ is a mapping learned by the model during training, enabling the conversion of input sequences to S -step-ahead predictions.

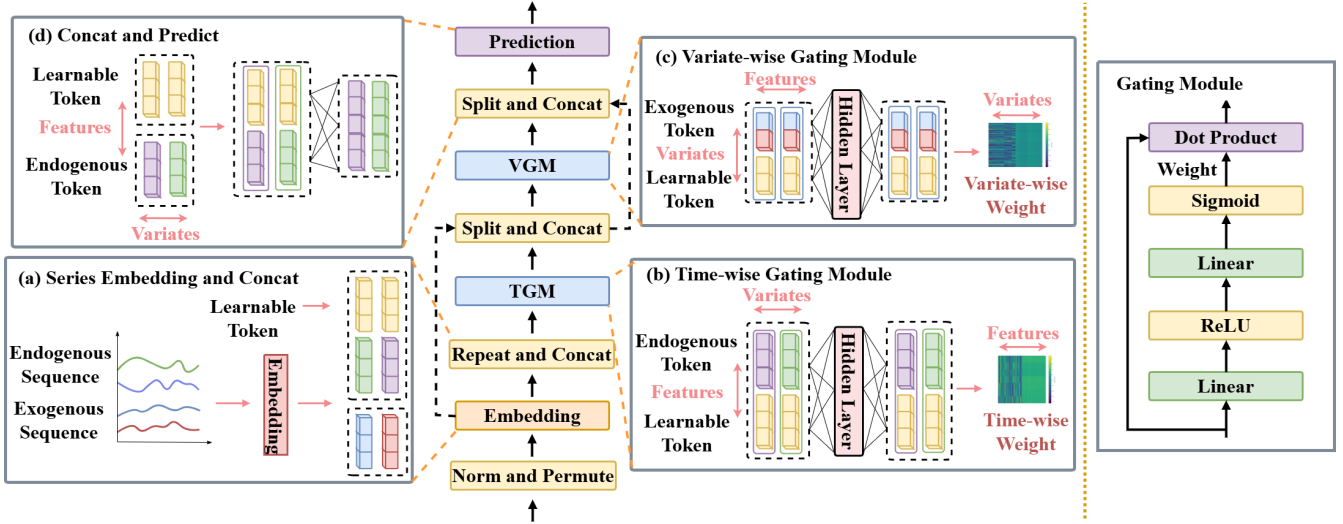


Figure 2: Architecture of XLinear. The left panel illustrates the core processing pipeline: (a) Synchronous processing of endogenous and exogenous sequences via a joint embedding layer, with learnable global representation tokens introduced for endogenous variables; (b) Time-wise Gating Module (TGM) enhances key temporal features and suppresses noise in the endogenous sequences while transferring critical temporal characteristics to the global tokens; (c) Variate-wise Gating Module (VGM) leverages the global tokens of the endogenous sequences to interact with the exogenous sequences across dimensions, extracting implicit cross-variable dependencies; (d) Cross-variable information carried by each global token is concatenated with its associated temporally enhanced endogenous sequence, followed by generating time-series predictions for endogenous variables via the prediction head. The right gating module illustrates the implementation method of TGM and VGM.

Embedding

To capture latent features in time-series data, we jointly embed the RevIN-processed (Kim et al. 2021) endogenous sequence $X_{1:T}$ and exogenous sequence $E_{1:T}$. The process is:

$$X_{\text{endo}}, E_{\text{exo}} = \text{Embedding}(X_{1:T}, E_{1:T}). \quad (2)$$

where $\text{Embedding}(\cdot)$ denotes a parameterized embedding function. To extract relevant information from the exogenous sequence, we introduce a learnable global context token X_{glob} , which is concatenated with the endogenous embedding along the temporal dimension to form an enhanced representation:

$$X_{\text{endo.tok}} = \text{Concat}_t(X_{\text{endo}}, X_{\text{glob}}). \quad (3)$$

where $\text{Concat}_t(\cdot)$ denotes concatenation along the time dimension.

Time-wise Gating Module (TGM)

TGM is designed to extract salient temporal patterns from historical endogenous sequences, as future values largely follow these dynamics. It employs an MLP to capture time-wise dependencies and map endogenous sequence information to the global token, generating time-wise gating weights that are subsequently passed through a sigmoid function to constrain them to a stable range, enabling selective feature filtering. TGM is formulated as:

$$[X'_{\text{endo}}, X'_{\text{glob}}] = \sigma(\text{Linear}_2(\phi(\text{Linear}_1(X_{\text{endo.tok}}))) \odot X_{\text{endo.tok}}. \quad (4)$$

where Linear_1 and Linear_2 are learnable linear layers, $\phi(\cdot)$ (ReLU) and $\sigma(\cdot)$ (sigmoid) are the activation functions, and \odot denotes element-wise multiplication.

Variate-wise Gating Module (VGM)

VGM is designed to capture cross-variable information. It constructs dependency associations between endogenous variables' global tokens (obtained in the previous step) and the exogenous sequences. By doing so, it maps relevant information from exogenous variables into the endogenous global tokens, thereby optimizing the prediction of the endogenous variables. This step is formulated as:

$$X_{\text{exo.tok}} = \text{Concat}_c(E_{\text{exo}}, X'_{\text{glob}}), \quad (5)$$

$$[E'_{\text{exo}}, X''_{\text{glob}}] = \sigma(\text{Linear}_4(\phi(\text{Linear}_3(X_{\text{exo.tok}})))) \odot X_{\text{exo.tok}}. \quad (6)$$

where $\text{Concat}_c(\cdot)$ denotes concatenation along the channel dimension.

Prediction Head

XLinear uses a fully connected layer as its final prediction head. It directly maps the integrated features to the target sequence by fusing the temporal characteristics of endogenous variables with key information extracted from exogenous variables. The computation follows:

$$\hat{X}_{(T+1):(T+S)} = \text{FC}(\text{Concat}_t(X'_{\text{endo}}, X''_{\text{glob}})). \quad (7)$$

where $\text{FC}(\cdot)$ denotes the fully connected layer.

Loss Function

We employ Mean Square Error (MSE) as the loss function to quantify the deviation between predicted values and ground truth. Specifically, losses generated from individual endogenous channels are aggregated and then averaged across all M endogenous variable sequences to obtain the overall loss. The calculation formula is as follows:

$$\mathcal{L} = \mathbb{E}_{\mathbf{X}} \frac{1}{M} \sum_{i=1}^M \left\| \hat{\mathbf{X}}_{(T+1):(T+S)}^{(i)} - \mathbf{X}_{(T+1):(T+S)}^{(i)} \right\|_2^2. \quad (8)$$

where $\mathbb{E}_{\mathbf{X}}$ denotes the expectation over time steps of \mathbf{X} , representing the average squared L_2 -norm error between predictions and true values for from time steps $T + 1$ to $T + S$.

Multivariate Forecasting

In multivariate forecasting, each variable serves a dual purpose: it acts as an endogenous variable to be predicted and concurrently as an exogenous variable, providing variate-wise information to help predict other variables. To handle complex interaction, the TGM is shared across all variables, the VGM across features.

Experimental Results and Comparisons

Experimental Settings

Datasets. We used 7 widely adopted datasets (ETT with 4 sub-datasets (Zhou et al. 2021), Weather, Electricity, Traffic (Wu et al. 2021)), along with five datasets with exogenous inputs: one for Crop yield and four for environmental monitoring. The Crop dataset came from Kaggle, two Dissolved Oxygen (DO) datasets, DO₄₂₅₀₁₂ and DO₄₀₉₂₁₅, were from Darling River (u/s Weir 32) and the Murray River (Tocumwa), Australia (Genova et al. 2025). Two water temperature datasets GTD_N and GTD_S were collected from the northern and southern monitoring sites of Grahamstown Dam, Australia, respectively¹. Their metadata are presented in Table 1.

Baselines. We compare against SOTA models spanning multiple architectures. We include top-performing Transformer variants – TimeXer (Wang et al. 2024b), which achieved the best results in the Long-term Forecasting Look-Back-96 scenario, and supports exogenous inputs, iTransformer (Liu et al. 2023b), PatchTST (Nie et al. 2023), and Crossformer (Zhang and Yan 2023), as well as the CNN-based TimesNet (Wu et al. 2022). For linear-based models, we benchmark TimeMixer (Wang et al. 2024a), RLinear (Li et al. 2023), TiDE (Das et al. 2023), and DLinear (Zeng et al. 2023). Finally, we compare with four leading GNN-based models (Huang et al. 2023; Cai et al. 2024; Wu et al. 2020; Yi et al. 2023).

Implementation Details. We reproduced the baselines based on the official repositories provided by each model to ensure a fair comparison. For TiDE, we followed the implementation provided in the Time Series Library² for

¹Please contact Warren Jin for data sets (Huang et al. 2025).

²<https://github.com/thuml/Time-Series-Library>

replication. All experiments were implemented using the PyTorch framework. The forecasting tasks on the PEMS07 and PEMS08 datasets (Table 5) were conducted on NVIDIA Quadro RTX 8000 (48 GB) GPU, while all other experiments were performed on NVIDIA Tesla V100-PCIE (32 GB) GPU. Following common practice in prior studies, we fixed the random seed to 2025 to ensure experimental reproducibility and control stochasticity across Python, NumPy, and PyTorch environments. For model optimization, the ADAM optimizer was employed. The initial learning rate was selected from the range $\{1 \times 10^{-4}, 2 \times 10^{-4}, 3 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$ based on the specific dataset used. The learning rate adjustment during training followed this rule: it remained constant for the first 3 epochs, and from the end of the 3rd epoch onwards, it was decayed by a factor of 0.9 for each subsequent epoch. Formally, the learning rate adjustment strategy can be expressed as:

$$\text{lr_adjust}(e) = \begin{cases} \lambda_{\text{init}}, & e < 3 \\ \lambda_{\text{init}} \times 0.9^{(e-3)}, & e \geq 3 \end{cases}$$

where e denotes the current training epoch and λ_{init} represents the initial learning rate.

Considering the variations across different datasets, the batch size for training each dataset is selected from $\{4, 8, 16, 32, 64, 128, 256\}$. During the training process, the number of training epochs is set to 30 in most cases, and an early stopping strategy is adopted. Detailed settings for each dataset can be found in the corresponding scripts.

Unless otherwise specified, the input length for all datasets was set to 96. The prediction horizon S was selected from $\{96, 192, 336, 720\}$, except for the Crop dataset, where it was chosen from $\{12, 24, 36, 48\}$. For a time series of length T (comprising d_x endogenous variables and d_z exogenous variables) with the batch size set as $batch_size$, the embedding layer projects both types of variable sequences into the same dimension d_{model} , where the length of the globally learnable tokens for endogenous variables is also d_{model} . In the temporal and channel gating modules, the hidden layer dimensions of the MLP are defined as t_{ff} and c_{ff} , respectively, and the feature dimension remains unchanged after sigmoid processing. After capturing temporal dependencies, both the endogenous variable sequences and their global tokens have a dimension of $(batch_size, d_x, d_{model})$. To mitigate overfitting and enhance generalization, dropout regularization is introduced in the embedding layer, TGM, VGM, and prediction head, with dropout rates denoted as $embed_dropout$, $t_dropout$, $c_dropout$, and $head_dropout$, respectively. The model’s prediction head maps the concatenated endogenous sequence tensor (shape $(batch_size, d_x, 2 * d_{model})$) to the next S time steps (shape $(batch_size, d_x, S)$) in a channel-independent manner.

Evaluation Metrics. For the seven public datasets, we follow the evaluation conventions in existing studies and adopt MSE and Mean Absolute Error (MAE) as evaluation metrics. For the Crop dataset and environmental datasets, we use additional evaluation metrics such as Nash-Sutcliffe Efficiency (NSE), Kling-Gupta Efficiency (KGE), and Mean

Datasets	Electricity	Weather	ETTh1	ETTh2	ETTm1	ETTm2	Traffic	Crop	DO ₄₀₉₂₁₅	DO ₄₂₅₀₁₂	GTD _N	GTD _S
Features	321	21	7	7	7	7	862	8	7	7	16	16
Timesteps	26304	52696	17420	17420	69680	69680	17544	8914	103827	77336	38593	28727
Granularity	1 hour	10 min	1 hour	1 hour	15 min	15 min	1 hour	5 min	15 min	15 min	1 hour	1 hour

Table 1: Dataset characteristics: number of variables and sequence length, and temporal granularity (the first 7 are benchmarks).

Models	XLiner		TimeXer		TimeMixer		iTransformer		RLinear		PatchTST		Crossformer		TiDE		TimesNet		DLinear		
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
Electricity	96	0.256	0.359	<u>0.261</u>	0.366	0.275	0.384	0.299	0.403	0.433	0.480	0.339	0.412	0.265	<u>0.364</u>	0.405	0.459	0.342	0.437	0.387	0.451
	192	0.292	0.380	<u>0.316</u>	0.397	<u>0.294</u>	<u>0.385</u>	0.321	0.413	0.407	0.461	0.361	0.425	0.313	0.390	0.383	0.442	0.384	0.461	0.365	0.436
	336	0.339	0.409	<u>0.367</u>	0.429	<u>0.344</u>	<u>0.419</u>	0.379	0.446	0.440	0.481	0.393	0.440	0.380	0.431	0.418	0.464	0.439	0.493	0.391	0.453
	720	<u>0.385</u>	<u>0.451</u>	0.365	0.439	0.417	0.465	0.461	0.504	0.495	0.523	0.482	0.507	0.418	0.463	0.471	0.507	0.473	0.514	0.428	0.487
Weather	96	0.001	<u>0.026</u>	<u>0.001</u>	0.027	<u>0.001</u>	0.028	<u>0.001</u>	<u>0.026</u>	0.001	0.025	<u>0.001</u>	0.027	0.004	0.048	0.001	0.025	0.002	0.029	0.006	0.062
	192	0.001	<u>0.029</u>	<u>0.002</u>	0.030	<u>0.002</u>	0.031	<u>0.002</u>	<u>0.029</u>	0.001	0.028	<u>0.002</u>	0.030	0.005	0.053	0.001	0.028	<u>0.002</u>	0.031	0.006	0.066
	336	0.002	<u>0.030</u>	<u>0.002</u>	0.031	<u>0.002</u>	0.031	<u>0.002</u>	0.031	0.002	0.029	<u>0.002</u>	0.032	0.004	0.051	0.002	0.029	<u>0.002</u>	0.031	0.006	0.068
	720	0.002	<u>0.034</u>	<u>0.002</u>	0.036	<u>0.002</u>	0.036	<u>0.002</u>	0.036	0.002	0.033	<u>0.002</u>	0.036	0.007	0.067	0.002	0.033	0.381	0.368	0.007	0.070
ETTh1	96	0.055	0.178	0.057	0.181	<u>0.056</u>	0.181	0.057	0.183	0.059	0.185	0.055	0.178	0.133	0.297	0.059	0.184	0.059	0.188	0.065	0.188
	192	0.071	0.202	0.071	<u>0.204</u>	<u>0.072</u>	<u>0.204</u>	0.074	0.209	0.078	0.214	<u>0.072</u>	0.206	0.232	0.409	0.078	0.214	0.080	0.217	0.088	0.222
	336	0.084	0.226	0.080	0.223	0.088	0.230	0.084	0.223	0.093	0.240	0.087	0.231	0.244	0.423	0.093	0.240	<u>0.083</u>	<u>0.224</u>	0.110	0.257
	720	0.083	0.227	<u>0.084</u>	<u>0.229</u>	0.086	0.230	<u>0.084</u>	<u>0.229</u>	0.106	0.256	0.098	0.247	0.530	0.660	0.104	0.255	0.083	0.231	0.202	0.371
ETTh2	96	0.130	0.277	<u>0.132</u>	<u>0.280</u>	0.134	0.282	0.137	0.287	0.136	0.286	0.136	0.285	0.261	0.413	0.136	0.285	0.159	0.310	0.135	0.282
	192	0.180	0.331	<u>0.181</u>	<u>0.333</u>	0.186	0.340	0.187	0.341	0.187	0.339	0.185	0.337	1.240	1.028	0.187	0.339	0.196	0.351	0.188	0.335
	336	0.209	0.365	0.223	0.377	0.222	0.378	0.221	0.376	0.231	0.384	<u>0.217</u>	<u>0.373</u>	0.974	0.874	0.231	0.384	0.232	0.385	0.238	0.385
	720	0.217	0.373	<u>0.220</u>	<u>0.376</u>	0.223	0.379	0.223	0.379	0.253	0.403	0.267	0.417	0.229	0.384	1.633	1.177	0.267	0.417	0.254	0.403
ETTm1	96	0.028	0.125	0.028	0.125	<u>0.029</u>	0.126	<u>0.029</u>	0.128	0.030	0.129	<u>0.029</u>	<u>0.126</u>	0.171	0.355	0.030	0.129	<u>0.029</u>	0.128	0.034	0.135
	192	0.043	0.158	0.043	0.158	<u>0.044</u>	<u>0.160</u>	0.045	0.163	<u>0.044</u>	<u>0.160</u>	<u>0.045</u>	<u>0.160</u>	0.293	0.474	<u>0.044</u>	<u>0.160</u>	<u>0.044</u>	<u>0.160</u>	0.055	0.173
	336	0.057	0.183	<u>0.058</u>	0.185	0.059	0.186	0.060	0.190	0.057	<u>0.184</u>	<u>0.058</u>	<u>0.184</u>	0.330	0.503	0.057	<u>0.184</u>	0.061	0.190	0.078	0.210
	720	0.079	0.216	0.079	<u>0.217</u>	0.082	0.218	0.079	0.218	<u>0.080</u>	<u>0.217</u>	0.082	0.221	0.852	0.861	<u>0.080</u>	<u>0.217</u>	0.083	0.223	0.098	0.234
ETTm2	96	0.064	0.182	<u>0.067</u>	<u>0.188</u>	0.068	0.189	0.071	0.194	0.074	0.199	0.068	0.188	0.149	0.309	0.073	0.199	0.073	0.200	0.072	0.195
	192	0.098	0.233	0.101	0.236	<u>0.100</u>	<u>0.234</u>	0.108	0.247	0.104	0.241	<u>0.100</u>	0.236	0.686	0.740	0.104	0.241	0.106	0.247	0.105	0.240
	336	<u>0.129</u>	<u>0.274</u>	0.130	0.275	0.132	0.278	0.140	0.288	0.131	0.276	0.128	0.271	0.546	0.602	0.131	0.276	0.150	0.296	0.136	0.280
	720	0.179	0.329	0.182	<u>0.332</u>	0.183	<u>0.332</u>	0.188	0.340	<u>0.180</u>	0.329	0.185	0.335	2.524	1.424	<u>0.180</u>	0.329	0.186	0.338	0.191	0.335
Traffic	96	0.139	0.216	<u>0.151</u>	<u>0.224</u>	0.170	0.256	0.156	0.236	0.350	0.431	0.176	0.253	0.154	0.230	0.350	0.430	0.154	0.249	0.268	0.351
	192	0.140	0.214	<u>0.152</u>	<u>0.229</u>	0.162	0.248	0.156	0.237	0.314	0.404	0.162	0.243	0.180	0.256	0.230	0.315	0.164	0.255	0.302	0.387
	336	0.143	0.221	<u>0.150</u>	<u>0.232</u>	0.159	0.251	0.154	0.243	0.305	0.399	0.164	0.248	-	-	0.220	0.208	0.167	0.259	0.298	0.384
	720	0.165	0.245	<u>0.172</u>	<u>0.253</u>	0.186	0.274	0.177	0.268	0.328	0.415	0.189	0.267	-	-	0.243	0.329	0.197	0.292	0.340	0.416
1st	25	21	6	4	0	0	1	1	5	5	2	2	0	0	5	5	0	0	0	0	

Table 2: Results of univariate forecasting with exogenous inputs. “-” denotes the occurrence of an Out of Memory (OOM) exception during model execution. The best results are in **bold** and the second best are underlined. All baseline results except those of TimeMixer are reported in (Wang et al. 2024b).

Absolute Percentage Error (MAPE) which are widely employed in environmental sciences, enabling a more holistic assessment of forecast model performance (Gupta et al. 2009). They are defined as: $NSE = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$ where y_i denotes the observed value, \hat{y}_i represents the predicted value, \bar{y} is the mean of observed values, and n is the number of samples; $KGE = 1 - \sqrt{(r - 1)^2 + (\alpha - 1)^2 + (\beta - 1)^2}$ where r is the Pearson correlation coefficient between observed and predicted values, variability ratio α is the ratio of the standard deviations of the predicted and the observed values, and bias ratio β indicates the ratio of the means of the predicted and the observed values; $MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100\%$. Higher NSE and KGE values indicate better performance, while lower values are preferred for the other metrics.

Main Results

Univariate Forecast with Exogenous inputs. In each dataset, for univariate forecast performance assessment, the final variable is designated as the endogenous variable, while the remaining variables are treated as exogenous variables, following common practice (Wang et al. 2024b). For the 5 environmental datasets, their last variable consistently serves as the endogenous one. In the Crop dataset, yield is driven by soil moisture, air temperature, humidity, and etc. In the water quality dataset, DO₄₂₅₀₁₂ and DO₄₀₉₂₁₅, dissolved oxygen content is guided by water temperature, mean discharge rate, and mean water level, among other variables. In the GTD_N and GTD_S datasets, the bottom water temperature at 9m depth is affected by shortwave radiation, air temperature, cloud cover, wind speed and direction, etc.

The forecast MSE and MAE results for the 7 benchmarks are listed in Table 2. XLiner achieves the lowest

Models	XLinear					TimeXer					iTransformer					PatchTST					
	Metric	MSE	MAE	NSE	KGE	MAPE	MSE	MAE	NSE	KGE	MAPE	MSE	MAE	NSE	KGE	MAPE	MSE	MAE	NSE	KGE	MAPE
DO ₄₂₅₀₁₂	96	0.113	0.207	0.862	0.930	0.451	<u>0.116</u>	<u>0.212</u>	<u>0.858</u>	<u>0.928</u>	0.460	0.117	<u>0.212</u>	0.857	0.927	<u>0.452</u>	0.117	<u>0.212</u>	<u>0.858</u>	0.925	0.466
	192	0.212	0.288	0.741	0.869	<u>0.653</u>	0.215	0.288	0.738	0.863	0.650	<u>0.214</u>	<u>0.289</u>	<u>0.740</u>	<u>0.867</u>	<u>0.653</u>	<u>0.214</u>	0.290	0.739	0.865	0.662
	336	<u>0.340</u>	0.379	<u>0.586</u>	0.789	0.935	0.341	0.379	0.585	<u>0.786</u>	0.924	0.337	0.380	0.590	0.789	<u>0.930</u>	0.345	0.382	0.580	0.783	0.948
	720	0.555	0.531	0.329	<u>0.649</u>	<u>1.526</u>	0.561	0.537	0.322	0.653	1.540	<u>0.556</u>	0.531	<u>0.328</u>	0.648	1.512	0.562	<u>0.536</u>	0.320	0.645	1.552
DO ₄₀₉₂₁₅	96	0.003	0.037	0.986	0.992	1.014	0.003	<u>0.038</u>	<u>0.985</u>	0.964	<u>1.022</u>	0.003	<u>0.038</u>	0.984	0.953	1.043	0.003	0.041	0.983	<u>0.976</u>	1.027
	192	0.006	0.054	0.970	0.982	<u>1.413</u>	0.006	0.054	<u>0.969</u>	0.930	1.418	0.006	<u>0.055</u>	0.968	0.926	1.428	<u>0.007</u>	0.059	0.965	<u>0.957</u>	1.404
	336	0.011	0.077	0.942	0.967	1.683	0.012	<u>0.079</u>	<u>0.939</u>	0.909	1.774	<u>0.012</u>	<u>0.079</u>	0.938	0.882	1.738	0.013	0.081	0.936	<u>0.953</u>	1.748
	720	0.023	0.111	0.880	0.913	2.397	0.027	<u>0.118</u>	0.861	0.844	<u>2.603</u>	0.027	0.121	0.859	0.818	2.854	<u>0.026</u>	0.119	<u>0.862</u>	<u>0.904</u>	2.820
GTD _N	96	0.008	0.062	0.991	<u>0.864</u>	0.455	<u>0.009</u>	0.064	<u>0.990</u>	0.781	<u>0.458</u>	<u>0.009</u>	<u>0.063</u>	<u>0.990</u>	0.851	0.455	0.010	0.068	0.989	0.898	0.481
	192	0.018	0.092	0.980	0.861	0.726	<u>0.019</u>	0.096	0.978	0.775	0.790	0.018	<u>0.094</u>	<u>0.979</u>	<u>0.792</u>	0.751	0.020	0.100	0.977	0.747	<u>0.740</u>
	336	0.033	0.127	0.962	0.824	<u>0.864</u>	0.036	0.133	0.958	0.741	0.917	<u>0.035</u>	<u>0.130</u>	<u>0.959</u>	<u>0.801</u>	0.855	0.037	0.138	0.957	0.661	0.933
	720	<u>0.093</u>	<u>0.221</u>	<u>0.889</u>	<u>0.567</u>	<u>1.651</u>	0.098	0.226	0.883	0.441	1.658	0.062	0.180	0.926	0.749	1.225	0.104	0.239	0.877	0.362	1.837
GTD _S	96	0.012	0.077	0.971	<u>0.983</u>	0.554	<u>0.013</u>	<u>0.079</u>	<u>0.970</u>	0.978	<u>0.558</u>	<u>0.013</u>	0.081	0.969	0.984	0.602	<u>0.013</u>	0.082	0.968	0.982	0.591
	192	0.025	0.112	0.936	<u>0.965</u>	0.729	0.025	<u>0.114</u>	0.936	0.966	<u>0.748</u>	<u>0.027</u>	0.118	0.933	0.964	0.786	<u>0.027</u>	0.118	0.931	0.961	0.755
	336	0.048	0.159	<u>0.869</u>	<u>0.926</u>	<u>0.941</u>	<u>0.049</u>	0.159	0.870	<u>0.928</u>	0.921	0.051	0.164	0.863	0.930	1.002	0.050	<u>0.162</u>	0.864	0.923	0.945
	720	<u>0.124</u>	<u>0.260</u>	<u>0.621</u>	<u>0.795</u>	<u>1.550</u>	0.121	0.258	0.634	0.816	<u>1.462</u>	0.125	0.268	<u>0.622</u>	0.782	1.370	0.131	0.268	0.600	0.785	1.631
Crop	12	0.147	0.236	<u>0.529</u>	<u>0.714</u>	0.982	0.143	0.245	0.541	0.689	<u>0.940</u>	0.153	0.251	0.510	0.719	0.991	<u>0.144</u>	<u>0.244</u>	<u>0.529</u>	0.705	0.920
	24	0.186	<u>0.284</u>	0.408	<u>0.612</u>	1.071	0.190	0.296	0.392	0.606	<u>1.019</u>	0.208	0.307	0.336	0.607	1.131	0.186	0.280	<u>0.394</u>	0.616	1.018
	36	0.219	<u>0.311</u>	0.303	<u>0.542</u>	1.137	0.224	0.327	<u>0.287</u>	0.547	<u>1.092</u>	0.247	0.343	0.212	0.448	1.205	<u>0.220</u>	0.310	<u>0.287</u>	0.539	1.077
	48	0.245	0.336	0.220	<u>0.433</u>	<u>1.175</u>	<u>0.251</u>	0.356	<u>0.203</u>	0.485	1.185	0.274	0.367	0.129	0.388	1.245	<u>0.252</u>	<u>0.338</u>	0.188	<u>0.462</u>	1.100
1 st Count	16	16	15	9	8	5	5	4	5	3	5	2	2	5	5	2	2	0	2	5	

Table 3: Results of single endogenous variable prediction based on exogenous variables for each model on environmental datasets. The best results are in **bold** and the second best are underlined.

errors in 89.3% (25/28) of experimental configurations for MSE, and 75.0% for MAE. It consistently ranks among the top two models, except at the 336-step lead time for ETTh1 – outperforming strong competitors such as TimeXer and TimeMixer in most cases. In addition to its accuracy, XLinear also demonstrates high efficiency. As illustrated in Fig 1, under relatively consistent parameter settings, among all baseline models, XLinear achieves at least 39.3% faster training, 13.9% lower memory usage, and 2.8% lower MSE than the others—except DLinear and RLinear, whose MSEs are around 23.5% and 39.6% higher than XLinear, respectively.

We further evaluate XLinear on the five real-world applications. Due to space constraints, we compare it with three representative models, as listed in Table 3. XLinear achieves the lowest MSE and MAE in over 75% of the cases. For the composite metric NSE, it ranks first in 75% of the rows, and second in 20%. Additionally, XLinear reaches the “Good” performance level (NSE > 0.65, per (Moriassi et al. 2007)) in 65% of the cases. For the KGE and MAPE metrics, XLinear ranks among the top two models in 90% and 70% of the cases, respectively. Overall, XLinear demonstrates consistent advantages across multiple evaluation metrics, confirming its effectiveness in real-world forecasting scenarios.

Long-term Multivariate Forecasting. The forecasting results for MSE and MAE across 7 benchmark datasets are presented in Table 4. On the first six datasets, XLinear achieves the best performance in over 91.7% of the cases for both metrics, and ranks second in the remaining cases.

The only exception is on the Traffic dataset with 862 variables, where iTransformer outperforms XLinear due to its variable-wise attention mechanism, which is more effective for high-dimensional forecasting tasks. In this scenario, each

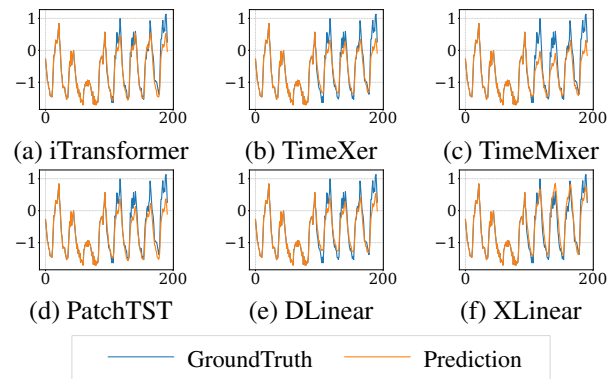


Figure 3: Typical forecast results of XLinear and 5 SOTA baselines on the last variable of the Electricity dataset (Input and prediction horizons are 96).

variable simultaneously serves as a prediction target and an exogenous driver, leading to a substantial expansion of the VGM input dimensionality. Extracting informative signals while suppressing noise in such high-dimensional spaces poses a significant challenge. Future work will focus on reducing the VGM input dimensionality while preserving essential inter-variable dependencies, thereby enhancing the model’s scalability and performance on high-dimensional datasets. Averaging over the four forecast horizons, XLinear has the lowest MSE and MAE on the first six data sets, and ranks second in MSE and third in MAE on Traffic. Across all 7 benchmarks, XLinear consistently outperforms the four GNN-based models (in the appendix).

Furthermore, we validate the effectiveness of the XLin-

Models		XLiner		TimeXer		TimeMixer		iTransformer		RLinear		PatchTST		Crossformer		TiDE		TimesNet		DLinear	
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Electricity	96	0.138	0.237	<u>0.140</u>	0.242	0.153	0.247	0.148	<u>0.240</u>	0.201	0.281	0.195	0.285	0.219	0.314	0.237	0.329	0.168	0.272	0.197	0.282
	192	0.156	0.251	<u>0.157</u>	0.256	0.166	0.256	0.162	<u>0.253</u>	0.201	0.283	0.199	0.289	0.231	0.322	0.236	0.330	0.184	0.289	0.196	0.285
	336	0.171	0.267	<u>0.176</u>	0.275	0.185	0.277	0.178	<u>0.269</u>	0.215	0.298	0.215	0.305	0.246	0.337	0.249	0.344	0.198	0.300	0.209	0.301
	720	0.206	0.297	<u>0.211</u>	<u>0.306</u>	0.225	0.310	0.225	0.317	0.257	0.331	0.256	0.337	0.280	0.363	0.284	0.373	0.220	0.320	0.245	0.333
Weather	96	0.149	0.198	<u>0.157</u>	<u>0.205</u>	0.163	0.209	0.174	0.214	0.192	0.232	0.177	0.218	0.158	0.230	0.202	0.261	0.172	0.220	0.196	0.255
	192	0.201	0.244	<u>0.204</u>	<u>0.247</u>	0.208	0.250	0.221	0.254	0.240	0.271	0.225	0.259	0.206	0.277	0.242	0.298	0.219	0.261	0.237	0.296
	336	0.259	0.288	0.261	0.290	0.251	0.287	0.278	0.296	0.292	0.307	0.278	0.297	0.272	0.335	0.287	0.335	0.280	0.306	0.283	0.335
	720	0.339	0.340	<u>0.340</u>	<u>0.341</u>	0.339	<u>0.341</u>	0.358	0.349	0.364	0.353	0.354	0.348	0.398	0.418	0.351	0.386	0.365	0.359	0.345	0.381
ETTh1	96	0.369	0.393	0.382	0.403	<u>0.375</u>	0.400	0.386	0.405	0.386	<u>0.395</u>	0.414	0.419	0.423	0.448	0.479	0.464	0.384	0.402	0.386	0.400
	192	0.421	<u>0.423</u>	<u>0.429</u>	0.435	<u>0.429</u>	0.421	0.441	0.436	0.437	0.424	0.460	0.445	0.471	0.474	0.525	0.492	0.436	0.429	0.437	0.432
	336	0.455	0.440	<u>0.468</u>	0.448	0.484	0.458	0.487	0.458	0.479	<u>0.446</u>	0.501	0.466	0.570	0.546	0.565	0.515	0.491	0.460	0.481	0.459
	720	0.453	0.456	<u>0.469</u>	<u>0.461</u>	0.498	0.482	0.503	0.491	0.481	0.470	0.500	0.488	0.653	0.621	0.594	0.558	0.521	0.509	0.519	0.516
ETTm2	96	0.286	0.337	0.286	<u>0.338</u>	0.289	<u>0.341</u>	0.297	0.349	<u>0.288</u>	<u>0.338</u>	0.302	0.348	0.745	0.584	0.400	0.440	0.340	0.374	0.333	0.387
	192	0.363	0.388	0.363	<u>0.389</u>	<u>0.372</u>	0.392	0.380	0.400	0.374	<u>0.390</u>	0.388	0.400	0.877	0.656	0.528	0.509	0.402	0.414	0.477	0.476
	336	0.378	0.407	0.414	0.423	<u>0.386</u>	<u>0.414</u>	0.428	0.432	0.415	0.426	0.426	0.433	1.043	0.731	0.643	0.571	0.452	0.452	0.594	0.541
	720	0.408	0.431	0.408	<u>0.432</u>	<u>0.412</u>	0.434	0.427	0.445	0.420	0.440	0.431	0.446	1.104	0.763	0.874	0.679	0.462	0.468	0.831	0.657
ETTm1	96	0.311	0.351	<u>0.318</u>	<u>0.356</u>	0.320	0.357	0.334	0.368	0.355	0.376	0.329	0.367	0.404	0.426	0.364	0.387	0.338	0.375	0.345	0.372
	192	0.353	0.376	0.362	0.383	<u>0.361</u>	<u>0.381</u>	0.387	0.391	0.391	0.392	0.367	0.385	0.450	0.451	0.398	0.404	0.374	0.387	0.380	0.389
	336	0.382	0.398	0.395	0.407	<u>0.390</u>	<u>0.404</u>	0.426	0.420	0.424	0.415	0.399	0.410	0.532	0.515	0.428	0.425	0.410	0.411	0.413	0.413
	720	0.444	0.436	<u>0.452</u>	0.441	0.454	0.441	0.491	0.459	0.487	0.450	0.454	<u>0.439</u>	0.666	0.589	0.487	0.461	0.478	0.450	0.474	0.453
ETTm2	96	0.167	0.250	<u>0.171</u>	<u>0.256</u>	0.175	0.258	0.180	0.264	0.182	0.265	0.175	0.259	0.287	0.366	0.207	0.305	0.187	0.267	0.193	0.292
	192	0.233	0.293	<u>0.237</u>	<u>0.299</u>	<u>0.237</u>	<u>0.299</u>	0.250	0.309	0.246	0.304	0.241	0.302	0.414	0.492	0.290	0.364	0.249	0.309	0.284	0.362
	336	0.290	0.332	<u>0.296</u>	<u>0.338</u>	0.298	0.340	0.311	0.348	0.307	0.342	0.305	0.343	0.597	0.542	0.377	0.422	0.321	0.351	0.369	0.427
	720	0.388	0.390	0.392	<u>0.394</u>	<u>0.391</u>	0.396	0.412	0.407	0.407	0.398	0.402	0.400	1.730	1.042	0.558	0.524	0.408	0.403	0.554	0.522
Traffic	96	0.432	0.281	<u>0.428</u>	<u>0.271</u>	0.462	0.285	0.395	0.268	0.649	0.389	0.462	0.295	0.522	0.290	0.805	0.493	0.593	0.321	0.650	0.396
	192	0.449	0.289	<u>0.448</u>	<u>0.282</u>	0.473	0.296	0.417	0.276	0.601	0.366	0.466	0.296	0.530	0.293	0.756	0.474	0.617	0.336	0.598	0.370
	336	<u>0.465</u>	0.295	0.473	<u>0.289</u>	0.498	0.296	0.433	0.283	0.609	0.369	0.482	0.304	0.558	0.305	0.762	0.477	0.629	0.336	0.605	0.373
	720	0.507	0.315	<u>0.516</u>	<u>0.307</u>	<u>0.506</u>	0.313	0.467	0.302	0.647	0.387	0.514	0.322	0.589	0.328	0.719	0.449	0.640	0.350	0.645	0.394
1 st Count	23	22	3	0	2	2	4	4	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 4: Multivariate time-series forecasting results. The best results are in **bold** and the second best are underlined. The results of TimeMixer are reported in (Wang et al. 2024a), while those of other baselines are reported in (Wang et al. 2024b).

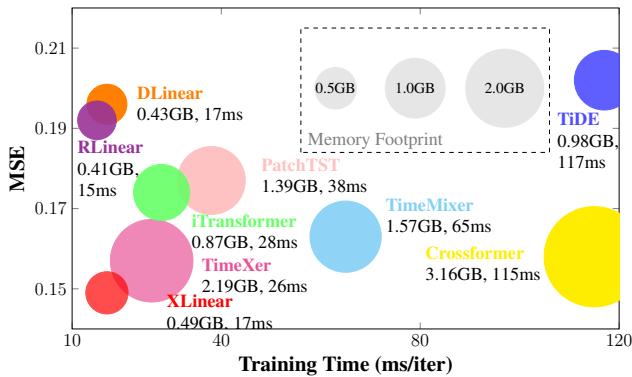


Figure 4: Model analysis of multivariate prediction on the Weather dataset.

ear using multivariate forecasting, rather than univariate, to more clearly visualise forecast differences. As illustrated in Fig. 3, the model’s predictions closely align with the ground truth, confirming its ability to effectively integrate multivariate information for accurate forecasting. We also assessed XLiner in a multivariate prediction setting using the Weather dataset. The experiment adopted a consistent training batch size of 128, with both the input and output win-

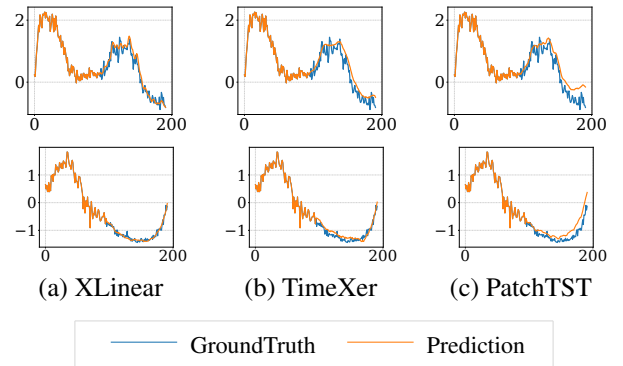


Figure 5: Comparison of prediction performance of various models on PEMS03 (Input and prediction horizons are 96).

dows set to 96, while the remaining parameters remained at their default values. As shown in Fig 4, XLiner demonstrates strong training efficiency and predictive accuracy. Except DLinear and RLinear, whose MSEs are 31.5% and 28.9% higher respectively, XLiner achieves at least 39.3% faster training, 43.7% lower memory usage, and 5.1% lower MSE than the other time series models in the figure.

Models		XLinear		TimeXer		PatchTST	
Metrics		MSE	MAE	MSE	MAE	MSE	MAE
PEMS03	12	0.030	0.128	<u>0.033</u>	<u>0.136</u>	0.036	0.139
	24	0.044	0.152	<u>0.050</u>	<u>0.168</u>	0.061	0.170
	48	0.075	0.190	<u>0.086</u>	<u>0.217</u>	0.125	0.233
	96	<u>0.133</u>	0.243	0.129	<u>0.264</u>	0.226	0.321
PEMS04	12	0.047	0.163	<u>0.051</u>	<u>0.170</u>	0.055	0.173
	24	0.056	0.178	<u>0.064</u>	<u>0.192</u>	0.069	0.194
	48	0.072	0.198	<u>0.086</u>	<u>0.219</u>	0.105	0.241
	96	0.108	0.246	<u>0.134</u>	<u>0.272</u>	0.180	0.309
PEMS07	12	0.075	0.200	<u>0.078</u>	<u>0.208</u>	0.093	0.220
	24	0.094	0.225	<u>0.096</u>	<u>0.234</u>	0.133	0.264
	48	<u>0.128</u>	<u>0.266</u>	0.124	0.264	0.208	0.332
	96	0.146	0.289	<u>0.158</u>	<u>0.304</u>	0.287	0.380
PEMS08	12	0.153	0.277	<u>0.158</u>	<u>0.281</u>	0.195	0.307
	24	0.189	0.305	<u>0.200</u>	<u>0.317</u>	0.262	0.357
	48	0.238	0.350	<u>0.258</u>	<u>0.361</u>	0.378	0.442
	96	0.287	0.402	<u>0.322</u>	<u>0.425</u>	0.490	0.580

Table 5: Main results of the PEMS forecasting task.

Model Analysis

Comparison with Patch-based Models. To mitigate the inherent limitations of the attention mechanism (Zeng et al. 2023), many efficient time series forecasting models have adopted patch-based attention mechanisms. While this approach effectively preserves local semantic information, it reduces sensitivity to fine-grained positional cues, making it challenging to model rapid fluctuations. To test this hypothesis, we conducted comparative experiments with TimeXer and PatchTST on the highly volatile PEMS datasets, under the setting of single endogenous variable forecasting with exogenous variables. Two representative examples are visualized in Fig. 5, and the detailed results are provided in Table 5, where XLinear demonstrates superior accuracy across most scenarios. All models can capture the overall trend reasonably well. However, in terms short-term fluctuations, XLinear more precisely forecasts fine-grained changes in the sequences. This advantage stems from its direct modeling of individual dimensions after data embedding, whereas the patching operations in patch-based models may hinder the learning of high-resolution temporal details.

Weight Visualization. To investigate the interpretability of XLinear, we selected a representative scenario dataset DE (a subset of the EPF dataset) (Lago et al. 2021) and used exogenous variables such as Wind power and Ampirion zonal load to predict Electricity price. Figure 6 presents a sample cases and their corresponding weights: (a)-(c) the original sequences, which show a significant correlation between Wind power and Electricity price across multiple time intervals; (e) the variable-wise weight distribution, where Channel 1 (representing Wind power) and Channel 3 (representing Electricity price) exhibit higher weights, consistent with the correlations observed in the input sequences; (d) Temporal weight distribution: In this process, XLinear identifies and enhances important features in the original sequences while establishing mapping relationships with global tokens.

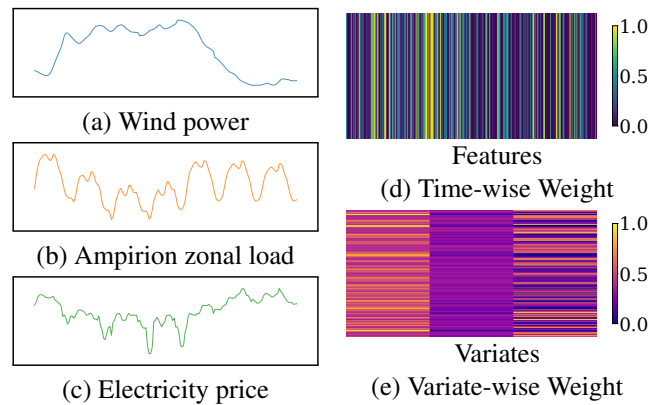


Figure 6: Visualization results of input samples and weights of XLinear on the DE dataset.

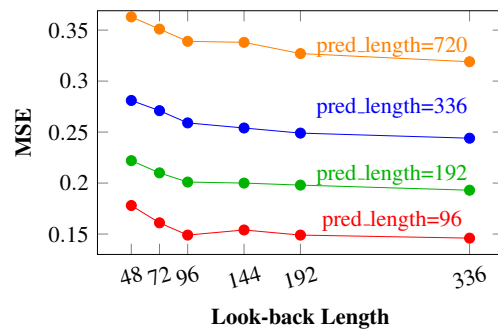


Figure 7: Visualization results of input samples and weights of XLinear on the DE dataset.

Overall, these observations demonstrate that XLinear can effectively highlight key variables and critical time steps, providing clear interpretability of its predictions.

Long lookback window. Our proposed MLP-based model, XLinear, with its “time-step-dependent” captured weights, can efficiently extract valid information from longer sequences without temporal attention. As shown in Fig. 7, its forecast MSEs for different prediction lengths drop rapidly as the lookback length increases to 96, and then decreases gradually with further increases. This indicates that XLinear, like most models, benefits from longer input sequences.

Furthermore, we investigate how model efficiency varies with the lookback window. On the Weather dataset, the experimental results are shown in Fig. 8, with the training batch size fixed at 32 and the forecast horizon set to 96. These demonstrate that XLinear outperforms all baseline models in terms of training time and memory usage, except being comparable to DLinear and RLinear. Notably, its training time and GPU memory consumption grow slowly and approximately linearly with the lookback window, whereas the resource demands of TimeMixer and PatchTST increase at least quadratically. This clearly demonstrates the high computational efficiency of XLinear when handling long input sequences.

To assess the effectiveness of XLinear in long-input sce-




Figure 8: Model efficiency vs. lookback window for multivariate forecasting on the Weather dataset.

Models	XLinear	TSMixer	PacthTST				
Metrics	MSE	MAE	MSE	MAE	MSE	MAE	
ETTh1	96	0.357	0.389	<u>0.361</u>	<u>0.392</u>	0.370	0.400
	192	0.393	0.411	<u>0.404</u>	<u>0.418</u>	0.413	0.429
	336	0.397	0.425	<u>0.420</u>	<u>0.431</u>	0.422	0.440
	720	<u>0.451</u>	<u>0.470</u>	0.463	0.472	0.447	0.468
	720	<u>0.451</u>	<u>0.470</u>	0.463	0.472	0.447	0.468
ETTh2	96	0.258	0.328	<u>0.274</u>	<u>0.341</u>	<u>0.274</u>	<u>0.337</u>
	192	0.314	0.367	<u>0.339</u>	<u>0.385</u>	<u>0.341</u>	<u>0.382</u>
	336	0.320	0.382	<u>0.361</u>	<u>0.406</u>	<u>0.329</u>	<u>0.384</u>
	720	<u>0.385</u>	<u>0.427</u>	0.445	0.470	0.379	0.422
	720	<u>0.385</u>	<u>0.427</u>	0.445	0.470	0.379	0.422
ETTm1	96	0.280	0.336	<u>0.285</u>	<u>0.339</u>	0.293	0.346
	192	0.325	0.364	<u>0.327</u>	<u>0.365</u>	0.333	0.370
	336	0.354	0.382	<u>0.356</u>	<u>0.382</u>	0.369	0.392
	720	0.415	0.414	0.419	0.414	<u>0.416</u>	<u>0.420</u>
	720	0.415	0.414	0.419	0.414	<u>0.416</u>	<u>0.420</u>
Weather	96	0.142	0.195	<u>0.145</u>	<u>0.198</u>	0.149	<u>0.198</u>
	192	0.185	0.235	<u>0.191</u>	<u>0.242</u>	0.194	<u>0.241</u>
	336	0.237	0.276	<u>0.242</u>	<u>0.280</u>	0.245	<u>0.282</u>
	720	0.311	0.330	<u>0.320</u>	<u>0.336</u>	<u>0.314</u>	<u>0.334</u>
	720	0.311	0.330	<u>0.320</u>	<u>0.336</u>	<u>0.314</u>	<u>0.334</u>

Table 6: Performance comparison with a lookback window length of 512.

narios, we conducted more comparative experiments with state-of-the-art models such as TSMixer and PatchTST, using a lookback window of length 512 for multivariate forecasting tasks. As presented in Table 6, for forecast horizons ranging from 96 to 336 steps, XLinear effectively leverages the long lookback window and consistently outperforms both TSMixer and PatchTST. At the longest horizon of 720 steps, XLinear performs comparably to PatchTST, which begins to show a general advantage over TSMixer. These results indicate that XLinear is better able to utilize long lookback windows than TSMixer.

Ablation Study

In the prediction phase, XLinear generates forecasts by synergistically integrating the temporal dependencies of the endogenous sequence and the key exogenous inputs embedded in the global tokens. To assess the contributions of the two information sources, we conducted ablation experiments under three configurations: (a) using both the endogenous sequence to capture temporal dependencies and the global tokens representing interaction from exogenous vari-

Models	XLinear	ES	GT				
Metrics	MSE	MAE	MSE	MAE	MSE	MAE	
ETTh1	96	0.369	0.393	0.370	0.396	0.382	0.406
	192	0.421	0.423	0.427	0.427	0.429	0.430
	336	0.455	0.440	0.462	0.444	0.467	0.449
	720	0.453	0.456	0.461	0.461	0.475	0.471
	720	0.453	0.456	0.461	0.461	0.475	0.471
ETTm1	96	0.311	0.351	0.328	0.366	0.326	0.364
	192	0.353	0.376	0.364	0.383	0.362	0.382
	336	0.382	0.398	0.392	0.404	0.390	0.404
	720	0.444	0.436	0.457	0.441	0.463	0.444
	720	0.444	0.436	0.457	0.441	0.463	0.444
Weather	96	0.149	0.198	0.175	0.216	0.153	0.200
	192	0.201	0.244	0.226	0.258	0.204	0.247
	336	0.259	0.288	0.282	0.298	0.263	0.291
	720	0.339	0.340	0.357	0.347	0.344	0.344
	720	0.339	0.340	0.357	0.347	0.344	0.344
GTD _s -MS	96	0.012	0.077	0.014	0.083	0.012	0.078
	192	0.025	0.112	0.027	0.118	0.025	0.114
	336	0.048	0.159	0.051	0.163	0.049	0.161
	720	0.124	0.260	0.130	0.267	0.127	0.266
	720	0.124	0.260	0.130	0.267	0.127	0.266

Table 7: Results of the ablation experiments, where the optimal results are highlighted in bold.

ables (XLinear); (b) using only the endogenous sequence (denoted as ES); (c) using only the global tokens (denoted as GT). As shown in Table 7, integrating both information sources yields the best performance. Notably, the endogenous global token alone generally outperforms the pure temporal sequence, as it captures both the target’s overall temporal features and exogenous information, highlighting their complementary roles.

Conclusion

Motivated by real-world time series forecasting applications, particularly those involving exogenous inputs, we designed XLinear. This lightweight model incorporates a gating mechanism, utilizing MLP with sigmoid activation functions. XLinear effectively captures both temporal patterns and variate-wise dependencies from exogenous variables to endogenous variables through their learnable global tokens via time-wise and variate-wise gating modules. These gated outputs are fused with the input feature maps to amplify salient signals and suppress noise. Extensive experiments on seven public benchmarks and five real-world application datasets demonstrated that XLinear achieved superior accuracy and efficiency for both multivariate forecasting and univariate forecasting influenced by exogenous variables. Notably, it matched the speed of other lightweight models while being significantly (at least 30%) faster than the most efficient Transformer-based alternatives.

For future work, we aim to enhance XLinear’s scalability to handle time series with hundreds of variables. Promising directions include further improving forecast accuracy by optimizing lookback windows and bolstering forecast reliability for operational deployment.

Acknowledgments

This research project was supported in part by National Key Research and Development Program of China under Grant 2023YFF1000100, and Hubei Key Research and Development Program of China under Grant 2024BBB055, 2024BAA008; and in part by the Major Science and Technology Project of Yunnan Province under Grant 202502AE090003, and in part by the Fundamental Research Funds for the Chinese Central Universities under Grant 2662025XXPY005. This work was partially supported by CSIRO Digital Water and Landscapes project.

References

- Bakar, K. S.; and Jin, H. 2018. Spatio-temporal quantitative links between climatic extremes and population flows: a case study in the Murray-Darling Basin, Australia. *Climatic Change*, 148(1): 139–153.
- Bakar, K. S.; Kokic, P.; and Jin, H. 2015. A spatiodynamic model for assessing frost risk in south-eastern Australia. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 64(5): 755–778.
- Bakar, K. S.; Kokic, P.; and Jin, H. 2016. Hierarchical spatially varying coefficient and temporal dynamic process models using spTDyn. *Journal of Statistical Computation and Simulation*, 86(4): 820–840.
- Cai, W.; Liang, Y.; Liu, X.; Feng, J.; and Wu, Y. 2024. MS-GNet: Learning multi-scale inter-series correlations for multivariate time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 11141–11149.
- Chen, S.-A.; Li, C.-L.; Arik, S. O.; Yoder, N. C.; and Pfister, T. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. *Transactions on Machine Learning Research*.
- Das, A.; Kong, W.; Leach, A.; Mathur, S.; Sen, R.; and Yu, R. 2023. Long-term Forecasting with TiDE: Time-series Dense Encoder. *Transactions on Machine Learning Research*.
- Genova, A.; Jin, H.; Taylor, J.; Dabrowski, J. J.; Toscas, P.; and Joehnk, K. 2025. Advancing Long-Term High-Frequency Dissolved Oxygen Forecasting for Australian Rivers. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 198–209. Springer.
- Gupta, H. V.; Kling, H.; Yilmaz, K. K.; and Martinez, G. F. 2009. Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. *Journal of hydrology*, 377(1-2): 80–91.
- Huang, Q.; Jin, H.; Nguyen, D.; Guo, Y.; Zheng, L.; Livingston, D.; Dabrowski, J. J.; Toscas, P.; and Joehnk, K. 2025. Enhanced Transformers for Long-Term, High-Resolution Lake Temperature Stratification Forecasting. *Engineering Applications of Artificial Intelligence*. (Under Preparation).
- Huang, Q.; Shen, L.; Zhang, R.; Ding, S.; Wang, B.; Zhou, Z.; and Wang, Y. 2023. Crossggnn: Confronting noisy multivariate time series via cross interaction refinement. *Advances in Neural Information Processing Systems*, 36: 46885–46902.
- Jin, H.; and Henderson, B. 2011. Towards a daily soil moisture product based on incomplete time series observations of two satellites. In *MODSIM*, volume 2011, 1959–1965.
- Jin, H.; Li, M.; Hopwood, G.; Hochman, Z.; and Bakar, K. S. 2022. Improving early-season wheat yield forecasts driven by probabilistic seasonal climate forecasts. *Agricultural and Forest Meteorology*, 315: 108832.
- Kim, T.; Kim, J.; Tae, Y.; Park, C.; Choi, J.-H.; and Choo, J. 2021. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International conference on learning representations*.
- Kokic, P.; Jin, H.; and Crimp, S. 2013. Improved point scale climate projections using a block bootstrap simulation and quantile matching method. *Climate dynamics*, 41(3): 853–866.
- Kottapalli, S. R. K.; Hubli, K.; Chandrashekhara, S.; Jain, G.; Hubli, S.; Botla, G.; and Doddaiiah, R. 2025. Foundation Models for Time Series: A Survey. *arXiv preprint arXiv:2504.04011*.
- Lago, J.; Marcjasz, G.; De Schutter, B.; and Weron, R. 2021. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293: 116983.
- Li, S.; Jin, X.; Xuan, Y.; Zhou, X.; Chen, W.; Wang, Y.-X.; and Yan, X. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in neural information processing systems*, 32.
- Li, Z.; Qi, S.; Li, Y.; and Xu, Z. 2023. Revisiting long-term time series forecasting: An investigation on linear mapping. *arXiv preprint arXiv:2305.10721*.
- Liu, S.; Wu, K.; Jiang, C.; Huang, B.; and Ma, D. 2023a. Financial time-series forecasting: Towards synergizing performance and interpretability within a hybrid machine learning approach. *arXiv preprint arXiv:2401.00534*.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2023b. iTransformer: Inverted transformers are effective for time series forecasting. In *The Twelfth International Conference on Learning Representations*.
- Luo, D.; and Wang, X. 2024. DeformableTST: Transformer for time series forecasting without over-reliance on patching. *Advances in Neural Information Processing Systems*, 37: 88003–88044.
- Moriasi, D. N.; Arnold, J. G.; Van Liew, M. W.; Bingner, R. L.; Harmel, R. D.; and Veith, T. L. 2007. Model evaluation guidelines for systematic quantification of accuracy in watershed simulations. *Transactions of the ASABE*, 50(3): 885–900.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.
- Shao, Q.; Bange, M.; Mahan, J.; Jin, H.; Jamali, H.; Zheng, B.; and Chapman, S. C. 2019. A new probabilistic forecasting model for canopy temperature with consideration of

- periodicity and parameter variation. *Agricultural and Forest Meteorology*, 265: 88–98.
- Shao, Q.; Roche, R.; Jamali, H.; Nunn, C.; Zheng, B.; Jin, H.; Chapman, S. C.; and Bange, M. 2025. Comprehensive Assessment of PeriodiCT Model for Canopy Temperature Forecasting. *Agronomy*, 15(7): 1665.
- Shi, X.; Wang, S.; Nie, Y.; Li, D.; Ye, Z.; Wen, Q.; and Jin, M. 2025. Time-MoE: Billion-Scale Time Series Foundation Models with Mixture of Experts. In *ICLR 2025: The Thirteenth International Conference on Learning Representations*. International Conference on Learning Representations.
- Stitsyuk, A.; and Choi, J. 2025. xPatch: Dual-Stream Time Series Forecasting with Exponential Seasonal-Trend Decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 20601–20609.
- Tang, P.; and Zhang, W. 2025. Unlocking the Power of Patch: Patch-Based MLP for Long-Term Time Series Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12640–12648.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and Zhou, J. 2024a. TimeMixer: Decomposable multiscale mixing for time series forecasting. In *International Conference on Learning Representations (ICLR)*.
- Wang, Y.; Wu, H.; Dong, J.; Qin, G.; Zhang, H.; Liu, Y.; Qiu, Y.; Wang, J.; and Long, M. 2024b. TimeXer: Empowering transformers for time series forecasting with exogenous variables. In *Advances in Neural Information Processing Systems*, volume 37, 469–498.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34: 22419–22430.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Chang, X.; and Zhang, C. 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 753–763.
- Yi, K.; Zhang, Q.; Fan, W.; He, H.; Hu, L.; Wang, P.; An, N.; Cao, L.; and Niu, Z. 2023. FourierGNN: Rethinking multivariate time series forecasting from a pure graph perspective. *Advances in neural information processing systems*, 36: 69638–69660.
- Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11121–11128.
- Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *The eleventh international conference on learning representations*.
- Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.
- Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, 27268–27286. PMLR.