

HN-MVTS: HyperNetwork-based Multivariate Time Series Forecasting

Andrey Savchenko^{1,2}, Oleg Kachan^{1,2}

¹Sber AI Lab, Moscow, Russia

²HSE University, Moscow, Russia

avsavchenko@hse.ru, okachan@hse.ru

Abstract

Accurate forecasting of multivariate time series data remains a formidable challenge, particularly due to the growing complexity of temporal dependencies in real-world scenarios. While neural network-based models have achieved notable success in this domain, complex channel-dependent models often suffer from performance degradation compared to channel-independent models that do not consider the relationship between components but provide high robustness due to small capacity. In this work, we propose HN-MVTS, a novel architecture that integrates the hypernetwork-based generative prior with an arbitrary neural network forecasting model. The input of this hypernetwork is a learnable embedding matrix of time series components. To restrict the number of new parameters, the hypernetwork learns to generate the weights of the last layer of the target forecasting networks, serving as a data-adaptive regularizer that improves generalization and long-range predictive accuracy. The hypernetwork is only used during training, so it does not increase the inference time compared to the base forecasting model. Extensive experiments on eight benchmark datasets demonstrate that application of HN-MVTS to the state-of-the-art models (DLinear, PatchTST, TSMixer, etc.) typically improves their performance. Our findings suggest that hypernetwork-driven parameterization offers a promising direction for enhancing existing forecasting techniques in complex scenarios.

Code — <https://github.com/av-savchenko/HN-MVTS>

Extended version — <https://arxiv.org/abs/2511.08340>

Introduction

Multivariate time series (MVTS) forecasting (Kostromina et al. 2025; Mendis, Wickramasinghe, and Marasinghe 2024) is critical to many real-world applications, including energy demand prediction, traffic management, financial forecasting, weather modeling, and health monitoring. In such scenarios, multiple interrelated time series, often referred to as channels or components, must be predicted simultaneously. Effective forecasting, therefore, depends not only on modeling temporal patterns within each series but also on capturing correlations across components.

Advancements in deep learning have led to powerful forecasting models, ranging from linear baselines (Zeng et al.

2023) to architectures based on convolutional networks (Luo and Wang 2024), multilayer perceptrons (MLPs) (Chen et al. 2023), and Transformers (Nie et al. 2023). Surprisingly, recent work has shown that channel-independent (CI) models, which treat each time series channel separately, often outperform channel-dependent (CD) approaches that attempt to model all channels jointly (Han, Ye, and Zhan 2024; Peiwen and Changsheng 2023). CI models reduce model complexity, but sacrifice the ability to leverage valuable inter-channel relationships, potentially limiting their accuracy in highly correlated settings.

On the other hand, CD models explicitly exploit inter-channel dependencies. Indeed, the improvement over baseline is often achieved by advanced modeling of temporal (Nie et al. 2023) and channel (Liu et al. 2024b) similarities, or both (Wu et al. 2023). However, CD models are often known to be less scalable due to their requirement for larger datasets (Peiwen and Changsheng 2023). Designing models that can balance the robustness of CI approaches with the expressiveness of CD strategies remains a key open challenge in MVTS forecasting (Wang et al. 2025; Zhao and Shen 2024). The difficulty lies in balancing the leverage of inter-channel correlations and maintaining robustness and scalability across diverse datasets. For example, a local model with individual parameters for each channel results in increased parameters and does not capture the channels' similarities (Montero-Manso and Hyndman 2021).

To address this, we propose HN-MVTS (Fig. 1), a novel architecture that bridges the gap between channel-independent and channel-dependent modeling by using a hypernetwork to generate channel-specific forecasting parameters. Specifically, a small MLP-based hypernetwork receives a learnable embedding for each channel and outputs the weights of the final prediction layer for that channel. This setup allows similar channels to share information through their embeddings while preserving the modularity and efficiency of CI models. A key strength of HN-MVTS is its architectural flexibility: it can be seamlessly integrated with a wide range of neural network-based forecasting models, including linear baselines, MLPs, ConvNets, and Transformer variants. By modifying only the final prediction layer, HN-MVTS enhances the forecasting performance of existing architectures with minimal overhead in training time and the number of learnable parameters and without any im-




Figure 1: **Top left:** Channel-dependent, **Bottom left:** channel-independent, and **Right:** proposed HN-MVTS forecasting approach, where a hypernetwork outputs parameters of the forecasting model’s last layer, taking a channel embedding as an input.

impact on the inference time. We demonstrate the effectiveness of HN-MVTS through comprehensive experiments on multiple benchmark datasets. Our approach typically substantially improves the performance of several state-of-the-art models, such as DLinear, PatchTST, and TSMixer, on multiple widely used datasets (e.g., ECL, ETTh), highlighting its robustness, adaptability, and ability to generalize across diverse settings.

Related Works

Models. Nowadays, the dominant architecture in MVTS forecasting is the Transformer. Modern papers illustrate an expanding design space where model efficacy increasingly depends on architectural alignment with the unique characteristics of time series (Su et al. 2025; Wang et al. 2024b). Inverse Transformer (iTransformer) (Liu et al. 2024b) processes entire time points as variate tokens rather than temporal tokens, enabling better capture of multivariate correlations while retaining Transformer components. It addresses key limitations of conventional temporal tokenization, particularly when handling large lookback windows and diverse variate relationships. PatchTST (Nie et al. 2023) demonstrates how patching from computer vision can adapt Transformers for time series through subseries-level embeddings and channel independence.

However, recent works have introduced various models that challenge traditional architectures by focusing on improved temporal and cross-variable dependencies. For example, the DLinear architecture (Zeng et al. 2023) has emerged as a compelling alternative by decomposing series into trend and seasonality components through linear layers, achieving superior performance over complex Transformer variants, such as FEDformer (Zhou et al. 2022), on key benchmarks while maintaining computational efficiency. This simplicity-first approach sparked reevaluations of model design priorities, with ModernTCN (Luo and Wang 2024) further advancing non-Transformer paradigms through modernized temporal convolutional networks that expand effective receptive fields while maintaining computational efficiency. TSMixer (Chen et al. 2023), TimesMixer (Wang et al. 2024a), and HD-Mixer (Huang et al. 2024) exemplify this through an all-MLP (multi-layer perceptron) architecture that integrates

historical data, future known inputs, and static covariates using conditional feature mixing layers, demonstrating competitive performance against Transformers.

Finally, techniques based on graph neural networks (Makarov et al. 2022), such as graph inter-series models like MSGNet (Cai et al. 2024a) and FourierGNN (Yi et al. 2023), or ForecastGrapher (Cai et al. 2024b), treat each series as a node and learn a graph; however, their scalability is usually relatively poor.

Channel-dependent vs channel-independent models.

MVTS forecasting strategies are broadly categorized into CI and CD (Montero-Manso and Hyndman 2021). The CI mode (Han, Ye, and Zhan 2024), which includes both global (single model applied to all univariate series) and local (dedicated models per series) options, treats each time series channel as independent, ignoring cross-channel correlations. In contrast, the CD strategy employs a single model that takes all components as input and forecasts all components simultaneously, explicitly modeling inter-channel dependencies (Hertel et al. 2023). The CI-CD performance gap stems from fundamental trade-offs between model capacity and robustness. CD methods theoretically possess a higher capacity for modeling complex channel interactions, but struggle with real-world challenges, such as limited training data. CI strategies overcome these issues by using larger effective datasets (through channel-wise replication) and avoiding error propagation from correlated channels. Hence, empirical studies (Hertel et al. 2023) demonstrate that CI methods outperform CD approaches across various domains, including energy load forecasting, financial data analysis, and epidemiological modeling. (Han, Ye, and Zhan 2024) benchmarks global univariate versus multivariate modeling for linear, autoregressive, MLP, gradient boosting, and transformer-based models. They show that CI (univariate) models generally significantly outperform multivariate ones, though most transformer-based MVTS forecasting models are trained using the multivariate strategy. Global CI methods achieve parameter efficiency comparable to local models while maintaining competitive accuracy (Montero-Manso and Hyndman 2021). Indeed, the top linear model (DLinear) (Zeng et al. 2023) trained using the global univariate approach outperforms multivariate-trained transformers. Similarly, (Hertel et al. 2023) reports

that global CI Transformers significantly reduce forecasting errors compared to CD counterparts in smart grid applications. The success of CI highlights the importance of robustness in real-world forecasting (Ashouri and Phoa 2022). Nevertheless, CD approaches retain promise for domains with strong inter-channel correlations, provided they incorporate mechanisms like temporal clustering or residual regularization (Han, Ye, and Zhan 2024). For example, DUET (Qiu et al. 2024) introduces dual temporal-channel clustering to enhance CD models, showing that structured dependency modeling combined with robustness mechanisms can match CI performance in scenarios with clear channel groupings. PatchTST (Nie et al. 2023) processes time series into patched “words” using CI Transformers to preserve univariate patterns while enabling cross-variate attention. (Chen et al. 2024) observes the boost in forecasting performance for multivariate models for correlated time series. These developments underscore the need for context-aware strategy selection, balancing the CI-CD trade-off against dataset characteristics and operational constraints. Similar challenges are addressed in COSCI-GAN (Seyfi, Rajotte, and Ng 2022) for another task, generative modeling, by coordinating single-channel generators from a common noise source and utilizing a central discriminator to enforce inter-channel relationships. However, it does not employ an embedding-based parameter generation approach.

Hypernetworks. Hypernetworks, neural networks that generate weights for another neural network (the target network), have emerged as a powerful approach for various machine learning tasks (Ha, Dai, and Le 2016), including time series forecasting. The foundational work (Duan et al. 2023) introduced Hyper Time Series Forecasting, which jointly learns time-varying distributions and corresponding forecasting models, enabling more accurate predictions when underlying data distributions change over time. Recent advances include HyperGPA (Lee et al. 2022), a hypernetwork that generates optimal model parameters for each time period through computational graph structures. It demonstrates the effectiveness in forecasting scenarios where future data patterns differ from historical observations. Finally, LPCNet employs a hypernetwork-based framework that dynamically updates neural network parameters in response to changing input distributions (Liu et al. 2024a). This adaptive parameter generation automatically enables the model to adjust to varying temporal patterns with acceptable computational complexity. These advancements highlight the ability of hypernetworks to model complex temporal dependencies while adapting to evolving data distributions.

Another research direction is the usage of a hypernetwork for meta-learning to handle multiple tasks. HyperTrees (März and Rasul 2024) generate parameters that adapt to local characteristics when applied to individual series, enabling effective handling of heterogeneous time series. Recent work (Staněk 2025) explored the hypernetwork’s meta-learning architectures, which are capable of constructing optimal parametric models for families of similar data-generating processes, achieving competitive performance in forecasting challenges. The HyperTime framework (Fons et al. 2022) provides an accurate and resolution-independent

encoding of time series data, facilitating accurate reconstruction and interpolation. It was effective for data augmentation, achieving competitive performance against state-of-the-art time series augmentation methods.

Thus, hypernetworks show promise for limited observations where nonparametric methods fail. However, their applications are still restricted to managing non-stationary time series, data augmentation, or training when multiple realizations of similar but not identical time series are available. To the best of our knowledge, they have not been applied to improve the performance metrics, such as mean squared error (MSE), of state-of-the-art models for MVTs forecasting. This paper fills this gap by utilizing hypernetworks to leverage the benefits of both CI and CD approaches.

Proposed Approach

Problem statement. Let the history be $\mathbf{X}_{1:T} \in \mathbb{R}^{N \times T}$, where $N \geq 1$ is the number of channels, and T is the length of the lookback window. The multivariate time series forecasting task is to predict $\mathbf{X}_{T+1:T+H} \in \mathbb{R}^{N \times H}$, where $H \geq 1$ is the prediction horizon.

We consider a supervised learning scenario, where time series $\mathbf{X}_{1:t} \in \mathbb{R}^{N \times t}$ with duration $t \gg T$ are available. It is typically used to create the training dataset of pairs $\{(\mathbf{X}_{i:T+i}, \mathbf{X}_{T+1+i:T+H+i})\}_{i=0}^M$, where the dataset size $M = t - (T + H)$. This dataset is used to learn the weights θ of the model f_θ that predicts future H steps for each of N channels:

$$\hat{\mathbf{X}}_{T+1+i:T+H+i} = f_\theta(\mathbf{X}_{i:T+i}). \quad (1)$$

The MSE loss function is typically minimized to obtain the optimal weights θ^* .

In this paper, we focus on neural network models f_θ with K layers, and *assume* that its last (K -th) layer linearly transforms an input hidden state $[\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(N)}]$, where $\mathbf{h}^{(n)} \in \mathbb{R}^D$ with hidden dimension D into the final output $\hat{\mathbf{X}} = [\hat{\mathbf{x}}^{(1)}, \dots, \hat{\mathbf{x}}^{(N)}] \in \mathbb{R}^{N \times H}$ using the matrix product for each channel $n = 1, \dots, N$:

$$\hat{\mathbf{x}}_n = \mathbf{W}_K^{(n)} \cdot \mathbf{h}^{(n)}, \quad (2)$$

where $\mathbf{W}_K^{(n)} \in \mathbb{R}^{H \times D}$ is a weight matrix for the n -th channel of the last layer.

Hypernetworks. Our core idea is to bridge the gap between CI and CD forecasting models by using a hypernetwork to generate channel-specific parameters based on learnable component embeddings. In contrast to conventional models that use a fixed set of parameters across all time series components, we dynamically adapt the final prediction layer for each channel, enabling personalized forecasting behavior while preserving parameter efficiency.

Let $f_{\theta_1, \dots, \theta_K} : \mathcal{X} \rightarrow \mathcal{Y}$ be a base neural network, parameterized by $\theta = \{\theta_k\}_{k=1}^K$, where θ_k are the parameters of the k -th layer. A hypernetwork $h_\phi : \mathcal{Z} \rightarrow \Theta$ predicts parameters of the base network f given a task embedding \mathbf{z} , resulting a network $f_{h_\phi} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ (Ha, Dai, and Le 2016).

The main idea of this paper is to parametrize the weights $\theta_K = \mathbf{W}_K = [\mathbf{W}_K^{(1)}, \dots, \mathbf{W}_K^{(N)}] \in \mathbb{R}^{N \times H \times D}$

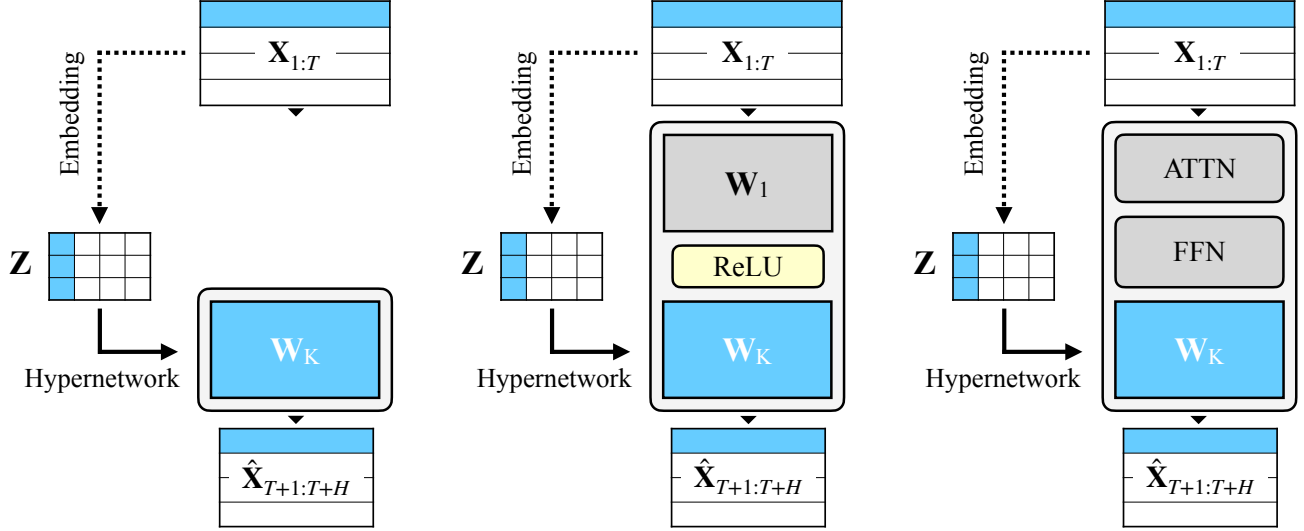


Figure 2: Examples of adding our HN-MVTS to various forecasting models. **Left:** Linear model. **Center:** Fully-connected neural network. **Right:** Transformer model with attention (ATTN) and feed-forward (FFN) layers.

of the last (prediction) layer (2) with a partial hypernetwork. In particular, each n -th component of a time series ($n \in \{1, 2, \dots, N\}$) is associated with a d -dimensional embedding vector $\mathbf{z}^{(n)} \in \mathbb{R}^d$, and the matrix $\mathbf{Z} = [\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N)}] \in \mathbb{R}^{N \times d}$ is fed into a hypernetwork:

$$\mathbf{W}_K = h_\phi(\mathbf{Z}). \quad (3)$$

In this paper, we implement a hypernetwork (3) as a MLP with one hidden layer that takes a learnable embedding vector for each time series component $\mathbf{z}^{(n)}$ (treating the number of channels N as the batch size) and outputs the corresponding weights for the final prediction layer (2) of an arbitrary base forecasting model. In the simplest case of an MLP without hidden layers, the weights of the base model are obtained using a linear transform:

$$\mathbf{W}_K^{(n)} = \mathbf{W}_\phi^{(n)} \cdot \mathbf{z}^{(n)}, \quad (4)$$

where $\mathbf{W}_\phi^{(n)} \in \mathbb{R}^{H \times D \times d}$ is the n -th component of the weights $\mathbf{W}_\phi^{(n)} = [\mathbf{W}_\phi^{(1)}, \dots, \mathbf{W}_\phi^{(N)}] \in \mathbb{R}^{N \times H \times D \times d}$ of the hypernetwork.

The central contribution of our hypernetwork-based approach (HN-MVTS) lies in utilizing a hypernetwork to generate part of the parameters of a forecasting model in a computationally efficient yet scalable and data-adaptive manner. In traditional MVTS models, all output channels typically share the same output projection layer (CD), or each channel has its own separate model (CI). Our approach aims to combine both benefits by introducing a shared neural forecasting model, while tailoring its final layer to each channel via a learnable hypernetwork. The difference between CD, CI, and our HN-MVTS is shown in Fig. 1. The former is the most common model, with a single transform f_θ of $N \times T$ input data into $N \times H$ outputs. The CI approach learns N separate models $f_{\theta_1}^{(1)}, \dots, f_{\theta_1}^{(N)}$, and processes each channel

independently. Our HN-MVTS saves the advantages of both types of models. If some components j_1 and j_2 of time series are similar, their embeddings will be close to each other $\mathbf{z}_{j_1} \approx \mathbf{z}_{j_2}$, and, hence, the training data for the j_1 component will have more influence to learn the weights for the j_2 component, and vice versa. In the ideal case of clustered embeddings $\mathbf{z}_{j_1} = \mathbf{z}_{j_2}$, a single (global) model (Montero-Manso and Hyndman 2021) is trained for the union of these components. In contrast, if the embeddings of j_1 and j_2 channels are significantly different, a CI mode is activated, in which the forecasting parameters are trained independently. This mechanism allows the model to adaptively interpolate between CD and CI behavior based on the learned embedding, without manual architecture changes or channel grouping.

Some examples of adding our HN-MVTS to various architectures are shown in Fig. 2. The above-mentioned motivation of mixing CD and CI models is especially clear for a forecasting model with one linear layer (left part of this figure), where the columns of the weight matrix W are learned based on the similarity between embeddings \mathbf{z} of time series components. However, our approach can be easily applied to MLP (Fig. 2b) or more complex architectures such as transformers (Fig. 2c), as only the weights of the last prediction layer are defined by the hypernetwork. It may be used for all layers, but the number of trained parameters will be significantly greater, making it more challenging to train the model. A practitioner can control the number of excess parameters by managing the complexity of the hypernetwork.

In all experiments, we use the following straightforward implementation of HN-MVTS. First, we initialize embeddings \mathbf{z} from the Pearson product-moment correlation coefficients across channels (Nguyen-Thai et al. 2024), which are projected onto the principal components of dimension d computed using the training-only splits. Our experiments indicate that alternative (random) initialization results in

slightly higher MSE. Next, we attach a simple MLP (4) and learn its weights and embeddings \mathbf{z} from given data together with the base model (1) by minimizing MSE. Our simplest implementation (4) adds only $N \cdot H \cdot D \cdot d$ parameters to learn (with addition of $N \cdot d$ if embeddings Z are learnable). As a result, the number of parameters can be even less than that of the CI approach, which applies the base model separately with its parameters for each channel. As the number of learnable weights is reasonably small, the training time is only slightly higher than that of the base forecasting model.

Moreover, since the trained hypernetwork outputs fixed weights that do not depend on the specific input time series, we can compute them only once after the training procedure is complete. Hence, we remove a hypernetwork and copy weights $\mathbf{W}_K^{(n)}$ (4) into the last linear layer of the base model (2). As a result, the inference time won't be affected (though deployment overheads, e.g., I/O, memory layout, may still cause minor variation), making our HN-MVTS as fast as the base model.

Thus, the proposed approach uses a learnable component embedding combined with a hypernetwork to produce channel-specific final-layer weights as a data-adaptive regularizer, integrates seamlessly with diverse backbone architectures, and preserves inference (Savchenko 2016). These points with empirical gains discussed in the next section are a substantive contribution beyond the straightforward reapplication of existing hypernetwork's ideas.

Experimental setup

Baseline models. We evaluated the proposed HN-MVTS with contemporary multivariate time series forecasting models of different architectures: 1) Linear model: DLinear (Zeng et al. 2023), 2) MLP-based model: TSMixer (Chen et al. 2023), 3) Convolutional neural network-based model: ModernTCN (Luo and Wang 2024), and 4) Transformer-based models: PatchTST (Nie et al. 2023) and inverse Transformer (Liu et al. 2024b). The source code for all models was taken from the original repositories of their authors.

Datasets. We leverage eight publicly available open datasets that are conventionally used in the MVTS literature (Shao et al. 2025), covering energy, traffic, and weather forecasting applications: `ECL` (UCI Electricity Consuming Load) with electricity consumption of hundreds of points/clients; 15-minute-level datasets (`ETTm1`, `ETTm2`) from the Electricity Transformer Temperature database (Zhou et al. 2021) with one oil and six load features of electricity transformers from 2016 to 2018; `Weather` (Max-Planck-Institute Weather Dataset for Long-term Time Series Forecasting) with weather indicators recorded every 10 min for 2020, and Traffic Flow from Performance Measurement System (`PEMS03`, `PEMS04`, `PEMS07`, `PEMS08`) collected by California Transportation Agencies (CalTrans) every 30 seconds and aggregated to 5 minutes.

Each dataset exhibits diverse temporal and structural properties, allowing us to assess the robustness and generalizability of our method. Their statistics are summarized in Table 1. We split the time series into train, validation, and test sets using a 7:2:1 ratio. For the `ETTm1` and `ETTm2`

Dataset	Timesteps t	Channels N	Split ratio	Granularity
<code>ECL</code>	26304	321	7:2:1	1h
<code>ETTm1</code>	57600	7	6:2:2	15m
<code>ETTm2</code>	57600	7	6:2:2	15m
<code>PEMS03</code>	26208	358	7:2:1	5m
<code>PEMS04</code>	16992	307	7:2:1	5m
<code>PEMS07</code>	28224	883	7:2:1	5m
<code>PEMS08</code>	17856	170	7:2:1	5m
<code>Weather</code>	52696	21	7:2:1	10m

Table 1: Datasets properties.

datasets, following standard testing protocols, we use a 6 : 2 : 2 split and take only the first 57600 timesteps. The validation set was used to obtain the best checkpoint of the model, and the results on the test set are reported.

Training and evaluation details. The training was performed on a server with 2x80Gb Nvidia A100 GPUs, Intel Xeon Gold 6326 CPUs (2.90GHz), and 512 GB of RAM. PyTorch 2.x was used to implement base models and our HN-MVTS. For the evaluation, we followed traditional settings (Nie et al. 2023; Wu et al. 2023): we set the input length T of the loopback window to 336, and evaluate the results for horizon lengths of $H = \{48, 96, 192, 336\}$. For a fair comparison, we apply reversible instance normalization (Kim et al. 2021) for all models. We minimize the MSE with the Adam optimizer with a learning rate of 0.0001 and a batch size of 64. These hyperparameters enable us to achieve approximately the same results as those reported in the original papers (Zeng et al. 2023; Nie et al. 2023). The dimensionality of embeddings for HN-MVTS is set to be less than or equal to the number of components in the time series.

Experimental Results

We report average MSE for four horizon lengths $H \in \{48, 96, 192, 336\}$ and 5 random seeds in Table 2. Here, the statistically significant (under the Wilcoxon signed-rank test with confidence 0.95 over 5 seeds) model among HN-MVTS and the baseline is shown in **bold**. The Mean Absolute Error (MAE) and detailed training curves, with the dependence of testing MSE on epoch, are presented in the arXiv version of this paper. As one can notice, in most cases, our HN-MVTS improves forecasting accuracy across all datasets and architectures, often by a substantial margin.

Regarding applications to various architectures, the efficient and straightforward linear baseline, namely DLinear (Zeng et al. 2023), benefits significantly from our method. For example, on the `Weather` dataset, MSE improves from 0.1369 to 0.1115 at horizon $H = 48$ and from 0.2641 to 0.2396 at $H = 336$. As a result, DLinear, combined with our HN-MVTS, became the top-performing solution for this dataset, although the vanilla DLinear was initially the worst model. Similar gains are observed on `ECL`, `PEMS04`, and `PEMS08`, i.e., improvements are dataset-dependent and often practically meaningful (particularly in high-dimensional or strongly correlated datasets). These results confirm that HN-MVTS can enhance lightweight models by incorporating channel relationships.

Dataset	H	DLinear	+HN-MVTS	TSMixer	+HN-MVTS	ModernTCN	+HN-MVTS	PatchTST	+HN-MVTS	iTransformer	+HN-MVTS
ECL	48	0.1255±0.0003	0.1184±0.0002	0.1377±0.0003	0.1220±0.0001	0.1270±0.0002	0.1191±0.0002	0.1162±0.0002	0.1157±0.0003	0.1091±0.0002	0.1076±0.0001
	96	0.1409±0.0004	0.1347±0.0004	0.1748±0.0011	0.1421±0.0003	0.1516±0.0007	0.1383±0.0005	0.1328±0.0003	0.1318±0.0002	0.1309±0.0002	0.1297±0.0002
	192	0.1577±0.0010	0.1533±0.0009	0.1972±0.0012	0.1666±0.0008	0.1783±0.0009	0.1708±0.0006	0.1523±0.0005	0.1532±0.0005	0.1575±0.0005	0.1536±0.0004
	336	0.1720±0.0010	0.1676±0.0008	0.1986±0.0014	0.1747±0.0010	0.1804±0.0015	0.1750±0.0011	0.1669±0.0012	0.1681±0.0012	0.1690±0.0007	0.1637±0.0006
ETTm1	48	0.2705±0.0022	0.2610±0.0017	0.2669±0.0044	0.2749±0.0049	0.2615±0.0017	0.2597±0.0013	0.2528±0.0063	0.2603±0.0049	0.2871±0.0005	0.2863±0.0004
	96	0.3019±0.0023	0.2899±0.0021	0.3002±0.0014	0.2969±0.0010	0.2929±0.0004	0.2899±0.0003	0.2899±0.0019	0.2898±0.0020	0.3034±0.0037	0.3171±0.0039
	192	0.3364±0.0010	0.3306±0.0009	0.3405±0.0008	0.3361±0.0008	0.3455±0.0010	0.3397±0.0011	0.3295±0.0034	0.3313±0.0029	0.3386±0.0087	0.3579±0.0105
	336	0.3717±0.0014	0.3721±0.0015	0.3820±0.0009	0.3759±0.0008	0.3824±0.0010	0.3760±0.0008	0.3657±0.0013	0.3669±0.0015	0.3767±0.0085	0.3946±0.0114
ETTm2	48	0.1279±0.0003	0.1246±0.0003	0.1373±0.0004	0.1326±0.0003	0.1251±0.0002	0.1230±0.0001	0.1302±0.0005	0.1310±0.0009	0.1363±0.0006	0.1376±0.0011
	96	0.1641±0.0003	0.1626±0.0004	0.1805±0.0005	0.1749±0.0004	0.1641±0.0007	0.1606±0.0003	0.1696±0.0009	0.1703±0.0010	0.1823±0.0013	0.1832±0.0015
	192	0.2183±0.0003	0.2165±0.0003	0.2375±0.0007	0.2308±0.0006	0.2285±0.0010	0.2297±0.0014	0.2339±0.0012	0.2357±0.0016	0.2538±0.0010	0.2448±0.0009
	336	0.2746±0.0003	0.2730±0.0003	0.2906±0.0006	0.2848±0.0005	0.2857±0.0071	0.3075±0.0113	0.2885±0.0049	0.2903±0.0061	0.3055±0.0062	0.3055±0.0060
PEMS03	48	0.1513±0.0002	0.1490±0.0001	0.1077±0.0003	0.1073±0.0003	0.1313±0.0007	0.1178±0.0003	0.1143±0.0004	0.1071±0.0002	0.0898±0.0001	0.0875±0.0001
	96	0.1899±0.0002	0.1871±0.0002	0.1369±0.0003	0.1376±0.0004	0.1828±0.0014	0.2063±0.0049	0.1461±0.0021	0.1364±0.0020	0.1118±0.0005	0.1087±0.0003
	192	0.2096±0.0011	0.2065±0.0009	0.1537±0.0021	0.1554±0.0028	0.2004±0.0034	0.1629±0.0029	0.1773±0.0008	0.1719±0.0006	0.1307±0.0003	0.1223±0.0002
	336	0.2265±0.0008	0.2227±0.0005	0.1654±0.0006	0.1645±0.0006	0.2505±0.0143	0.2135±0.0106	0.1899±0.0017	0.1862±0.0013	0.1504±0.0010	0.1481±0.0008
PEMS04	48	0.1639±0.0008	0.1572±0.0005	0.1185±0.0006	0.1101±0.0003	0.1048±0.0004	0.1056±0.0005	0.1306±0.0014	0.1130±0.0011	0.0960±0.0004	0.0878±0.0003
	96	0.2016±0.0012	0.1943±0.0009	0.1312±0.0007	0.1270±0.0007	0.1744±0.0009	0.1682±0.0010	0.1627±0.0132	0.1461±0.0066	0.1136±0.0015	0.1030±0.0011
	192	0.2208±0.0014	0.2125±0.0012	0.1566±0.0043	0.1395±0.0039	0.1874±0.0075	0.1662±0.0034	0.1873±0.0068	0.1638±0.0029	0.1267±0.0024	0.1189±0.0018
	336	0.2444±0.0016	0.2390±0.0012	0.1760±0.0024	0.1519±0.0019	0.1985±0.0037	0.1786±0.0022	0.2005±0.0039	0.1741±0.0018	0.1533±0.0013	0.1333±0.0016
PEMS07	48	0.1479±0.0005	0.1411±0.0005	0.0959±0.0002	0.0904±0.0001	0.1328±0.0004	0.1070±0.0003	0.0992±0.0001	0.0888±0.0002	0.0681±0.0001	0.0637±0.0001
	96	0.1871±0.0020	0.1766±0.0016	0.1117±0.0006	0.1062±0.0005	0.1253±0.0067	0.1474±0.0106	0.1240±0.0063	0.1092±0.0024	0.0790±0.0002	0.0721±0.0001
	192	0.2057±0.0018	0.1994±0.0011	0.1308±0.0008	0.1212±0.0010	0.1399±0.0023	0.1489±0.0026	0.1478±0.0012	0.1286±0.0010	0.0908±0.0003	0.0818±0.0002
	336	0.2287±0.0020	0.2200±0.0013	0.1533±0.0009	0.1326±0.0011	0.1442±0.0034	0.1545±0.0078	0.1619±0.0016	0.1415±0.0013	0.1030±0.0004	0.0928±0.0003
PEMS08	48	0.1889±0.0035	0.1651±0.0027	0.1120±0.0014	0.1027±0.0009	0.1268±0.0001	0.1250±0.0013	0.1221±0.0038	0.1052±0.0015	0.0870±0.0003	0.0799±0.0002
	96	0.2618±0.0105	0.2257±0.0074	0.1482±0.0040	0.1268±0.0028	0.2284±0.0035	0.2356±0.0042	0.1562±0.0014	0.1348±0.0008	0.1113±0.0003	0.0957±0.0003
	192	0.3067±0.0063	0.2780±0.0026	0.2217±0.0027	0.1884±0.0013	0.2327±0.0029	0.1894±0.0015	0.2128±0.0019	0.1964±0.0008	0.1459±0.0016	0.1187±0.0009
	336	0.3273±0.0024	0.3252±0.0026	0.2496±0.0012	0.2477±0.0008	0.2694±0.0007	0.2640±0.0007	0.2343±0.0013	0.2231±0.0010	0.1770±0.0008	0.1769±0.0009
Weather	48	0.1369±0.0025	0.1115±0.0007	0.1168±0.0008	0.1162±0.0007	0.1151±0.0004	0.1127±0.0005	0.1136±0.0007	0.1143±0.0008	0.1235±0.0004	0.1195±0.0003
	96	0.1733±0.0032	0.1425±0.0014	0.1490±0.0008	0.1475±0.0009	0.1480±0.0005	0.1429±0.0004	0.1480±0.0014	0.1448±0.0023	0.1559±0.0006	0.1505±0.0003
	192	0.2167±0.0063	0.1857±0.0012	0.2000±0.0011	0.1948±0.0009	0.1941±0.0007	0.1914±0.0006	0.1865±0.0024	0.1903±0.0027	0.1992±0.0031	0.1995±0.0029
	336	0.2641±0.0042	0.2396±0.0028	0.2491±0.0010	0.2418±0.0011	0.2500±0.0009	0.2476±0.0006	0.2427±0.0003	0.2419±0.0002	0.2507±0.0006	0.2483±0.0003

Table 2: Mean and standard deviation (over 5 runs) of MSE for multivariate time series forecasting for different horizons H .

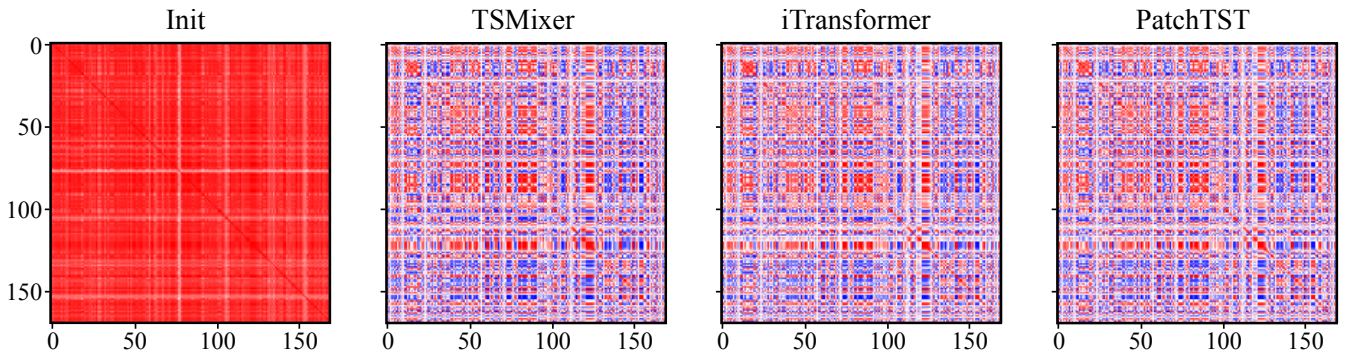


Figure 3: Embeddings for our HN-MVTS, PEMS08 dataset: (a) Initial embeddings before training, (b)-(d) learned embeddings for TSMixer, iTransformer and PatchTST.

Second, as an MLP-based model, TSMixer (Chen et al. 2023) benefits from the added expressivity and channel specificity provided by our hypernetwork. For example, on the ECL dataset, MAE drops from 0.2453 to 0.2350 at $H = 48$, and from 0.3040 to 0.2876 at $H = 336$. Notably, TSMixer + HN-MVTS shows significant gains on PEMS07 and PEMS08, which involve high-dimensional traffic sensor data with strong inter-channel dependencies. Third, while already competitive, ModernTCN (Luo and Wang 2024) exhibits modest yet significant improvements when paired with HN-MVTS. On ETTm2 and PEMS datasets, it becomes more robust at longer, typically more challenging horizons.

Finally, even advanced Transformer-based models benefit from our approach. On the PEMS07 dataset at $H = 336$, PatchTST’s MSE decreases from 0.1619 to 0.1415, and iTransformer from 0.1030 to 0.0928. This demonstrates that our approach complements sophisticated attention mechanisms by generating adaptive, channel-aware output param-

eters. The visualization of HN-MVTS embeddings (Fig. 3) shows that channel embeddings learned for various base models are very similar and reflect the specifics of the dataset rather than the architecture of the MVTS model.

Regarding the dataset-specific observations, the Weather dataset shows the most significant improvement. For DLinear, the MSE at $H = 96$ drops from 0.1733 to 0.1425 (an 18% relative improvement). ECL also sees noticeable improvement across horizons and architectures. For instance, TSMixer+HN-MVTS improves MSE from 0.1972 to 0.1666 at $H = 192$, showing that even a compact MLP benefits from our parameter-sharing scheme. Traffic datasets contain hundreds of correlated sensor channels, making them particularly suitable for channel-specific modeling. At $H = 336$, HN-MVTS improves PatchTST on PEMS08 from 0.2343 to 0.2231, showing how our approach enhances scalability in high-dimensional input settings.

Thus, our experiments demonstrate that while the abso-

Dataset	H	DLinear	+HN-MVTS	TSMixer	+HN-MVTS	ModernTCN	+HN-MVTS	PatchTST	+HN-MVTS	iTransformer	+HN-MVTS
ECL	48	15.31±0.23	17.24±0.18	20.53±0.27	20.32±0.32	90.10±0.54	91.23±0.57	46.83±0.36	48.74±0.38	20.71±0.19	24.85±0.25
	336	23.04±0.29	33.29±0.33	25.40±0.24	25.75±0.22	95.64±0.48	111.55±0.67	52.38±0.45	66.26±0.53	27.75±0.22	39.13±0.48
ETTm1	48	4.71±0.01	5.61±0.05	2.87±0.02	3.36±0.04	15.87±0.09	11.01±0.11	8.71±0.08	11.37±0.08	7.18±0.12	9.13±0.11
	336	5.39±0.05	8.12±0.03	3.06±0.07	4.76±0.06	10.22±0.10	17.27±0.13	6.16±0.04	12.90±0.08	7.87±0.03	10.82±0.14
ETTm2	48	4.79±0.06	5.63±0.07	2.83±0.02	3.38±0.01	15.03±0.18	11.47±0.15	9.12±0.09	8.29±0.10	7.41±0.06	9.28±0.13
	336	5.56±0.04	7.98±0.03	2.97±0.06	4.73±0.05	10.14±0.09	17.44±0.12	6.16±0.08	12.88±0.07	7.28±0.07	10.76±0.11
PEMS03	48	16.67±0.16	18.83±0.13	19.42±0.19	20.13±0.17	99.71±0.39	102.28±0.50	51.58±0.28	54.14±0.34	22.14±0.23	27.26±0.26
	336	25.31±0.28	35.84±0.40	26.57±0.37	26.17±0.34	107.65±0.54	121.93±0.62	58.46±0.29	73.27±0.61	29.60±0.28	43.39±0.31
PEMS04	48	9.73±0.10	10.61±0.09	11.37±0.12	11.71±0.12	51.23±0.38	53.37±0.41	28.96±0.23	30.27±0.25	13.02±0.17	15.40±0.16
	336	14.52±0.12	19.39±0.21	16.44±0.18	15.33±0.18	55.87±0.35	63.75±0.49	32.63±0.30	40.38±0.33	16.85±0.13	21.20±0.16
PEMS07	48	40.87±0.34	44.62±0.35	33.85±0.39	34.86±0.33	298.28±0.78	300.66±0.82	133.25±0.54	139.22±0.61	57.74±0.47	77.60±0.85
	336	59.21±0.37	83.22±0.44	53.67±0.41	58.45±0.50	315.42±0.92	372.10±1.05	153.61±0.64	188.44±0.77	84.84±0.59	107.37±0.63
PEMS08	48	6.43±0.06	7.14±0.06	12.22±0.10	11.27±0.12	29.62±0.19	31.27±0.22	17.44±0.25	18.66±0.26	8.59±0.17	9.87±0.17
	336	9.36±0.05	12.73±0.12	13.88±0.13	14.90±0.14	32.28±0.36	38.38±0.33	19.24±0.25	25.09±0.21	11.35±0.12	13.04±0.14
Weather	48	5.61±0.01	6.54±0.01	3.08±0.02	3.54±0.01	24.42±0.27	15.46±0.31	9.25±0.08	10.73±0.10	8.15±0.08	10.35±0.07
	336	7.10±0.19	9.72±0.17	3.34±0.03	5.10±0.04	14.52±0.12	22.86±0.45	9.34±0.14	17.53±0.36	9.28±0.13	12.89±0.12

Table 3: Mean and standard deviation of training time (seconds) per epoch.

lute gains vary across datasets and backbones, the improvements are consistent in direction and sometimes substantial (e.g., *Weather*, *traffic datasets*), reflecting HN-MVTS’s role as a general plug-and-play mechanism rather than a dataset-specific optimization. Our HN-MVTS improves all models across nearly every dataset and horizon length, confirming the value of channel-aware, dynamic parameter generation. Moreover, our approach is model-agnostic, yielding gains across all model types without requiring architectural overhauls, thereby underscoring the plug-and-play nature of our framework. Furthermore, we observe a resilience to horizon length H . Performance gains persist across short and long horizons. Particularly for longer-term forecasts (e.g., $H = 336$), where models typically degrade, HN-MVTS mitigates accuracy loss.

To assess the computational overhead introduced by HN-MVTS, we measured the average training time per epoch for each model on all datasets and forecast horizons. Results are presented in Table 3. Despite introducing hypernetwork and learnable embeddings, the additional training cost remains modest. Across all models and datasets, the training time increase ranges from approximately 5% to 25%, depending on the model complexity and input size. The most lightweight models, such as DLinear and TSMixer, experience the most minor increases (e.g., DLinear on *ECL* at $H = 48$: from 15.31s to 17.24s, or 12.6%). The relative increase is moderate given their baseline costs, even for more complex Transformer-based models like PatchTST and iTransformer. Even on high-dimensional datasets (e.g., *PEMS07* with 883 channels), the added parameter cost is modest, and the increase in training time is negligible, demonstrating scalability. Thus, the HN-MVTS adds minimal training overhead, making it a practical enhancement for research and production scenarios. These efficiency gains and performance improvements further strengthen the case for integrating HN-MVTS into a wide range of forecasting pipelines.

Conclusion

In this work, we introduced HN-MVTS, a novel general-purpose framework for MVTS forecasting that combines the strengths of CI and CD modeling (Fig. 1). Unlike traditional CI models, which ignore inter-channel relationships,

or CD models, which can be overly complex, our method leverages the concept of hypernetworks to generate channel-specific parameters in a lightweight and scalable manner. We enable the model to condition its predictions on channel similarity, allowing shared statistical strength across similar components while still preserving the robustness and efficiency of CI models. HN-MVTS contributes a training-time, architecture-agnostic generative prior for per-channel final-layer weights. Importantly, our approach is modular and architecture-agnostic, allowing it to be seamlessly applied to a wide range of existing neural forecasting models (Fig. 2), including both CI (DLinear, PatchTST) and CD (TSMixer, iTransformer). Through extensive empirical evaluations across multiple datasets, we demonstrated that HN-MVTS, in most cases, improves the accuracy of state-of-the-art forecasting models (Table 2). These gains come at minimal computational cost, with only a modest increase in training time and parameter count (Table 3). The inference time is not affected, as the hypernetwork can be removed during the deployment phase after the model has been trained.

Despite its flexibility and strong empirical results, HN-MVTS has several limitations. First, we assume that the last layer of a base forecasting model linearly transforms an input hidden state into the final output; therefore, multi-layer hyperparameterization (e.g., generating weights beyond the last layer) may offer deeper insights into representation sharing. Second, although the method is architecture-agnostic, integrating it into highly specialized, non-standard forecasting architectures may require additional tuning or design adaptation. We did not consider the application of hypernetworks to non-neural network models (gradient boosting, statistical models, etc.), which are still widely used.

Beyond immediate performance improvements, HN-MVTS opens up promising avenues for future research in adaptive and structured forecasting. Its ability to learn flexible representations of channel relationships makes it a natural fit for settings involving heterogeneous or evolving multivariate time series. Moreover, it provides a principled mechanism for incorporating prior knowledge, such as known similarity structures or hierarchical groupings, through initializing or designing component embeddings.

Acknowledgments

The work was supported by the grant for research centers in the field of AI provided by the Ministry of Economic Development of the Russian Federation in accordance with the agreement 000000C313925P4E0002 and the agreement with HSE University No. 139-15-2025-009.

The research was supported in part through the computational resources of HPC facilities at HSE University.

References

- Ashouri, M.; and Phoa, F. K. H. 2022. Interactive tool for clustering and forecasting patterns of Taiwan COVID-19 spread. *Plos One*, 17(6): e0265477.
- Cai, W.; Liang, Y.; Liu, X.; Feng, J.; and Wu, Y. 2024a. 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.
- Cai, W.; Wang, K.; Wu, H.; Chen, X.; and Wu, Y. 2024b. ForecastGrapher: Redefining multivariate time series forecasting with graph neural networks. *arXiv preprint arXiv:2405.18036*.
- Chen, J.; Lenssen, J. E.; Feng, A.; Hu, W.; Fey, M.; Tassulas, L.; Leskovec, J.; and Ying, R. 2024. From Similarity to Superiority: Channel Clustering for Time Series Forecasting. *arXiv preprint arXiv:2404.01340*.
- Chen, S.-A.; Li, C.-L.; Arik, S.; Yoder, N. C.; and Pfister, T. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. *Transactions on Machine Learning Research*.
- Duan, W.; He, X.; Zhou, L.; Thiele, L.; and Rao, H. 2023. Combating distribution shift for accurate time series forecasting via hypernetworks. In *28th International Conference on Parallel and Distributed Systems (ICPADS)*, 900–907. IEEE.
- Fons, E.; Sztrajman, A.; El-Laham, Y.; Iosifidis, A.; and Vyetrenko, S. 2022. HyperTime: Implicit neural representation for time series. *arXiv preprint arXiv:2208.05836*.
- Ha, D.; Dai, A.; and Le, Q. V. 2016. Hypernetworks. *arXiv preprint arXiv:1609.09106*.
- Han, L.; Ye, H.-J.; and Zhan, D.-C. 2024. The Capacity and Robustness Trade-Off: Revisiting the Channel Independent Strategy for Multivariate Time Series Forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 7129–7142.
- Hertel, M.; Beichter, M.; Heidrich, B.; Neumann, O.; Schäfer, B.; Mikut, R.; and Hagenmeyer, V. 2023. Transformer training strategies for forecasting multiple load time series. *Energy Informatics*, 6(1): 20.
- Huang, Q.; Shen, L.; Zhang, R.; Cheng, J.; Ding, S.; Zhou, Z.; and Wang, Y. 2024. HDMixer: Hierarchical dependency with extendable patch for multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12608–12616.
- 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 (ICLR)*, volume 9.
- Kostromina, A.; Kuvshinova, K.; Yugay, A.; Savchenko, A.; and Simakov, D. 2025. Tsururu: A Python-based Time Series Forecasting Strategies Library. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, 11077–11081. Demo Track.
- Lee, J.; Kim, C.; Lee, G.; Lim, H.; Choi, J.; Lee, K.; Lee, D.; Hong, S.; and Park, N. 2022. Time series forecasting with hypernetworks generating parameters in advance. *arXiv preprint arXiv:2211.12034*.
- Liu, G.; Hu, Z.; Wang, L.; Zhang, H.; Xue, J.; and Matthaïou, M. 2024a. A hypernetwork based framework for non-stationary channel prediction. *IEEE Transactions on Vehicular Technology*, 73(6): 8338–8351.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024b. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *International Conference on Learning Representations (ICLR)*, volume 12.
- Luo, D.; and Wang, X. 2024. ModernTCN: A modern pure convolution structure for general time series analysis. In *International Conference on Learning Representations (ICLR)*, volume 12.
- Makarov, I.; Savchenko, A.; Korovko, A.; Sherstyuk, L.; Severin, N.; Kiselev, D.; Mikheev, A.; and Babaev, D. 2022. Temporal network embedding framework with causal anonymous walks representations. *PeerJ Computer Science*, 8: e858.
- März, A.; and Rasul, K. 2024. Forecasting with Hyper-Trees. *arXiv preprint arXiv:2405.07836*.
- Mendis, K.; Wickramasinghe, M.; and Marasinghe, P. 2024. Multivariate time series forecasting: A review. In *Proceedings of the 2nd Asia Conference on Computer Vision, Image Processing and Pattern Recognition*, 1–9.
- Montero-Manso, P.; and Hyndman, R. J. 2021. Principles and algorithms for forecasting groups of time series: Locality and globality. *International Journal of Forecasting*, 37(4): 1632–1653.
- Nguyen-Thai, B.; Le, V.; Tieu, N.-D. T.; Tran, T.; Venkatesh, S.; and Ramzan, N. 2024. Learning evolving relations for multivariate time series forecasting. *Applied Intelligence*, 54(5): 3918–3932.
- Nie, Y.; Nguyen, N.; Sinthong, P.; and Kalagnanam, J. 2023. A time series is worth 64 words: Long-term forecasting with transformers. In *International Conference on Learning Representations (ICLR)*, volume 11.
- Peiwen, Y.; and Changsheng, Z. 2023. Is Channel Independent strategy optimal for Time Series Forecasting? *arXiv preprint arXiv:2310.17658*.
- Qiu, X.; Wu, X.; Lin, Y.; Guo, C.; Hu, J.; and Yang, B. 2024. Duet: Dual clustering enhanced multivariate time series forecasting. *arXiv preprint arXiv:2412.10859*.
- Savchenko, A. V. 2016. *Search techniques in intelligent classification systems*. Springer.
- Seyfi, A.; Rajotte, J.-F.; and Ng, R. 2022. Generating multivariate time series with COMmon Source COordInated GAN (COSCI-GAN). *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 32777–32788.

Shao, Z.; Wang, F.; Xu, Y.; Wei, W.; Yu, C.; Zhang, Z.; Yao, D.; Sun, T.; Jin, G.; Cao, X.; Cong, G.; Jensen, C. S.; and Cheng, X. 2025. Exploring Progress in Multivariate Time Series Forecasting: Comprehensive Benchmarking and Heterogeneity Analysis. *IEEE Transactions on Knowledge and Data Engineering*, 37(1): 291–305.

Staněk, F. 2025. Designing time-series models with hypernetworks and adversarial portfolios. *International Journal of Forecasting*, 41(4): 1461–1476.

Su, L.; Zuo, X.; Li, R.; Wang, X.; Zhao, H.; and Huang, B. 2025. A systematic review for transformer-based long-term series forecasting. *Artificial Intelligence Review*, 58(3): 80.

Wang, H.; Mo, Y.; Xiang, K.; Yin, N.; Dai, H.; Li, B.; and Fan, S. 2025. CSformer: Combining Channel Independence and Mixing for Robust Multivariate Time Series Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 21090–21098.

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.; Liu, Y.; Long, M.; and Wang, J. 2024b. Deep time series models: A comprehensive survey and benchmark. *arXiv preprint arXiv:2407.13278*.

Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *International Conference on Learning Representations (ICLR)*, volume 11.

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 (NeurIPS)*, 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.

Zhao, L.; and Shen, Y. 2024. Rethinking Channel Dependence for Multivariate Time Series Forecasting: Learning from Leading Indicators. *arXiv preprint arXiv:2401.17548*.

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 (ICML)*, 27268–27286. PMLR.