

Take an Irregular Route: Enhance the Decoder of Time-Series Forecasting Transformer

Li Shen¹, Yuning Wei¹, Yangzhu Wang, and Huaxin Qiu¹

Abstract—With the development of Internet of Things (IoT) systems, precise long-term forecasting method is requisite for decision makers to evaluate current statuses and formulate future policies. Currently, transformer and MLP are two paradigms for deep time-series forecasting and the former one is more prevailing in virtue of its exquisite attention mechanism and encoder-decoder architecture. However, data scientists seem to be more willing to dive into the research of encoder, leaving decoder unconcerned. Some researchers even adopt linear projections in lieu of the decoder to reduce the complexity. We argue that both extracting the features of input sequence and seeking the relations of input and prediction sequence, which are respective functions of encoder and decoder, are of paramount significance. Motivated from the success of FPN in CV field, we propose FPPformer to utilize bottom-up and top-down architectures, respectively, in encoder and decoder to build the full and rational hierarchy. The cutting-edge patchwise attention is exploited and further developed with the combination, whose format is also different in encoder and decoder, of revamped elementwise attention in this work. Extensive experiments with six state-of-the-art baselines on twelve benchmarks verify the promising performances of FPPformer and the importance of elaborately devising decoder in time-series forecasting transformer. The source code is released in <https://github.com/OrigamiSL/FPPformer>.

Index Terms—Deep learning, neural network, time-series forecasting, transformer.

I. INTRODUCTION

A. Background

THE ADVENT of Big Data era has brought immense volume and variety of data in the 21st century, especially in Internet of Things (IoT) systems with tons of sensors [1]. Consequently, it necessitates long-term time-series forecasting methods with demanding accuracy and efficiency to assist decision makers and engineers in the appraisal of sensor statuses and future plans. Since traditional forecasting methods based on statistics [2], [3] are no longer sufficient for such sophisticated situations, more and more data scientists pay their attention to deep time-series forecasting [4]. After decades of development and competition, time-series forecasting MLP (TSFM) [5], [6], [7]

and time-series forecasting transformer (TSFT) [8], [9], [10], [11] become the mainstream.

B. Problems

TSFM and TSFT have different pros and cons. TSFM is known for its parsimonious but efficient architecture so that forecasting models based on TSFM excel in resisting nonstationarity brought by distribution shifts [12] and concept drifts [13]. Conversely, forecasting models based on TSFT own more complicated architecture and better capability of capturing long-term dependencies of time-series at the expense of being more vulnerable to over-fitting problem caused by nonstationarity [5]. Fortunately, pioneers have striven to get around plenty of problems of TSFT. Direct forecasting strategy [14] reduces the time complexity and alleviates the error accumulation problem [15]. RevIN [12] solves the problem of distribution shifts among windows with distinct time spans. The channel-independent [16] forecasting method renders TSFT refraining from extracting vague inter-relationships of different variables. Patchwise attention mechanism [10], [16] further attenuates the space complexity and brings the capability of local feature extraction to TSFT. Indeed, recent works have proven that TSFT models [9], [17] can also be stable and robust in forecasting. Evidently, the majority of these enhancement focus on improving the encoder architecture and tackling input sequence features. It cannot be denied that they are very important, but not solely. The connections of input and prediction sequences, manifested by decoder in TSFT, are also of paramount significance, especially for pursuing precise forecasting in IoT. However, its significance is frequently omitted and itself is inadequately explored. Normally, the decoder architectures of existing TSFT models are simply duplicates of their encoder architectures, barring with little indispensable modifications, such like changing self-attentions to cross-attentions [8], [15]. Furthermore, some researchers have gone so far as to substitute decoder in TSFT with simple linear projection [16], [18], which is analogous to TSFM, for the sake of enhancing their efficiency. Now it is time to enhance the decoder of TSFT to fully develop its potential and push its forecasting performances to a new altitude.

C. Contributions

Different from existing TSFT models, we Fully develop the tried-and-tested Patchwise attention mechanism and Pyramid architecture in *both encoder and decoder* and thereby propose FPPformer. Alike FPN [19] and PAN [20] architectures which are prevalent in CV fields, FPPformer hierarchically extracts

Manuscript received 2 October 2023; revised 24 November 2023; accepted 6 December 2023. Date of publication 11 December 2023; date of current version 9 April 2024. This work was supported in part by the National Natural Science Foundation of China under Grant U19B2033 and Grant 62103414. (Corresponding author: Li Shen.)

Li Shen and Yuning Wei are with the School of Automation Science and Electrical Engineering, Beihang University, Beijing 100191, China (e-mail: shenli@buaa.edu.cn; yuning@buaa.edu.cn).

Yangzhu Wang and Huaxin Qiu are with the Flying College, Beijing University, Beijing 100191, China (e-mail: wangyangzhu@buaa.edu.cn; qiuhuaxin@buaa.edu.cn).

Digital Object Identifier 10.1109/JIOT.2023.3341099

input sequence features *from fine to coarse* and constructs prediction sequence *from coarse to fine*. To strengthen the feature extraction capability of patchwise attention, we further insert an elementwise attention block into each patchwise attention block to extract fine-grained inner-relations of each patch in encoder and decoder with merely linear complexity. A *channel-independent and temporal-independent embedding method* is utilized to modify the size of feature maps in FPPformer to meet the needs of elementwise attention and patchwise attention. Within each attention block in encoder, *the diagonal line of query-key matching matrix is masked* to ensure the generality of features extracted from input sequences. Primary contributions of this work are fivefold.

- 1) We propose a novel TSFT, i.e., FPPformer, which uncommonly and efficaciously improves the decoder architecture of TSFT to break its fetters and excavate its potential.
- 2) We renovate the decoder architecture of TSFT and change it into top-down architecture for the sake of rationally constructing the prediction sequence in a hierarchical manner.
- 3) Motivated by a pioneer anomaly detection method, we propose diagonal-masked (DM) self-attention to mitigate the negative impacts of the outliers in input sequences.
- 4) A new combination of elementwise attention and patchwise attention is proposed by us to compensate the weakness of conventional patch-attention in extracting the inner-features of each patch, with only additional linear complexity.
- 5) Extensive experiments under diverse settings validate that FPPformer is capable of reaching state-of-the-art on twelve benchmarks with peerless accuracy and robustness.

II. RELATED WORKS

The past few years have witnessed the development and the success of deep-learning-based forecasting methods. Thanks to the help of neural network, the long-term multivariate forecasting is no more a dream so that even IoT systems with plenty of sensors and explosive data can be predicted [1], [21], [22]. Researchers have developed deep forecasting methods built upon diverse networks and transformer is a hot topic among corresponding literature.

A. Time-Series Forecasting Transformer

Traditionally, TSFT executes the forecasting via encoder-decoder architecture. The transformer encoder is used to extract the features of input sequence, then the transformer decoder is able to construct the prediction sequence by the extracted features of encoder and the prediction sequence, which is initialized with a certain number since it is unknown at the beginning. These two processes are completed predominantly by *attention mechanism*, thereby researchers always keep an eye on it. LogSparse transformer [23] and Informer [15] discover the sparsity of query-key matching matrix and they force the elements of query to attach to the partial elements of key for the sake of reducing the complexity. Autoformer [24], FEDformer [25] and ETSformer [26]

combine the TSFT with seasonal-trend decomposition and signal processing method, e.g., fourier analysis, in attention mechanism to enhance their interpretability. Patchwise attention is more popular and proven to be more useful recently. TSFTs with patchwise attention, including Triformer [10], Crossformer [8], and PatchTST [16], achieve more promising performances than preceding models. However, whichever TSFT always emphasizes that the modified architecture is intended for *more efficient or effective feature extraction for input sequence*. Hardly ever can statements *involved with the profits of decoder* be found. Indeed, their decoders seem to play the role of requisite appendages in the entire transformer architecture. Once some parts of encoders are changed by their proposed methods, mirrored changes are made to their decoders. Some researches [16], [18] even abandon the decoder to circumvent these changes. Contrary to them, studying and figuring out the correct way of designing decoders in TSFT is exactly what this work is supposed to do.

B. Other Miscellaneous Deep Forecasting Methods

Barring TSFT, there are plenty of other types of deep forecasting methods. Forecasting methods based on RNN and CNN are feasible ones. Their respective representatives LSTMNet [27] and SCINet [28] both achieved shiny performances during their periods. However, compared with the foregoing two types of deep forecasting methods, TSFM relatively receives more attentions. These forecasting networks are solely composed of linear-projection layers, whereas they still achieve promising performances. Due to their simple architectures, it is convenient for them to combine with statistics models for the objective of improving their interpretability and forecasting capability. NBEATS [29] and DLinear [5] adopt seasonal-trend decomposition methods in their networks more concisely than FEDformer [25] but achieve better results in general. Challu et al. [7] further presented N-HITS that employs sampling and interpolation strategies on the basis of NBEATS for more precise and hierarchical prediction. The reconstruction method motivated from Legendre Polynomials is taken into account by Zhou et al. [6] to come out with FiLM. TSMixer proposed by Ekambaram et al. [30] considered the temporal patterns, cross-variate information and additional auxiliary information to render TSFM ready for more complicated forecasting cases. They are challenging competitors for TSFT and we chiefly compare FPPformer with other TSFTs and these TSFMs in forthcoming experiments.

III. PRELIMINARY

A. Problem Statement

This work primarily concentrates on multivariate forecasting problem. As the term suggests, a multivariate forecasting problem is to predict a certain window $\{x_{t_2 : t_3}\}^{1:V}$ with time duration of $(t_3 - t_2)$ and variable number of V with its anterior window $\{x_{t_1 : t_2}\}^{1:V}$. Each $x_t^v \in \mathbb{R}$, where $t \in [t_1 : t_3)$ and $v \in [1, V]$, denotes an element at timestamp t and stemming from variable v . There are quite a few nomenclature style to name the dimension of t and v . In this work, the dimension of t is termed the temporal dimension and the dimension of v is

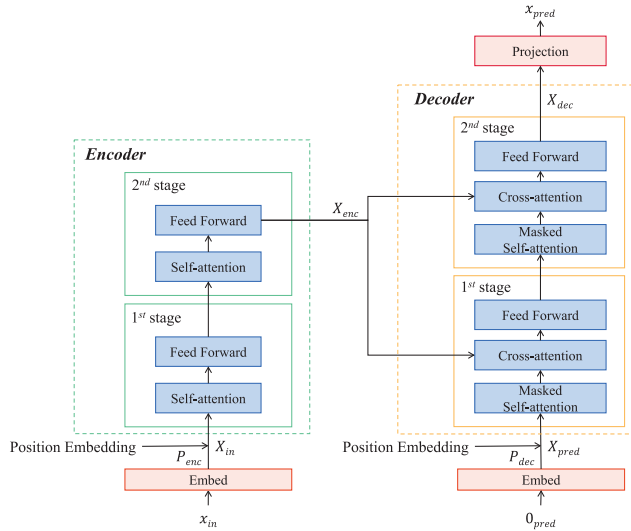


Fig. 1. Schematic of a vanilla TSFT with two-stage encoder (green dashed box containing green solid boxes in the left) and two-stage decoder (orange dashed box containing orange solid boxes in the right).

termed the variable dimension. Note that the word “channel-independent” mentioned above refers to the independence at the variable dimension. Moreover, we still discuss several univariate forecasting cases where $V = 1$ in our experiments since the variable treatment strategies of different forecasting methods can be very distinctive, making multivariate forecast comparison, albeit prevailing, not persuasive enough.

B. Vanilla TSFT Architecture

The architecture of vanilla TSFT is chiefly composed of an *encoder*, a *decoder*, and a *projection layer*. An example of a vanilla TSFT with 2-stage encoder and 2-stage decoder is sketched in Fig. 1. We can see that input sequence x_{in} passes through the *encoder* embedding layer, then the superimposition of the embedded input sequence X_{in} and its position embedding P_{enc} , which covers the input time span, is sent to the encoder. The encoder processes $X_{in} + P_{enc}$ with M ($M = 2$ in Fig. 1) stages, each consisting of a self-attention block and a feed forward layer, and the ultimate output feature map of the encoder is X_{enc} . Typically, the canonical self-attention conducts scaled dot-products with the formula of

$$\text{Attn}(q, k, v) = \text{Softmax}\left(\frac{qk^T}{\sqrt{D}}v\right) \quad (1)$$

where $q, k, v \in \mathbb{R}^{L \times D}$ are the linear projections of the identical sequence tensor, L is the number of tokens (or sequence tensor length), and D is the hidden dimension. Readers can refer to [31] to be familiar with attention mechanism and feed forward layer. Correspondingly, the *decoder* receives both X_{enc} and the zero-initialized prediction sequence 0_{pred} . 0_{pred} propagates through the decoder embedding layer to obtain X_{pred} and its position embedding P_{dec} , which covers the prediction time span, is superimposed on it. Afterward, they are sequentially sent into N ($N = 2$ in Fig. 1) stages, each consisting of a masked self-attention block, a cross-attention block and a feed forward layer. Causality is essential

as the prediction sequence is unknown, so that the masked self-attention, rather than normal self-attention, is utilized in decoder. The cross-attention block is intended to construct the prediction sequence via the encoder feature map X_{enc} . Eventually, a *projection layer* maps the output feature map of decoder X_{dec} to the prediction sequence x_{pred} .

C. Employed Mechanisms

Barring our proposed methods, which will be introduced in the upcoming section, we also employ several advanced time-series forecasting mechanisms in FPPformer.

- 1) *Direct forecasting method* [14], which is widely employed by recent deep forecasting method, performs the prediction of the entire sequence with only one forward process to alleviate the error accumulation.
- 2) *Channel-independent forecasting method*, which has been mentioned in the foregoing sections, treats the sequences of different variables as different instances. The sequences of different variables are parallel sent into the network without interfering with each other so that the network can seek shared characteristics of different variable sequences without imposing any inductive bias to the correlations of different variables.
- 3) *RevIN* [12], which is devised for nonstationary time-series forecasting, normalizes each input sequence with its own statistics before sent into the network and restores the original statistics to the prediction sequence via the reverse instance normalization to handle the distribution shifts of real-world long time-series.
- 4) *Patchwise attention* [16], which segments the sequence into patches of the same length, treats each single patch, rather than each single element, as a token in (1) and treats the elements inside each patch or their latent representations as the hidden features of each token (patch) for better efficiency and generality.

Note that the majority of recent deep forecasting models [5], [6], [7], [8], [10], [16], [28], [30], [32], including those employed in our experiments, at least adopt two of above mechanisms so that they are not something that distinguish our methods from others.

IV. METHODOLOGY

A. Analysis of Decoder in TSFT

Before commencing the introduction to our proposed FPPformer, we point out some deficiencies of current decoder architecture in TSFT to clarify the necessity and rationality of the decoder improvement in FPPformer.

We first discover the *redundant self-attention problem in decoder*. To elaborate, we notice that the input to decoder in Fig. 1 is a zero-initialized prediction sequence 0_{pred} owing to the unknown future. Consequently, the first (masked) self-attention is performed only on the position embedding of the prediction sequence. No matter it is fixed [15] or learnable [8], it is completely independent of input sequence. *Keep in mind that time-series forecasting is an auto-regressive problem*. It makes no sense to perform (masked) self-attention only on position embedding with the attempt of deducing some

groundless relations. Moreover, position embedding is always static after training while the input sequence can be dynamic and nonstationary, shattering the last hope that the position embedding can fit the statistics of time-series sequences due to some sort of assumptions with respect to homogeneity [33], [34]. Start token [15] can be a solution, however short start token is not enough for long-term forecasting and longer start token brings about excessive complexity.

Besides, the connection of encoder and decoder is unitary, leading to the *multiscale insufficiency* problem. As shown in Fig. 1, the encoder with M stages can produce M feature maps of input sequence, however merely the last one is sent to decoder. This problem is more noteworthy when it comes to some modified TSFT with *hierarchical architecture in encoder*. For instance, Informer [15] employs convolution layers between every two stages in encoder and FEDformer [25] keeps decomposing input sequences to acquire more precise seasonal features but neither of them apply the same operations to decoder and only the feature map of the last stage in encoder is sent to decoder. Crossformer [8] merges adjacent segments to obtain bigger patches in deeper stages in both encoder and decoder. However, just as what we claim in foregoing sections, the architecture of decoder is merely the replica of encoder with an additional cross-attention in Crossformer. By merging patches from small to big in decoder, Crossformer attempts to construct the unknown prediction sequence from fine to coarse, whose irrationality is self-evident.

B. Model Architecture

The overview of our proposed FPPformer is illustrated in Fig. 2 and its major enhancement on vanilla TSFT concentrates on addressing the preceding two problems of decoder. Comparing the schematics in Figs. 1 and 2, the differences with respect to the overall architecture can be readily noticed. To handle the first redundant self-attention problem in decoder, we *change the order of the self-attention and cross-attention* in decoder. Thereby, before embarking upon deducing any relations within unknown prediction sequence, the prediction sequence receives *the auto-regressive parts from the deepest encoder feature map*, which serves as a better role for prediction sequence initialization before the first self-attention in decoder than simple zero-initialization with start token [15], randomly generated parameters [8], the trend decomposition of raw input sequence [25], and so forth. It is evident that the latter initialization formats of other TSFTs are either relatively simple or inefficient.

We employ the hierarchical pyramids both in encoder and decoder with lateral connections to tackle the second multiscale insufficiency problem. As we adopt the patchwise attention in FPPformer, the patches are merged before sent to the next stages in the bottom-up architecture of encoder and opposite operations, i.e., the splitting, are performed in the top-down architecture of decoder. The feature map of input sequence gets deeper and more coarse-grained within later stages, which is also the property shared by encoders of many TSFTs. However, things get different when we attempt to construct the prediction sequence from the position embedding

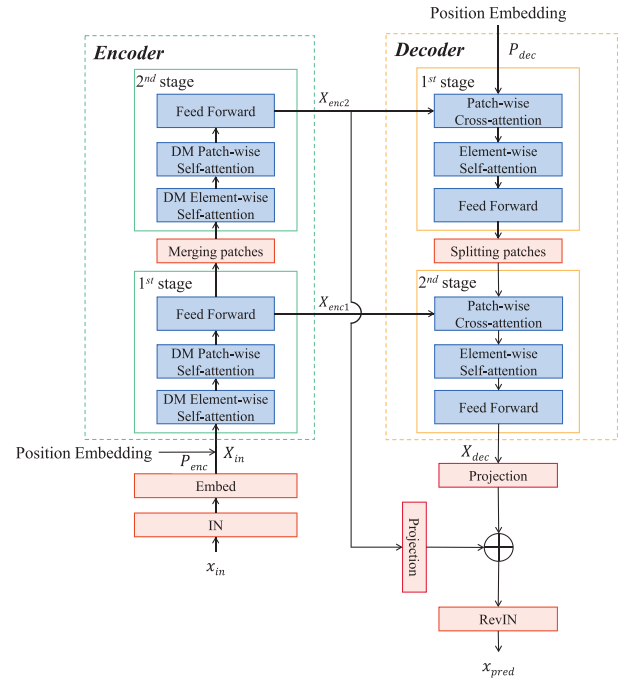


Fig. 2. Overview of FPPformer’s hierarchical architecture with two-stage encoder and two-stage decoder. Different from the vanilla one in Fig. 1, the encoder owns bottom-up structure while the decoder owns top-down structure. Note that the direction of the propagation flow in decoder is opposite to that in Fig. 1 to highlight the top-down structure. “DM” in the stages of encoder means “diagonal masked.”

and encoder feature maps. Recall how we decompose and reconstruct an arbitrary sequence from a certain multiresolution analysis $\{V_j\}_{j \in \mathbb{Z}}$ of $L^2(\mathbb{R})$ and wavelet spaces $\{W_j\}_{j \in \mathbb{Z}}$ in wavelet theory [35], which owns a transcendent position in signal processing. When we decompose certain sequence $f_{i+1}(t) \in V_{j+1}$, we decompose it into coarser spaces V_j and W_j . Whereas the reconstruction is opposite, i.e., we recover the sequence in finer space $f_{i+2}(t) \in V_{j+2}$ from V_{j+1} and W_{j+1} . Omitting the existence of wavelet space W_j , which contains the information of details or noises, we can find that the encoder and decoder processes separately correspond to the decomposition and (re)construction processes in multiresolution analysis. From another perspective, the unknown prediction sequence is initialized with zero or other parameters not pertaining to the ground truths at first. Thereby, when we strive to construct it from input sequence features, it is natural to commence with the most universal features to ensure the exactitude of general characteristics of prediction sequence features, then we can prudently take steps to seek finer features of prediction sequence to avoid over-fitting. The success of FPN [19], which also employs bottom-up and top-down architectures, in CV fields further confirms the preceding idea. Therefore, we keep splitting the patches in decoder and commence the hierarchical prediction sequence constructions with the feature maps from encoder, separately with identical resolutions, via lateral connections. The encoder in FPPformer presents a bottom-up architecture while the decoder presents a top-down architecture. Differences between the hierarchical design in Crossformer [8], whose decoder architecture is




Fig. 3. Comparison of hierarchical architecture of (a) FPPformer and (b) Crossformer. The discrepancies are highlighted with red. Obviously, the decoder structure of Crossformer is nearly a duplicate of encoder's bottom-up structure whereas the decoder of FPPformer owns a different "top-down" structure.

merely a replication of that encoder.¹ Shabani et al. [9] also noticed the analogous thing but they neither expound the reasons of doing so nor they carry out any change to decoder.

C. Combined Elementwise Attention and Patchwise Attentions

The preceding two problems are shared by the majority of TSFTs and we would like to mention another specific problem of TSFT with patchwise attention. We name this type of transformer PTSFT for brevity. Different from the elementwise attention of vanilla TSFT, which seeks the correlations of sequence elements, the patchwise attention seeks the correlations of different patches or segments of sequences to improve the efficiency and reduce the risk of over-fitting. PatchTST [16] and Crossformer [8] are such PTSFTs and their experiments have proven the superiority of patchwise attention. However, they neglect the inner-relations of the elements inside the patches or only employ simple linear projections to mix them up. Therefore, we make different changes in employed patchwise attention in encoder and decoder for the sake of extracting more fine-grained features in encoder or pursuing finer prediction sequence construction in decoder.

As shown in the schematic of Fig. 2, a elementwise attention block is arranged before each patchwise block in every encoder stage to extract the inner-relationships of all patches before seeking their inter-relationship. This elementwise self-attention is patch-independent so that the additional complexity is $O(P^2 \times L/P) = O(L \times P)$, which is linear with input sequence length. P is the patch size and L is the input sequence length in the last complexity expression. Observing that the elementwise attention requires the preservation of independent sequence elements information, therefore we cannot directly map the initially segmented patches into the latent space like other PTSFTs otherwise the elementwise information is no longer preserved and elementwise attention cannot be implemented. To address this issue, we adopt a channel-independent and element-independent embedding method. As illustrated in Fig. 4 and Table I, the input sequence elements of different timestamps do not interfere with each other during the embedding process and reshaping operations

¹In effect, the decoder of Crossformer even does not own a pyramid architecture. We use this statement since the hierarchical process of constructing the unknown prediction sequence is determined by how the model hierarchically uses the encoder features and for the more vivid comparison between FPPformer and Crossformer in Fig. 3.




Fig. 4. Changes in the size of a single input sequence when propagating through the first encoder stage. The batch size and the variable dimension are omitted. The red and blue letters in the last two sizes separately refer to the token dimension and its latent representation dimension. The reshaping operation is used to treat the features of all elements in a single patch as a unity for the sake of connecting elementwise self-attention and patchwise attention.

TABLE I
ARCHITECTURE OF THE FIRST STAGE IN ENCODER

Layer/Operation Name	Output Size
Input	L
Embedding	$L \times D$
Segmentation	$S \times P \times D, P = L/S$
Element-wise self-attention	$S \times P \times D$
Reshape	$S \times PD$
Patch-wise self-attention	$S \times PD$

are performed for the different needs of tensor shapes of elementwise attention and patchwise attention.

The changes to decoder's attention are analogous but not completely. Since decoder itself has already owned two attention blocks, we maintain the patchwise attention in cross-attention block to ensure the general construction of prediction sequence via auto-regressive process. Simultaneously, the masked self-attention block of vanilla TSFT is transformed to elementwise self-attention block, which is also patch-independent, in FPPformer. Just as we mentioned in the preceding sections, the prediction sequence is unknown so that we need to foremost guarantee the correctness of its general characteristics, manifested by placing patchwise attention before the elementwise attention in decoder, then we can pursue the fine-grained features of prediction sequences without over-fitting. As the patchwise cross-attention treats each patch as a unity, the respect to causality within prediction sequence, i.e., the masking to the upper triangular parts of query-key match matrix, is superfluous in the decoder of FPPformer.

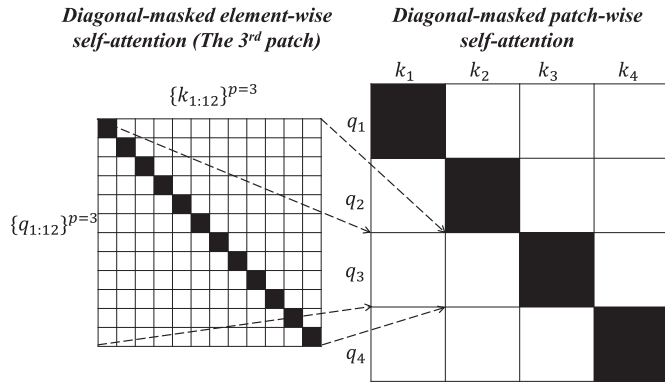


Fig. 5. Example of query-key matching matrices in DM elementwise self-attention and patchwise self-attention. The total sequence length is 48 and the sequence is divided into four patches, each with the length of 12, in this example. The white hollow boxes denote the normal unchanged matrix elements while the black solid boxes, i.e., the matrix elements at the diagonal, denote the masked matrix elements.

D. Diagonal-Masked Self-Attention

It is known that outliers always occur in real-world systems, especially for IoT systems owning immense and diverse data. These anomalies sometimes exist in the form of small patches [36] so that patch-attention cannot be immune to them. Compared with smoothing with filters [6], [7], which is natural but not flexible enough, it is better to devise mechanisms inside the networks to circumvent the negative effects of outliers in the latent space. Representation learning is a fair answer [16] but needs heavy parameter tuning and not very stable. Directly masking the input sequence [37] gives rise to another problem analogous to the preceding position embedding problem in decoder since the fixed parameters cannot sufficiently represent the dynamic sequence features. Enlightened by [38], we mask the diagonal of query-key matching matrix of both elementwise and patchwise self-attention blocks in encoder, as sketched in Fig. 5. Thereby, any element (patch) during the attention can merely be expressed by the values of the rest of the elements (patches). Those elements or patches whose characteristics confine with the general ones are scarcely affected but the outliers are impossible to be expressed by normal elements (patches), hence their values are restored to approach the general level and the negative effects of them are mitigated.

E. Projection

The prediction sequence is acquired through the summation of the linear projections of encoder output and decoder output. The first linear projection is supposed to represent the linear correlations of input and prediction sequence while the second linear projection, together with the entire decoder, is supposed to represent the nonlinear correlations. The loss function (2) is the summation of MSE function (3) and MAE function (4) according to [39], [40]

$$\text{Loss} = \text{MSE}(x_{t_2:t_3}^{1:V}, y_{t_2:t_3}^{1:V}) + \text{MAE}(x_{t_2:t_3}^{1:V}, y_{t_2:t_3}^{1:V}) \quad (2)$$

$$\text{MSE}(x, y) = \frac{1}{n} \sum_{i=1}^n (x_i - y_i)^2 \quad (3)$$

TABLE II
NUMERICAL DETAILS OF EIGHT MULTIVARIATE DATA SETS

Datasets	Sizes	Dimensions	Frequencies
ETTh ₁	17420	7	1h
ETTh ₂	17420	7	1h
ETTM ₁	69680	7	15min
ETTM ₂	69680	7	15min
ECL	26304	321	1h
Traffic	17544	862	1h
Weather	52696	21	10min
Solar	52560	137	10min

$$\text{MAE}(x, y) = \frac{1}{n} \sum_{i=1}^n |x_i - y_i|. \quad (4)$$

V. EXPERIMENTS

We attempt to answer three questions via the experiments on FPPformer.

- 1) Can FPPformer outperform temporarily state-of-the-art TSFTs and TSFMs on commonly used benchmarks with the settings of both short input sequence length and long input sequence length (Section V-C)?
- 2) Are the unique mechanisms proposed to be applied in FPPformer literally effective or useful (Section V-D) and what is about their parameter sensitivity (Section V-E)?
- 3) Why does FPPformer own better or worse performances than other baselines? Can we figure it out via visualization (Section V-G)?

A. Baselines and Data Sets

To unveil the empirical forecasting capability of FPPformer, we perform multivariate forecasting experiments on eight benchmarks involved in four types of IoT systems, including electricity consumption (ETTh₁, ETTh₂, ETTm₁, ETTm₂ [15], and ECL [41]), traffic flow (traffic [42]), meteorological conditions (weather [43]), and solar power production (solar [44]). Their numerical details are presented in Table II. Eight temporarily state-of-the-art forecasting baselines, including four TSFTs (Triformer [10], Crossformer [8], Scaleformer [9], and PatchTST [16]) and two TSFMs (FiLM [6] and TSMixer [30]), are employed to make comparison with FPPformer. It is worth mentioning that they are all superb forecasting methods proposed in the recent two years. Especially, besides Scaleformer, the other three of four TSFTs are PTSFTs so that FPPformer does not have an edge on pure attention mechanism design. Furthermore, we notice that these six forecasting baselines own different variable treatment strategies, e.g., some of them are channel-independent while some of them are not, which means that multivariate forecasting results cannot fully typify their forecasting capabilities. Therefore, we additionally perform univariate forecasting results on M4 data set [45], which is a competition data set qualified for univariate forecasting, rather than deliberately choosing a variate within the above multivariate forecasting data sets to perform univariate forecasting experiments like many other researches [6], [8], [9], [10]. Its details are elaborated in Table III.

TABLE III
NUMERICAL DETAILS OF FOUR M4 SUBDATA SETS

Sub-datasets	Prediction Lengths	Periodicities	Frequencies	Instances
M4-Monthly	18	12	1month	48,000
M4-Weekly	13	1	1week	359
M4-Daily	14	1	1day	4,227
M4-Hourly	48	24	1hour	414

B. Implementation Details

We would like to present a persuasive and fair comparison of FPPformer and other baselines, therefore we set the hyperparameters of FPPformer identical to the commonly used ones. The input sequence lengths are different in different sub-experiments but are kept identical for all baselines. The number of stages are 3 in both encoder and decoder of FPPformer, and the size of the initial segmented patch is 6, which are in accordance with those of Crossformer [8]. The embedding dimension D (in Fig. 4) is 32. As for hyperparameters with respect to the training process, FPPformer is trained via an Adam optimizer with the learning rate of $1e-4$, which decays by half per epoch with totally ten epochs and the patience of one. The batch size is 16 and the dropout rate is 0.1. These are all commonly employed settings. All experiments, which are conducted on a single NVIDIA GeForce RTX 3090 24-GB GPU, are repeated for five times with casual seeds and the average results are presented. The source codes are implemented by Python 3.8.8 and Pytorch 1.11.0 in <https://github.com/OrigamiSL/FPPformer>. Correspondingly, the other baselines used in this work also merely employ the fixed hyperparameters and settings, which are chosen after referencing their default ones. As for those with multiple choices and versions, we choose the one that owns the best general performance. How we use the other baselines for experiment can all be found in our provided GitHub repository. The best results in each table are highlighted with *bold and italic* and the second best are highlighted with *underline and italic*, barring a special table in Section V-D.

C. Quantitative Results

We commence with the multivariate forecasting experiments, whose results are shown in Tables IV and V. Under many real-world occasions, training samples are limited so that long input sequence length is not always available for some deep forecasting methods needing it for satisfactory performances. Therefore, we first measure the performances of FPPformer and its six competitors in eight multivariate benchmarks with input sequence length of 96, which is ascribable to well-known Autoformer [24] in Table IV. The prediction lengths are commonly agreed-upon {96, 192, 336, 720}. Then we evaluate the performances of the same seven models and data sets using longer input sequence length within {192, 384, 576}, whose results are shown in Table V. MSE (3) and MAE (4) are utilized as the evaluation metrics. The average results of eight benchmarks with prediction length of 720 are given in Table V to refrain from tedious data stacking. Full results are given in our released repository provided in

Section V-B. “—” refers to the fact that certain model is out of the memory (24GB) even batch size is set to 1.

As Tables IV and V show, FPPformer outperforms other baselines in most of situations with both short and long input sequence lengths. When input sequence length is set to 96, FPPformer obtains 31.7%/60.0%/10.5%/6.4%/12.8%/14.7% MSE reduction compared with Triformer/Crossformer/Scaleformer/PatchTST/FiLM/TSMixer, which illustrates the superb forecasting capability of FPPformer with the setting of short input sequence length. Though it seems that FPPformer fails to own a superior performance when experimenting on Solar data set, FPPformer reconquers its leading position with longer input length when handling the same data set (concrete results are available at github repository provided in Section V-B). Furthermore, if also equipped with cross-variable attention² like Crossformer, which means that a cross-variable attention module proposed by Crossformer is arranged at the end of each stage of the encoder and decoder in FPPformer, the modified FPPformer, denoted by FPPformer-Cross in Table VI, is capable of completely outperforming Crossformer and itransformer [46], which is another state-of-the-art model employing cross-variable attention, under Solar data set.

When prolonging the input sequence length within {192, 384, 576}, FPPformer achieves better forecasting performances and, respectively, obtains 9.5%/12.3%/13.0% general MSE reduction when compared with the FPPformer with short input length in Table IV. Moreover, FPPformer obtains 67.4%/72.3%/40.2%/3.5%/9.5%/7.5% MSE reduction compared with Triformer/Crossformer/Scaleformer/PatchTST/FiLM/TSMixer with longer input lengths in general. These phenomena illustrate the superb forecasting performances of FPPformer especially handling longer input sequence.

Then we compare the univariate forecasting capability of FPPformer with other six baselines on M4. We omit the first two subsets with sampling frequencies of a year and a quarter since many of their instance lengths are too short, and only perform experiments on the rest of four subsets. The prediction lengths, which are regulated by [45], are {18, 13, 14, 48} for {M4-monthly, M4-weekly, M4-daily, M4-hourly}. The input sequence lengths for them are separately {72, 65, 84, 336} after consulting [29]. We change these four input sequence lengths a little to {72, 72, 96, 384} for the sake of rendering them fitting the patchwise attention in FPPformer. The M4-specific metrics SMAPE (5) and OWA (7) are used for measurement. m refers to the periodicity of series and naïve2 refers to the results of a seasonally adjusted forecast model by [45] for scaling in OWA

$$\text{SMAPE} = \frac{200}{t_3 - t_2} \sum_{t \in [t_2, t_3)} \frac{|y_t - x_t|}{|y_t| + |x_t|} \quad (5)$$

$$\text{MASE} = \frac{1}{t_3 - t_2} \sum_{t \in [t_2, t_3)} \frac{|y_t - x_t|}{\frac{1}{t_3 - t_1 - m} \sum_{j \in [t_1 + m + 1, t_3)} |x_j - x_{j-m}|} \quad (6)$$

²We use the statement of “cross-variable,” rather than “cross-dimension” in Crossformer, to maintain the identical description of the variable dimension throughout this work.

TABLE IV
MULTIVARIATE FORECASTING RESULTS WITH SHORT INPUT LENGTH

Methods	Datasets	ETTh ₁				ETTh ₂				ETTm ₁				ETTm ₂			
		Metrics/Prediction Lengths	96	192	336	720	96	192	336	720	96	192	336	720	96	192	336
FPPformer	MSE	0.373	0.425	0.470	0.479	0.296	0.372	0.418	0.422	0.313	0.362	0.393	0.448	0.176	0.243	0.302	0.398
	MAE	0.391	0.421	0.442	0.463	0.343	0.392	0.427	0.435	0.350	0.377	0.401	0.437	0.258	0.301	0.340	0.396
Triformer	MSE	0.419	0.484	0.513	0.711	0.742	1.028	1.049	1.223	0.362	0.419	0.466	0.531	0.240	0.387	0.545	1.928
	MAE	0.446	0.486	0.489	0.638	0.585	0.708	0.745	0.836	0.402	0.443	0.484	0.513	0.326	0.449	0.532	0.924
Crossformer	MSE	0.472	0.533	0.698	0.847	0.884	2.835	2.339	4.387	0.444	0.522	0.649	0.818	0.350	0.643	1.051	3.988
	MAE	0.493	0.516	0.602	0.713	0.690	1.381	1.241	1.810	0.467	0.524	0.592	0.672	0.421	0.596	0.774	1.426
Scaleformer	MSE	0.443	<u>0.430</u>	0.457	0.498	0.328	0.434	0.489	0.474	<u>0.327</u>	0.394	0.416	<u>0.467</u>	0.175	0.251	0.316	0.438
	MAE	0.460	<u>0.451</u>	0.472	<u>0.503</u>	0.380	0.448	0.490	0.495	0.381	0.422	0.441	0.474	0.268	0.319	0.361	0.436
PatchTST	MSE	<u>0.393</u>	0.438	0.482	<u>0.492</u>	<u>0.296</u>	<u>0.380</u>	0.417	<u>0.425</u>	0.332	<u>0.377</u>	<u>0.407</u>	0.473	0.177	<u>0.243</u>	<u>0.305</u>	<u>0.401</u>
	MAE	0.427	0.459	0.494	0.519	0.346	<u>0.395</u>	<u>0.428</u>	<u>0.443</u>	<u>0.366</u>	<u>0.385</u>	<u>0.405</u>	0.442	0.274	0.324	0.363	0.420
FiLM	MSE	0.423	0.474	0.510	0.557	0.299	0.384	0.425	0.443	0.354	0.390	0.422	0.481	0.183	0.249	0.310	0.409
	MAE	<u>0.423</u>	0.452	<u>0.468</u>	0.521	<u>0.345</u>	0.397	0.436	0.454	0.371	0.387	0.408	0.441	<u>0.266</u>	0.307	<u>0.344</u>	0.399
TSMixer	MSE	0.441	0.491	0.531	0.536	0.320	0.406	0.438	0.455	0.339	0.380	0.410	0.472	0.184	0.248	0.308	0.406
	MAE	0.441	0.468	0.488	0.509	0.365	0.415	0.446	0.463	0.370	0.388	0.407	<u>0.441</u>	0.267	<u>0.307</u>	0.345	<u>0.399</u>

Methods	Datasets	ECL				Traffic				Weather				Solar			
		Metrics/Prediction Lengths	96	192	336	720	96	192	336	720	96	192	336	720	96	192	336
FPPformer	MSE	0.179	0.186	0.203	0.246	0.481	0.479	0.486	0.518	<u>0.174</u>	<u>0.219</u>	0.275	0.352	0.234	0.263	0.270	0.273
	MAE	0.254	0.264	0.281	0.317	0.295	0.290	0.291	0.309	0.210	0.251	0.291	0.342	0.261	0.276	<u>0.277</u>	0.281
Triformer	MSE	<u>0.179</u>	<u>0.191</u>	<u>0.209</u>	<u>0.253</u>	—	—	—	—	0.174	0.219	<u>0.272</u>	0.357	<u>0.185</u>	<u>0.203</u>	<u>0.241</u>	<u>0.220</u>
	MAE	0.277	0.289	0.308	0.344	—	—	—	—	0.242	0.290	0.323	0.378	<u>0.247</u>	<u>0.253</u>	0.285	<u>0.269</u>
Crossformer	MSE	0.234	0.285	0.342	0.543	<u>0.532</u>	<u>0.543</u>	<u>0.562</u>	<u>0.571</u>	0.156	0.210	0.260	0.355	0.167	0.203	0.216	0.216
	MAE	0.323	0.351	0.381	0.471	<u>0.295</u>	<u>0.296</u>	<u>0.303</u>	<u>0.321</u>	<u>0.227</u>	0.278	0.321	0.385	0.214	0.239	0.251	0.249
Scaleformer	MSE	0.180	0.195	0.215	0.257	0.551	0.578	0.587	0.611	0.198	0.313	0.395	0.595	0.206	0.218	0.266	0.283
	MAE	0.297	0.308	0.328	0.363	0.341	0.350	0.354	0.364	0.273	0.372	0.420	0.549	0.305	0.303	0.330	0.333
PatchTST	MSE	0.196	0.199	0.214	0.256	0.589	0.579	0.589	0.631	0.188	0.231	0.285	0.359	0.240	0.274	0.295	0.295
	MAE	0.279	0.283	<u>0.299</u>	<u>0.331</u>	0.379	0.368	0.371	0.391	0.228	<u>0.263</u>	<u>0.301</u>	<u>0.349</u>	0.285	0.308	0.317	0.317
FiLM	MSE	0.198	0.199	0.217	0.279	0.648	0.600	0.610	0.694	0.197	0.240	0.290	0.361	0.314	0.355	0.395	0.397
	MAE	<u>0.274</u>	<u>0.278</u>	0.300	0.357	0.385	0.362	0.367	0.427	0.239	0.272	0.307	0.352	0.334	0.356	0.371	0.364
TSMixer	MSE	0.222	0.226	0.244	0.287	0.725	0.706	0.727	0.775	0.191	0.235	0.286	0.358	0.263	0.302	0.332	0.335
	MAE	0.308	0.315	0.332	0.363	0.453	0.449	0.458	0.477	0.236	0.271	0.307	0.352	0.305	0.328	0.345	0.345

TABLE V
AVERAGE MULTIVARIATE RESULTS WITH LONG INPUT LENGTHS

Methods	Metrics/Input Lengths	192	384	576
FPPformer	MSE (Avg. of all)	0.358	0.347	0.345
	MAE (Avg. of all)	0.354	0.352	0.350
Triformer	MSE (Avg. of all)	0.932	0.966	1.546
	MAE (Avg. of all)	0.621	0.639	0.828
Crossformer	MSE (Avg. of all)	1.319	1.192	1.290
	MAE (Avg. of all)	0.729	0.707	0.747
Scaleformer	MSE (Avg. of all)	0.605	0.555	0.598
	MAE (Avg. of all)	0.539	0.517	0.520
PatchTST	MSE (Avg. of all)	<u>0.377</u>	<u>0.357</u>	<u>0.354</u>
	MAE (Avg. of all)	<u>0.378</u>	<u>0.368</u>	<u>0.366</u>
FiLM	MSE (Avg. of all)	0.397	0.381	0.383
	MAE (Avg. of all)	0.391	0.385	0.390
TSMixer	MSE (Avg. of all)	0.397	0.372	0.367
	MAE (Avg. of all)	0.392	0.379	0.377

$$\text{OWA} = \frac{1}{2} \left(\frac{\text{SMAPE}}{\text{SMAPE}_{\text{Naive2}}} + \frac{\text{MASE}}{\text{MASE}_{\text{Naive2}}} \right). \quad (7)$$

The univariate forecasting results are shown in Table VII. It could be observed that FPPformer gains 88.6%/47.8%/21.0%/54.1%/50.8%/20.8% MSE reduction in general when compared with Triformer/Crossformer/Scaleformer/PatchTST/FiLM/TSMixer, which is more persuasive to verify the forecasting capability of FPPformer.

TABLE VI
COMPARISON OF MODELS EMPLOYING CROSS-VARIABLE ATTENTION

Methods (Dataset: Solar)	Metrics/Prediction Length	96	192	336	720
FPPformer-Cross	MSE	0.162	0.174	0.191	0.199
	MAE	0.209	0.232	0.237	0.238
iTransformer	MSE	0.215	0.247	0.267	0.268
	MAE	0.255	0.279	0.294	0.297
Crossformer	MSE	<u>0.167</u>	<u>0.203</u>	<u>0.216</u>	<u>0.216</u>
	MAE	<u>0.214</u>	<u>0.239</u>	<u>0.251</u>	<u>0.249</u>

D. Ablation Study

We conduct ablation studies to validate the functions of the architecture of FPPformer and its components. All variants are experimented with multivariate benchmarks with prediction sequence length of 720. The results of all eight benchmarks are presented in Table VIII. Each value is the average of four subexperiment results with input sequence lengths of {96, 192, 384, 576}. As expected, only using pointwise attention like canonical TSFT gives rise to a losses increasing of 84.6% (Average loss: 0.345→0.637). Meanwhile, only using patchwise attention like other PTSFTs does not suffer a severe performance degradation but still owns apparently worse performance than our proposed combined patchwise attention and pointwise attention (Average loss: 0.380 versus 0.345). As for the decoder architecture design, FPPformer




Fig. 6. Results of training time per epoch (1), GPU memory computation (2), and inference time per instance (3) in two experiments (a) and (b) of Complexity Analysis Experiment. FPPformer-Enc in (b) refers to FPPformer without decoder.

TABLE VII
UNIVARIATE FORECASTING RESULTS

Methods	Metrics	M4-Monthly	M4-Weekly	M4-Daily	M4-Hourly
FPPformer	SMAPE	8.674	8.711	3.214	17.468
	OWA	0.891	1.464	1.057	1.373
Triformer	SMAPE	65.161	106.905	82.288	78.169
	OWA	11.068	19.181	41.565	5.756
Crossformer	SMAPE	24.801	12.702	6.901	28.471
	OWA	2.704	3.021	3.077	2.706
Scaleformer	SMAPE	14.275	<u>10.591</u>	<u>3.362</u>	19.939
	OWA	1.163	1.409	1.163	2.017
PatchTST	SMAPE	15.127	15.687	4.457	47.657
	OWA	1.237	2.679	1.519	7.526
FiLM	SMAPE	32.329	15.029	8.376	21.681
	OWA	3.489	3.966	3.899	2.240
TSMixer	SMAPE	<u>13.886</u>	10.962	3.429	<u>19.769</u>
	OWA	<u>0.947</u>	1.515	<u>1.131</u>	<u>1.516</u>

surpasses the same models but replacing the decoder architecture of FPPformer with that of Crossformer (Average loss: 0.345 versus 0.389) or simply substituting the entire decoder with a linear projection like PatchTST (Average loss: 0.345 versus 0.376). Besides, removing DM mechanisms in the self-attention blocks of encoder also results in worse performances (Average loss: 0.345→0.379). Conclusively, the efficiency and necessity of all unique parts and architectures of FPPformer are verified.

E. Parameter Sensitivity Analysis

It is well-known that TSFT cannot own too many layers or stages, otherwise the risk of over-fitting substantially rises. Thereby, we perform parameter analysis on the stage number of FPPformer in this section to check out whether FPPformer is capable of handling this problem. The parameter analysis of patch size is no longer tested as it has been well studied by [8], [10], and [16]. The input sequence length is chosen as 576 for using more stages for FPPformer and other models use their default input sequence lengths, which are supposed to be best for them (96 for {Scaleformer, Triformer}; 512 for PatchTST; 336 for Crossformer). The number of stages are chosen in {1, 2, 3, 4}. The prediction sequence length is set as 720 and the average results (MSEs) of all eight

multivariate benchmarks are presented in Table IX. The result of each model with stage number of one is utilized as the normalization factor to further measure their absolute performance deviations with more stage numbers which are manifested by

$$\text{Dev} = \sum_{i=2}^4 |\tilde{M}_i - \tilde{M}_1|, \tilde{M}_i = M_i/M_1. \quad (8)$$

$M_i (i = 1, 2, 3, 4)$ are the original average MSE results and $\tilde{M}_i (i = 1, 2, 3, 4)$ are the normalized ones. When comparing the performances and performance deviations of five TSFTs, it is evident that FPPformer not only keeps its leading position (smaller errors) among all TSFTs with different stage numbers but also maintains its robustness ascendancy (smaller deviations) over other TSFTs.

F. Complexity Analysis

We compare the training time per epoch/the GPU memory computation/the inference time per instance on GPU of FPPformer and Crossformer [8], and the identical four measuring criteria of FPPformer without decoder and PatchTST [16] during multivariate forecasting under Solar data set. Solar data set is selected in that it owns the intermediate variable number among all data sets. Crossformer and PatchTST are chosen as they are also patchwise attention-based models. The decoder of FPPformer is removed when compared with PatchTST since PatchTST does not employ the decoder architecture. The input sequence lengths are chosen within {96, 192, 384, 576} and the prediction sequence length is 96. The other hyperparameters and settings are identical with those used in the quantitative multivariate results, barring the size of the hidden(embedding) dimension. The size of the hidden dimension, which is the one that exceedingly affects the model complexity, is identical for all baselines in this experiment so that the model architecture design can determine the model complexity to the utmost extent. These two experiment results are shown in Fig. 6(a) and (b), respectively.

As Fig. 6(a) shows, the full FPPformer owns linear computation and space complexity with input sequence length. Moreover, Fig. 6(b) illustrates that the encoder of FPPformer (FPPformer-Enc for short) also only owns linear complexity

TABLE VIII
ABLATION RESULTS WITH PREDICTION LENGTHS OF 720

Architecture	Attention	Decoder Structure	DM	ETTh ₁	ETTh ₂	ETM ₁	ETM ₂	ECL	Traffic	Weather	Solar	Avg. MSE(Increase rate)
FPPformer	Point-wise	top-down	✓	0.706	0.446	0.708	0.404	0.818	1.298	0.343	0.371	0.637(+84.5%)
	Patch-wise	top-down	✓	0.479	0.422	0.441	0.384	0.227	0.501	0.336	0.248	0.380(+10.0%)
	Point-wise & Patch-wise	bottom-up (e.g. Crossformer)	✓	0.471	0.427	0.446	0.405	0.235	0.522	0.340	0.264	0.389(+12.6%)
	Point-wise & Patch-wise	Linear (e.g. PatchTST)	✓	0.473	0.425	0.437	0.385	0.229	0.488	0.238	0.333	0.376(+9.0%)
(Complete FPPformer)	Point-wise & Patch-wise	top-down	×	0.469	0.421	0.439	0.385	0.227	0.500	0.337	0.254	0.379(+9.8%)
	Point-wise & Patch-wise	top-down	✓	0.449	0.394	0.411	0.366	0.199	0.432	0.316	0.194	0.345

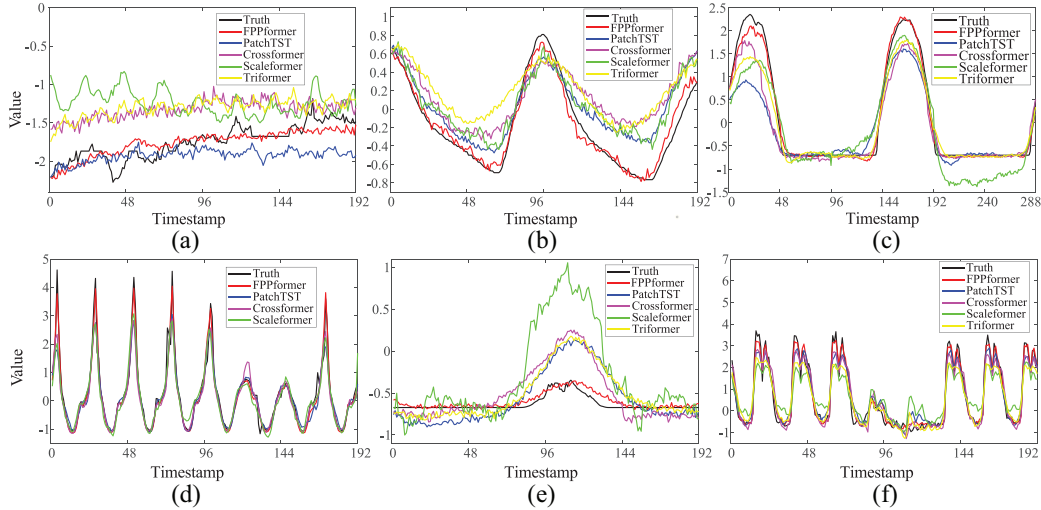


Fig. 7. Forecasting windows of FPPformer and other four TSFTs from six data sets. The black line is the ground truth, the red line is the forecasting curve of FPPformer, and the rest are the forecasting curves of other TSFTs. (a) ETTh₁. (b) ETM₂. (c) Solar. (d) Traffic. (e) Weather. (f) ECL.

TABLE IX
MSE RESULTS OF PARAMETER SENSITIVITY ON STAGE NUMBERS

Methods/Stages	$M_1(\bar{M}_1)$	$M_2(\bar{M}_2)$	$M_4(\bar{M}_3)$	$M_4(\bar{M}_4)$	Dev
FPPformer	0.349(1.00)	0.344(0.99)	0.345(0.99)	0.351(1.01)	0.03
Scaleformer	0.417(1.00)	0.425(1.02)	0.434(1.04)	0.433(1.04)	0.10
Triformer	0.787(1.00)	0.797(1.01)	0.964(1.23)	0.814(1.03)	0.27
PatchTST	0.370(1.00)	0.369(1.00)	0.385(1.04)	0.382(1.03)	0.07
Crossformer	1.475(1.00)	1.339(0.91)	1.466(0.99)	1.498(1.02)	0.12

and the complexity of FPPformer and PatchTST are analogous, demonstrating that the elementwise attention, which is almost the only difference between PatchTST and FPPformer-Enc, merely brings minuscule additional complexity.

G. Case Study

To more vividly illustrate the outstanding forecasting performance of FPPformer, we visualize several forecasting windows of FPPformer and other TSFTs from different data sets in Fig. 7. Benefitting from exploiting better decoder and attention mechanism, FPPformer excels in capturing the features of trends [Fig. 7(a)], seasons [Fig. 7(b)] and their hybrids [Fig. 7(c)] so that smaller forecasting errors than others can be obtained. Moreover, its preponderance of robustness and immunity against outliers are revealed in Fig. 7(d)–(f) where certain distribution shifts occur in partial distincts. In

addition, we present some visualizations of the feature maps of FPPformer and several competitors in the latent space to validate the functions of its unique modules.

- 1) We visualize the attention score distribution of the first DM patchwise and DM elementwise attention in FPPformer with a certain input window of length 96 in ETTh₁ data set, via heat map in Fig. 8. As illustrated in Fig. 8(b1), the attention score distribution is uniformly distributed if applying the DM patchwise attention during the training phase. Even substituting the DM patchwise attention with the normal attention solely in the testing phase [Fig. 8(b2)] will not lead to the self-matches with high scores, demonstrating the enhancement of DM attention mechanism on universal feature extraction. However, it can be observed in Fig. 8(b3) that the highest attention score chiefly lies in the fifth patch, which corresponds to an outlier patch with a exorbitant dip. The visualization [Fig. 8(c)] of the three different elementwise attention score heat maps of that outlier patch in Fig. 8(b) also manifests the same phenomenon, i.e., the excessive high or low values can result in the excessive high-attention scores which burden the universal feature extraction.
- 2) As the additional elementwise attention is employed in FPPformer to reinforce the inner-patch feature extraction, we visualize the feature maps outputted by each

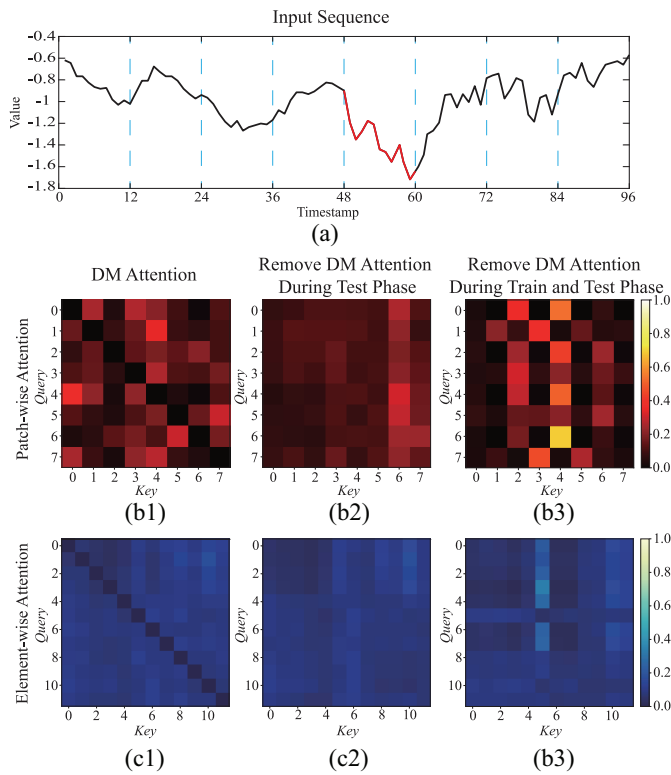


Fig. 8. Visualization of DM patchwise (b) and elementwise (c) self-attention score distributions via heat map. All figures are obtained by FPPformer or its variants with the identical input sequence (a). Especially, the elementwise attention score heat maps stem from the fifth patch, which is supposed to be anomalous and marked with red in (a).

patchwise attention in the encoders of FPPformer, PatchTST and Crossformer by T-SNE [47] in Fig. 9. To get rid of the influences from other different modules, this experiment is performed as a reconstruction experiment, which reconstructs the input sequences via the ultimate encoder features, and we remove the DM technique in the patchwise attentions of FPPformer but keep the DM elementwise attention. The input sequence length is set as 576 to provide enough data points in the T-SNE figures. Apparently, the distances of the data points in three FPPformer-related figures are the smallest among three models, indicating that the DM elementwise attentions can literally assist in better representing the features of each patch in the latent space so that it is easier for the following patchwise attentions in FPPformer encoder to extract the universal features among different patches and FPPformer can have an edge on extracting universal features compared with other two baselines.

- 3) To vividly illustrate that the top-down architecture in decoder can genuinely render the construction of prediction sequence feature maps more general in the latent space, we visualize the patchwise cross-attention score distributions of different decoder layers in FPPformer and Crossformer via heat maps in Fig. 10. Obviously, with the stage number grows, the attention score matrix size is getting bigger in FPPformer and smaller in Crossformer, illustrating the differences

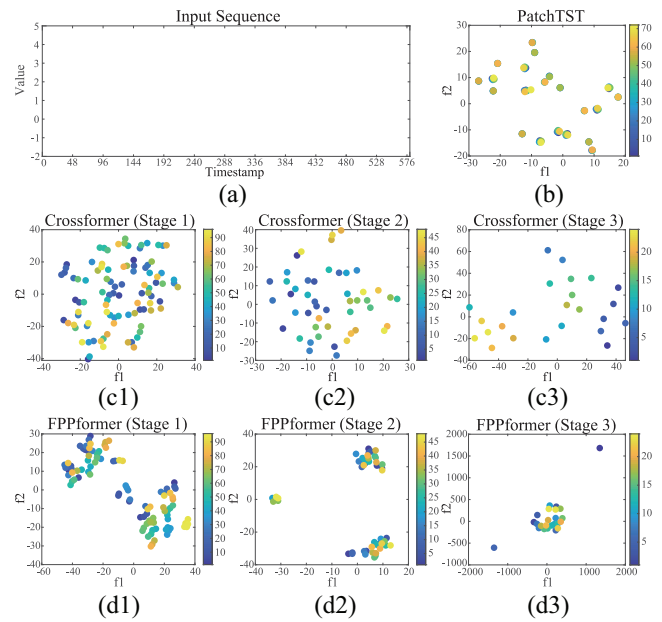


Fig. 9. (a) Visualization of the feature maps of the patches in different encoder stages (b) in PatchTST, (c) Crossformer, and (d) FPPformer via T-SNE. The points with different colors denote different patches. Only the feature map of the last encoder stage of PatchTST is shown since it is not hierarchical.

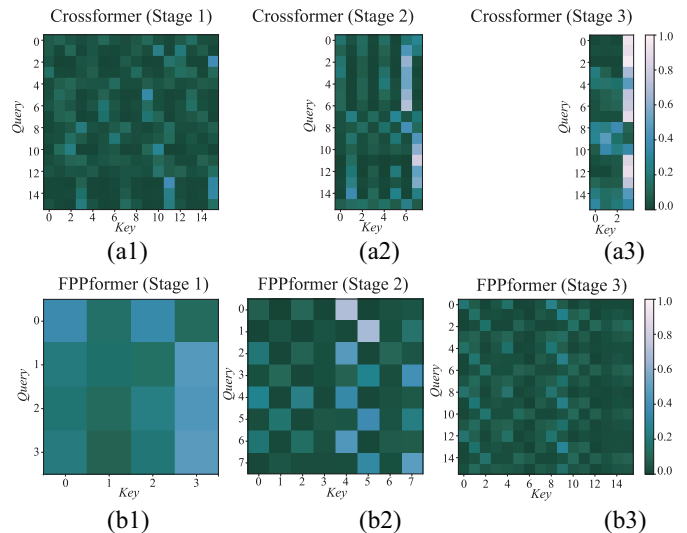


Fig. 10. Visualization of the cross-attention score distribution of different decoder stages in Crossformer (a) and FPPformer (b) via heat map.

between “bottom-up” and “top-down” decoder architecture. Meanwhile, it can be observed that the highest attention scores in Crossformer primarily locate at the last patch, especially the third stage, indicating that the construction of the prediction sequence features heavily rely on the rear end of input sequence features and it fails to build up the prediction sequence in an universal manner. On contrary, the cross-attention scores in the decoder uniformly distribute along the temporal dimension, which implies the preponderance of the top-down architecture in FPPformer decoder. In effect, the instance in Fig. 10. is not a particular situation. We collect the




Fig. 11. Distribution of the highest cross-attention score positions of the third stage in Crossformer (a) and the (b) first stage in FPPformer.

highest attention score positions of the third stage in Crossformer and the first stage in FPPformer, where the most coarse-grained features lie in, when handling the whole ETTh₁ data set, i.e., with over 100 000 instances. The result is shown in Fig. 11. Apparently, the highest attention score distribution of FPPformer is much more uniform than that of Crossformer, indicating the success of the top-down decoder design in FPPformer.

VI. DISCUSSION

Though FPPformer has achieved state-of-the-art performances, it still owns at least two limits.

- 1) The hierarchy in FPPformer can be more exquisitely devised. The “merging” operation in the encoder of FPPformer is too simple to well represent the feature map of the bigger patch via the two smaller patch ingredients. So does the “splitting” operation in the decoder. The cutting-edge methods to handle the combination or the split of patches, e.g., SwinTransformer [48], in CV field, where patchwise attention is also prevailing, can be learned, imitated and modified in TSFT.
- 2) Currently, the outlier is tackled via DM self-attention, which roughly mask the entire diagonal of the self-attention score matrix, in FPPformer. Notice that the outliers shall be fewer than the normal segments of time-series sequences, which implies that the majority of masked patches are indeed normal and the masking behavior can negatively influences the feature extraction of input sequences. We believe that applying a prior anomaly detection method to each input sequence before forecasting and then only masking the detected anomalous patches can be a better format of utilizing the DM self-attention.

Both of the foregoing two limits and potential solutions will be our future research directions.

VII. CONCLUSION

In this work, we attempt to further develop the TSFT from the perspective of decoder. We lucubrate the existing decoder designs, point out their drawbacks and propose our solutions. The ultimate product, i.e., FPPformer, achieves state-of-the-art performances in multiple benchmarks, including multivariate

and univariate ones, leveraging from refined attention mechanism and enhanced encoder–decoder architecture proposed by us.

REFERENCES

- [1] J. Han, G. H. Lee, S. Park, J. Lee, and J. K. Choi, “A multivariate-time-series-prediction-based adaptive data transmission period control algorithm for IoT networks,” *IEEE Internet Things J.*, vol. 9, no. 1, pp. 419–436, Jan. 2022.
- [2] G. E. P. Box, G. M. Jenkins, and J. F. MacGregor, “Some recent advances in forecasting and control,” *J. Roy. Statist. Soc.*, vol. 23, no. 2, pp. 158–179, 1974. [Online]. Available: <https://riss.onlinelibrary.wiley.com/doi/abs/10.2307/2346997>
- [3] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*. Hoboken, NJ, USA: Wiley, 2015.
- [4] K. Benidis et al., “Deep learning for time series forecasting: Tutorial and literature survey,” *ACM Comput. Surv.*, vol. 55, no. 6, pp. 1–36, Dec. 2022. [Online]. Available: <https://doi.org/10.1145/3533382>
- [5] A. Zeng, M. Chen, L. Zhang, and Q. Xu, “Are transformers effective for time series forecasting?” in *Proc. 37th AAAI Conf. Artif. Intell.*, 2023, pp. 1–8.
- [6] T. Zhou et al., “FiLM: Frequency improved legendre memory model for long-term time series forecasting,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 12677–12690. [Online]. Available: <https://openreview.net/forum?id=zTQdHSQUWc>
- [7] C. Challu, K. G. Olivares, B. N. Oreshkin, F. Garza, M. Mergenthaler-Canseco, and A. Dubrawski, “NHITS: Neural hierarchical interpolation for time series forecasting,” in *Proc. AAAI Conf. Artif. Intell.*, 2023, pp. 6989–6997.
- [8] Y. Zhang and J. Yan, “Crossformer: Transformer Utilizing cross-dimension dependency for multivariate time series forecasting,” in *Proc. 11th Int. Conf. Learn. Represent.*, 2023, pp. 1–21. [Online]. Available: <https://openreview.net/forum?id=vSVLM2j9eie>
- [9] M. A. Shabani, A. H. Abdi, L. Meng, and T. Sylvain, “Scaleformer: Iterative multi-scale refining transformers for time series forecasting,” in *Proc. 11th Int. Conf. Learn. Represent.*, 2023, pp. 1–23. [Online]. Available: <https://openreview.net/forum?id=sCnlllCtjoe>
- [10] R.-G. Cirstea, C. Guo, B. Yang, T. Kieu, X. Dong, and S. Pan, “Triformer: Triangular, variable-specific attentions for long sequence multivariate time series forecasting,” in *Proc. 31 Int. Joint Conf. Artif. Intell.*, (IJCAI), 2022, pp. 1994–2001. [Online]. Available: <https://doi.org/10.24963/ijcai.2022/277>
- [11] H. Wang, J. Peng, F. Huang, J. Wang, J. Chen, and Y. Xiao, “MICN: Multi-scale local and global context modeling for long-term series forecasting,” in *Proc. 11th Int. Conf. Learn. Represent.*, 2023, pp. 1–22. [Online]. Available: <https://openreview.net/forum?id=zt53IDUR1U>
- [12] T. Kim, J. Kim, Y. Tae, C. Park, J.-H. Choi, and J. Choo, “Reversible instance Normalization for accurate time-series forecasting against distribution shift,” in *Proc. Int. Conf. Learn. Represent.*, 2022, pp. 1–25. [Online]. Available: <https://openreview.net/forum?id=cGDAkQo1C0p>
- [13] Z. Liu, R. Godahewa, K. Bandara, and C. Bergmeir, “Handling concept drift in global time series forecasting,” 2023, *arXiv:2304.01512*. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257921602>
- [14] S. B. Taieb and A. F. Atiya, “A bias and variance analysis for multistep-ahead time series forecasting,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 1, pp. 62–76, Jan. 2016.
- [15] H. Zhou et al., “Informer: Beyond efficient transformer for long sequence time-series forecasting,” in *Proc. 35th AAAI Conf. Artif. Intell.*, 2021, pp. 11106–11115.
- [16] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, “A time series is worth 64 words: Long-term forecasting with transformers,” in *Proc. 11th Int. Conf. Learn. Represent.*, 2023, pp. 1–24. [Online]. Available: <https://openreview.net/forum?id=Jbdc0vTOcol>
- [17] Y. Liu, H. Wu, J. Wang, and M. Long, “Non-stationary transformers: Exploring the stationarity in time series forecasting,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 1–21. [Online]. Available: <https://openreview.net/forum?id=ucNDIDRNjiv>
- [18] S. Liu et al., “Pyrformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting,” in *Proc. Int. Conf. Learn. Represent.*, 2022, pp. 1–20. [Online]. Available: <https://openreview.net/forum?id=0EXmFzUn5I>
- [19] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature pyramid networks for object detection,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 936–944.

- [20] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, "Path aggregation network for instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8759–8768.
- [21] B. Huang, H. Dou, Y. Luo, J. Li, J. Wang, and T. Zhou, "Adaptive spatiotemporal transformer graph network for traffic flow forecasting by IoT loop detectors," *IEEE Internet Things J.*, vol. 10, no. 2, pp. 1642–1653, Jan. 2023.
- [22] Y. Jiang et al., "Spatial-temporal graph data mining for IoT-enabled air mobility prediction," *IEEE Internet Things J.*, vol. 9, no. 12, pp. 9232–9240, Jun. 2022.
- [23] S. Li et al., "Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 1–11. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/6775a0635c302542da2c32aa19d86be0-Paper.pdf>
- [24] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 22419–22430. [Online]. Available: <https://proceedings.neurips.cc/paper/2021/file/bcc0d400288793e8bdcd7c19a8ac0c2b-Paper.pdf>
- [25] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting," in *Proc. 39th Int. Conf. Mach. Learn.*, 2022, pp. 27268–27286. [Online]. Available: <https://proceedings.mlr.press/v162/zhou22g.html>
- [26] G. Woo, C. Liu, D. Sahoo, A. Kumar, and S. C. H. Hoi, "ETSformer: Exponential smoothing transformers for time-series forecasting," 2022, *arXiv:2202.01381*.
- [27] G. Lai, W.-C. Chang, Y. Yang, and H. Liu, "Modeling long-and short-term temporal patterns with deep neural networks," in *Proc. 41st Int. ACM SIGIR Conf. Res. amp; Develop. Inf. Retrieval*, 2018, pp. 95–104. [Online]. Available: <https://doi.org/10.1145/3209978.3210006>
- [28] M. Liu et al., "SCINet: Time series modeling and forecasting with sample convolution and interaction," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 1–18. [Online]. Available: <https://openreview.net/forum?id=AyajSjTAzmg>
- [29] B. N. Oreshkin, D. Carпов, N. Chapados, and Y. Bengio, "N-BEATS: Neural basis expansion analysis for interpretable time series forecasting," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–31. [Online]. Available: <https://openreview.net/forum?id=r1ecqn4YwB>
- [30] V. Ekambaram, A. Jati, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "TSMixer: Lightweight MLP-mixer model for multivariate time series forecasting," in *Proc. 29th ACM SIGKDD Conf. Knowl. Discovery Data Min.*, 2023, pp. 459–469. [Online]. Available: <https://api.semanticscholar.org/CorpusID:259187817>
- [31] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 1–15. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [32] W. Fan, P. Wang, D. Wang, D. Wang, Y. Zhou, and Y. Fu, "Dish-TS: A general paradigm for alleviating distribution shift in time series forecasting," in *Proc. AAAI Conf. Artif. Intell.*, 2023, pp. 7522–7529.
- [33] S. Chen, G. Long, T. Shen, and J. Jiang, "Prompt federated learning for weather forecasting: Toward foundation models on meteorological data," in *Proc. 32nd Int. Joint Conf. Artif. Intell. (IJCAI)*, 2023, pp. 3532–3540. [Online]. Available: <https://doi.org/10.24963/ijcai.2023/393>
- [34] H. Xue and F. D. Salim, "PromptCast: A new prompt-based learning paradigm for time series forecasting," 2022, *arXiv:2210.08964*.
- [35] D. K. Ruch and P. J. Van Fleet, *Wavelet Theory an Elementary Approach with Applications*. Hoboken, NJ, USA: Wiley, 2009. [Online]. Available: <https://search.library.wisc.edu/catalog/9911067192202121>
- [36] K.-H. Lai, D. Zha, J. Xu, Y. Zhao, G. Wang, and X. Hu, "Revisiting time series outlier detection: Definitions and benchmarks," in *Proc. 35th Conf. Neural Inf. Process. Syst.*, 2021, pp. 1–13. [Online]. Available: <https://openreview.net/forum?id=r8IvOsnHchr>
- [37] G. Zerveas, S. Jayaraman, D. Patel, A. Bhamidipaty, and C. Eickhoff, "A transformer-based framework for multivariate time series representation learning," in *Proc. 27th ACM SIGKDD Conf. Knowl. Discovery Data Min.*, 2021, pp. 2114–2124. [Online]. Available: <https://doi.org/10.1145/3447548.3467401>
- [38] Z. You et al., "A unified model for multi-class anomaly detection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 1–14. [Online]. Available: https://openreview.net/forum?id=bMYU8_qD8PW
- [39] Z. Yue et al., "TS2Vec: Towards universal representation of time series," in *Proc. 36th AAAI Conf. Artif. Intell.*, 2022, pp. 1–8.
- [40] G. Woo, C. Liu, D. Sahoo, A. Kumar, and S. Hoi, "CoST: Contrastive learning of disentangled seasonal-trend representations for time series forecasting," in *Proc. Int. Conf. Learn. Represent.*, 2022, pp. 1–18. [Online]. Available: <https://openreview.net/forum?id=PilZY3omXV2>
- [41] A. Trindade, "ElectricityLoadDiagrams20112014," Dataset, UCI Machine Learning Repository, Aigle, Switzerland, 2015, doi: [10.24432/C58C86](https://doi.org/10.24432/C58C86).
- [42] "Caltrans pems (n.d.)." Dec. 2023. [Online]. Available: <http://pems.dot.ca.gov/>
- [43] "Max-planck-institut fuer biogeochemie - wetterdaten (n.d.)." Dec. 2023. [Online]. Available: <https://www.bgc-jena.mpg.de/wetter/>
- [44] "Solar power data for integration studies (n.d.)." Dec. 2023. [Online]. Available: <https://www.nrel.gov/grid/solar-power-data.html>
- [45] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The M4 competition: Results, findings, conclusion and way forward," *Int. J. Forecasting*, vol. 34, no. 4, pp. 802–808, Oct.–Dec. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169207018300785>
- [46] Y. Liu et al., "ITransformer: Inverted transformers are effective for time series forecasting," 2023, *arXiv:2310.06625*. [Online]. Available: <https://api.semanticscholar.org/CorpusID:263830644>
- [47] L. Van Der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: <http://jmlr.org/papers/v9/vandermaaten08a.html>
- [48] Z. Liu et al., "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 9992–10002.



Li Shen received the B.S. degree in navigation and control from the School of Automation Science and Electrical Engineering, Beihang University, Beijing, China, in 2020. He is currently pursuing the Doctoral degree in electronic information engineering with Beihang Institute of Unmanned System, Beijing.

His research interests include computer vision and time series forecasting.



Yuning Wei received the B.S. degree in agricultural electrification from the School of Engineering, Nanjing Agricultural University, Nanjing, China, in 2021. She is currently pursuing the MA.Eng. degree in electronic information engineering with Beihang Institute of Unmanned System, Beijing, China.

Her research interests include deep learning and time series forecasting.



Yangzhu Wang received the B.S. and Ph.D. degrees in measurement and control technology and instrumentation program from the Institute of Unmanned System, Beihang University, Beijing, China, in 1994 and 2001, respectively.

He is currently a Research Fellow of the Flying College, Institute of Unmanned System, Beihang University. His research interests include computer vision, unmanned system, and measurement and control technology.



Huaxin Qiu received the B.S. degree in automation and the Ph.D. degree in guidance, navigation, and control from Beihang University, Beijing, China, in 2014 and 2019, respectively.

From 2019 to 2022, she was a Postdoctoral Researcher with Qian Xuesen Laboratory of Space Technology, China Academy of Space Technology, Beijing. Since 2022, she has been an Associate Professor with the Flying College, Beihang University. Her main research interests include swarm intelligence with its application to distributed decision making and optimization in multiagent systems.