

DP-LET : An Efficient Spatio-Temporal Network Traffic Prediction Framework

Xintong Wang*, Haihan Nan[†], Ruidong Li* and Huaming Wu[‡]

*Kanazawa University, Kanazawa, Japan

[†]George Mason University, Virginia, US

[‡]Tianjin University, Tianjin, China

Email: *wangxt@stu.kanazawa-u.ac.jp, [†]hnan@gmu.edu, *liruidong@ieee.org, [‡]whming@tju.edu.

Abstract—Accurately predicting spatio-temporal network traffic is essential for dynamically managing computing resources in modern communication systems and minimizing energy consumption. Although spatio-temporal traffic prediction has received extensive research attention, further improvements in prediction accuracy and computational efficiency remain necessary. In particular, existing decomposition-based methods or hybrid architectures often incur heavy overhead when capturing local and global feature correlations, necessitating novel approaches that optimize accuracy and complexity. In this paper, we propose an efficient spatio-temporal network traffic prediction framework, DP-LET, which consists of a data processing module, a local feature enhancement module, and a Transformer-based prediction module. The data processing module is designed for high-efficiency denoising of network data and spatial decoupling. In contrast, the local feature enhancement module leverages multiple Temporal Convolutional Networks (TCNs) to capture fine-grained local features. Meanwhile, the prediction module utilizes a Transformer encoder to model long-term dependencies and assess feature relevance. A case study on real-world cellular traffic prediction demonstrates the practicality of DP-LET, which maintains low computational complexity while achieving state-of-the-art performance, significantly reducing MSE by 31.8% and MAE by 23.1% compared to baseline models.

Index Terms—Traffic Prediction, Deep Learning, Transformer, Performance Evaluation.

I. INTRODUCTION

The rapid proliferation of smart devices has led to the widespread adoption of the Internet and an explosion in data traffic and the number of applications. This exponential growth significantly enlarges the volume of data that must be collected and managed, increasing network complexity. In the Internet of Things (IoT) era and advanced wireless communications, these trends have further complicated network environments, making efficient management and optimization critical [1].

Network traffic prediction is a vital technology designed to forecast the overall traffic volume based on historical data. By accurately predicting future traffic, network operators can preemptively mitigate network congestion and ensure high-quality service, which is crucial for devising effective resource allocation strategies and maintaining robust network performance. Traditionally, statistical methods—such as the Auto-Regressive Integrated Moving Average (ARIMA) model—have been used for traffic prediction. However, while

these linear models can perform reasonably well in short-term scenarios, they often struggle with the complex and dynamic traffic patterns observed in large-scale IoT networks and modern communication systems [2]. As a result, deep learning techniques have emerged as a promising alternative to capture better the inherent non-linearities and long-range dependencies present in today’s network traffic.

In deep learning, some studies have treated network traffic as images and leveraged Convolutional Neural Networks (CNNs) to capture and predict spatio-temporal dependencies [3]. In contrast, others have employed Graph Neural Networks (GNNs) because they can learn topological information [4]. However, CNNs rely on fixed receptive fields, making it challenging to model long-range temporal dependencies. Meanwhile, GNNs depend on predefined adjacency matrices, which poses challenges when accommodating dynamic node additions or removals [5]. By contrast, Transformer-based methods can address these issues effectively, offering a more flexible approach to spatio-temporal traffic prediction through self-attention mechanisms that excel at handling long-range dependencies and adapting to evolving network topologies. Nevertheless, existing Transformer architectures suffer from two critical deficiencies that hinder them from achieving higher accuracy. Specifically, they exhibit insufficient sensitivity to local features—such as burst traffic patterns and periodic fluctuations in the temporal domain. Furthermore, they lack explicit spatial modeling mechanisms, leading to spatial confusion when predicting large-scale network topologies.

To further address Transformers’ insufficient sensitivity to local features, current research primarily focuses on improving temporal feature extraction in the following two ways. The first approach decomposes raw data into multi-scale, seasonally informed components for separate predictions [6], which inevitably increases computational overhead. The second approach adopts hybrid models that integrate additional networks with Transformers [7], but still faces an inherent trade-off between effective local feature capture and computational efficiency. Meanwhile, although techniques such as channel-independence [8] and modified attention mechanisms [9] have been proposed to handle spatial dependencies, they cannot bring significant breakthroughs in enhancing temporal feature extraction.

To address these challenges, we propose **DP-LET**, a spatio-

temporal traffic prediction framework composed of a **Data Processing** module, a **Local Enhancement** module, and a **Transformer-based prediction** module. Specifically, the data processing module leverages Truncated Singular Value Decomposition and Reconstruction (TSVDR) combined with a spatial decoupling strategy, circumventing complex spatial modeling while reinforcing critical temporal features. Furthermore, the local feature enhancement module projects data into high-dimensional space. It employs Temporal Convolutional Networks (TCNs) with a dense layer to effectively capture multi-scale local patterns, compensating for the Transformer's limitations in local temporal dependency modeling. Finally, we employ Transformer as the backbone of the prediction module to capture global dependencies. DP-LET enhances prediction accuracy while maintaining computational efficiency. The main contribution of this study can be summarized as follows:

- We design a data processing module based on TSVDR with a spatial decoupling strategy. Specifically, TSVDR enables noise removal while preserving key temporal correlations, and explicit spatial decoupling eliminates the need for complex spatial modeling. Compared with the seasonal decomposition-based method, this module reduces GPU memory consumption by 38.4%.
- We develop a local feature enhancement module that maps network data into high-dimensional space and employs TCNs with dense layers to capture multi-scale local patterns.
- We integrate the data processing module, local feature enhancement module, and prediction module into the DP-LET framework. This architecture explicitly addresses spatial dependencies while modeling local and global interactions to capture spatio-temporal correlations. Evaluations of real-world cellular network traffic data demonstrate that DP-LET achieves significant performance gains, reducing the average MSE by 31.8% and the average MAE by 23.1% compared to state-of-the-art baselines.

The remainder of this paper is organized as follows. Section II introduces the recent literature on network traffic prediction. Section III presents the structure of DP-LET and describes each module in detail. In Section IV, a case study is presented. Section V concludes the paper by summarizing the main findings.

II. LITERATURE REVIEW

Recent research has addressed two critical challenges in Transformer-based spatio-temporal traffic prediction: local feature extraction, for example, detecting traffic bursts, and spatial dependency modeling, such as establishing node relationships.

To enhance temporal feature capture, Autoformer [6] pioneered seasonal decomposition to explicitly separate data into seasonal and trend components for prediction. This approach has been widely adopted by subsequent works [5], [10]. However, since the seasonal and trend components must be

predicted separately, the computational cost increases substantially. Some studies have proposed using simpler predictors—such as a linear layer [11] or multilayer perceptrons (MLPs) [12]—as a replacement for the Transformer to address the resulting parameter explosion. Although these approaches reduce some complexity, they do not fundamentally resolve the parameter explosion caused by decomposition.

A promising direction emerges from natural language processing (NLP), where Singularformer [13] leverages TSVDR to reconstruct attention mechanisms, achieving significant complexity reduction. However, directly applying TSVDR to network traffic prediction faces a unique challenge. While the denoised results are remarkably effective in the temporal dimension, they disrupt the interrelationships among different spatial dimensions, making subsequent spatial modeling more challenging.

Model enhancement approaches have also emerged to mitigate the computational overhead introduced by decomposition-based strategies. Some researchers have explored hybrid architectures to leverage multiple neural network techniques within a single framework. For example, the CNN-LSTM-Transformer structure proposed in [14] uses sliding convolutional layers to reinforce local feature extraction. However, the depth of this network inevitably adds substantial computational overhead. One recent variant further replaces the CNN component with a TCN to refine local feature capture [7]. However, these approaches still overlook explicit spatial relationships during spatio-temporal prediction.

For spatial dependency challenges, PatchTST [8] attempts spatial decoupling by treating variables as independent time series, while Crossformer [9] designs attention mechanisms to capture cross-variable dependencies explicitly. Nevertheless, these methods cannot address attention mechanisms' inherent temporal feature extraction limitations.

While existing methods either focus on improving temporal feature extraction or spatial dependency modeling, they fail to address both challenges simultaneously. Therefore, we propose DP-LET, an efficient spatio-temporal traffic prediction framework that simultaneously addresses the traditional Transformers' feature extraction deficiencies and spatial relationship modeling challenges while maintaining computational efficiency.

III. DP-LET: AN EFFICIENT FRAMEWORK FOR SPATIO-TEMPORAL NETWORK TRAFFIC PREDICTION

DP-LET consists of three main modules: the data processing module, the local feature enhancement module, and the prediction module. The data processing module is subdivided into preprocessing and postprocessing parts. As shown in Fig. 1, in the preprocessing stage, the raw data first undergoes TSVDR-based denoising and is then decoupled into multiple univariate time series. Afterward, these time series are normalized and segmented. Next, the local feature enhancement module refines local temporal dependencies within the resulting segments. The prediction module then employs a Transformer encoder to capture global dependencies and

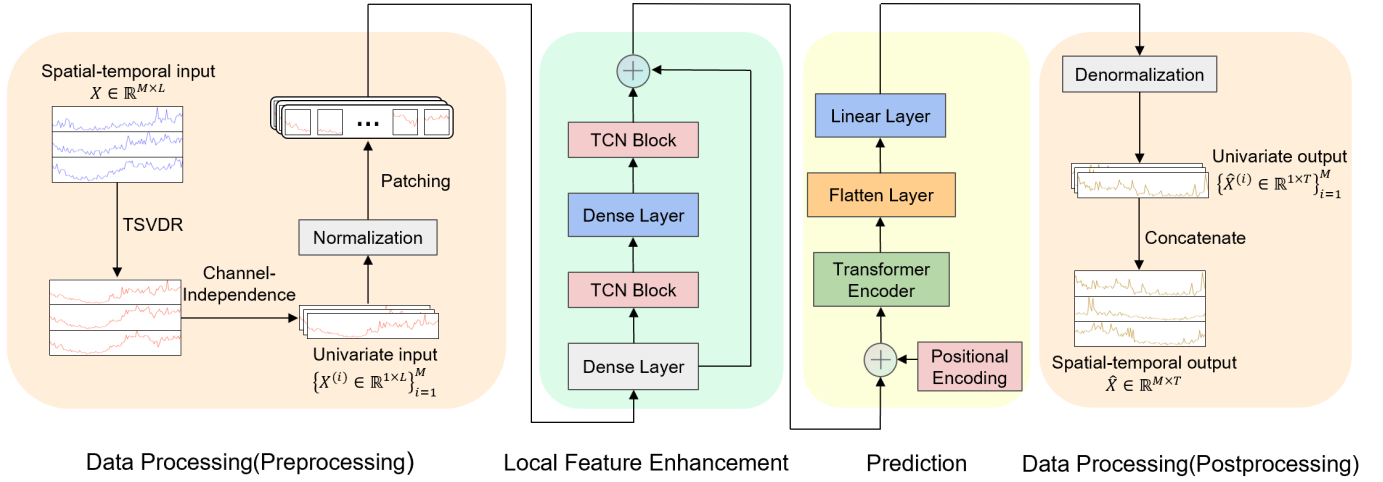


Fig. 1. DP-LET framework architecture

generate channel-wise forecasts. Finally, the postprocessing part restores the forecasts to their original scale and concatenates them into the final spatio-temporal predictions. In the following subsections, we describe each module in detail.

A. Data Processing Module (Preprocessing Part)

In the preprocessing stage, the input spatio-temporal data $X \in \mathbb{R}^{M \times L}$ is first denoised via TSVDR, where M denotes the number of channels and L denotes the input time horizon. Channel independence is applied to decouple the spatio-temporal data into M univariate time series $\{X^{(i)} \in \mathbb{R}^{1 \times L}\}_{i=1}^M$ to address challenges in spatial relationship modeling that arise after TSVDR denoising. This approach circumvents the pitfalls of imposing artificial spatial relationships while preserving critical temporal dynamics. Reversible instance normalization and temporal patching mechanisms are incorporated to enhance data stationarity and capture multi-scale patterns, as illustrated in Algorithm 1.

After the preprocessing stage, a total of N temporal segments of length ℓ are obtained for each univariate channel. These segments can be expressed as $\{X_p^{(i)}[n] \in \mathbb{R}^{1 \times \ell}\}_{n=0, \dots, N-1, i=1, \dots, M}$, where each segment captures local temporal dynamics within a fixed window of length ℓ . The resulting representation $\{X_p^{(i)}[n]\}_{n=0, \dots, N-1, i=1, \dots, M}$ not only encompasses the denoised, normalized, and segmented data but also preserves the essential characteristics of the original time series, thereby laying a solid foundation for efficient multi-scale feature extraction and forecasting in subsequent stages.

B. Local Feature Enhancement Module

Multi-branch TCNs serve as the backbone of this module. Compared to CNNs, TCNs can effectively capture local features via dilated convolutions, reducing the required network depth. Initially, a dense layer projects

Algorithm 1: Data Processing Module (Preprocessing Part)

Data: $X \in \mathbb{R}^{M \times L}$, threshold $c \in \mathbb{R}_+$, segment length ℓ , stride S

Result: $\{X_p^{(i)}[n] \in \mathbb{R}^{1 \times \ell}\}_{n=0, \dots, N-1, i=1, \dots, M}$, $\{\mu_i, \sigma_i\}_{i=1}^M$

```

1  $U \Sigma V^T \leftarrow \text{SVD}(X)$ ;
2 for  $j \leftarrow 1$  to  $\min(M, L)$  do
3   if  $\Sigma_{j,j} < c$  then
4      $\Sigma_{j,j} \leftarrow 0$ ;
5   end
6 end
7  $X' \leftarrow U \Sigma V^T$ ;
8 for  $i \leftarrow 1$  to  $M$  do
9    $X^{(i)} \leftarrow X'[i, :]$ ;
10   $\mu_i \leftarrow \frac{1}{L} \sum_{t=1}^L X^{(i)}(t)$ ;
11   $\sigma_i \leftarrow \sqrt{\frac{1}{L} \sum_{t=1}^L (X^{(i)}(t) - \mu_i)^2}$ ;
12   $X_{\text{norm}}^{(i)} \leftarrow \frac{X^{(i)} - \mu_i}{\sigma_i}$ ;
13   $X_{\text{pad}}^{(i)} \leftarrow [X_{\text{norm}}^{(i)}, 0, 0, \dots, 0]$ ;
14   $N \leftarrow \lfloor \frac{L-\ell}{S} \rfloor + 2$ ;
15  for  $n \leftarrow 0$  to  $N-1$  do
16     $X_p^{(i)}[n+1] \leftarrow X_{\text{pad}}^{(i)}[n \cdot S + 1 : n \cdot S + \ell]$ ;
17  end
18 end
19 return  $\{X_p^{(i)}[n]\}_{n=0, \dots, N-1, i=1, \dots, M}$ ;

```

each segment $\{X_p^{(i)}[n]\}_{n=0, \dots, N-1, i=1, \dots, M}$ to a new representation $\{\tilde{X}_p^{(i)}[n] \in \mathbb{R}^{1 \times d_{\text{model}}}\}_{n=0, \dots, N-1, i=1, \dots, M}$, thus enhancing the representational capacity of the input. The projected segment $\{\tilde{X}_p^{(i)}[n]\}_{n=0, \dots, N-1, i=1, \dots, M}$ is then processed by a TCN block to

strengthen local feature extraction further, followed by a dense layer that refines and recalibrates the learned features. Subsequently, a second TCN block is applied to capture additional temporal dynamics. Finally, a residual connection is introduced to mitigate issues such as gradient explosion and to ensure stable training. The final output of this module is $\{X_e^{(i)}[n] \in \mathbb{R}^{1 \times d_{\text{model}}}\}_{n=0, \dots, N-1, i=1, \dots, M}$.

The framework incorporates this module as an integral component, enabling it to seamlessly replace the conventional Transformer data embedding layer and be readily generalized to other Transformer-based architectures.

C. Prediction Module

A classical Transformer encoder is employed in the prediction module to model long-range temporal dependencies and generate final forecasts. A learnable positional encoding is then added to $\{X_e^{(i)}[n]\}_{n=0, \dots, N-1, i=1, \dots, M}$, and the sequence is fed into the Transformer encoder, retaining the same shape while capturing global temporal dependencies across all segments. Next, the encoder output is flattened into $\{X_{\text{flat}}^{(i)} \in \mathbb{R}^{1 \times (N \cdot d_{\text{model}})}\}_{i=1}^M$, thereby consolidating information from all segments. Finally, a linear layer projects $\{X_{\text{flat}}^{(i)}\}_{i=1}^M$ to $\{X_{\text{linear}}^{(i)} \in \mathbb{R}^{1 \times T}\}_{i=1}^M$, where T is the prediction horizon.

D. Data Processing Module (Postprocessing Part)

Since the preprocessing transformations alter the original data structure, corresponding inverse operations must project the model predictions back to the original data space. In the postprocessing phase, a denormalization step is applied to restore the predictions to their original scale. Subsequently, the univariate prediction results for all channels, $\{\hat{X}^{(i)} \in \mathbb{R}^{1 \times T}\}_{i=1}^M$, are concatenated to reconstruct the final spatio-temporal prediction result $\hat{X} \in \mathbb{R}^{M \times T}$. The detailed process is described in Algorithm 2.

Algorithm 2: Data Processing Module (Postprocessing Part)

Data: $\{X_{\text{linear}}^{(i)} \in \mathbb{R}^{1 \times T}\}_{i=1}^M, \{\mu_i, \sigma_i\}_{i=1}^M$
Result: $\hat{X} \in \mathbb{R}^{M \times T}$

```

1 for  $i \leftarrow 1$  to  $M$  do
2    $\hat{X}^{(i)} \leftarrow X_{\text{linear}}^{(i)} \cdot \sigma_i + \mu_i$ 
3 end
4  $\hat{X} \leftarrow \mathbf{0}^{M \times T}$ 
5 for  $i \leftarrow 1$  to  $M$  do
6    $\hat{X}[i, :] \leftarrow \hat{X}^{(i)}$ 
7 end
8 return  $\hat{X}$ ;
```

IV. A CASE STUDY: REAL-WORLD CELLULAR TRAFFIC PREDICTION

Cellular network traffic is a critical data type underpinning modern communication systems, serving billions of mobile

TABLE I
SPATIO-TEMPORAL CHARACTERISTICS OF THE MILAN CDRs DATASET

Dimension	Data Size	Range	Minimum unit
Spatial	100×100	10,000 grids	Per-grid network traffic
Temporal	8,928	62 days	10 minutes

users and devices. The evolution of cellular networks—from the first-generation (1G) analog systems to the sophisticated 5G networks—has led to dramatic increases in data transmission rates and service complexity. In this context, accurate traffic prediction becomes essential for optimizing network performance, ensuring reliable service, and enabling strategic operations such as failure detection, adaptive base station sleep scheduling, and efficient resource allocation [15].

In this case study, we leverage a real-world cellular network traffic dataset to validate the effectiveness of our proposed framework. We begin by outlining the fundamental characteristics of cellular traffic data. Next, we compare DP-LET with other state-of-the-art methods regarding prediction accuracy and computational complexity. To further assess the contribution of each module, we conduct a series of ablation experiments analyzing their impact on accuracy and complexity. Finally, we evaluate the transferability of the local enhancement module, demonstrating its potential for seamless integration into other Transformer-based architectures.

A. Experimental Setting

1) *Dataset:* In this study, we utilized the Call Detail Records (CDRs) dataset [16] collected by Telecom Italia in Milan, which represents the mobile network traffic data for the city of Milan. The data were acquired at a 10-minute interval, yielding 8,928 time stamps that chronicle the evolution of mobile network traffic from November 1, 2013, to January 1, 2014 (62 consecutive days). Spatially, Milan is divided into a 100×100 grid, resulting in 10,000 equally-sized regions. The spatio-temporal characteristics of the dataset are detailed in Table I.

2) *Evaluation criteria:* In prediction tasks, Mean Squared Error (MSE) and Mean Absolute Error (MAE) are commonly used metrics because they assess regression performance from different perspectives. These metrics are adopted in this paper to evaluate prediction accuracy comprehensively. Their formulas are defined as follows:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2, \quad (1)$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad (2)$$

where N denotes the total number of samples, y_i represents the actual network traffic, and \hat{y}_i represents the predicted network traffic. In addition, the number of parameters is reported to provide insight into the computational complexity of the proposed framework.

TABLE II
PERFORMANCE COMPARISON OF DP-LET AND BASELINE METHODS

Methods	DP-LET	TimeMixer	PatchTST	iTransformer	TimesNet	DLinear	FEDformer	Autoformer	Crossformer	Informer
Metrics	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE	MSE MAE
$T = 72$	0.225 0.324	<u>0.230</u> <u>0.329</u>	0.234 0.337	0.237 0.333	0.304 0.385	0.311 0.428	0.387 0.468	0.429 0.506	0.510 0.517	0.653 0.615
$T = 144$	0.238 0.333	0.255 0.347	<u>0.249</u> 0.350	0.256 <u>0.346</u>	0.341 0.410	0.335 0.447	0.408 0.485	0.443 0.513	0.646 0.611	0.749 0.654




Fig. 2. Efficiency analysis of DP-LET and baseline methods

3) *Baselines*: We compare our framework with nine state-of-the-art models, including six leading Transformer-based methods: iTransformer [17], Informer [18], PatchTST, Autoformer, FEDformer and Crossformer; in addition, we consider the CNN-based model TimesNet [19], the linear model DLinear, and the MLP-based model TimeMixer. All of these models have demonstrated strong performance in spatio-temporal prediction tasks.

4) *Implementation details*: All experiments were conducted on an NVIDIA RTX 5000 GPU. We randomly selected 100 spatial grids from Milan's 100×100 cellular network to ensure a representative subset of the overall traffic distribution. The temporal configuration adopts a 10-minute resolution, with input sequences spanning three consecutive days ($L = 432$ time steps) to predict two distinct horizons: 12-hour ($T = 72$) and 24-hour ($T = 144$) network traffic flows.

All baseline models were implemented using their officially reported hyperparameters. The maximum number of training epochs was set to 100, with an early stopping patience of 20 epochs.

B. Experimental Result

1) *Main results*: We conduct experiments on a cellular network traffic dataset and compare our proposed framework with multiple baseline methods under two prediction horizons, $T \in \{72, 144\}$. As shown in Table II, the best performance is highlighted in **bold**, and the second-best result is underlined. DP-LET consistently outperforms the baselines at both $T = 72$ and $T = 144$. On average, it achieves a 31.8% reduction in MSE and a 23.1% reduction in MAE compared to the other methods, underscoring its superior predictive capabilities.

Beyond predictive performance, we also evaluate the model complexity of each method by comparing the number of trainable parameters and time per training iteration, as shown in Fig. 2. Among Transformer-based approaches, DP-LET has the third-smallest parameter count. It substantially reduces both parameters and iteration time relative to seasonal-decomposition Transformers such as Autoformer and FEDformer, and its complexity is close to that of the MLP-based TimeMixer. These results demonstrate that our approach achieves high forecasting accuracy while maintaining a relatively light model footprint, effectively balancing efficiency and performance.

2) *Ablation study*: We perform an ablation study to evaluate the contributions of the local feature enhancement module and the data processing module in our proposed framework for a prediction horizon of $T = 144$. Specifically, we compare three configurations: (i) the prediction module combined with the data processing module, (ii) the prediction module combined with the local feature enhancement module, and (iii) the complete framework that integrates all modules. As shown in Table III, the complete framework achieves the highest accuracy. Furthermore, compared with the local feature enhancement module, the data processing module contributes more substantially to prediction accuracy, as reflected by the lower error metrics when integrated with the prediction module.

3) *Effect of the data processing module*: To validate the effectiveness of our proposed data processing module, we integrated both the proposed module and a conventional seasonal decomposition approach into the prediction module and assessed their performance in terms of prediction accuracy

TABLE III
RESULTS OF ABLATION EXPERIMENTS

Methods	MSE	MAE
Prediction+Data Processing	0.254	0.349
Prediction+Local Feature Enhancement	0.343	0.423
Proposed Framework	0.238	0.333

TABLE IV
EFFECT OF THE DATA PROCESSING MODULE

Metrics	Complexity		Accuracy	
	Para.	Mem. (MiB)	MSE	MAE
Seasonal Decomposition	3,694,456	13,485	0.263	0.349
Proposed Data Processing	1,847,128	8,308	0.254	0.349

TABLE V
EFFECT OF THE LOCAL FEATURE ENHANCEMENT MODULE

Methods	Transformer		Non-stationary		FEDformer		Reformer	
Metrics	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Original	0.689	0.624	0.414	0.457	0.408	0.485	0.632	0.594
+ Local Feature Enhancement	0.606	0.588	0.373	0.437	0.388	0.471	0.620	0.586
Improvement (%)	12.05%	5.77%	9.90%	4.37%	4.90%	2.89%	1.90%	1.35%

and computational complexity. As shown in Table IV, our method achieves a 3.42% reduction in MSE relative to the baseline. Meanwhile, the proposed module requires only 50% of the parameters (Para.) and 61.6% of the GPU memory (Mem.) compared to seasonal decomposition, underscoring its superior efficiency.

4) *Effect of the local feature enhancement module:* We applied the proposed local feature enhancement module to four Transformer-based models at a prediction horizon of $T = 144$. Table V reports the MSE and MAE values before and after integrating our module. Notably, each method sees a clear performance boost, with the Transformer experiencing the largest gains, followed by Nonstationary-Transformer [20], Reformer [21], and FEDformer. These findings confirm the effectiveness of capturing local dependencies for improved spatio-temporal forecasting.

V. CONCLUSION

This paper introduced DP-LET as a novel and efficient spatio-temporal network traffic prediction framework. The proposed approach integrates a robust data processing module for noise reduction and spatio-temporal data decoupling, a local feature enhancement module for capturing multi-scale temporal dependencies, and a Transformer-based prediction module for extracting global temporal patterns. This unified design strikes an effective balance between computational complexity and predictive accuracy.

Extensive experiments on a real-world cellular network traffic dataset demonstrate that DP-LET consistently outperforms state-of-the-art baselines across various forecasting horizons. Moreover, DP-LET's modular architecture not only boosts prediction performance but also underscores the exceptional portability of the local feature enhancement module, enabling seamless integration into other Transformer-based models and broadening its applicability to diverse prediction frameworks.

ACKNOWLEDGMENT

This work was partly supported by Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Number 23K28070.

REFERENCES

- [1] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, D. Niyato, O. Dobre, and H. V. Poor, "6g internet of things: A comprehensive survey," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 359–383, 2021.
- [2] S. Siarni-Namini, N. Tavakoli, and A. S. Namin, "The performance of lstm and bilstm in forecasting time series," in *2019 IEEE International conference on big data (Big Data)*. IEEE, 2019, pp. 3285–3292.
- [3] C. Zhang, H. Zhang, D. Yuan, and M. Zhang, "Citywide cellular traffic prediction based on densely connected convolutional neural networks," *IEEE Communications Letters*, vol. 22, no. 8, pp. 1656–1659, 2018.
- [4] X. Wang, Z. Zhou, F. Xiao, K. Xing, Z. Yang, Y. Liu, and C. Peng, "Spatio-temporal analysis and prediction of cellular traffic in metropolis," *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2190–2202, 2018.
- [5] H. Nan, R. Li, X. Zhu, J. Ma, and D. Niyato, "An efficient data-driven traffic prediction framework for network digital twin," *IEEE Network*, vol. 38, no. 1, pp. 22–29, 2023.
- [6] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," *Advances in neural information processing systems*, vol. 34, pp. 22 419–22 430, 2021.
- [7] Z. Zhang, S. Gong, Z. Liu, and D. Chen, "A novel hybrid framework based on temporal convolution network and transformer for network traffic prediction," *Plos one*, vol. 18, no. 9, p. e0288935, 2023.
- [8] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," *arXiv preprint arXiv:2211.14730*, 2022.
- [9] Y. Zhang and J. Yan, "Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting," in *The eleventh international conference on learning representations*, 2023.
- [10] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting," in *International conference on machine learning*. PMLR, 2022, pp. 27 268–27 286.
- [11] M. Chen, Q. Xu, A. Zeng, L. Zhang *et al.*, "Are transformers effective for time series forecasting?" *arXiv preprint, arXiv: 2205.13504*, 2022.
- [12] S. Wang, H. Wu, X. Shi, T. Hu, H. Luo, L. Ma, J. Y. Zhang, and J. Zhou, "Timemixer: Decomposable multiscale mixing for time series forecasting," *arXiv preprint arXiv:2405.14616*, 2024.
- [13] Y. Wu, S. Kan, M. Zeng, and M. Li, "Singularformer: Learning to decompose self-attention to linearize the complexity of transformer," in *IJCAI*, 2023, pp. 4433–4441.
- [14] M. Gong, H. Cui, S. Jing, J. Huang, L. Ma, Y. Wang, and W. Xu, "Network traffic prediction model based on convolutional neural networks-long short-term memory and itransformer," *Internet Technology Letters*, p. e658, 2025.
- [15] W. Jiang, "Cellular traffic prediction with machine learning: A survey," *Expert Systems with Applications*, vol. 201, p. 117163, 2022.
- [16] T. Italia, "Telecommunications - SMS, Call, Internet - MI," 2015. [Online]. Available: <https://doi.org/10.7910/DVN/EGZHFV>
- [17] Y. Liu, T. Hu, H. Zhang, H. Wu, S. Wang, L. Ma, and M. Long, "itranformer: Inverted transformers are effective for time series forecasting," *arXiv preprint arXiv:2310.06625*, 2023.
- [18] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 12, 2021, pp. 11 106–11 115.
- [19] H. Wu, T. Hu, Y. Liu, H. Zhou, J. Wang, and M. Long, "Timesnet: Temporal 2d-variation modeling for general time series analysis," *arXiv preprint arXiv:2210.02186*, 2022.
- [20] Y. Liu, H. Wu, J. Wang, and M. Long, "Non-stationary transformers: Exploring the stationarity in time series forecasting," *Advances in neural information processing systems*, vol. 35, pp. 9881–9893, 2022.
- [21] N. Kitaev, L. Kaiser, and A. Levskaya, "Reformer: The efficient transformer," *arXiv preprint arXiv:2001.04451*, 2020.