# FDALLM+: A Functional Data Analysis-Driven Large Language Model Framework for Network Traffic Prediction

Yujie Sun, *Graduate Student Member, IEEE*, Xuyu Wang, *Member, IEEE*, Guanqun Cao, and Shiwen Mao, *Fellow, IEEE*

***Abstract*—In communication network management, prediction of mobile network traffic is essential to ensure efficient system operation. Although significant progress has been made in the application of neural networks to traffic prediction tasks, traditional models still face considerable challenges when handling high-dimensional and highly time-dependent data. To address these issues, this paper proposes a new prediction framework that leverages large language models (LLMs), by constructing efficient prompts to enhance the ability of large language models (LLMs) in traffic prediction and improve their understanding of complex traffic patterns. Specifically, we introduce functional data analysis (FDA), a technique that offers superior capabilities compared to traditional methods in processing continuous and high-dimensional data structures, to preprocess traffic data and extract key features. Extensive experiments conducted on multiple LLMs using a real-world dataset validate the effectiveness and scalability of the proposed method. The experimental results demonstrate that the framework achieves significant improvements in predictive performance, providing a promising and efficient solution for traffic data analysis in future communication networks.**

***Index Terms*—Large Language Models (LLM), Functional Data Analysis (FDA), Traffic Prediction, Demonstration Prompts, Generative AI.**

## I. Introduction

In existing and emerging communication networks, prediction of mobile traffic is a critical task for ensuring efficient system operation. Accurately forecasting changes in traffic data remains a complex challenge due to the high dimensionality and strong temporal correlations inherent in traffic datasets. These characteristics pose significant obstacles to traditional prediction methods. Although machine learning and deep learning approaches, such as the autoregressive integrated moving average (ARIMA) model [1] and recurrent neural networks (RNN) [2], have achieved notable success in certain scenarios, they typically rely on large volumes of training data, require complex parameter tuning, and exhibit heavy dependence on contextual information. These limitations hinder their flexibility and generalization capabilities in practical mobile traffic prediction applications.

Recently, the powerful information processing capabilities of LLMs have opened new opportunities for addressing these challenges [3]–[6]. LLMs, through extensive pretraining, possess broad language comprehension abilities and can handle complex tasks without requiring additional training on specific test sets [7]. For instance, LLM4TS [8] and TTMs [9] employ pretrained LLMs for time series prediction and enhance their understanding of temporal data through fine-tuning. MultiCast [10] and Wav2Prompt [11] further demonstrate that LLMs excel in zero-shot tasks. However, several studies have unveiled LLMs' limitations in this area. For example, Spathis and Kawsar in [12] and Time-LLM [13] reveal that LLMs encounter difficulties in directly interpreting time series data, necessitating the preprocessing of raw data into formats more compatible with LLM architectures.

Beyond these developments, LLMs also exhibit strong learning capacity and transferability, offering new possibilities in predictive tasks. For example, UmiTime [14] shows that LLMs can significantly enhance time series prediction, particularly in contexts involving limited data or cross-domain challenges. Commercial models such as GPT-4 [15] contribute to overcoming the resource constraints traditionally associated with time series modeling, providing a more accessible prediction platform. Gruver et al. [16] further demonstrate the potential of fine-tuned LLMs in achieving zero-shot time series prediction. Despite these advances, predicting complex, high-dimensional data remains a significant challenge. Enhancements in prompt design are crucial to improving LLMs' ability to understand the intricate structures and dependencies of traffic data. Recent studies, including PrompCast [17], TEMPO [18], and LSTPrompt [19], have made progress in this area by incorporating additional feature information from raw data into prompts, thereby improving LLMs' understanding of traffic patterns. However, a noticeable performance gap persists, as many existing approaches still struggle to capture the dependencies in high-dimensional data, and LLMs remain limited in modeling these intricate structures effectively.

To overcome these challenges, this paper introduces an enhanced framework that leverages Functional Data Analysis (FDA) [20] to strengthen the feature extraction and interpretability of network traffic data for LLM-based predictions.

This work is supported by the NSF under Grants CNS-2319342 and CNS-2319343, and by the Wireless Engineering Research and Education Center at Auburn University.

Y. Sun and S. Mao are with the Department of Electrical and Computer Engineering, Auburn University, Auburn, AL 36849-5201, USA. X. Wang is with the Knight Foundation School of Computing and Information Sciences at Florida International University, Miami, FL 33199, USA. G. Cao is with the Department of Statistics & Probability, Michigan State University, East Lansing, MI 48824. Email: yzs0124@auburn.edu, caoguanq@msu.edu, xuywang@fiu.edu, smao@ieee.org

A preliminary version of this work was presented in part at the 2025 IEEE International Conference on Communications, June 2025, Montreal, Canada.

FDA is a robust statistical approach known for its superior performance in analyzing continuous and complex data structures, making it well-suited to the characteristics of traffic datasets. Moreover, FDA has proven particularly effective in handling high-dimensional data [21], further supporting its application in traffic prediction tasks. However, existing state-of-the-art research on FDA primarily focuses on deep learning models rather than LLMs [22], leaving a gap in the integration of FDA techniques with LLMs.

To bridge this gap, we propose the FDALLM+ framework, an updated and enhanced version of our previous FDALLM method [23]. Our approach integrates FDA with LLMs, enabling the extraction of essential features from high-dimensional traffic data and enhancing the model's ability to capture complex temporal dependencies. Unlike existing methods that rely heavily on extensive training data or intricate fine-tuning procedures, FDALLM+ achieves high prediction accuracy in both zero-shot and few-shot scenarios, demonstrating greatly improved flexibility and generalization. Experimental results confirm that FDALLM+ not only outperforms prior approaches in predictive accuracy, but also reduces reliance on large annotated datasets, offering a more scalable and effective solution for mobile traffic prediction in modern communication networks.

The main contributions of this study are summarized in the following.

- Enhancing previous work with FDA for traffic data analysis. This paper improves previous research by incorporating FDA to analyze traffic data, emphasizing its periodicity features. The enhanced approach enables a more precise characterization of traffic flow variations.
- Leveraging periodicity features to optimize prompt tuning and generation. We exploit the periodicity features in traffic data to optimize prompt tuning and generation methods, enhancing the ability of LLMs to capture and predict traffic trends more effectively.
- Extensive experiments on real-world datasets to improve zero-shot and few-shot prediction. We extend our experiments to real-world datasets to validate our prompt generation method. Our results demonstrate that this approach helps LLMs improve zero-shot and few-shot prediction accuracy, showcasing its practical applicability and effectiveness.

In the remainder of this paper, we present the related work in Section II, motivation and problem formulation in Section III, followed by our system design in Section IV. We discuss our experimental setup and analyze the experimental results in Section V. Section VI concludes this paper.

## II. Related Works

### A. Traffic Data Prediction and Periodicity Analysis

Traffic data prediction plays a crucial role in various applications, including network performance optimization and efficient data flow management. Accurate forecasting of traffic patterns enables proactive measures against potential threats and ensures effective resource allocation. One of the key challenges in traffic prediction is capturing the periodicity inherent in traffic data, which significantly influences the accuracy of the prediction.

*1) Traditional Methods in Traffic Prediction:* Early research primarily relied on statistical models to predict traffic data. For instance, Hasegawa, et al. [24] explored nonlinear prediction methods to analyze traffic trends, while Chen, et al. [25] applied the ARIMA model, a classical time series forecasting approach widely used to capture linear temporal dependencies. Additionally, Xu, et al. [26] introduced the Kalman Filter model, which has been widely adopted for real-time traffic state estimation due to its recursive update mechanism. While these traditional methods provided a strong foundation for traffic prediction, they struggle to capture the complex periodicity of traffic data, especially in dynamic and nonlinear environments. Traffic patterns are often influenced by daily, weekly, and seasonal cycles, and failing to incorporate such periodic features can lead to suboptimal forecasting results.

*2) The Emergence of Machine Learning and Periodicity-Aware Models:* With the advent of machine learning, researchers have developed more flexible models capable of capturing nonlinear relationships and periodic patterns in traffic data. Machine learning approaches such as Support Vector Regression [27], Random Forest [28], and Gradient Boosting Machines [29] have demonstrated improved performance in traffic forecasting compared to traditional statistical methods. However, these models still require explicit feature engineering to extract periodicity-related insights. More recently, deep learning methods such as Recurrent Neural Networks [30], Long Short-Term Memory (LSTM) [31] networks, and Spatiotemporal Graph Neural Networks [32] have shown promising results in modeling complex periodic traffic patterns. These models leverage the temporal recurrence of traffic flows and dynamically learn periodic features without requiring handcrafted input features. Additionally, integrating FDA techniques [20] has further enhanced the ability to analyze and predict periodic traffic behaviors.

*3) Significance of Periodicity Analysis in Traffic Prediction:* Some studies [33] have shown that the incorporation of periodicity-sensitive mechanisms can significantly enhance traffic forecast performance. Traffic patterns often exhibit recurring trends, such as daily, weekly, and seasonal cycles, which, if effectively captured, can improve the accuracy and robustness performance of the model. Leveraging periodicity helps models distinguish between structured patterns and random fluctuations, leading to more reliable predictions. It also plays a crucial role in optimizing prompt generation for LLMs, improving their zero-shot and few-shot learning capabilities in traffic forecasting. Furthermore, integrating cyclical information from historical data reduces forecast errors by correcting deviations from expected trends.

Despite these advancements, challenges remain to effectively embed periodicity in machine learning models and refining zero-shot prediction strategies. Traditional statistical approaches struggle to capture complex periodic behaviors, and existing machine learning models often overlook long-term cyclical dependencies. To address these issues, this study further explores FDA as a novel mathematical approach to extract and analyze periodic structures in traffic data. By optimizing previous methods, we refine previous experimental results, enhancing the ability of the model to learn and adapt to periodic variations. This not only improves forecast accuracy, but also strengthens zero-shot learning by structuring prompts with more informative periodic insights, ultimately advancing traffic prediction methodologies.

### B. FDA for Traffic Data

FDA is a powerful mathematical framework particularly suited for analyzing high-dimensional and complex time-series data. In the context of traffic data, FDA provides an effective approach to capturing and modeling the continuous variation of traffic flow, addressing the limitations of traditional statistical and machine learning methods, which often rely on discrete observations and stationarity assumptions [21].

Unlike conventional time-series methods that treat traffic data as discrete and independent observations, FDA models traffic data as continuous functions [34]. This approach provides a more natural representation of periodic traffic patterns, allowing models to capture long-term trends and gain deeper insights into the underlying dynamics of traffic flow. By employing smoothing techniques, FDA reduces noise, enhances the structural integrity of data, and ultimately improves data quality. Additionally, the functional representation enables smooth interpolation, making it effective in handling missing data and mitigating errors caused by data gaps. For example, S. Wang [35] applied FDA to enhance deep neural networks, demonstrating its potential to improve predictive performance in complex modeling tasks.

Research has demonstrated the efficacy of FDA in traffic data analysis. For instance, Functional Principal Component Analysis (FPCA) has been utilized to extract the primary periodic patterns in traffic flow, thereby optimizing both short-term and long-term forecasts. Capturing long-term periodic trends allows models to better identify variations in traffic flow [36].

In this study, we further explore periodicity analysis based on FDA to enhance traffic data prediction. Using FDA, we can gain deeper insights into periodic dependencies within traffic data and refine existing experimental results. This approach not only improves forecasting performance but also provides a more interpretable framework for understanding the temporal variations in traffic flow.

### C. LLMs for Traffic Prediction: Challenges and Zero-Shot Learning

*1) LLMs for Data-Driven Prediction:* LLMs have demonstrated remarkable capabilities in data-driven prediction tasks, particularly in processing large-scale textual data and integrating structured information. Unlike traditional machine learning models that rely heavily on large amounts of labeled training data, LLMs utilize pretrained knowledge and contextual reasoning to generate predictions through structured prompts. This makes them potentially useful for traffic flow prediction, even in situations where training data is scarce.

Recent studies have shown that LLMs, when provided with well-structured prompts, can effectively generate traffic forecasts without the need for extensive retraining on numerical datasets. For instance, Liu, et al. leveraged the Spatial-Temporal Large Language Model in [37], while Seyed, et al. adopted the Graph LLM [38] for traffic prediction tasks. Additionally, pretrained models, such as TPLLM [39] and Traj-LLM [40], have been developed specifically for traffic forecasting.

However, despite their ability to process complex text-based reasoning tasks, LLMs still encounter significant challenges in zero-shot prediction, particularly when dealing with numerical time-series data, which fundamentally differs from the textual information they are primarily trained on.

*2) Challenges of Zero-Shot Prediction for LLMs and Existing Methods:* Zero-shot learning refers to the model's ability to make predictions without being explicitly trained on specific datasets. In the context of traffic forecasting, LLMs face several unique challenges, primarily due to the nature of traffic data and the differences in how LLMs are traditionally trained [41].

One of the main challenges is the strong temporal dependency in traffic data. Traffic flow is inherently sequential (usually self-similar [42]), with past patterns significantly influencing future trends. However, LLMs are primarily designed for text-based reasoning and lack the ability to natively model long-term dependencies in time series data [7]. Unlike dedicated time-series models such as ARIMA and LSTMs, which are highly capable in modeling sequential dependencies, LLMs struggle with raw numerical sequences unless explicitly structured within prompts [43].

Another major limitation is that LLMs inherently lack the ability to model periodicity [44]. Traditional time-series forecasting models can automatically recognize daily, weekly, and seasonal periodic patterns, which are crucial for accurate traffic prediction. In contrast, LLMs do not inherently learn or leverage periodicity unless explicitly provided with structured context. This makes zero-shot forecasting of traffic data particularly difficult, as LLMs may fail to capture recurring congestion patterns without proper guidance.

Current approaches to applying LLMs in traffic prediction heavily rely on prompt engineering, where structured prompts are designed to help models interpret time-series data. Recent studies have proposed various strategies to enhance LLM-based forecasting. For instance, xTP-LLM [45] converted multimodal traffic data into natural language descriptions, incorporating them into prompts to improve contextual understanding. R2T-LLM [46] captured complex spatiotemporal patterns and external factors from comprehensive traffic data and integrates them into prompt design. The authors in [47] processed textual information and extracted embeddings that were added to prompts, further refining the model's ability to interpret traffic-related content. S2IPLLM [48] optimized prompt learning by leveraging LLMs for semantic space-informed prompt tuning, improving their adaptability to diverse traffic scenarios.

However, LLMs are highly sensitive to prompt phrasing and structure, making their predictions susceptible to inconsistencies caused by slight variations in input formatting. This presents a significant challenge in ensuring interpretability and stability of LLM-based traffic forecasting. A key issue that remains to be addressed is how to better integrate periodicity into LLM prompts, enabling models to recognize and utilize the cyclic nature of traffic data more effectively in zero-shot prediction scenarios.

## III. Motivation and Problem Formulation

### A. Motivation

Building upon recent advancements in LLMs and FDA, we propose leveraging FDA for preprocessing traffic data to enhance prompt quality and optimize the predictive performance of LLMs in traffic forecasting tasks. Unlike traditional data processing methods, FDA exhibits distinct advantages when handling continuous time-series data. By treating discrete data points as samples of continuous functions, FDA effectively captures both global trends and local variations, reducing data complexity while preserving essential features. Moreover, FDA's inherent smoothing and denoising capabilities improve data robustness without the need for additional assumptions, which is particularly valuable when dealing with complex and dynamic traffic patterns.

Compared to traditional models such as LSTM networks, LLMs demonstrate superior predictive capabilities in traffic forecasting. While LSTM models are proficient in capturing temporal dependencies, they often encounter limitations when addressing high-dimensional data, long-range dependencies, and cross-domain tasks. LLMs, on the other hand, possess powerful contextual understanding and reasoning abilities, enabling them to recognize intricate patterns in traffic data and maintain high predictive accuracy, even with limited samples. Furthermore, LLMs exhibit strong cross-domain generalization, allowing them to adapt efficiently to various traffic scenarios and enhancing their flexibility and applicability in real-world contexts.

This paper focuses on the periodic characteristics of traffic data, which often exhibit daily peak hours, weekly travel fluctuations, and seasonal variations. Incorporating periodicity analysis during the preprocessing phase allows LLMs to capture these regular patterns more accurately, thereby improving their ability to model complex traffic behaviors and enhancing prediction accuracy. Effectively modeling periodic features is especially critical for predicting traffic surges during holidays or responding to unexpected disruptions, ultimately leading to more robust and precise forecasts.

Additionally, this paper will integrate few-shot learning strategies to further exploit the potential of LLMs. In many traffic prediction scenarios, data scarcity poses a considerable challenge, such as when dealing with newly constructed roads, rare events, or underrepresented regions lacking historical data. The few-shot learning capabilities of LLMs enable them to extract valuable information from minimal training samples, maintaining reliable prediction performance. This reduces the dependency on large-scale labeled datasets and enhances the model's adaptability and generalization across diverse traffic prediction tasks.

Notably, the synergy between periodic feature analysis and few-shot learning plays a crucial role in improving overall model performance. In scenarios with limited data, periodic patterns provide structured prior information that compensates for the lack of extensive training data. For instance, by learning traffic flow variations between peak and off-peak hours from historical patterns, the model can accurately predict future traffic conditions even with sparse data. Additionally, capturing these periodic characteristics aids in the efficient identification of anomalies, enabling the model to respond more accurately to irregular traffic patterns and unexpected events.

In conclusion, by combining FDA's capability to analyze periodic features in traffic data with the few-shot learning abilities of LLMs, this approach offers significant improvements in traffic forecasting performance. It enables the model to capture long-term dependencies and complex spatiotemporal patterns while maintaining high accuracy and robustness, even under data-limited conditions. Future research will further refine this framework across various traffic scenarios and validate its effectiveness in large-scale traffic prediction systems. Ultimately, this methodology holds the potential to advance intelligent transportation systems, support urban traffic management, and facilitate real-time traffic forecasting, contributing to the broader development and application of smart transportation technologies.

### B. Problem Formulation

In this study, we present a formulation of the network traffic data prediction problem. The dataset consists of hourly traffic data, with 24 samples collected per day. Each traffic data sample is represented by a multi-dimensional vector $x(t)$ accompanied by a corresponding timestamp $t$. Thus, a single sample can be expressed as $s[t] = [x(t), t]$. To process the raw traffic data, we adopt a sliding window approach for segmentation. The window size is denoted by $L$, meaning that at any given time, the window contains $L$ consecutive samples. The sliding window, represented as $w(t)$, captures a sequence of $L$ data points starting at time $t$.

To define the prediction task, we introduce a prediction horizon $N$, which specifies the number of future time steps to be forecasted, resulting in a predicted set $\hat{W}(t)$ of elements $\hat{x}(t)$. Correspondingly, the true values from the original dataset are represented by the set $W(t)$ with elements $x(t)$. The sliding window of input data is represented as:

$$w(t_1) = \{x(t_1), x(t_2), ..., x(t_L)\}, \tag{1}$$

where $t_1$ represents the starting time of the window. The prediction target generated by the model is given by:

$$\hat{W}(t_1) = \{\hat{x}(t_1), \hat{x}(t_2), ..., \hat{x}(t_N)\}. \tag{2}$$

The corresponding ground truth values are expressed as:

$$W(t_1) = \{x(t_1), x(t_2), ..., x(t_N)\}, \tag{3}$$

where $x(t_i)$ represents the actual traffic data corresponding to the predicted time points $\hat{x}(t_i)$, for $i = 1, 2, ...N$.

To enable continuous prediction, we define a sliding step size $M$, which determines the number of samples the window advances after processing the current window. After completing the prediction for the first window, the next window shifts forward by $M$ samples, making the first element of the new window $x[t_1+M]$. This sliding mechanism allows the model to iteratively generate new input sequences for subsequent prediction rounds by moving forward along the time axis.

The primary objective of this study is to design effective prompts for LLMs based on the input samples $w(t)$ generated by the sliding window, enabling the models to predict future traffic data $\hat{W}(t)$. The predictive performance of the model will be evaluated by comparing the predicted values in $\hat{W}(t)$ with the corresponding ground truth values in $W(t)$. This approach facilitates the modeling of temporal dependencies and complex dynamics within traffic data, ultimately leading to more accurate and robust traffic flow predictions.

## IV. System Design

### A. Overview

This study presents a greatly enhanced system design that builds upon our previous framework FDALLM [23], termed FDALLM+ as shown in Fig. 1, which integrated FDA with LLMs for traffic data prediction. The original framework primarily used FDA to preprocess high-dimensional traffic data, extracting essential features that captured both global trends and local variations. Extracted features, combined with raw data, form structured prompts that allow LLMs to achieve high-accuracy predictions in zero-shot and few-shot scenarios, even with limited domain-specific data.

The enhanced system focuses on the critical periodic characteristics of traffic data. Traffic data typically exhibit significant periodic patterns, such as daily peak hours, weekly commuting routines, and seasonal fluctuations. These features are crucial for capturing long-term dependencies and traffic flow dynamics. To capture these periodic characteristics, FDA is applied for functional fitting and periodicity analysis, refining feature representations. This process highlights recurring patterns and short-term anomalies, enabling the model to better understand underlying temporal structures.

The extracted periodic features are then embedded into structured prompts, enhancing the LLMs' ability to discern deep traffic patterns beyond surface-level temporal fluctuations. By integrating periodic characteristics into the prompts, the model is guided to leverage its contextual understanding more effectively, enabling it to capture deeper temporal dynamics and long-term dependencies within the traffic data. This results in more robust and accurate traffic predictions, especially when dealing with complex traffic scenarios.

Additionally, the system integrates few-shot learning techniques to further improve predictive performance in data-scarce scenarios. Using periodic feature-driven prompts, LLMs can extract valuable information from minimal training samples, allowing for rapid generalization across different traffic patterns. This approach significantly reduces reliance on large-scale labeled datasets while maintaining high prediction accuracy, making it particularly suitable for cross-domain traffic prediction tasks or cases with limited historical data.

Overall, the redesigned system architecture strategically combines FDA-based periodic feature extraction with the contextual understanding and few-shot learning capabilities of LLMs. This integration enhances LLM-based traffic forecasting by improving predictive accuracy, adaptability, and generalization across diverse traffic scenarios.

### B. Data Simplification and Extraction

In traffic data analysis, raw traffic data streams often contain redundancy and irrelevant information, increasing computational complexity and hindering downstream analysis. To address this, we first simplify and extract data to provide high-quality input for periodic feature analysis and modeling. Specifically, data points corresponding to each exact hour are selected, while noncritical time intervals are removed. This process reduces data redundancy, ensures temporal consistency, and highlights variations in traffic flow at key temporal nodes, thereby facilitating the capture of long-term trends and periodic fluctuations.

After temporal filtering, the dataset is segmented into seven subsets, each corresponding to a day of the week (Monday to Sunday). This classification strategy enables the identification of distinctive traffic patterns associated with each weekday, such as the differences in traffic flow distribution between weekdays and weekends. By adopting this data simplification and extraction approach—structured around a weekly cycle and analyzed on a daily basis—this study establishes a robust foundation for the extraction of daily periodic features. This step strengthens the model's ability to capture complex temporal patterns in traffic data, improving both prediction accuracy and interpretability.

### C. FDA for Data Processing

After completing the preprocessing of raw data, this study further transforms the high-dimensional traffic data into an FDA framework by representing it in functional form. In this approach, discrete traffic data is treated as sampled observations from a continuous function, enabling a precise characterization of its temporal dynamics. This transformation provides a robust foundation for capturing smooth trends, periodic fluctuations, and complex temporal dependencies within time series data.

Specifically, let the original data be observed at a set of discrete time instances $t_1$, $t_2$, ..., $t_n$ with corresponding values $x(t_1)$, $x(t_2)$, ..., $x(t_n)$. Using the FDA approach, the data can be represented as a continuous function $X(t)$, which can be expanded with respect to a set of basis functions as follows:

$$X(t)=\sum_{k=1}^{K} c_k \phi_k(t), \tag{4}$$

where $\phi_k(t)$ denotes the $k$th basis function, and $c_k$ represents the corresponding coefficient, indicating the contribution of each basis function in the functional representation.

In this manner, the original discrete data points are mapped into a smooth curve across the continuous-time domain, where each sampled data point corresponds to the function's evaluation at the respective time. Functional data representation provides a continuous analytical perspective, reducing noise while preserving trends and periodic structures, an advantage

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2025.3604332



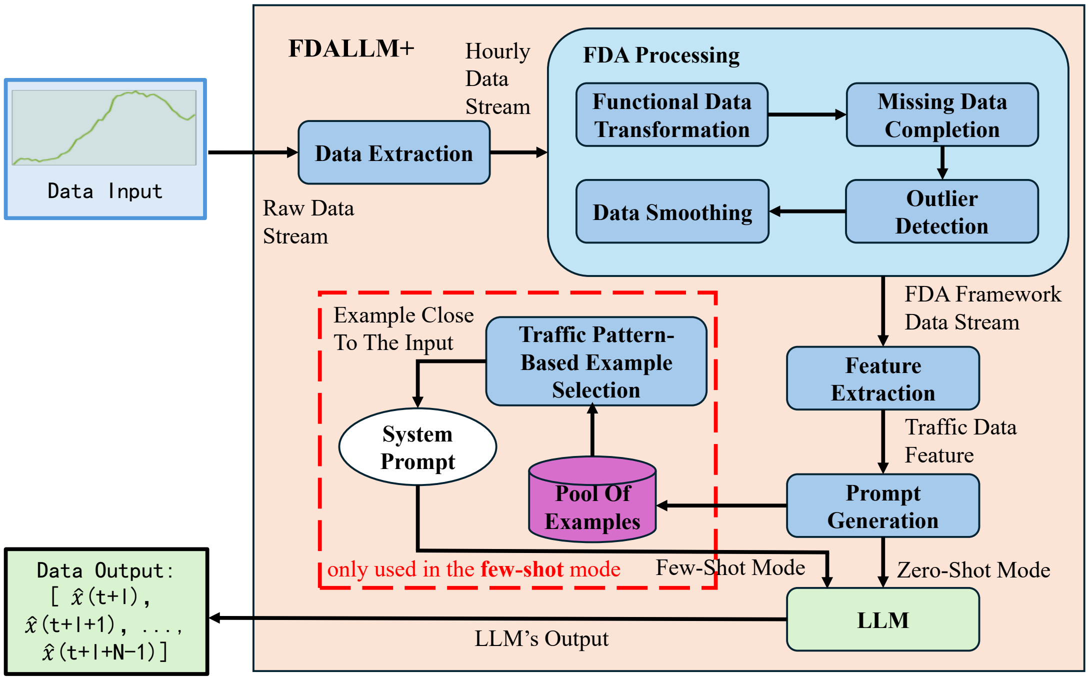


Figure 1: Overview of the architecture of the proposed FDALLM+ system: This system integrates FDA with LLMs to enhance traffic data prediction. Initially, raw traffic data is collected and processed into an hourly data stream. Through FDA processing, the data undergoes functional transformation, smoothing, missing data completion, and outlier detection to improve quality and consistency. The refined data is then used for feature extraction, generating structured prompts that serve as input to LLMs. In the few-shot mode, a traffic pattern-based example selection mechanism retrieves relevant past data from a pool of examples to enhance model performance, while in the zero-shot mode, the LLMs processes the generated prompt without additional reference data. Finally, the LLMs produces traffic predictions, providing valuable insights for downstream traffic forecasting applications.

over discrete time-series methods. Second, this representation facilitates the identification of long-term trends, seasonal patterns, and other complex temporal dependencies. Lastly, functional representation allows for an efficient description of the data through a limited number of basis functions and their corresponding coefficients, enabling dimensionality reduction and enhancing computational efficiency. Overall, functional data representation establishes a robust and efficient foundation for subsequent periodic feature extraction, trend analysis, and predictive modeling, significantly improving the model's ability to interpret and forecast complex temporal data patterns.

To further enhance data continuity, B-spline basis functions are applied for missing value imputation. B-splines are widely used in data fitting and interpolation due to their strong local control properties, smoothness, and flexibility, enabling seamless interpolation while preserving overall trends. B-spline basis functions are defined recursively, starting from the simplest first-order spline and progressively building higher-order splines to achieve varying levels of smoothness. The first-order B-spline basis function is defined as:

$$B_{i,1}(x) = \begin{cases} 1 & \text{if } t_i \le x < t_{i+1} \\ 0 & \text{otherwise.} \end{cases}$$

This function takes a value of 1 within the interval $[t_i, t_{i+1})$ and 0 elsewhere, acting as a piecewise step function and forming the basis for higher-order splines. Higher-order B-spline basis functions are recursively defined as:

$$B_{i,k}(x) = \frac{x - t_i}{t_{i+k-1} - t_i} B_{i,k-1}(x) + \frac{t_{i+k} - x}{t_{i+k} - t_{i+1}} B_{i+1,k-1}(x). \quad (5)$$

This formula generates higher-order basis functions through a linear combination of two lower-order basis functions $B_{i,k-1}(x)$ and $B_{i+1,k-1}(x)$. The weighting factors $\frac{x-t_i}{t_{i+k-1}-t_i}$ and $\frac{t_{i+k}-x}{t_{i+k}-t_{i+1}}$ ensure smooth transitions and continuity between the spline segments. Through this recursive definition, B-splines achieve high-order continuity between piecewise polynomials, providing smooth and coherent fitting results for missing data completion. Ultimately, the original data sequence $X(t)$ can be globally represented as a linear combination of B-spline basis functions:

$$X(t) = \sum_{k=1}^{K} c_k B_k(t), \quad (6)$$

where $B_k(t)$ represents the $k$th B-spline basis function, and $c_k$ is the corresponding coefficient.

In the process of missing value completion, the B-spline basis function library is first constructed using the recursive definition. The coefficients $c_k$ are then estimated using least squares fitting or similar methods to ensure that the fitted curve accurately matches the observed data points while providing smooth predictions in regions with missing data. The missing values at specific time points $t_{\text{missing}}$ are subsequently estimated by evaluating the fitted spline function, as

$$X(t_{\text{missing}}) = \sum_{k=1}^{K} c_k B_k(t_{\text{missing}}). \tag{7}$$

The primary advantage of the B-spline approach in missing value completion lies in its localized control property; each control knot influences only its neighboring interval, preventing drastic global fluctuations when completing missing values. Moreover, B-splines offer flexibility in adjusting the complexity of the fit by modifying the number and order of basis functions, thereby achieving an optimal balance between fitting accuracy and smoothness. By employing this method, all missing data points are successfully completed while ensuring continuity and smoothness, providing a robust and efficient foundation for subsequent periodic feature extraction, trend analysis, and predictive modeling.

To ensure robust smoothing, outlier detection and correction must first be applied to maintain data quality and reliability. Outliers, if left untreated, can significantly distort the smoothing process and adversely affect the performance of subsequent predictive models. In this study, the Interquartile Range (IQR) method is employed for outlier detection due to its robustness and simplicity. The IQR method is particularly effective because it does not assume any specific data distribution and is less sensitive to extreme values. The $IQR$ is calculated using the first quartile ($Q_1$) and the third quartile ($Q_3$) of the data, where $Q_1$ represents the 25th percentile and $Q_3$ represents the 75th percentile. The $IQR$ is calculated as $IQR = Q3 - Q1$. Any data point $x$ that lies outside the following range is considered an outlier:

$$x \notin [Q1 - 1.5 \times IQR,\ Q3 + 1.5 \times IQR]. \tag{8}$$

This criterion identifies both lower and upper outliers in the dataset. Once detected, the outliers are handled by replacing them with the nearest boundary values of the specified range. This replacement ensures that the continuity of the data distribution is preserved while mitigating the impact of extreme values. By correcting outliers in this manner, the data becomes cleaner and more consistent, providing a stable foundation for the smoothing process.

Following outlier detection and correction, the next step involves data smoothing to reduce noise and highlight underlying trends and periodic features. This study employs kernel smoothing, implemented through the Kernel Smoother function. Kernel smoothing is a non-parametric technique that reduces random fluctuations by computing weighted averages of neighboring data points, thereby revealing the overall trend more clearly.

The fundamental concept of kernel smoothing is that for each time point $t$, a smoothed value is calculated by taking a weighted average of the surrounding data points. The smoothing process can be mathematically represented as:

$$\hat{y}(x) = \left( \sum_{i=1}^{n} K\left( \frac{x - x_i}{h} \right) y_i \right) / \left( \sum_{i=1}^{n} K\left( \frac{x - x_i}{h} \right) \right), \tag{9}$$

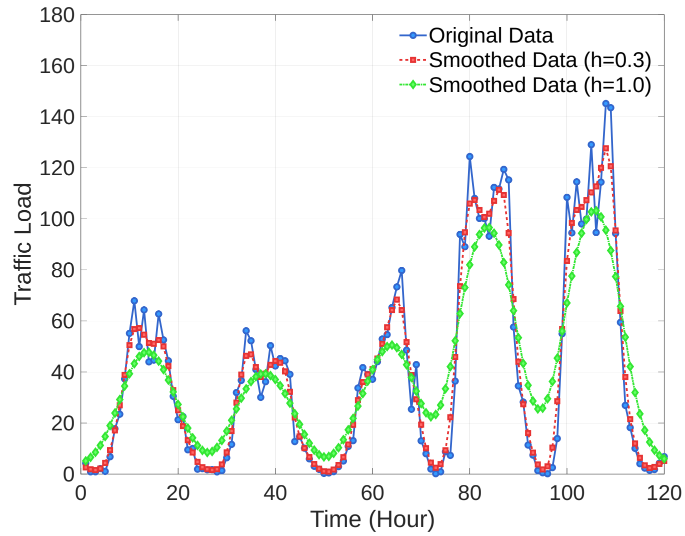


Figure 2: The smoothing method exhibits varying levels of effectiveness across different bandwidths $h$.

where $\hat{y}(t)$ represents the smoothed estimate at time $t$, $y_i$ denotes the observed value at time $t_i$, $K(\cdot)$ is the kernel function that determines the weights assigned to neighboring points, and $h$ is the bandwidth parameter that controls the degree of smoothing. A smaller $h$ value results in a curve that fits the data more closely, potentially leading to overfitting, whereas a larger $h$ value generates a smoother curve that may risk underfitting. As shown in Fig. 2, the smoothing effect varies across different bandwidths, demonstrating its adaptability to capture different data characteristics.

Kernel smoothing allows for the selection of different kernel functions—Gaussian, uniform, or Epanechnikov—while adjusting the bandwidth parameter to balance smoothing and fidelity. By selecting an appropriate kernel function and tuning the bandwidth parameter, a balance between smoothness and fitting accuracy can be achieved, preserving critical trends and periodic patterns in the data while minimizing noise.

After applying kernel smoothing, random fluctuations in the data are effectively suppressed, and long-term trends and periodic variations are more prominently highlighted. This refined dataset provides more reliable and higher-quality input for subsequent periodic feature extraction, dimensionality reduction, and predictive modeling. Ultimately, kernel smoothing refines temporal pattern recognition, enhances prediction accuracy, and preserves data integrity.

### *D. Feature Extraction: Weekday-Specific Analysis for Enhanced Temporal Understanding*

In the next stage of data processing, the focus shifts to feature extraction, where a new approach is adopted compared to previous methodologies. In earlier approaches, feature extraction was typically performed by first applying dimensionality reduction, followed by decomposing the entire dataset into trend, seasonality, and residual components [23]. These extracted features were subsequently integrated into the prompt for predictive modeling. In contrast, the current method introduces a more nuanced strategy by concentrating

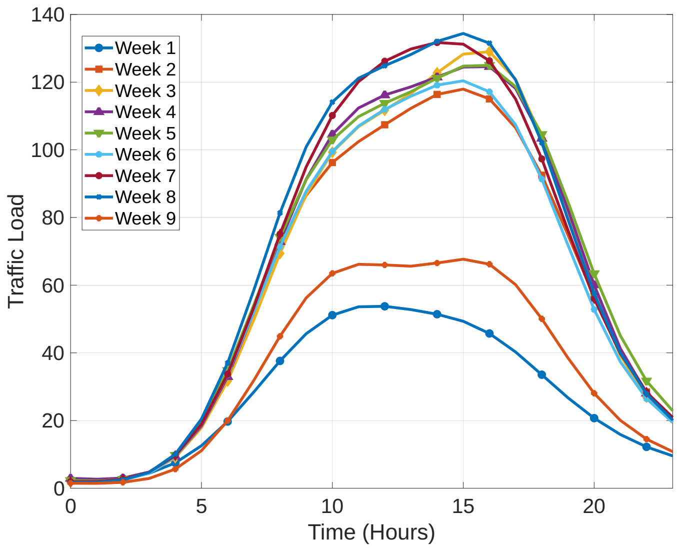


Figure 3: The periodic traffic pattern obtained by superimposing the original traffic data from all Mondays across weeks in the dataset.

on weekday-specific data aggregation and analysis, rather than mixing all data into a single analysis. This refined approach emphasizes medium-term temporal patterns and prioritizes critical weekday-level variations.

Data is first aggregated for each weekday, creating high-dimensional weekday-specific representations. Rather than a uniform dataset approach, each weekday's data (e.g., all Mondays, all Tuesdays, etc.) is analyzed separately. As shown in Fig. 3, the overlaid data from all Mondays reveal a clear periodic pattern, indicating the presence of recurring trends within the dataset. This strategy facilitates focused extraction of weekday-specific temporal features, improving the model's ability to capture recurring patterns that would otherwise be diluted. For each weekday, the average traffic flow is computed, representing the baseline flow characteristics of that day. This averaging process filters out random fluctuations, highlighting the essential temporal patterns while minimizing the influence of outliers.

A key aspect of this analysis is identifying anomalous dates, such as public holidays or pre/post-holiday periods. These anomalous periods typically exhibit traffic patterns that deviate significantly from the norm. By isolating such dates and conducting separate analyses, the method ensures that holiday-induced anomalies do not skew the standard weekday patterns. This targeted handling of exceptional dates further refines the accuracy of feature extraction, allowing for a clearer understanding of typical versus atypical temporal behaviors.

The extracted features encompass several critical dimensions of traffic flow dynamics. These include peak characteristics (e.g., peak flow values and the corresponding times they occur), valley characteristics (lowest flow points and their timings), traffic flow trends (e.g., periods of rapid flow increase or gradual decrease), and intra-day fluctuation amplitudes (quantifying daily variations in flow). Additionally, the analysis captures periodicity features, emphasizing the recurring 24-hour patterns typical in traffic data. The average weekday flow serves as a stable reference for distinguishing normal and abnormal traffic patterns.

This weekday-specific feature extraction emphasizes medium-term temporal patterns, aligning extracted features with key traffic dynamics. The methodology not only provides richer temporal context for each weekday, but also enhances the model's ability to recognize patterns critical for accurate forecasting. By isolating anomalous periods and focusing on weekday-level temporal dependencies, this approach refines feature extraction, enhancing model interpretability and robustness.

### E. Prompt Building

In traffic data forecasting tasks, the design of the prompt plays a critical role in enabling LLMs to understand data effectively and generate accurate predictions. In our previous work FDALLM [23], prompt construction primarily relied on the FDA approach, where global data was subject to dimensionality reduction and decomposed into three core components: trend $T(t)$, seasonality $S(t)$, and residual $R(t)$. These features captured long-term patterns and periodic behaviors, helping the model learn the data structure. However, this global, mixed-data analysis approach often overlooked subtle variations in traffic flow patterns across different temporal dimensions, especially between distinct weekdays. Given that traffic patterns can differ significantly depending on the day of week (e.g., weekends vs. weekdays), a unified global analysis may fail to capture these medium-term temporal dynamics effectively.

To address this limitation, this study proposes a novel prompt construction methodology that incorporates weekday-specific feature analysis. Unlike the previous approach, which extracted global features through unified analysis, this new method focuses on aggregating and analyzing data independently for each weekday to uncover distinct temporal patterns unique to individual days. By doing so, the new approach provides a refined analytical perspective that emphasizes medium-term temporal patterns, allowing the model to capture behavioral differences between weekdays and weekends more accurately.

The process begins by aggregating high-dimensional data corresponding to each weekday, forming weekday-specific datasets. For instance, all historical data corresponding to Mondays are aggregated into a single dataset, with similar datasets created for Tuesdays through Sundays. This aggregation strategy enables clearer identification of typical traffic flow patterns for each weekday, patterns that might otherwise be obscured in a globally mixed dataset. Once these weekday-specific datasets are established, the feature extraction process focuses on multiple critical dimensions. The average traffic flow level for each weekday is calculated to provide a baseline representation of the typical traffic intensity for that day. This averaging process not only reduces noise but also preserves the unique flow dynamics characteristic of each weekday.

Additionally, the analysis extracts peak flow characteristics, including the peak value and the corresponding time it occurs. These peaks typically represent rush hours or periods of heightened traffic demand triggered by special events. The valley characteristics, representing the lowest traffic flow levels and their corresponding times—often observed during late-night or early-morning hours—are also extracted. Beyond identifying peaks and valleys, the traffic flow trend is analyzed by examining how traffic volumes rise and fall throughout the day. For example, understanding how traffic surges from early morning lows to morning rush-hour peaks and subsequently declines after evening rush hours enables the model to better capture temporal flow dynamics.

To provide a more comprehensive description of intra-day traffic dynamics, the analysis includes the intra-day fluctuation amplitude, which measures the difference between maximum and minimum traffic volumes in a day. This metric reflects the overall intensity and variability of daily traffic. Explicitly handling anomalous dates prevents irregular patterns from distorting standard traffic behaviors, strengthening model reliability and forecast accuracy. These features collectively provide the LLM with rich temporal information, enabling it to recognize daily traffic cycles, detect behavioral trends, and predict future traffic flows more accurately.

A notable innovation in this new approach is the identification and handling of anomalous dates. Traditional analyses often mix public holidays or dates adjacent to holidays with regular dates, potentially distorting the underlying temporal patterns and introducing biases into the predictive model. To mitigate this, the proposed feature extraction process isolates these anomalous dates and analyzes them separately, ensuring that holiday-induced irregularities do not interfere with standard weekday patterns. For example, traffic patterns on days leading up to holidays often differ markedly due to increased travel demand. By constructing independent feature sets for these special dates, the model can account for their unique characteristics during forecasting, thereby enhancing its adaptability and robustness.

As shown in Fig. 4, in the newly designed prompt structure, the feature module has been completely redefined to reflect the comprehensive nature of weekday-specific analysis. Formally, the new feature module is expressed as:

$$\begin{aligned} Features_{\text{Weekday}} = \{&P_{\text{peak}}, T_{\text{trend}}, V_{\text{valley}}, \\ &F_{\text{fluctuation}}, A_{\text{average}}, C_{\text{cyclic}}, E_{\text{exception}}\}. \end{aligned} \tag{10}$$

The newly designed feature module incorporates several critical components that provide a comprehensive understanding of traffic flow dynamics. The peak feature ($P_{\text{peak}}$) captures the peak traffic flow values along with their corresponding occurrence times, offering insights into the timing and intensity of high-demand periods. The trend feature ($T_{\text{trend}}$) describes the rising and falling patterns of traffic throughout the day, highlighting key intervals of traffic build-up and decline. The valley feature ($V_{\text{valley}}$) identifies the lowest traffic volumes and their corresponding times, which typically represent off-peak hours when traffic demand is minimal. To reflect the overall daily variability in traffic, the fluctuation feature ($F_{\text{fluctuation}}$) measures the intra-day fluctuation amplitude, representing the range between maximum and minimum traffic flow. The average feature ($A_{\text{average}}$) provides the average daily traffic flow level, serving as a baseline for understanding typical traffic

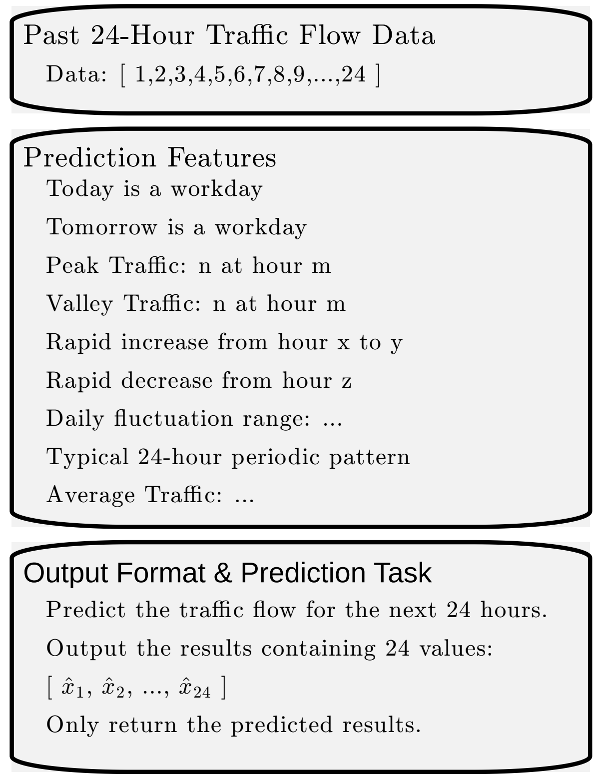


Figure 4: The structured design of the prompt, organizing input data, extracted features, and the expected output format.

conditions. Additionally, the cyclic feature ($C_{\text{cyclic}}$) highlights the periodicity of traffic patterns, capturing the typical 24-hour cycles that characterize daily traffic flows. Lastly, the exception feature ($E_{\text{exception}}$) focuses on anomalous dates, such as public holidays or adjacent days, which exhibit traffic patterns that deviate significantly from regular weekday behaviors. Together, these features provide a multi-dimensional perspective on traffic data, allowing the model to capture complex temporal dependencies and produce more accurate and context-aware forecasts.

With these new features incorporated, the overall prompt structure is summarized as follows:

$$\begin{aligned} Prompt = \{&Task\ Details,\quad Raw\ Data, \\ &Output\ Format,\quad Features_{\text{Weekday}}\}. \end{aligned} \tag{11}$$

This updated prompt structure enhances the traditional design by incorporating weekday-specific advanced features that provide richer temporal context. By focusing on the unique traffic dynamics of individual weekdays and isolating anomalous periods, the model gains a deeper understanding of the temporal dependencies critical to accurate traffic forecasting. Moreover, the explicit handling of anomalous dates ensures that irregular patterns do not skew the model's interpretation of standard traffic behaviors, resulting in more robust and more reliable forecasts.

## V. Experimental Evaluation

### A. Dataset Description

Our proposed FDALLM+ model is evaluated using the Milan dataset [49], which contains comprehensive information about telecommunication activities in the city of Milan, Italy. The dataset comprises two key components. The first component is SMS (Short Message Service Data, which captures the density of SMS reception and transmission across various regions. SMS data helps analyze distribution patterns, offering insights into network coverage and stability. Examining SMS density and frequency can reveal communication trends, helping to understand regional behaviors and network bottlenecks.

The second component, Call Data, captures hourly incoming and outgoing call densities across regions. This data is critical for identifying patterns in mobile communication traffic, enabling an understanding of how call activities vary throughout the day and across locations. Additionally, the call data offers valuable information regarding regional activity levels, highlighting areas with high telecommunication demand and providing insights into overall network usage. Analyzing these patterns helps predict peak network demand, assess performance, and optimize telecommunication infrastructure.

Together, these components form a robust foundation for evaluating FDALLM+ system in traffic prediction. FDALLM+ utilizes this dataset to model spatial-temporal dependencies in telecommunication traffic, improving urban traffic forecasting.

### B. Experiment Setup

We comprehensively evaluate FDALLM+ in traffic prediction tasks. Compared to previous experiments [23], FDALLM+ optimizes data preprocessing and introduces complex experimental scenarios for a more comprehensive evaluation. Specifically, the new FDALLM+ method, while maintaining a deep understanding of traffic data, incorporates both zero-shot and few-shot experimental modes. This design enables the assessment of the model's generalization ability under conditions with no example inputs and limited example support. By simulating varying data availability, this approach further validates FDALLM+'s applicability and robustness in complex traffic environments.

All prediction tasks were conducted using the ChatGPT client via prompt-based interactions. No local training or fine-tuning was performed, and the model weights remained fixed throughout the experiments. Therefore, our evaluation does not rely on specific hardware setups or traditional hyperparameter configurations such as learning rate or batch size. This setup reflects a black-box inference scenario, emphasizing the practical applicability of LLMs in resource-constrained or deployment-ready environments.

In terms of foundational model selection, this experiment eliminates the previously used PCA preprocessing method and instead focuses on periodic information extraction. This shift reduces the interference of traditional preprocessing steps and highlights the advantages of FDALLM+ in data comprehension. Specifically, we selected three state-of-the-art LLMs: GPT-4 [15], GPT-4o1, and Gemini [50]. GPT-4, developed by OpenAI, is a multimodal large language model known for its exceptional text comprehension and generation capabilities. GPT-4o1, an optimized version of GPT-4, offers enhanced performance for specific tasks. Gemini specializes in code generation and task automation, demonstrating strong generative capabilities. These models establish a strong foundation for evaluating FDALLM+'s adaptability across architectures and objectives.

To evaluate model performance, we use Mean Squared Error (MSE) and Mean Absolute Error (MAE) as primary metrics. These two metrics effectively reflect the accuracy and robustness of the models in traffic pattern prediction tasks. To demonstrate FDALLM+'s advantages in traffic prediction accuracy, we compare its performance against LSTM across various models and conditions. As a classical deep learning model for time series prediction, LSTM is widely used in traffic prediction tasks due to its ability to capture temporal dependencies and long-term trends in traffic data. Comparing FDALLM+ and LSTM results highlights the framework's effectiveness in improving accuracy and stability.

### C. Zero-Shot Prediction Perference Evaluation

In the zero-shot experiments, we systematically evaluated the performance of various foundational models combined with different preprocessing methods in traffic pattern prediction tasks. The experiment consisted of two parts: Table I presents the baseline prediction performance of different models across multiple traffic data types, while Table II further explores how periodic features contribute to improving prediction accuracy under various temporal scenarios. By integrating the analysis of these two parts, we can comprehensively understand the crucial role of periodic features in enhancing prediction accuracy, especially when there are significant changes in periodic information.

From Table I, it is evident that there are significant differences in prediction accuracy among the models under zero-shot conditions. Overall, the GPT-4o1 based FDALLM+ configuration demonstrates the best performance across all data types. Specifically, in the prediction of Call-out and SMS-out data, this configuration achieves the lowest MSE and MAE values, with 0.176 and 0.095 for Call-out, and 0.241 and 0.205 for SMS-out, respectively. These results highlight the model's exceptional ability to capture fluctuations in traffic flow, accurately identifying peak and off-peak traffic trends. Compared to the FDALLM and raw data methods, FDALLM+ achieves significant error reductions across all foundational models. For example, in the Call-out data, GPT-4o1 based FDALLM+ achieves a 34.3% reduction in error compared to FDALLM, while GPT-4 based FDALLM+ achieves a 20.2% reduction in the SMS-in data compared to the raw data method. These results demonstrate that FDALLM+, through deep modeling of traffic flow features during the preprocessing stage, significantly enhances the foundational models' understanding of the data, thereby improving overall prediction accuracy and robustness.

Furthermore, Table II provides an in-depth analysis of the impact of periodic features on prediction accuracy under different temporal scenarios. The experiment categorizes these scenarios into Workday-to-Workday (W-to-W), Workday-to-Holiday (W-to-H), Holiday-to-Holiday (H-to-H), and Holiday-to-Workday (H-to-W) to evaluate model performance under varying periodic conditions. The results highlight the critical

Table I: Zero-shot Traffic Prediction Performance (MSE and MAE) Across Different Models and Preprocessing Methods

| Model & Method | MSE | | | | MAE | | | |
|---|---|---|---|---|---|---|---|---|
| | Call-in | Call-out | SMS-in | SMS-out | Call-in | Call-out | SMS-in | SMS-out |
| LSTM (Baseline) | 0.422 | 0.464 | 0.512 | 0.449 | 0.355 | 0.242 | 0.400 | 0.353 |
| GPT-4 Based FDALLM | 0.287 | 0.247 | 0.445 | 0.368 | 0.305 | 0.153 | 0.347 | 0.289 |
| GPT-4 Based FDALLM+ | **0.242** | **0.226** | **0.402** | **0.333** | **0.297** | **0.127** | **0.304** | **0.247** |
| GPT-4 Based Raw Data | 0.322 | 0.265 | 0.504 | 0.407 | 0.335 | 0.182 | 0.393 | 0.309 |
| GPT-4o1 Based FDALLM | 0.243 | 0.182 | 0.404 | 0.317 | 0.232 | 0.134 | 0.304 | 0.256 |
| GPT-4o1 Based FDALLM+ | **0.221** | **0.176** | **0.370** | **0.241** | **0.163** | **0.095** | **0.287** | **0.205** |
| GPT-4o1 Based Raw Data | 0.303 | 0.217 | 0.491 | 0.372 | 0.244 | 0.157 | 0.362 | 0.270 |
| Gemini Based FDALLM | 0.332 | 0.303 | 0.456 | 0.403 | 0.265 | 0.173 | 0.279 | 0.304 |
| Gemini Based FDALLM+ | **0.291** | **0.242** | **0.431** | **0.381** | **0.225** | **0.142** | **0.356** | **0.272** |
| Gemini Based Raw Data | 0.352 | 0.314 | 0.543 | 0.442 | 0.287 | 0.202 | 0.403 | 0.335 |

Table II: Prediction Performance (MSE and MAE) Across Different Temporal Scenarios (W: workday; H: holiday)

| Model & Method | MSE | | | | MAE | | | |
|---|---|---|---|---|---|---|---|---|
| | W-to-W | W-to-H | H-to-H | H-to-W | W-to-W | W-to-H | H-to-H | H-to-W |
| LSTM (Baseline) | 0.332 | 0.453 | 0.345 | 0.508 | 0.255 | 0.382 | 0.277 | 0.368 |
| GPT-4 Based FDALLM | 0.243 | 0.344 | 0.224 | 0.329 | 0.249 | 0.332 | 0.222 | 0.341 |
| GPT-4 Based FDALLM+ | 0.258 | 0.291 | 0.258 | 0.334 | 0.247 | 0.329 | 0.275 | 0.353 |
| GPT-4 Based Raw Data | 0.281 | 0.378 | 0.263 | 0.415 | 0.239 | 0.402 | 0.257 | 0.339 |
| GPT-4o1 Based FDALLM | 0.188 | 0.296 | 0.176 | 0.307 | 0.141 | 0.288 | 0.151 | 0.261 |
| GPT-4o1 Based FDALLM+ | 0.176 | 0.235 | 0.153 | 0.254 | 0.132 | 0.203 | 0.142 | 0.199 |
| GPT-4o1 Based Raw Data | 0.245 | 0.362 | 0.253 | 0.357 | 0.153 | 0.292 | 0.162 | 0.308 |
| Gemini Based FDALLM | 0.281 | 0.378 | 0.256 | 0.339 | 0.229 | 0.304 | 0.209 | 0.342 |
| Gemini Based FDALLM+ | 0.254 | 0.328 | 0.218 | 0.328 | 0.203 | 0.241 | 0.209 | 0.251 |
| Gemini Based Raw Data | 0.321 | 0.423 | 0.331 | 0.398 | 0.235 | 0.312 | 0.267 | 0.305 |

role of periodic features in improving prediction accuracy, particularly in scenarios where significant periodic shifts occur.

For instance, in the Workday-to-Holiday scenario, GPT-4o1 based FDALLM+ outperforms all other models and methods, achieving the lowest MSE (0.235) and MAE (0.203). Compared to the LSTM baseline, which records MSE (0.453) and MAE (0.382), this represents reductions of 48.1% in MSE and 46.9% in MAE, demonstrating the significant impact of periodic feature integration. Similarly, in the Holiday-to-Workday scenario, GPT-4o1 based FDALLM+ achieves MSE (0.254) and MAE (0.199), substantially outperforming the LSTM baseline (MSE: 0.508, MAE: 0.368), with reductions of 50.0% and 45.9%, respectively.

These findings emphasize that leveraging periodic features allows the model to better capture shifts in traffic flow patterns between workdays and holidays, significantly enhancing prediction accuracy. The consistent improvements across different scenarios further underscore the essential role of periodic information in traffic forecasting tasks.

Notably, in the Holiday-to-Holiday and Workday-to-Workday scenarios, where traffic patterns are relatively stable, all models exhibit generally lower prediction errors. For example, GPT-4o1 based FDALLM+ achieves MSE and MAE values of 0.241 and 0.205, respectively, in the Workday-to-Workday scenario, and 0.241 and 0.177 in the Holiday-to-Holiday scenario. These results further indicate that when periodic variations are minimal, models can achieve high prediction accuracy even without relying heavily on complex periodic features.

However, in scenarios involving significant periodic fluctuations, such as transitions between workdays and holidays, periodic features become critical for enhancing prediction accuracy. For example, Fig. 5 provides a detailed comparison of the performance of FDALLM and FDALLM+ Zero-Shot across different time periods, which is an example of the Workday-to-Holiday scenario. It is evident that FDALLM exhibits significant deviations during peak hours, with its predicted peak position much earlier than the Ground Truth and the peak value being considerably overestimated. This discrepancy may stem from the model's insufficient capture of periodic features, particularly because FDALLM does not specifically extract features for each week day. When the periodicity of the target prediction day differs significantly from the current day—for example, when today is a weekday and tomorrow is a weekend—the prediction error of FDALLM tends to increase substantially, leading to either an exaggerated response to traffic surges or a temporal misalignment. In contrast, FDALLM+ Zero-Shot aligns more closely with the Ground Truth in both peak positioning and overall trend, indicating that this method more effectively utilizes periodic features, enabling the model to accurately capture traffic pattern shifts between weekdays and holidays, while avoiding overestimation of peak traffic. Additionally, FDALLM may have overfitted certain historical patterns during training, whereas FDALLM+ Zero-Shot, through the incorporation of additional feature guidance, enhances the model's general-

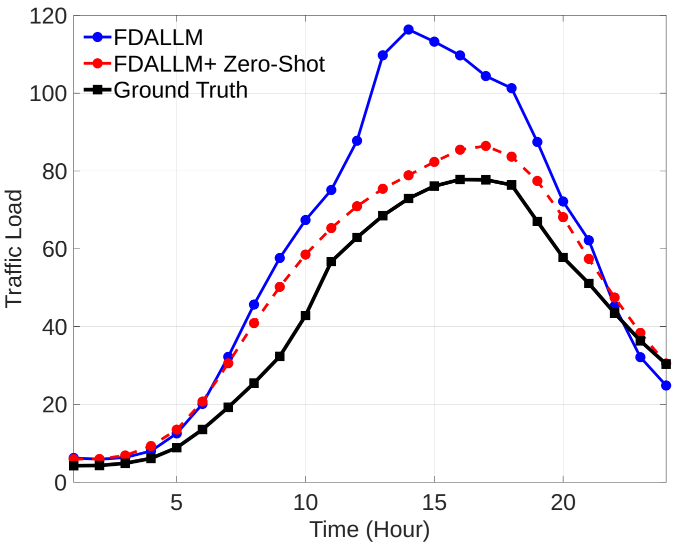


Figure 5: Comparison of FDALLM, FDALLM+ Zero-Shot, and Ground Truth over a 24-hour period from Weekday to Weekend.

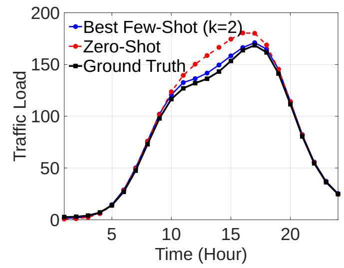


(a) Call-in ($k$=2).

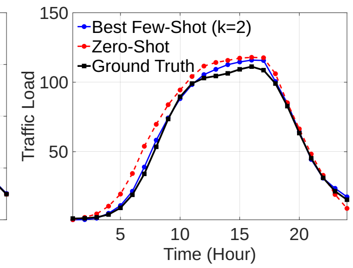


(b) Call-out ($k$=2).

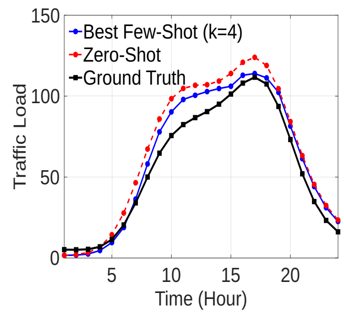


(c) SMS-in ($k$=4).

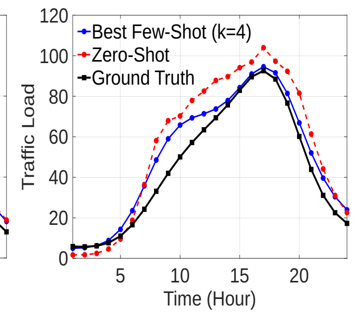


(d) SMS-out ($k$=4).

Figure 6: Comparison of Ground Truth, Zero-Shot Prediction, and Best Few-Shot Prediction in all data types.

ization ability across different time periods. Overall, these experimental results further confirm the critical role of periodic features in traffic prediction, particularly for peak-hour traffic forecasting, where properly integrating periodic information can significantly improve predictive accuracy.

Combining the findings from Table I and Table II, it is clear that foundation models exhibit certain limitations in prediction performance without the support of periodic features, especially when significant changes in periodic information occur. By incorporating periodic features, the FDALLM+ framework significantly improves the prediction accuracy of foundation models in these complex scenarios. For example, GPT-4o1 based FDALLM+ achieves the lowest MSE and MAE values across all temporal scenarios, demonstrating its superior ability to capture periodic information. Moreover, compared to other foundation models such as Gemini and GPT-4, GPT-4o1 consistently achieves higher prediction accuracy in temporal scenarios, further validating its advantages in complex traffic flow prediction tasks.

In summary, the zero-shot experiment results strongly confirm the critical role of periodic features in enhancing traffic pattern prediction accuracy. In particular, periodic features substantially improve model performance when there are significant changes in periodic information. Combined with FDALLM+, foundation models can effectively capture and interpret periodic traffic flow characteristics, achieving high accuracy and robustness in complex prediction tasks.

### D. *Few-Shot Prediction Perference Evaluation*

Few-shot experiments were conducted to evaluate the impact of providing a limited number of examples ($k$=1 to $k$=5) on the performance of the GPT-4o1 based FDALLM+ model in traffic pattern prediction tasks. In our few-shot experiments, we do not fine-tune or update any model parameters. Instead, we adopt a prompt-based in-context learning approach, where we directly input a small number of prompt-label pairs into the ChatGPT client as part of the prompt. The model uses these examples to infer patterns and make predictions on new data, but no gradient updates or weight changes occur. Table III presents the results, showcasing both MSE and MAE across four data types (Call-in, Call-out, SMS-in, and SMS-out) under varying few-shot conditions. The zero-shot mode ($k$=0) is included for comparison, allowing a direct examination of the performance improvements achieved through incremental example provision.

As shown in Fig. 6, the results clearly demonstrate that introducing even a small number of examples significantly enhances prediction accuracy across all data types. For instance, for the Call-in data, the MSE decreased from 0.221 in the zero-shot scenario ($k$=0) to 0.131 when $k$=2, representing a 40.7% reduction. Similarly, MAE for Call-in dropped from 0.163 when $k$=0 to 0.084 when $k$=2, indicating a 48.5% improvement. These findings suggest that the model greatly benefits from a small number of training examples, with the largest improvements typically occurring between $k$=0 and $k$=2. After $k$=2, the rate of improvement diminishes, and performance slightly declines at higher $k$ values ($k$=4 and $k$=5), potentially due to overfitting or redundancy in additional few-shot examples.

The Call-out data further supports this trend, with MSE decreasing from 0.176 at $k$=0 to 0.151 at $k$=2, and MAE dropping from 0.095 to 0.060. Notably, Call-out consistently exhibits the lowest error rates among all data types across different $k$ values, suggesting that traffic patterns in Call-out data are inherently easier to model and predict. This may be attributed to its lower volatility and more stable periodic trends compared to SMS data.

For the SMS-in and SMS-out datasets, although the error rates are generally higher than those for Call data, a similar

Table III: Prediction Performance (MSE and MAE) of GPT-4o1 Based FDALLM+ Across Zero-shot ($k$=0) and Few-shot ($k$=1∼5) Scenarios.

| **Data Type** | **MSE** | | | | | | **MAE** | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | $k$=0 | $k$=1 | $k$=2 | $k$=3 | $k$=4 | $k$=5 | $k$=0 | $k$=1 | $k$=2 | $k$=3 | $k$=4 | $k$=5 |
| Call-in | 0.221 | 0.157 | **0.131** | 0.148 | 0.167 | 0.192 | 0.163 | 0.107 | **0.084** | 0.116 | 0.130 | 0.168 |
| Call-out | 0.176 | 0.162 | **0.151** | 0.159 | 0.163 | 0.170 | 0.095 | 0.068 | **0.060** | 0.067 | 0.086 | 0.093 |
| SMS-in | 0.370 | 0.253 | 0.233 | 0.202 | **0.193** | 0.214 | 0.287 | 0.259 | 0.205 | 0.163 | **0.155** | 0.186 |
| SMS-out | 0.241 | 0.227 | 0.201 | 0.190 | **0.179** | 0.192 | 0.205 | 0.216 | 0.198 | 0.191 | **0.162** | 0.175 |

pattern of improvements with increased $k$ values is observed. In the SMS-in scenario, MSE dropped from 0.370 at $k$=0 to 0.233 at $k$=2, while MAE decreased from 0.287 to 0.205 over the same range. This suggests that while SMS data introduce more variability, the GPT-4o1 based FDALLM+ model can still leverage few-shot examples effectively to improve its predictive performance. However, at $k$=5, both MSE and MAE slightly increased, indicating diminishing returns and possible overfitting. Similar observations can be made for the SMS-out data results.

These results collectively indicate that the GPT-4o1 based FDALLM+ model achieves substantial performance gains with as few as two examples ($k$=2), making it highly effective in low-data scenarios. However, after $k$=3 or $k$=4, the improvements plateau, and additional examples may introduce noise or redundancy. The initial sharp improvements highlight the model's ability to generalize quickly from limited examples, a crucial advantage for real-world applications where large labeled datasets are often unavailable.

Moreover, the differences in performance across data types emphasize the varying complexity of traffic patterns. Call-out data appears more predictable, likely due to its regular structure, whereas SMS-in and SMS-out exhibit higher variability, making prediction more challenging. Nonetheless, the few-shot learning approach effectively narrows the performance gap across these data types, confirming its effectiveness.

The few-shot experiments confirm that the GPT-4o1 based FDALLM+ framework significantly improves traffic pattern prediction accuracy with minimal training data. The ability to achieve substantial performance improvements with as few as two examples ($k$=2) highlights the model's strong generalization capabilities, making it a promising solution for scenarios where annotated data are scarce. The observed performance trends also provide practical guidance for selecting the optimal number of few-shot examples, balancing between accuracy gains and the risk of overfitting.

### E. Bandwidth Sensitivity Analysis

To evaluate the influence of the kernel bandwidth parameter $h$ in the functional smoothing process, we conducted a sensitivity analysis by varying $h$ from 0.0 to 1.0 and measuring its impact on prediction performance. As shown in Table IV, both MSE and MAE initially decrease with increasing $h$, reaching their minimum values at $h = 0.3$, and then gradually increase as $h$ becomes larger. This trend is consistent across both Call and SMS data.

Table IV: Avg Impact of Bandwidth $h$ on MSE and MAE

| $h$ | **Call** | | **SMS** | |
|---|---|---|---|---|
| | MSE | MAE | MSE | MAE |
| 0.0 | 0.210 | 0.185 | 0.205 | 0.242 |
| 0.1 | 0.195 | 0.174 | 0.192 | 0.219 |
| 0.2 | 0.173 | 0.166 | 0.183 | 0.189 |
| **0.3** | **0.141** | **0.143** | **0.176** | **0.145** |
| 0.4 | 0.160 | 0.165 | 0.181 | 0.169 |
| 0.5 | 0.182 | 0.179 | 0.190 | 0.184 |
| 0.8 | 0.201 | 0.191 | 0.208 | 0.217 |
| 1.0 | 0.221 | 0.202 | 0.220 | 0.255 |

The results indicate that an overly small bandwidth (e.g., $h = 0.0$) retains too much noise, while an overly large bandwidth (e.g., $h = 1.0$) over smooths the signal and leads to underfitting. The optimal performance at $h = 0.3$ reflects a desirable balance between smoothing and signal fidelity. This validates the importance of proper kernel selection when applying functional data analysis to real-world traffic prediction tasks.

## VI. Conclusions

This paper introduced the FDALLM+ framework, leveraging FDA and LLMs for traffic pattern prediction, whose effectiveness and robustness were validated by an extensive experimental study. The results demonstrated FDALLM+'s superior predictive capabilities in both zero-shot and few-shot scenarios, with significantly reduced MSE and MAE compared to the baselines. Notably, the FDALLM+ framework effectively captured periodic traffic variations, particularly in workday-to-holiday transitions, highlighting the crucial role of periodic features in improving traffic prediction accuracy. Few-shot experiments further confirmed that FDALLM+ achieves substantial performance gains with as few as two examples, though diminishing returns were observed when more examples are used, emphasizing the need for better sample control.

While our experiments were based solely on the Milan dataset, we believe the proposed FDALLM+ framework is generalizable to other traffic datasets that exhibit similar temporal dynamics. For future work, we plan to validate the framework on additional datasets across different geographic and network contexts. In addition,x it is worth augmenting FDA with anomaly detection mechanisms or adaptive models, to deal with abrupt or irregular traffic behaviors caused by unforeseen events.

## References


[1] S. V. Kumar and L. Vanajakshi, "Short-term traffic flow prediction using seasonal ARIMA model with limited input data," *European Transport Research Review*, vol. 7, no. 1, p. 21, June 2015.

[2] R. Vinayakumar, K. P. Soman, and P. Poornachandran, "Applying deep learning approaches for network traffic prediction," in *Proc. 2017 Int. Conf. Advances Comput., Commun. Inform.*, Udupi, India, Sept. 2017, pp. 2353–2358.

[3] S. Long, J. Tan, B. Mao, F. Tang, Y. Li, M. Zhao, and N. Kato, "A survey on intelligent network operations and performance optimization based on Large Language Models," *IEEE Communications Surveys & Tutorials*, to appear. DOI: 10.1109/COMST.2025.3526606.

[4] S. Long, F. Tang, Y. Li, T. Tan, Z. Jin, M. Zhao, and N. Kato, "6G comprehensive intelligence: network operations and optimization based on Large Language Models," *IEEE Network*, to appear. DOI: 10.1109/MNET.2024.3470774.

[5] Y. Cheng, W. Zhang, Z. Zhang, C. Zhang, S. Wang, and S. Mao, "Towards federated large language models: Motivations, methods, and future direction," *IEEE Communications Surveys and Tutorials*, vol. 27, no. 4, pp. 2733–2764, Aug. 2025.

[6] M. Xu, N. Dusit, J. Kang, Z. Xiong, S. Mao, Z. Han, D. Kim, and K. Lataief, "When large language model agents meet 6G networks: Perception, grounding, and alignment," *IEEE Wireless Communications*, vol. 23, no. 6, pp. 63–71, Dec. 2024.

[7] M. A. Merrill, M. Tan, V. Gupta, T. Hartvigsen, and T. Althoff, "Language models still struggle to zero-shot reason about time series," arXiv:2404.11757, Apr. 2024. [Online]. Available: https://arxiv.org/abs/2404.11757

[8] C. Chang, W.-Y. Wang, W.-C. Peng, and T.-F. Chen, "LLM4TS: Aligning pre-trained LLMs as data-efficient time-series forecasters," arXiv:2308.08469, Aug. 2023. [Online]. Available: https://arxiv.org/abs/2308.08469

[9] V. Ekambaram, et al., "Tiny time mixers (TTMs): Fast pre-trained models for enhanced zero/few-shot forecasting of multivariate time series," arXiv:2401.03955, June 2024. [Online]. Available: https://arxiv.org/abs/2401.03955

[10] G. Chatzigeorgakidis, K. Lentzos, and D. Skoutas, "MultiCast: Zero-shot multivariate time series forecasting using LLMs," in *Proc. 2024 IEEE Int. Conf. Data Engineering Workshops (ICDEW)*, Utrecht, Netherlands, May 2024, pp. 119–127.

[11] K. Deng, G. Sun, and P. C. Woodland, "Wav2Prompt: End-to-end speech prompt generation and tuning for LLM in zero and few-shot learning," arXiv:2406.00522, June 2024. [Online]. Available: https://arxiv.org/abs/2406.00522

[12] D. Spathis and F. Kawsar, "The first step is the hardest: Pitfalls of representing and tokenizing temporal data for Large Language Models," arXiv:2309.06236, Sept. 2023. [Online]. Available: https://arxiv.org/abs/2309.06236

[13] M. Jin, et al., "Time-LLM: Time series forecasting by reprogramming Large Language Models," arXiv:2310.01728, Jan. 2024. [Online]. Available: https://arxiv.org/abs/2310.01728

[14] X. Liu, J. Hu, Y. Li, S. Diao, Y. Liang, B. Hooi, and R. Zimmermann, "UniTime: A language-empowered unified model for cross-domain time series forecasting," arXiv:2310.09751, Oct. 2023. [Online]. Available: https://arxiv.org/abs/2310.09751

[15] OpenAI, "GPT-4 Technical Report," Mar. 2024. [Online]. Available: https://cdn.openai.com/papers/gpt-4.pdf

[16] N. Gruver, M. Finzi, S. Qiu, and A. G. Wilson, "Large language models are zero-shot time series forecasters," arXiv:2310.07820, May 2024. [Online]. Available: https://arxiv.org/abs/2310.07820

[17] H. Xue and F. D. Salim, "PromptCast: A new prompt-based learning paradigm for time series forecasting," arXiv:2210.08964, Dec. 2023. [Online]. Available: https://arxiv.org/abs/2210.08964

[18] D. Cao, et al., "TEMPO: Prompt-based generative pre-trained transformer for time series forecasting," arXiv:2310.04948, Oct. 2023. [Online]. Available: https://arxiv.org/abs/2310.04948

[19] H. Liu, et al., "LSTPrompt: Large language models as zero-shot time series forecasters by long-short-term prompting," arXiv:2402.16132, Feb. 2024. [Online]. Available: https://arxiv.org/abs/2402.16132

[20] J.-L. Wang, J.-M. Chiou, and H.-G. Müller, "Functional data analysis," *Annual Review of Statistics and Its Application*, vol. 3, pp. 257–295, June 2016.

[21] S. Wang, Y. Huang, and G. Cao, "Review on functional data classification," *WIREs Computational Statistics*, vol. 16, no. 1, p. e1638, Nov. 2023.

[22] T. Zhao, N. Wang, G. Cao, S. Mao, and X. Wang, "Functional data analysis assisted cross-domain Wi-Fi sensing using few-shot learning," in *Proc. IEEE ICC 2024*, Denver, CO, June 2024, pp. 4780–4785.

[23] Y. Sun, X. Wang, G. Cao, and S. Mao, "FDALLM: Traffic data prediction with functional data analysis and Large Language Models," in *Proc. IEEE ICC 2025*, Montreal, Canada, June 2025, pp. 845–850.

[24] M. Hasegawa, G. Wu, and M. Mizuni, "Applications of nonlinear prediction methods to the internet traffic," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Sydney, Australia, May 2001, pp. 169–172.

[25] C. Chen, J. Hu, Q. Meng, and Y. Zhang, "Short-time traffic flow prediction with ARIMA-GARCH model," Baden-Baden, Germany, pp. 607–612, June 2011.

[26] D. Xu, Y. Wang, L. Jia *et al.*, "Real-time road traffic state prediction based on ARIMA and Kalman filter," *Frontiers of Information Technology & Electronic Engineering*, vol. 18, pp. 287–302, Feb. 2017.

[27] X. Zheng, W. Lai, H. Chen, and S. Fang, "Data prediction of mobile network traffic in public scenes by SOS-vSVR method," *MDPI Sensors*, vol. 20, no. 3, p. 603, Jan. 2020.

[28] B. W. J. Evans and A. Hamilton, "Forecasting road traffic conditions using a context-based random forest algorithm," *Transportation Planning and Technology*, vol. 42, no. 6, pp. 554–572, June 2019.

[29] M. Lopez-Martin, B. Carro, and A. Sanchez-Esguevillas, "IoT type-of-traffic forecasting method based on gradient boosting neural networks," *Elsevier Future Generation Computer Systems*, vol. 105, pp. 331–345, Apr. 2020.

[30] Y. He, P. Huang, W. Hong, Q. Luo, L. Li, and K.-L. Tsui, "In-depth insights into the application of recurrent neural networks (rnns) in traffic prediction: A comprehensive review," *Algorithms*, vol. 17, no. 9, Sept. 2024.

[31] H. D. Trinh, L. Giupponi, and P. Dini, "Mobile traffic prediction from raw data using lstm networks," in *Proc. IEEE PIMRC'18*, Bologna, Italy, Sept. 2018, pp. 1827–1832.

[32] X. Ta, Z. Liu, X. Hu, L. Yu, L. Sun, and B. Du, "Adaptive spatio-temporal graph neural network for traffic forecasting," *Elsevier Knowledge-Based Systems*, vol. 242, p. 108199, Apr. 2022.

[33] J. Tang, F. Liu, Y. Zou, W. Zhang, and Y. Wang, "An improved fuzzy neural network for traffic speed prediction considering periodic characteristic," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 9, pp. 2340–2350, Sept. 2017.

[34] J. Gertheiss, D. Rügamer, B. X. W. Liew, and S. Greven, "Functional data analysis: An introduction and recent developments," *Biometrical Journal*, Sept. 2024.

[35] S. Wang, W. Zhang, G. Cao, and Y. Huang, "Functional data analysis using deep neural networks," *WIREs Computational Statistics*, vol. 16, no. 4, p. e70001, Aug. 2024.

[36] I. M. Wagner-Muns, I. G. Guardiola, V. A. Samaranayke, and W. I. Kayani, "A functional data analysis approach to traffic volume forecasting," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 878–888, Mar. 2018.

[37] C. Liu, S. Yang, Q. Xu, Z. Li, C. Long, Z. Li, and R. Zhao, "Spatial-temporal large language model for traffic prediction," *arXiv preprint arXiv:2401.10134*, 2024. [Online]. Available: https://arxiv.org/abs/2401.10134

[38] S. M. Moghadas, Y. Lyu, B. Cornelis, A. Alahi, and A. Munteanu, "Strada-LLM: Graph LLM for traffic prediction," *arXiv preprint arXiv:2410.20856*, 2024. [Online]. Available: https://arxiv.org/abs/2410.20856

[39] Y. Ren, Y. Chen, S. Liu, B. Wang, H. Yu, and Z. Cui, "TPLLM: A traffic prediction framework based on pretrained Large Language Models," *arXiv preprint arXiv:2403.02221*, 2024. [Online]. Available: https://arxiv.org/abs/2403.02221

[40] Z. Lan, L. Liu, B. Fan, Y. Lv, Y. Ren, and Z. Cui, "Traj-LLM: A new exploration for empowering trajectory prediction with pre-trained large language models," *IEEE Transactions on Intelligent Vehicles*.

[41] I. Chamieh, T. Zesch, and K. Giebermann, "LLMs in short answer scoring: Limitations and promise of zero-shot and few-shot approaches," in *Proc. Workshop Innov. Use NLP Educ. Appl. (BEA)*, Mexico City, Mexico, June 2024, pp. 309–315.

[42] S. Mao and S. S. Panwar, "A survey on envelope processes and their applications in quality of service provisioning," *IEEE Communications Surveys and Tutorials*, vol. 8, no. 3, pp. 2–20, Third Quarter 2006.

[43] X. Ding and L. Wang, "Do language models understand time?" *arXiv preprint arXiv:2412.13845*, 2024. [Online]. Available: https://arxiv.org/abs/2412.13845

[44] C. Yuan, Q. Xie, and S. Ananiadou, "Zero-shot temporal relation extraction with chatgpt," *arXiv preprint arXiv:2304.05454*, 2023. [Online]. Available: https://arxiv.org/abs/2304.05454

[45] X. Guo, Q. Zhang, J. Jiang, M. Peng, M. Zhu, and H. F. Yang, "Towards explainable traffic flow prediction with large language models," *Elsevier Communications in Transportation Research*, vol. 4, p. 100150, Dec. 2024.

[46] X. Guo, Q. Zhang, J. Jiang, M. Peng, H. F. Yang, and M. Zhu, "Towards responsible and reliable traffic flow prediction with large language models," *SSRN Electronic Journal*, 2024. [Online]. Available: https://ssrn.com/abstract=4805901

[47] X. Huang, "Enhancing traffic prediction with textual data using large language models," *arXiv preprint arXiv:2405.06719*, 2024. [Online]. Available: https://arxiv.org/abs/2405.06719

[48] Z. Pan, Y. Jiang, S. Garg, A. Schneider, Y. Nevmyvaka, and D. Song, "$S^2$IP-LLM: Semantic space informed prompt learning with LLM for time series forecasting," in *Proc. ICML 2024*, Vienna Austria, July 2024, pp. 39 135–39 153.

[49] G. Barlacchi, M. De Nadai, R. Larcher, and et al., "A multi-source dataset of urban life in the city of Milan and the province of Trentino," *Scientific Data*, vol. 2, p. 150055, Oct. 2015.

[50] Gemini Team, "Gemini: A family of highly capable multimodal models," arXiv:2312.11805, Dec. 2023. [Online]. Available: https://arxiv.org/abs/2312.11805

**Yujie Sun** (Graduate Student Member, IEEE) [S' 24] received the B.S. degree in Information Security from Civil Aviation University of China, Tianjin, China, in 2021. He received the M.S. degree in Applied Machine Intelligence from Northeastern University, Boston, MA, USA, in 2024. He has been pursuing a Ph.D. degree in the Department of Electrical and Computer Engineering at Auburn University since 2024. His current research focuses on Artificial Intelligence of Things (AIoT) and wireless sensing.

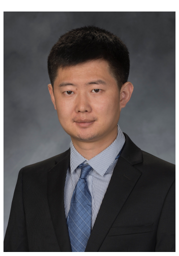

**Xuyu Wang** [S'13-M'18] received the M.S. in Signal and Information Processing in 2012 and B.S. in Electronic Information Engineering in 2009, both from Xidian University, Xi'an, China. He received a Ph.D. in Electrical and Computer Engineering from Auburn University, Auburn, AL, USA in Aug. 2018. He is currently an Assistant Professor in the Knight Foundation School of Computing and Information Sciences, Florida International University, Miami, FL, USA. His research interests include wireless sensing, Internet of Things, wireless localization, smart health, wireless networks, and deep learning. He received the NSF CRII Award in 2021. He was a co-recipient of the ACM FAcct 2023 Best Paper Award, the 2022 Best Journal Paper Award of IEEE ComSoc eHealth Technical Committee, the IEEE INFOCOM 2022 Best Demo Award, the IEEE ICC 2022 Best Paper Award, the IEEE Vehicular Technology Society 2020 Jack Neubauer Memorial Award, the IEEE GLOBECOM 2019 Best Paper Award, the IEEE ComSoc MMTC Best Journal Paper Award in 2018, the IEEE PIMRC 2017 Best Student Paper Award, and the IEEE SECON 2017 Best Demo Award. He is an associate editor for the IEEE Transactions on Mobile Computing.

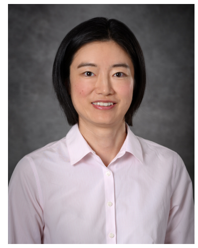

**Guanqun Cao** received a Ph.D. degree in Statistics from Michigan State University in 2012. She is an associate professor of Statistics and Probability with a joint appointment in Computational Mathematics, Science and Engineering at Michigan State University. Prior joining Michigan State University in 2023, she was on the faculty of Auburn University (2012-2023). Her work is at the interface of statistics, mathematics, and computer science, developing cutting-edge statistical methods for solving issues related to data science and big data analytics. She is an Elected Member of the International Statistical Institute (2022).

**Shiwen Mao** (Fellow, IEEE) received a Ph.D. in Electrical Engineering from Polytechnic University in 2004. He is a Professor and Earle C. Williams Eminent Scholar, and Director of the Wireless Engineering Research and Education Center at Auburn University. His research interests include wireless networks, multimedia communications, and smart grid. He is the editor-in-chief of IEEE Transactions on Cognitive Communications and Networking, a member-at-large on the Board of Governors of IEEE Communications Society, and Vice President of Technical Activities of IEEE Council on Radio Frequency Identification (CRFID). He received the IEEE ComSoc MMTC Outstanding Researcher Award in 2023, the SEC 2023 Faculty Achievement Award for Auburn, the IEEE ComSoc TC-CSR Distinguished Technical Achievement Award in 2019, the Auburn University Creative Research & Scholarship Award in 2018, the NSF CAREER Award in 2010, and several service awards from IEEE ComSoc. He is also a co-recipient of the 2022 Best Journal Paper Award of IEEE ComSoc eHealth Technical Committee, the 2021 Best Paper Award of Elsevier/KeAi Digital Communications and Networks Journal, the 2021 IEEE Internet of Things Journal Best Paper Award, the 2021 IEEE Communications Society Outstanding Paper Award, the IEEE Vehicular Technology Society 2020 Jack Neubauer Memorial Award, the 2018 Best Journal Paper Award and the 2017 Best Conference Paper Award from IEEE ComSoc MMTC, and the 2004 IEEE Communications Society Leonard G. Abraham Prize in the Field of Communications Systems. He is a co-recipient of the Best Paper/Demo Awards of 12 conferences.