



Full Length Article

Towards explainable traffic flow prediction with large language models

Xusen Guo^{a,1}, Qiming Zhang^{a,1}, Junyue Jiang^b, Mingxing Peng^a, Meixin Zhu^{a,c,*}, Hao Frank Yang^{b,**}

^a Intelligent Transportation Thrust, System Hub, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, 511400, China

^b Department of Civil and System Engineering, Johns Hopkins University, Baltimore, 21218, USA

^c Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things, Guangzhou, 511400, China

ARTICLE INFO

Keywords:

Traffic flow prediction
Large language models
Spatial-temporal prediction
Explainability

ABSTRACT

Traffic forecasting is crucial for intelligent transportation systems. It has experienced significant advancements thanks to the power of deep learning in capturing latent patterns of traffic data. However, recent deep-learning architectures require intricate model designs and lack an intuitive understanding of the mapping from input data to predicted results. Achieving both accuracy and explainability in traffic prediction models remains a challenge due to the complexity of traffic data and the inherent opacity of deep learning models. To tackle these challenges, we propose a traffic flow prediction model based on large language models (LLMs) to generate explainable traffic predictions, named xTP-LLM. By transferring multi-modal traffic data into natural language descriptions, xTP-LLM captures complex time-series patterns and external factors from comprehensive traffic data. The LLM framework is fine-tuned using language-based instructions to align with spatial-temporal traffic flow data. Empirically, xTP-LLM shows competitive accuracy compared with deep learning baselines, while providing an intuitive and reliable explanation for predictions. This study contributes to advancing explainable traffic prediction models and lays a foundation for future exploration of LLM applications in transportation.

1. Introduction

Traffic network prediction is a critical component of transportation management systems, aiming to forecast future traffic conditions such as congestion, traffic volume, and travel time. It plays a vital role in various applications, including route planning, traffic management, and intelligent transportation systems (ITS) (Dimitrakopoulos and Demestichas, 2010). Accurate predictions are crucial for providing valuable insights to stakeholders in traffic systems, aiding in informed decision-making. However, achieving reliability and precision in predictions is challenging due to the inherent nonlinear dynamics of traffic, spatial and temporal variability, and dynamic nature influenced by factors like accidents, weather conditions, and other events, necessitating models capable of capturing complex dependencies and rapid changes. As large-scale traffic data becomes more available and deep learning techniques advance, data-driven methods are increasingly favored for modeling complex traffic flow systems.

Deep learning-based traffic analysis (Wang et al., 2022) has been well researched, covering human mobility research (Jin et al., 2018; Zhang et al., 2017), traffic management (Du et al., 2020; Ranjan et al., 2020), and accident analysis (Yannis et al., 2017). Typically, these issues are treated as spatio-temporal deep learning problems. Deep learning methods consistently learn hierarchical feature representations from spatial-temporal data, understand historical trends (Fan et al., 2018; Yang et al., 2016), and employ graphs to illustrate the spatial relationships between locations (Gao et al., 2018; Marblestone et al., 2016). The spatio-temporal-graph learning paradigm is the primary methodology for learning representations and capturing potential trends and relationships from traffic data.

Despite the comprehensive model architectures, this domain format still faces challenges. First, deep learning methods require specific network structure designs (Asanjan et al., 2018; Li et al., 2017), to consider the multi-modal dynamic nature and spatio-temporal complexity of traffic data. Although these designs can help the model

* Corresponding author. Intelligent Transportation Thrust, System Hub, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, 511400, China.

** Corresponding author.

E-mail addresses: meixin@ust.hk (M. Zhu), haofrankyang@jhu.edu (H.F. Yang).

¹ Xusen Guo and Qiming Zhang contributed equally to this work.

improve the accuracy of predictions, they are generally artificially crafted, and abstract representations further obscure generalization ability. Additionally, some traffic forecasting studies (Chen et al., 2022; Dai and Wang, 2020; Zhou et al., 2019) related to the reliability and robustness of models prioritize the degree of fit to actual values. They consider a variety of dynamic factors affecting traffic flow to improve the reliability of traffic flow predictions. However, these methodologies encounter difficulties in providing credentials for prediction and taking accountability for outcomes in real-world traffic prediction tasks that lack labeled data.

Recently, with the popularity of foundation models (Chang, 2023; Touvron et al., 2023), large language models (LLMs) provide more intuitive explainable trials in spatial-temporal domain tasks, benefiting from textual paradigms. Spatial-temporal learning tasks can be refined into language format for exploring the potential of LLMs in various application fields, including forecasting (Jin et al., 2023; Shi et al., 2023), data imputation (Chen et al., 2023), and anomaly detection (Wan et al., 2024), as depicted in Fig. 1. Considering multi-modalities in urban big data, the LLM framework transfers original data into natural language description and is capable of capturing latent relationships between inputs from complicated contexts. Also, LLMs can generate explanations of the reasoning process, which provides available supplements for prediction and decision-making. However, though LLMs have expertise in language-based tasks, how to formulate specific spatial-temporal tasks into LLM frameworks for accurate prediction remains challenging.

In response to the challenges mentioned, we present xTP-LLM, a novel traffic flow prediction framework utilizing LLMs. This framework effectively predicts future traffic flow using multi-modal data while providing interpretable explanations. By converting traffic data and external factors—such as points of interest, weather conditions, dates, and holidays—into a structured prompt, xTP-LLM captures domain-specific information more effectively. Our experiments reveal that xTP-LLM not only outperforms state-of-the-art deep learning models in accuracy but also exhibits robust generalization across various scenarios. The main contributions of this work are as follows.

- We design a structured textual prompt that incorporates multi-modal traffic flow information, facilitating LLMs to capture traffic patterns better. Additionally, we reorganized the multi-modal traffic flow data in California to create a comprehensive text-based dataset, CATraffic, for future exploration in traffic prediction with LLMs.
- We propose xTP-LLM, a traffic flow prediction framework based on LLMs. This framework exhibits competitive accuracy against state-of-the-art deep learning models and demonstrates effective

generalization abilities across different traffic flow prediction scenarios without additional training.

- xTP-LLM can provide insightful explanations for its prediction results, offering greater interpretability than traditional deep learning approaches, which facilitates rational decision-making in traffic management and planning.

For the rest parts of this paper, the related works are involved in Section 2, and details about methodology are illustrated in Section 3. Comparison results, ablation, explanation, and what-if studies are included in experimental results and analysis, as Section 4. Section 5 summarizes this paper and gives insights for future exploration of LLM applications in transportation.

2. Related works

In this section, we will first explore advancements in spatial-temporal prediction, emphasizing the integration of deep learning methodologies. Subsequently, we will delve into the importance of reliable prediction, discussing methods for enhancing explainability in spatial-temporal learning. Finally, we will highlight the transformative role of LLMs across diverse domains, elucidating their pre-training and fine-tuning practices for domain-specific tasks.

2.1. Spatial-temporal prediction

Recently, the field of spatial-temporal learning has witnessed significant advancements, especially in the traffic, environment, and society field, due to the emergence of deep learning methodologies. These approaches have enabled the modeling of latent relationships among various features of urban data in diverse formats. These architectural designs are meticulously crafted to comprehend and represent the intricate interplay between spatial and temporal dimensions within datasets. Convolutional neural networks (CNNs) (Liang et al., 2018; Zhang et al., 2017), quite renowned for their efficacy in computer vision, are employed to discern the spatial relations among grid regions, by filtering the input data. Moreover, recurrent neural networks (RNNs) (Hopfield, 1982) are usually leveraged to adeptly capture temporal dependencies, through the maintenance of a memory state, facilitating the reusing of information over time. Notably, more spatial-temporal learning frameworks introduce graph neural networks (GNNs) (Guo et al., 2019; Wu et al., 2019; Yu et al., 2018), advanced in the representation of complex spatial relationships inherent in data structured as

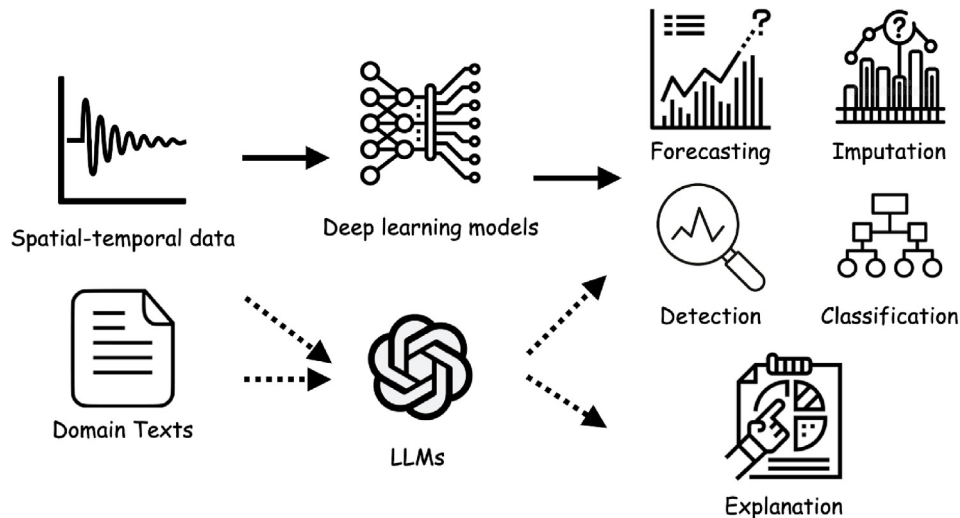


Fig. 1. Spatial-temporal learning formats for deep learning models and LLMs: Compared with well-developed deep learning models, LLMs for spatial-temporal learning advance in adapting domain knowledge with urban multi-modal data and generate reasonable explanations.

graphs, wherein nodes correspond to spatial locations and edges encapsulate the connections between them. Additionally, the adaptation of transformers (Vaswani et al., 2017), originally proposed for natural language processing, has proven effective in long-sequence modeling to capture comprehensive information. Various transformer blocks (Cai et al., 2020; Yu et al., 2020) have been tailored for different dependencies among spatial-temporal features, enabling the modeling of intricate relationships. A notable trend in this domain is the combination of different model architectures, leveraging various modules for spatial or temporal features (Jiang et al., 2023; Xu et al., 2020). This amalgamation gradually becomes the prevailing paradigm, showcasing promising performance in prediction tasks. However, it is worth noting that while these methods perform excellent in prediction accuracy, they often fall short in terms of explainability and generalization.

2.2. Explainable prediction

The explainability of spatial-temporal learning is also worthy of consideration for reliable prediction, which provides abundant views beyond prediction accuracy. Most recent works studied which features mostly affect decisions generated by models (Barredo-Arrieta et al., 2019) focus on the dependency on latent variables of road forecasting based on black-box machine learning methods, including RNNs and Random Forests. The spatial-temporal causal graph inference, as presented in Zhang et al. (2022), offers an approximation of the Granger causality test, thereby enhancing the accessibility of forecasting. Counterfactual explanations for time series (Ates et al., 2021; Yan and Wang, 2023) are also highly regarded, as they concentrate on generating alternative prediction outcomes by selecting time series data points from the training set and substituting them into the sample under analysis. This method allows for the illustration of results by examining a limited number of variables. Besides these methods focusing on model transparency, LLMs offer an alternative approach to generate convincing explanations along with prediction results with greater intuitiveness (Gruver et al., 2024; Huang et al., 2023; Peng et al., 2024). As a bridge between complex systems and humans, LLMs have the potential to convert input-output mappings into natural, easy-to-understand narratives, helping humans make more informed and reliable decisions (Cambria et al., 2024; Zytek et al., 2024).

2.3. Large language models

Large language models (LLMs) have achieved remarkable success across a wide range of tasks and fields, including natural language processing (Ray, 2023), vision-language tasks (Liu et al., 2023), and various other interdisciplinary domains (Thirunavukarasu et al., 2023; Wu et al., 2023). Originally designed as pre-trained language foundation models for addressing various natural language tasks, LLMs have exhibited the capacity to acquire intricate semantic and knowledge representations from extensive text corpora over time. This newfound ability has been a profound source of inspiration within the community for addressing a variety of tasks. The success of models like GPT-4 (Chang, 2023) in natural language understanding and generation tasks has spurred interest in exploring their potential for handling complex, multi-modal datasets beyond traditional linguistic domains. They can extract valuable information and relationships from complex textual contexts, thereby enhancing the learning of city data. With the popularity of decoder-architecture LLMs, domain tasks are normally formulated into the next token generation, which provides a unified formulation to learn the map from the input to the output. To acquire large models for specific fields, the practices of pre-training and fine-tuning have become widely accepted in the model training process. Pre-training a foundation model from scratch necessitates substantial computing resources and domain-specific datasets, resulting in its superior performance within professional domains compared to baseline models. On the other hand, fine-tuning based on foundation models offers a more accessible

approach, involving adjustments to only a few parameters (Hu et al., 2021). This method preserves most general knowledge while targeting expertise in domain-specific tasks. In some cases, researchers freeze all parameters of LLMs and focus solely on training the extended encoders and decoders (Jin et al., 2023; Zhu et al., 2023). This strategy aims to extend the learning capabilities of LLMs to domain-specific tasks while leveraging their existing knowledge base.

3. Methodology

This study presents a novel approach utilizing LLMs for traffic flow prediction. Our objective is to develop a predictive model that not only forecasts traffic flow patterns but also provides explanations for the predicted trends. In the following sections, we will provide a comprehensive overview of our approach. Firstly, we will describe the problem formulation and the predictive framework. Next, we will discuss the construction of prompts, which is crucial for fine-tuning LLMs. Finally, we will talk about the fine-tuning technique and how to generate the explanation corresponding to the predicted result.

3.1. Problem description

The traffic flow prediction problem as a part of time-series prediction problems, can be formulated as forecasting future values according to the historical data. In our framework, the goal is to predict future-step values and generate explanations based on historical values and external factors. We can represent this using Eq. (1):

$$X_{T:T+H}, \mathbb{I}_i = P_\theta(X_{T-H-1:T-1}, E_i) \quad (1)$$

where the function P_θ is the predictive model, that learns the mapping relationship between input and output. $X_{T-H-1:T-1}$ and $X_{T:T+H}$ represent continuous historical values and future predicted values with H steps. E_i is the external factor for data sample i , which includes date, holiday information, meteorological data and PoI data, etc. \mathbb{I}_i is the explanation corresponding to the predicted output $X_{T:T+H}$. In this work, we treat the language model (LLM) as the predictive model P_θ , and utilize the language tokenizer to transform the input and output data as a sequence of tokens, thus reframing the task of traffic flow prediction as a language modeling problem. Specifically, given the tokenized input sequence $S_{T-H-1:T-1}$ and tokenized external factors SE_i for data sample i , the tokenized output sequence $S_{T:T+H}$ is reconstructed autoregressively:

$$\hat{s}_{T+i} = \underset{s_{T+i}}{\operatorname{argmax}} P_\theta(s_{T+i} | S_{T:T+i-1}, S_{T-H-1:T-1}, SE_i) \quad (2)$$

where \hat{s}_{T+i} represents the i -th predicted next token in output sequence. Similarly, the explanation can also be generated following the output sequence in this way. By incorporating the next token prediction into the framework, we enhance the language model's ability to generate coherent and contextually relevant explanations alongside traffic flow predictions.

3.2. Prompt construction

This section demonstrates how to textualize multi-modal traffic data and construct instructions to inspire the reasoning and predicting abilities of LLMs. The instruction part of Fig. 2 showcases a meticulously crafted prompt template, designed to capture essential details consistently and comprehensively, including system prompts, spatial-temporal information, and external factors. Empirically, this structured format is tailored to convey diverse data modalities, boost the model's comprehension, and refine predictive accuracy.

- **System prompts.** To enhance the integration of LLMs into traffic prediction tasks, system prompts are structured with clear task descriptions and domain-specific knowledge, as shown in Fig. 3. Task

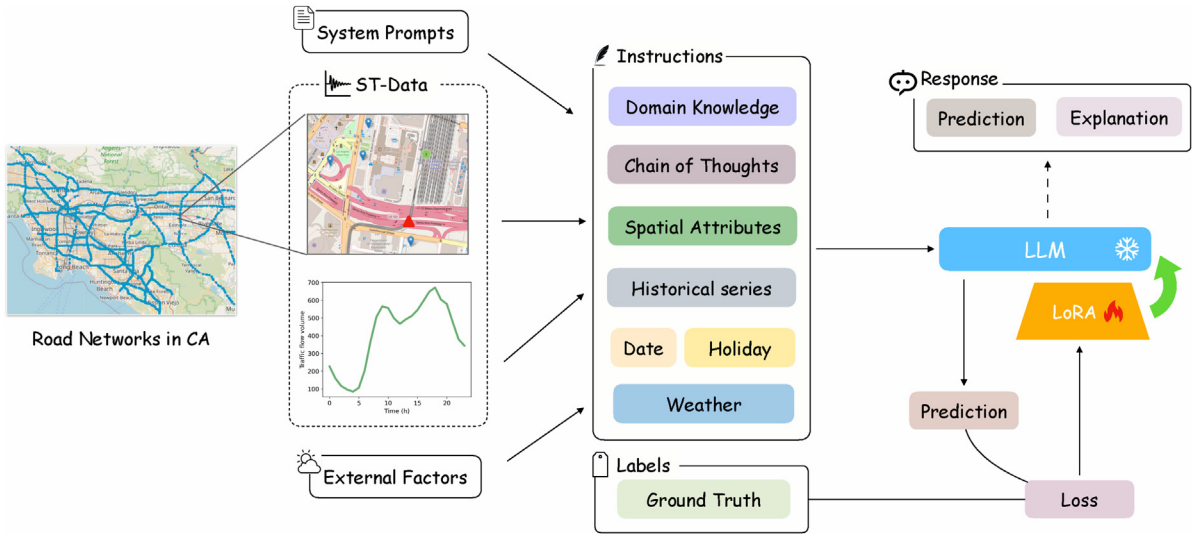


Fig. 2. Framework of xTP-LLM: Multi-modal traffic flow data is converted into text-based prompts, leveraging task settings and domain knowledge to enable a LLM to discern latent relationships across diverse inputs and traffic patterns. Through fine-tuning, the LLM gains the ability to predict future values and furnish pertinent explanations. This refined model adeptly captures regional, traffic patterns and dependencies on input factors, facilitating generalized predictions even in unseen datasets, thus enhancing its predictive capabilities.

settings explicitly outline the role of LLMs and leverage their pre-existing understanding of the traffic domain. Domain knowledge about traffic flow is embedded within the text description, incorporating few-shot examples to prompt LLMs to consider spatial factors, temporal fluctuation, and their interdependence. Spatial-specific traffic pattern cases are presented to assist models in establishing connections between geographical characteristics and traffic flow variations. For instance, residential areas may experience heightened traffic volume and increased travel activities during morning and evening rush hours.

- **Chain of thoughts.** Chain of thought (CoT) (Wei et al., 2022) prompts have been proven to improve LLMs' reasoning capabilities in intricate problems, which promote models to think from shallow to deep. Inspired by zero-shot CoT (Kojima et al., 2022), spatial-temporal CoT prompts in Fig. 4 are strategically crafted to enhance the LLMs' inference abilities. Initially, prompt questions guide LLMs to contemplate the spatial characteristics and potential traffic dynamics within a given area. Subsequently, LLMs assess whether the predicted time slot corresponds to special periods such as rush hours, weekdays, weekends, or holidays. After capturing spatial-temporal information, LLMs are directed to explore the profound connection between spatial data and the fluctuation of historical traffic flow, facilitating more precise predictions.
- **Spatial attributes.** Multi-modal information prompts serve as the cornerstone of our approach. In this context, spatial attributes are derived from nearby points of interest (PoIs). We preprocess PoI category data within different proximity ranges (1, 3, and 5 km), aligning with the locations of traffic volume sensors. To reduce redundant input and avoid the LLM's over-reliance on PoI information, we summarized the PoI distribution into descriptions of regional attributes, such as transportation hubs, commercial zones, and residential areas, effectively representing the key characteristics of each geographical area. This approach facilitates the effective representation of spatial factors, enabling the model to grasp the intricate interplay between various factors impacting traffic flow dynamics. Also, regional data comprises details such as the city, and road location. Our xTP-LLM framework leverages this multifaceted information to recognize and integrate spatial and temporal patterns across diverse regions and periods.
- **Historical time series.** Historical series are transcribed into textual descriptions with direct numerical representation. Data from the past

12 h are illustrated for each hour's time slot, enabling LLMs to accurately perceive the temporal change of traffic flow, compared with encoding temporal features. This approach empowers LLMs to generate insights, understanding, and explanations for emerging trends and patterns.

- **External factors.** External factors influencing traffic flow, such as dates, holidays, weather conditions, temperature, and vehicle visibility, are systematically considered. These diverse data points are uniformly transformed into textual information for comprehensive analysis.

A detailed prompt design can be found in Table A1 in Appendix A. Overall, this integration ensures a holistic understanding of the surrounding environment, historical trends, and external influences, enriching the model's contextual reasoning and enhancing its predictive capabilities.

3.3. Supervised fine-tuning

Supervised fine-tuning is a key technique for adapting LLMs to specific tasks (Ding et al., 2023; Ziegler et al., 2019). Initially, LLMs are trained on extensive, diverse datasets to develop broad linguistic knowledge. However, to excel in particular applications or domains, these models need additional training on task-specific data. In supervised fine-tuning, the pre-trained model is further trained on a dataset with labeled examples that are directly relevant to the target task. For instance, if the task is traffic prediction, the fine-tuning dataset would include historical traffic data paired with accurate forecasts. This process involves adjusting the model's parameters to improve its performance by optimizing a loss function that measures the discrepancy between the model's predictions and the true outputs. Fine-tuning enables the model to leverage its general knowledge while becoming more proficient in the specific domain, enhancing both accuracy and relevance.

During the fine-tuning phase, the traffic forecasting task is framed as a next-token generation task. The model predictions future traffic volumes autoregressively by optimizing the following objective function for the entire annotated dataset \mathbb{D} :

$$\mathbb{L}^{FT}(\mathbb{D}) = - \sum_{j=1}^n \log P(Y_j | Y_{1:j-1}; X_{1:n}) \quad (3)$$

where Y_j is the ground truth of the next token, $X_{1:n}$ is the sequence of input historical tokens, and $Y_{1:j-1}$ is the previously generated tokens. Through this process, the model learns to make more accurate predictions tailored to the specific task or domain it is being fine-tuned for. This approach maximizes the learning potential of the LLM, enabling it to develop specialized proficiency in the target domain.

Our xTP-LLM model is built upon the renowned open-source large language model, Llama2-7B-chat (Touvron et al., 2023), and is fine-tuned with LoRA technique (Hu et al., 2021). LoRA is a technique designed to efficiently adapt large pre-trained models to new tasks by introducing low-rank updates to the model's parameters. In LoRA, the adapted model's weight matrix W is expressed as $W = W_0 + \alpha \Delta W$, where W_0 represents the original pre-trained weights, α is the scaling factor that determines the magnitude of the low-rank update, and ΔW is the update applied during fine-tuning. Instead of updating the entire matrix ΔW directly, LoRA approximates it using the product of two lower-dimensional matrices: $\Delta W = B \times A$. Here, $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, with r being the rank of the approximation, chosen to be much smaller than the dimensions of W_0 (i.e., $r \ll \min(d, k)$). This low-rank factorization significantly reduces the number of parameters involved in the update from $d \times k$ to $r \times (d + k)$, leading to substantial computational and memory savings.

3.4. Explanation generation

The chat models in Llama2 are tailored to excel in understanding and generating text in conversational contexts. Therefore, by incorporating explanation requirements into the prompts, our model can not only generate prediction results but also provide explanations simultaneously. Initially, we add instructions for generating explanatory text directly within the input prompt. Although xTP-LLM was able to produce such explanations, they often lacked coherence with the prediction results. We attribute this issue to the model's insufficient alignment between explanations and predictions, as only the predictions are used in loss calculations during the fine-tuning phase. We addressed this misalignment by employing few-shot learning (Dai and Wang, 2020; Radford et al., 2019). By including a few carefully selected examples in the input prompt—examples that illustrate how the explanation should align with the traffic flow prediction—the LLM is able to learn this alignment dynamically during inference. To conduct few-shot learning in our context, we follow the steps below.

- **Create few-shot examples:** We selected 2–3 examples where the traffic flow prediction and the corresponding explanations were correctly aligned. These examples are generated through ChatGPT with ground truth predictions in input.
- **Incorporate into the input prompt:** These few-shot examples were then included directly in the input before the instruction fine-tuning prompt.
- **Model inference:** The LLM uses these examples as a guide during inference, enabling it to generate explanations that align well with the predicted sequences.

By structuring the few-shot examples in this way, we effectively leveraged the LLM's in-context learning capability to resolve the misalignment issue without requiring extensive new data or additional fine-tuning. This approach not only improved the coherence of the output explanations but also demonstrated the adaptability of LLMs in handling complex tasks with minimal examples.

4. Experimental settings and results

In this section, we will start by introducing our experimental setups, which include the dataset we used, the evaluation methods we employed, the baseline models we compared against, and the parameter settings we

used in the fine-tuning process. Subsequently, we will present the main results of our experiments, which involve comparing xTP-LLM with the baseline models, discussing how the performance of our model varied in the spatio-temporal domain, conducting ablation studies, and exploring the generalization capabilities of our model. Finally, we will discuss the explanations generated by xTP-LLM in corresponding to the traffic flow predictions. This discussion will elucidate how the robust reasoning capability inherent in LLMs enhances traffic forecasting tasks.

4.1. Dataset description

Our experiments were conducted based on our proposed multi-modal traffic prediction dataset, named CATraffic. This dataset consists of traffic volume data from various regions in California, as well as meteorological information, nearby point of interests (PoIs) data, and holiday information. The traffic volume data is sourced from the LargeST dataset (Liu et al., 2024), which comprises five years (2017–2021) of traffic flow data in California, encompassing 8600 traffic sensors sampled at a 15-min interval. We constructed CATraffic by selecting a subset of the LargeST dataset, focusing on 1,000 sensors from the greater los angeles (GLA) and greater bay area (GBA). Our dataset spans two years from January 1, 2018 to December 30, 2019, with data sampled hourly. The choice of a 1-h granularity represents a strategic compromise that balances the need for temporal resolution with the computational constraints as well as the input context length constraints of LLMs. Our PoIs data is obtained through OpenStreetMap² with the help of the Overpass Turbo API.³ For the meteorological data, we collected information from the National Oceanic and Atmospheric Administration (NOAA⁴). These include factors such as reported weather events, temperature, and visibility, which are considered as they have direct impacts on traffic patterns.

During the sensor selection process, for a diverse representation of traffic patterns, we clustered all 8,600 sensors into 1,000 categories based on their nearby PoIs features. These PoIs features are formalized as $\mathbb{F}_{poi} = [(f_1^e, f_2^e, \dots, f_n^e), (f_1^w, f_2^w, \dots, f_n^w), (f_1^n, f_2^n, \dots, f_n^n), (f_1^s, f_2^s, \dots, f_n^s)]$, where f^e , f^w , f^n , and f^s represent the PoIs features in the four directions (East, West, North, and South), respectively. We selected the top- n PoI categories, normalized their counts in a range of 5 km, and expressed them as f_i for i -th category. In our experiments, we set n to 20 for convenience. These PoIs feature vectors were then used as input for clustering via the K-means algorithm, which grouped the sensors into 1,000 clusters. For each cluster, we selected one representative sensor, resulting in a final set of 1,000 sensors. This approach effectively reduces the dataset to a manageable size while preserving a diverse representation of the various spatial and feature characteristics present in the original dataset.

We split the collected data into two parts: data from 2018 was used as the training set, while data from 2019 was reserved for model validation. We believe an entire year of data for testing provides a more comprehensive assessment of the model's performance across a full spectrum of temporal variations, including different weather conditions and holiday periods. All experiments were configured to predict traffic flows for the next 12 h based on historical 12-h traffic flow data. During the data preprocessing stage, we filtered out samples with zero values for 24 consecutive hours, which likely resulted from malfunctioning sensors. Additionally, to assess the model's generalization capability, we created a zero-shot dataset derived from LargeST (Liu et al., 2024). This dataset comprises data from 100 sensors (also selected through clustering) in San Diego (SD), covering the period from November 1, 2019 to December 31, 2019. These data were not utilized in the model fine-tuning process and served to the evaluation of the model's generalization performance.

² OpenStreetMap: <https://openmaptiles.org>.

³ Overpass Turbo: <https://overpass-turbo.eu>.

⁴ NOAA: <https://www.ncei.noaa.gov>.

4.2. Evaluation metrics

In time series forecasting tasks, researchers commonly employ root mean square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE), to evaluate the accuracy of forecasting results. These metrics are defined as Eqs. (4)–(6):

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (4)$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (5)$$

$$\text{MAPE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100\% \quad (6)$$

where y_i represents the ground truth value of the i -th data point, \hat{y}_i denotes the corresponding prediction value, and n stands for the total number of samples. RMSE is advantageous for emphasizing larger errors due to its square term, while MAE provides a straightforward interpretation by averaging absolute errors, treating all errors equally. On the other hand, MAPE measures the average percentage difference between predictions and ground truths, offering interpretability in terms of relative accuracy but being sensitive to zero values in the denominator. Researchers typically employ a combination of these metrics to comprehensively assess model performance.

4.3. Baseline models

We extensively compared our proposed xTP-LLM with 9 advanced baseline models. Among these, LSTM (Hochreiter, 1997) stands as a temporal-only deep model based on recurrent neural networks (RNNs), disregarding spatial correlations. Additionally, we select DCRNN (Li et al., 2017) and AGCRN (Bai et al., 2020) as the representation of RNNs-based methods. We also choose TCN-based methods such as STGCN (Yu et al., 2018) and GWNEN (Wu et al., 2019), along with attention-based methods ASTGCN (Guo et al., 2019) and STTN (Xu et al., 2020). These models were proposed between 2018 and 2020, reflecting the prevalent research direction in time series forecasting during those years. Furthermore, we integrated three representative methods from recent years, including STGODE (Fang et al., 2021) and DSTAGNN (Lan et al., 2022). STGODE adeptly utilizes neural ordinary differential equations to capture the continuous changes of traffic signals, while

DSTAGNN is specifically designed to capture the dynamic correlations among traffic sensors.

4.4. Experiment settings

Our model is fine-tuned based on the well-known open-source LLM, Llama2 (Touvron et al., 2023), specifically utilizing the chat version with a size of 7 B. We load the base model in 8 bits for fine-tuning, with training parameters including a batch size of 8, a learning rate of 5e-4, a warm-up step of 400, gradient accumulation steps of 8, and a training epoch for 2. The LoRA parameters are configured with a rank of 16 and an alpha value of 32. During the inference phase, a temperature of 0.95 was applied.

4.5. Overall performance

We report the performance comparison results between xTP-LLM and the baseline models in Table 1. All models are trained and evaluated based on our CATraffic dataset, with the same dataset settings. The task is to utilize traffic flow data from the historical 12 h to forecast future traffic flows in the next 12 h. We represent the results of horizons 3, 6, 9, and 12, as well as the average performance over all 12 steps in the table. The results demonstrate that the overall performance of our proposed xTP-LLM exceeds that of the baseline models by a large margin, especially in the MAE and MAPE. For example, In terms of average performance of all 12 horizons, our model outperforms the best two baseline models, GWNEN (Wu et al., 2019) by 18.37% in MAE, and STGCN (Yu et al., 2018) by 34.00% in MAPE, which shows the impressive capability in traffic flow forecasting of our model.

We further demonstrate the results with different prediction horizons of the baseline models and xTP-LLM, depicted in Fig. 6. From left to right, RMSE, MAE, and MAPE of compared models at different prediction steps are shown. The results yielded the following observations (Figs. 3–5).

- As the prediction horizon increases, performance generally declines across all models, as longer-term forecasts inherently entail greater uncertainty and complexity. However, several models exhibit improved performance in longer-term forecasting, such as LSTM, ASTGCN, and DCRNN. This phenomenon may be attributed to these models' ability to capture and leverage the periodicity within the data, allowing them to make more accurate predictions over extended time horizons.
- Our proposed model consistently outperforms the comparison methods at each time step, showing significant advantages in both

Table 1

Performance comparison between our proposed xTP-LLM and baseline models on our CATraffic dataset. The horizon for traffic flow prediction is 12, and we show the results in Steps 3, 6, 9, and 12. The last three rows are the average performance for all 12 steps. We highlight the best results in bold and the second best in underline.

Step	Metrics	LSTM	DCRNN	STGCN	ASTGCN	GWNEN	AGCRN	STTN	STGODE	DSTAGNN	xTP-LLM
3	RMSE	62.65	61.12	43.43	66.60	<u>42.59</u>	43.92	43.41	52.30	51.79	40.34
	MAE	40.00	37.32	25.83	44.47	<u>25.56</u>	25.67	26.00	29.53	34.06	20.20
	MAPE (%)	24.75	24.52	16.09	28.36	15.96	<u>15.60</u>	16.22	19.69	23.90	9.94
6	RMSE	74.69	74.21	46.64	75.07	<u>45.91</u>	48.92	51.63	55.79	53.44	44.53
	MAE	48.10	45.83	<u>27.94</u>	49.47	28.27	28.94	30.79	32.97	33.72	22.70
	MAPE (%)	33.98	31.58	<u>15.75</u>	31.10	18.00	17.31	18.99	21.62	21.27	11.39
9	RMSE	77.61	78.22	50.04	81.76	<u>46.61</u>	52.06	56.19	59.97	56.21	45.89
	MAE	49.60	47.97	30.16	53.43	<u>28.70</u>	30.74	33.11	35.23	36.45	23.89
	MAPE (%)	39.26	32.22	<u>17.52</u>	35.35	19.88	18.38	20.69	22.44	24.00	12.09
12	RMSE	69.86	71.79	54.91	69.75	<u>49.79</u>	56.13	61.38	60.47	60.33	47.82
	MAE	42.05	43.08	33.20	43.40	<u>30.47</u>	32.70	38.27	36.22	37.91	24.99
	MAPE (%)	25.99	25.05	<u>19.27</u>	28.03	20.46	19.50	26.09	23.16	24.23	12.30
Avg.	RMSE	68.14	67.89	46.69	70.48	<u>43.97</u>	47.97	50.03	54.41	52.82	42.81
	MAE	43.23	41.45	28.02	46.17	<u>26.84</u>	28.22	30.04	31.95	34.02	21.91
	MAPE (%)	30.21	27.14	<u>16.68</u>	29.99	17.44	16.94	19.04	20.72	22.45	11.01

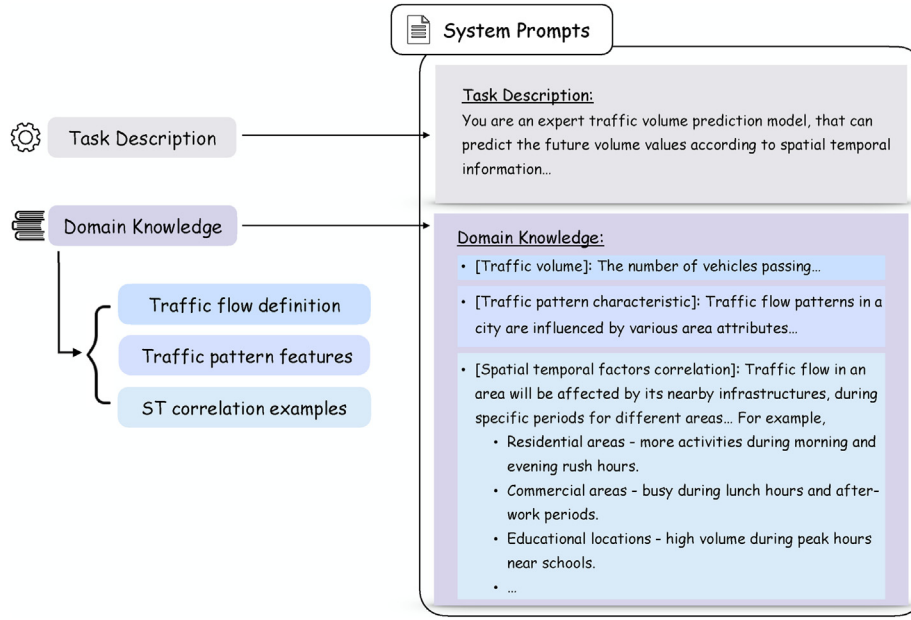


Fig. 3. System prompts. They encompass task descriptions of traffic flow prediction and domain knowledge, including definitions, traffic pattern features, and examples of regional characteristics and traffic-pattern correlations. Equipped with these foundational settings and backgrounds, language models can tailor their knowledge and capabilities to the specific domain.

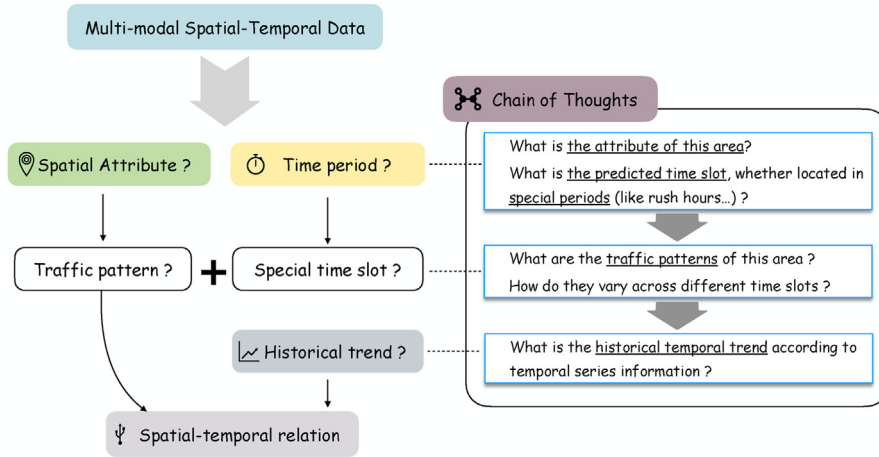


Fig. 4. Chain of thoughts. Chain of thought (CoT) prompts encourage LLMs to extract relevant factors from provided information, prompting deeper consideration of their potential causal relationships with related knowledge. This process facilitates a comprehensive analysis and interpretation of the data.

short-term and long-term traffic flow forecasting. This indicates the robustness of our model in various prediction horizons.

- Our proposed model shows a more significant advantage in MAE and MAPE than RMSE, which may be attributed to the amplification effect of RMSE on outliers, leading to inaccurate assessments. Compared to RMSE, MAE and MAPE are less sensitive to extreme errors because they measure the mean absolute error and the percentage error, respectively. The excellent performance of our model on MAE and MAPE suggests that it is effective in mitigating the effects of outliers and provides more accurate and stable predictions, especially in cases where extreme values may occasionally occur.

These findings underscore the effectiveness of our method in capturing complex temporal patterns in traffic flow data, leading to more accurate and reliable predictions.

4.6. Spatial and temporal homogeneity

Evaluating spatial and temporal homogeneity in traffic prediction is vital for evaluating model performance, generalizability, and robustness. It ensures that traffic flow prediction models can effectively adapt to diverse real-world conditions. Therefore, in this section, we will thoroughly analyze the performance of our proposed model in terms of spatio-temporal consistency.

Spatial homogeneity helps to assess the ability of different models to learn traffic patterns at different spatial locations. To evaluate this kind of capability, we tested our xTP-LLM as well as four baseline models (LSTM, AGCRN, STGCN, and GWNENET) at different locations with varying urban characteristics in the Greater Los Angeles Area (GLA). The result, depicted in Fig. 7a, illustrates MAPE values for four prediction horizons of 3, 6, 9, and 12 h, with darker colors indicating poorer performance.

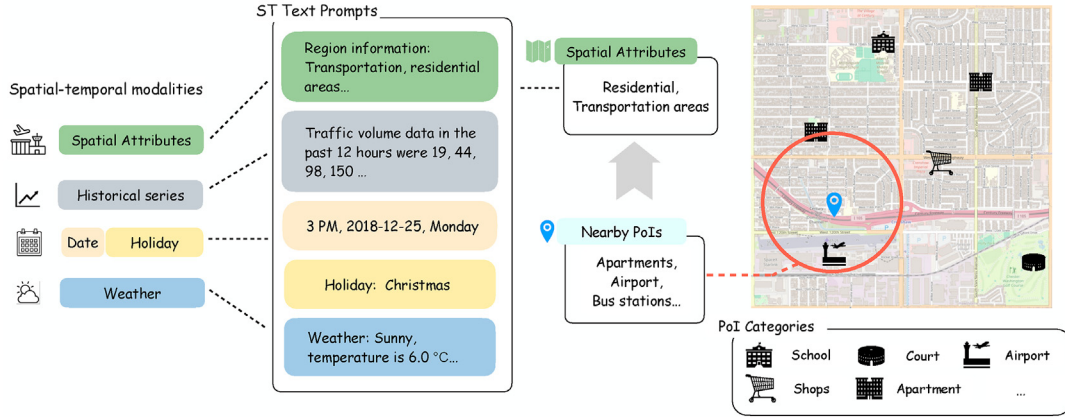


Fig. 5. Multi-modal spatial-temporal (ST) text prompts. Spatial-temporal data covers spatial attributes and time series of historical traffic flow. To consider the spatial aspect of a region, nearby PoIs within a defined range are utilized to characterize its location features. For instance, consider a sensor positioned along the Century Freeway, surrounded by residential complexes and an airport, as the blue symbol illustrated in the right part of the figure. This location can be classified as a blend of residential and transportation zones, reflecting the area's geographical attributes. To incorporate historical flow data and external factors, we represent these datasets textually as the rest components.

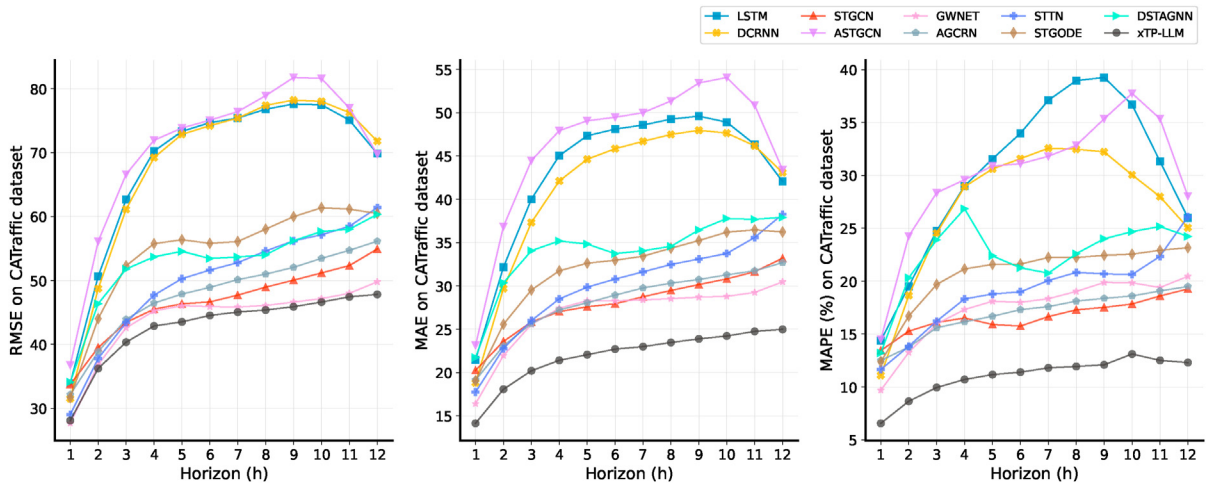


Fig. 6. Time step-based prediction comparison with baseline models and our xTP-LLM. All experiments are conducted based on our CATraffic dataset. Results show that the performance of all models decreases with increasing prediction time, whereas our model significantly outperforms the others in all prediction steps.

Overall, our xTP-LLM model exhibits relatively consistent prediction performance across different locations, effectively capturing traffic flow trends irrespective of varied spatial characteristics. In contrast, other models demonstrate inferior homogeneity, particularly in areas featuring complex road network intersections and intricate facility distributions. These findings underscore the robustness and adaptability of our proposed xTP-LLM model in effectively learning and predicting traffic patterns across diverse spatial contexts, thus highlighting its potential for real-world application in urban traffic management and planning scenarios.

In the analysis of temporal homogeneity, we evaluated models on a subset of the test dataset (traffic data in November 2019) and reported the daily averaged MAPE values. Illustrated in Fig. 7b as calendar heat maps, the results showcase the comparative performance across the month. Our model consistently exhibits lower daily average MAPE, demonstrating robustness in capturing temporal nuances of traffic flow patterns. Notably, it consistently outperforms others throughout the whole month, emphasizing its superior capability in handling diverse temporal dynamics inherent in real-world traffic scenarios.

4.7. Ablation studies

4.7.1. Impact of input components on model performance

To assess the effect of varying prompt input on model performance, we conducted a thorough ablation analysis. This approach enabled us to systematically evaluate the contribution of each component to the overall performance of the model. Specifically, we investigated the effects of including date, weather, PoIs information, domain knowledge, and CoT prompting on model performance, and the results presented in Table 2. We observed that date, weather, and points of interest (PoIs) each independently enhance model performance, as evidenced by comparisons between input settings B, C, D, and A. Notably, date information yields the most significant improvement, with a 37.62% reduction in MAE, underscoring its critical role in enhancing predictive accuracy. Weather and PoIs also contribute positively, with MAE reductions of 4.53% and 14.58%, respectively, though their impact is less pronounced compared to date information. Moreover, The interaction between date, weather, and PoIs significantly enhances model performance, as demonstrated by input settings E, F, and G. Any combination of these

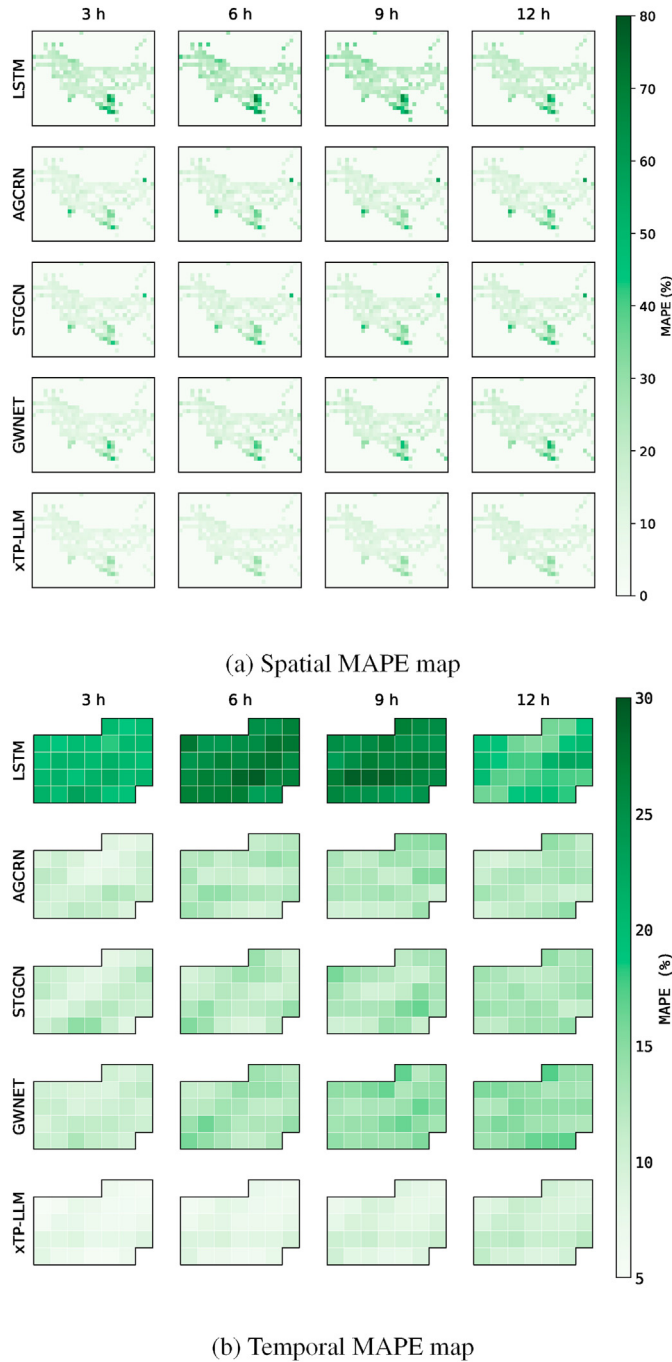


Fig. 7. Comparison of the spatial and temporal prediction error distributions among different models: (a) presents the distribution of MAPE for our proposed xTP-LLM and four baseline models in various locations across the Greater Los Angeles (GLA), considering four different prediction horizons; (b) illustrates the distribution of daily average MAPE over the whole November in 2019, with four different prediction horizons.

three components can further improve the model's predictive performance, illustrating their synergistic effect. For example, setting F results in an MAE of 24.31, shows a 21.25% improvement compared to setting B, which shows the synergistic effect of date and PoIs on model performance. Similarly, the including weather and PoIs (setting G) shows a 14.41% improvement in MAE compared to setting C. These results underscore how the integration of date, weather, and PoIs interactively boosts model performance, with each component amplifying the benefits provided by the others.

Additionally, an interesting observation emerges when comparing input settings H, I, J, and K. We found that adding domain knowledge or CoT prompting alone did not significantly improve model performance and sometimes even degraded it (as seen in settings I and J compared to H). However, when both were combined (setting K), significant performance improvements were observed, with a 4.57% reduction in MAE. This finding contrasts with our initial expectations and suggests that domain knowledge and CoT prompting work synergistically. We attribute this to CoT prompting effectively guiding the model in utilizing domain knowledge, extracting relevant information for the prediction task. Conversely, CoT prompting also benefits from the inclusion of domain knowledge, enabling it to achieve its full potential.

4.7.2. Impact of different horizons on model performance

To explore the impact of different horizons on model performance, we conducted additional experiments with historical horizons of 4 and 8 h. The evaluation results in the test dataset are presented in Fig. 8. We observed that as the historical horizon decreases, the model's prediction error, measured by MAE, increases significantly. Specifically, the 12-h horizon consistently yielded the lowest MAE values, indicating superior performance compared to shorter horizons. In contrast, the 4-h horizon exhibited the highest MAE, suggesting that a shorter historical horizon substantially impairs the model's ability to accurately predict traffic flow. This is because a longer historical horizon provides more information, allowing the model to more accurately identify traffic flow patterns. Moreover, examining different prediction steps provides a direct view of how varying prediction horizons impact the model's predictions. As expected, the performance deteriorates as the prediction horizon increases in all three historical horizons. Two primary factors contribute to this phenomenon. Firstly, the increasing temporal distance introduces greater uncertainty and variability into the forecasting process, thereby exacerbating the challenge of accurately predicting future events. Secondly, the iterative nature of large language model (function as a next-token predictor) means that any error in early prediction is propagated and magnified in subsequent steps, leading to a cumulative effect that significantly impairs the model's overall performance on the long range.

4.8. Generalization studies

4.8.1. Robustness analysis on the test dataset

To ensure the robustness of our xTP-LLM model, we conducted a comprehensive analysis of its performance on the test dataset, focusing on various time periods such as peak hours, off-peak hours, weekdays, weekends, and different weather scenarios. The results, detailed in Tables 3 and 4, demonstrate the model's resilience and consistency across these diverse scenarios.

Table 3 illustrates the model's performance across different time periods, including peak hours, off-peak hours, weekdays, and weekends. The average MAE values range from 21.25 to 22.57, indicating that the model maintains a stable level of accuracy regardless of temporal variations. This consistency underscores the robustness of the xTP-LLM model in handling different time-related scenarios. Table 4 presents the model's performance under various weather conditions, such as sunny, rainy, foggy, and stormy situations. Despite the imbalanced representation in the dataset, where sunny conditions dominate and extreme weather like sleet and snow are underrepresented, the model delivers reliable predictions. The MAE values remain relatively stable across all weather conditions, demonstrating the model's robustness even in less frequent and challenging scenarios. Notably, Table 4 also shows that even when certain weather conditions, such as storm and snow, are absent from the training dataset, the model still achieves good results on the corresponding test data. This highlights the strong generalization ability of the xTP-LLM model.

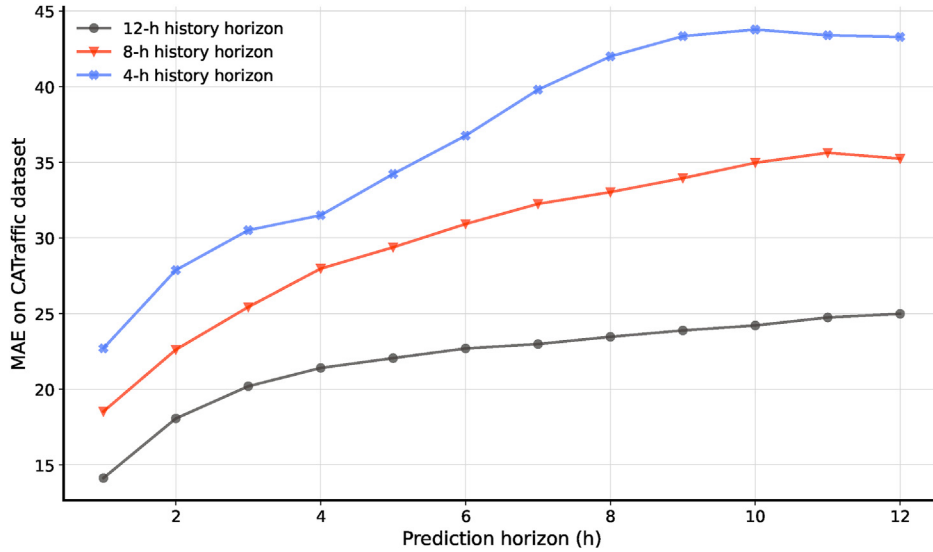
4.8.2. Zero-shot testing analysis

LLMs are well known for their excellent zero-shot capabilities. In this

Table 2

Ablation study results showing the impact of different input settings (date, weather, PoIs, domain knowledge, and CoT prompting) on model performance.

Input setting	Date	Weather	PoIs	Domain knowledge	CoT prompting	RMSE	MAE	MAPE (%)
A						89.17	49.49	30.08
B	✓					56.43	30.87	17.05
C		✓				82.43	47.25	28.52
D			✓			75.14	42.28	26.39
E	✓	✓				51.51	28.35	14.98
F	✓		✓			48.50	24.31	12.88
G		✓	✓			69.97	40.44	24.90
H	✓	✓	✓			46.20	22.36	12.05
I	✓	✓	✓	✓		44.22	22.64	11.23
J	✓	✓	✓		✓	45.49	23.55	12.38
K	✓	✓	✓	✓	✓	42.81	21.91	11.01

**Fig. 8.** Model performances with different history horizons. We show that as the history horizon shortens, the MAE increases significantly, indicating that the model performs better with longer historical data.**Table 3**

Model's performance across different time periods, including peak and off-peak hours, weekdays, and weekends.

Different period	Peak hours (7 a.m.–9 a.m., 4 p.m.–6 p.m.)	Off-peak hours (10 a.m.–3 p.m., 7 p.m.–7 a.m.)	Weekday	Weekend
Avg. MAE	21.25	22.57	22.05	21.56

section, we delve into the generalization of our proposed xTP-LLM and compare it with other LLMs. We conducted zero-shot experiments using two datasets: the first is our proposed CATraffic-based zero-shot dataset, and the other is the taxiBJ dataset (Zhang et al., 2017). The TaxiBJ dataset comprises taxicab GPS data and meteorology data in Beijing from four time intervals within 2013–2016, focusing on the inflow/outflow prediction task, which is different from the traffic volume prediction task in CATraffic. To ensure a fair comparison, we reorganized the taxiBJ dataset into the same format as our CATraffic dataset. The overall results

are presented in Table 5, which indicate that our proposed model exhibits superior performance across all three tasks compared to the original Llama2 series models, as well as GPT-3.5-turbo and GPT-4. For the CATraffic zero-shot dataset, our xTP-LLM model exhibits notable improvements over the best-performing comparative model (GPT-4) by 61.53%, 71.22%, and 85.11% in RMSE, MAE, and MAPE, respectively. On the TaxiBJ dataset, our model achieves substantial performance enhancements in both inflow and outflow tasks. Compared to GPT-4, xTP-LLM shows improvements of 40.32% in RMSE, 44.01% in MAE and 80.75% in MAPE in inflow prediction task, and improvements of 44.89% in RMSE, 47.97% in MAE and 77.26% in MAPE in outflow prediction task. These results demonstrate superior zero-shot learning capability of our xTP-LLM on the cross-domain datasets, even with a smaller scale of parameters (7 billion for xTP-LLM, compared to 70 billion for Llama2-70B-chat, 173 billion for GPT-3.5, and approximately 1.76 trillion for GPT-4). While LLMs like GPT-4 and Llama2-70B-chat exhibit strong generalization capabilities in general domains such as conversation and translation, they generally do not perform well in the

Table 4

Distribution of various weather conditions in the training and test datasets, as well as the model's performance across these different weather conditions.

Weather	Sunny	Rain	Foggy	Thunderstorm	Sleet	Storm	Snow
Percentage on training dataset	83.300%	11.600%	4.700%	0.430%	0.016%	0.000%	0.000%
Percentage on test dataset	76.400%	19.800%	3.000%	0.740%	0.000%	0.028%	0.011%
Avg. MAE	21.88	21.42	20.05	22.11	—	20.17	19.93

Table 5

Zero-shot performance comparison of our proposed xTP-LLM with other LLMs without fine-tuning. We conducted tests on both the CATraffic zero-shot dataset and the TaxiBJ dataset to compare the zero-shot capabilities of different models. The results show that our model outperforms other LLMs significantly on both unseen datasets.

Model	CATraffic zero-shot			TaxiBJ (inflow)			TaxiBJ (outflow)		
	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)
Llama2-7B-chat	273.02	226.64	126.00	255.00	151.57	172.60	218.44	135.39	151.96
Llama2-13B-chat	234.73	187.37	146.08	230.55	129.39	184.90	234.14	144.24	143.47
Llama2-70B-chat	222.02	168.04	122.72	205.67	123.46	136.83	209.05	125.39	139.89
GPT-3.5-turbo	168.93	135.34	74.49	150.71	95.08	196.00	150.91	94.77	195.31
GPT-4	121.02	103.32	61.85	91.86	57.05	165.40	99.05	59.50	137.21
xTP-LLM	46.56	29.73	9.21	54.82	31.94	31.84	54.59	30.96	31.20

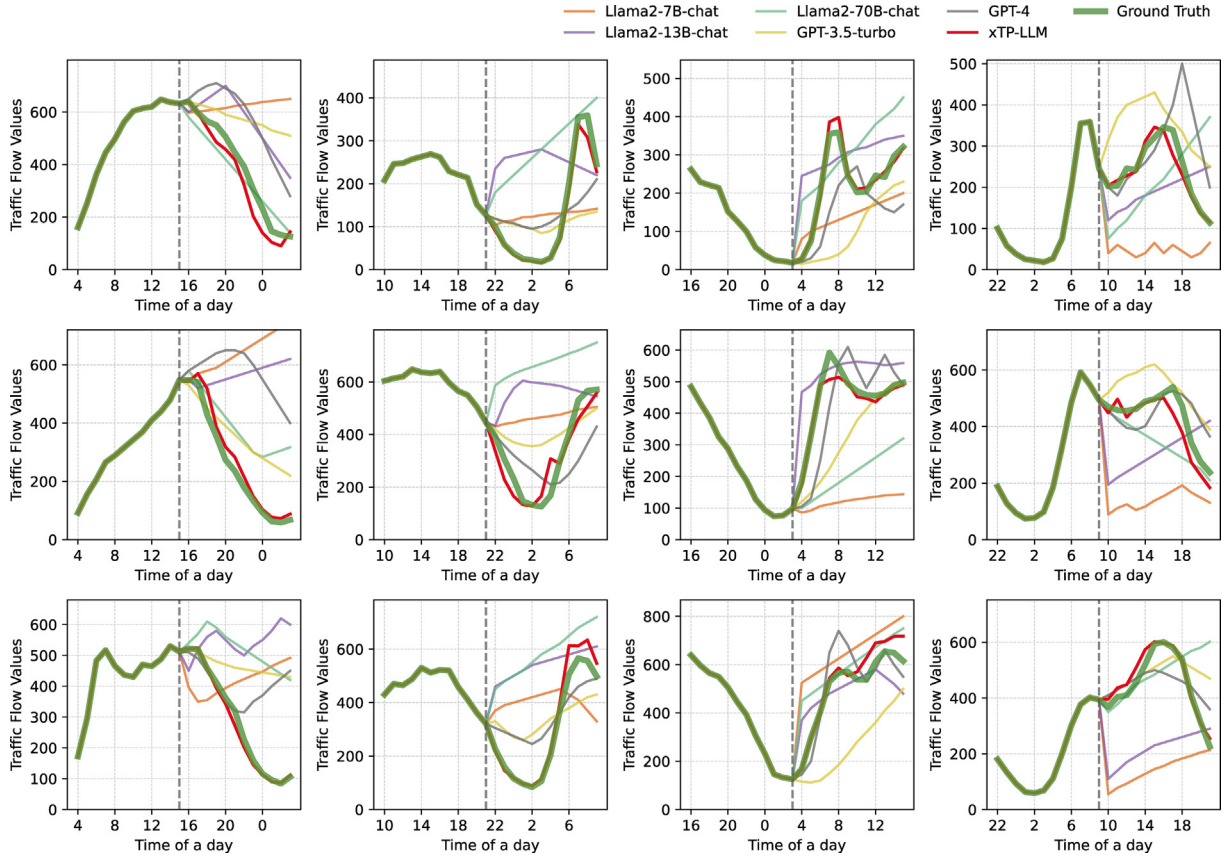


Fig. 9. Visualization of traffic flow prediction results on CATraffic zero-shot dataset. We randomly select test samples of three sensors with four time periods and visualized the prediction results of our proposed xTP-LLM, Llama2-chat models with three different sizes, GPT-3.5-turbo and GPT-4. The results show that our fine-tuned model performs significantly better than other LLMs.

traffic prediction domain when applied in a zero-shot learning setting.

We also provide visualization results in Fig. 9, showing 12 test samples from three different sensors in four different time periods. The results highlight that our model not only captures the traffic trends of new scenarios effectively but also delivers accurate prediction values. In contrast, Llama2 series of LLMs struggle to capture the dynamic pattern of traffic flow; while GPT-3.5-turbo and GPT-4 demonstrate the ability to describe some general trends in traffic, they perform inadequately in capturing nuances of variation, hindering their ability to provide accurate predictions. We attribute these discrepancies to our systematic prompt design and effective model fine-tuning. Prompt inputs effectively describe the task context and align information from different modalities into a unified representation, while fine-tuning infuses crucial expert knowledge into the model, thereby substantially enhancing its performance in traffic prediction. Our experiments further validate that instruction fine-tuning allows smaller models to excel in specific domains, often outperforming larger models.

4.9. Case studies for explainable predictions

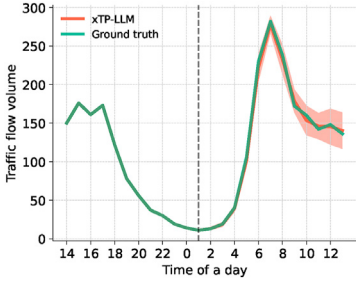
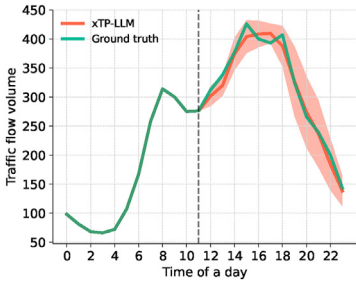
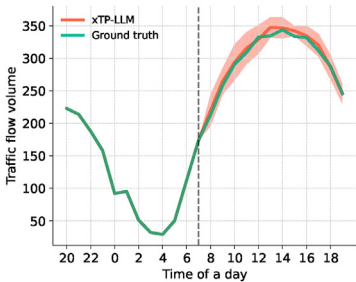
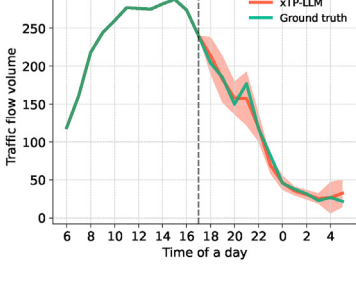
4.9.1. Result-explainable studies

To illustrate the reliability and accountability of our model's predictions, we report four cases with different prediction time slots and external factors in Table 6. In each example, the ground truths and predictions are displayed in the figure on the left (light red areas indicate the 95% confidence interval), with the corresponding explanatory texts on the right. The original result explanations are too long, so we used chatGPT to summarise them into a brief paragraph for presentation purposes. A complete output can be found in Table B1 in Appendix B (corresponding to Example (d) in Table 6).

Our proposed xTP-LLM consistently delivers accurate traffic flow predictions across various time periods, including weekdays, weekends, and holidays like Christmas. Notably, our model not only provides precise forecasts but also offers insightful explanations, considering a wide range of external factors. For instance, in Example (a) in Table 6, our

Table 6

Explanation examples. Light red shows the 95% confidence interval, and yellow highlights traffic trends.

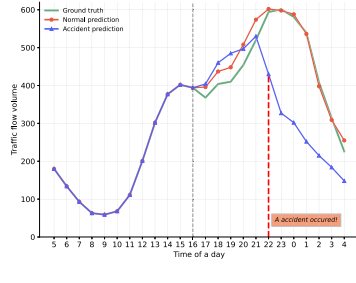
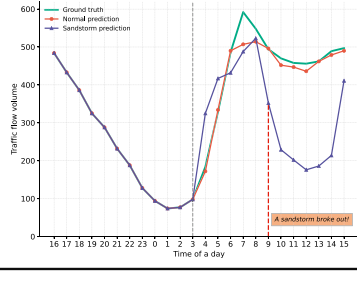
Prediction	Explanation
	(a): Traffic volume is expected to be low before dawn, then it is expected to increase gradually throughout the day. Since today is Wednesday, weekday patterns suggest higher traffic during rush hours and lower traffic during non-peak times. Nearby transportation areas and commercial districts will contribute to an increase in traffic volume during certain periods. Sunny weather can also contribute to increased traffic.
	(b): Traffic volume is expected to increase and then decrease in the next 12 h. Sunny weather with a comfortable temperature will attract more day travelers. The onset of Sunday afternoon rush hour will further contribute to increased traffic. With no known holidays or events, historical data suggest that traffic volume will resemble previous Sundays, with a slight increase during the afternoon rush hour followed by a gradual decline at night.
	(c): Traffic volume is expected to steady increase in the rainy morning and then decrease in the afternoon. Activities at the morning rush hour in residential and commercial areas are likely to contribute to the initial rise. The rain might cause congestion, leading to a peak in volumes around midday. As the day progresses, the volume is expected to gradually decrease as people return home.
	(d): Traffic volume is expected to decrease significantly over the next 12 h considering a rainy Christmas Day. As it is a holiday, commuter and commercial activities are minimal, leading to reduced traffic. Additionally, the rainy weather may further deter people from traveling or driving, contributing to the decline in traffic volumes. The focus on indoor celebrations and festivities is likely to keep the roads quiet at night.

model predicts an initial increase followed by a decrease in traffic flow, aligning with typical weekday patterns. Moreover, in Example (c) in Table 6, our model identifies rainy weather as an impact factor for traffic congestion, leading to a delayed peak traffic flow. Furthermore, Example (d) in Table 6 highlights our model's ability to factor in holiday impacts on traffic flow.

While our xTP-LLM excels in leveraging diverse external data for robust predictions, it faces challenges in capturing nuanced flow fluctuations, such as those occurring between 14 and 18 h in Example (b) and 20–22 h in Example (d) in Table 6. These complexities may stem from the

Table 7

Examples of what-if analysis.

Prediction	Explanation
	Sudden accident: ... Important! A serious traffic accident occurred on this road at 10 p.m.! ... Explanation: ... The traffic volume in the next 12h will increase due to the post-accident recovery. The accident had a significant impact on traffic flow causing delays and changes in traffic patterns. After the accident, the number of cars decreased but still remained at a relatively high level.
	Sudden weather event: ... Important! A severe sandstorm broke out at 9 a.m.! ... Explanation: ... The sandstorm causes a sharp drop in traffic after 9 a.m., with volumes significantly lower than normal due to hazardous conditions, recovering slightly in the afternoon as the storm eases ...

dynamic nature of human activities and the intricacies of transportation systems. Nevertheless, although not good at capturing the full details of the traffic flow changes, our xTP-LLM provides textual rationales for its forecasts, which shows a high transparency and interpretability.

4.9.2. What-if analysis

To investigate the generalization and conditional reasoning abilities of the framework, this section examines how prediction results change when influenced by external factors, such as traffic accidents and inclement weather. For the situational analysis, we consider two hypothetical scenarios: Case 1 involves a nearby traffic accident occurring at 10 p.m., and Case 2 involves a severe sandstorm breaking out at 9 a.m. These conditional descriptions are incorporated into the input prompt, while the other sections of the prompt remain consistent with the original prediction tasks. We then compare the normal predictions, what-if predictions, and the ground truth, as shown in Table 7.

The traffic system is highly dynamic and complex, interacting with human factors and real-time external conditions. From the experimental results, we observe that when a traffic accident occurs at a specific moment, the nearby traffic volume decreases due to the congestion caused by the accident, dropping from 530 to 430. The model can interact with real-time incidents, avoiding prediction delays and incorrect analysis. Similarly, when severe weather events like a sandstorm occur, the reduced visibility and road passability lead to a sharp decline in the number of vehicles passing through the area, from 532 to 352. However, over time, traffic flow gradually recovers to normal levels in the region. This indicates that our model provides timely feedback to sudden weather changes, and its inherent reasoning and analytical capabilities help make more reasonable predictions.

5. Conclusions and future work

In conclusion, our research introduces xTP-LLM, a novel traffic prediction model designed for both accuracy and interpretability. By incorporating multi-modal inputs and employing language-based representations, xTP-LLM achieves competitive performance compared to state-of-the-art models while offering insightful explanations into its

predictions. xTP-LLM's language-based framework, coupled with spatial-temporal alignment instructions, provides a transparent and adaptable approach suitable for various urban prediction tasks. In general, our work contributes to the advancement of effective and reliable traffic prediction methods, essential for informed decision-making in urban transportation planning and management.

In the future, we aim to delve into methods that enable LLMs to harness spatial information more effectively and to grasp how different sensors are related spatially. This will help models to make better predictions by considering data from nearby sensors. Moreover, further external factors can be considered, such as traffic accidents, human activities, and big events, contributing to more accurate predictions. Additionally, exploring the development of LLM systems tailored for urban brains is also a very interesting but challenging topic. This entails integrating city-level data into LLMs to tackle various downstream tasks like urban planning, traffic management, and pollution control, etc. Achieving these involves the challenge of enabling LLMs to efficiently utilize city-level multi-modal data, alongside the need for substantial computational resources and exceptional engineering capabilities.

CRediT authorship contribution statement

Xusen Guo: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Qiming Zhang:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Junyue Jiang:** Validation, Investigation, Data curation. **Mingxing Peng:** Methodology, Formal analysis, Conceptualization. **Meixin Zhu:** Writing – review & editing,

Supervision, Resources, Project administration, Funding acquisition, Formal analysis, Conceptualization. **Hao Frank Yang:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Replication and data sharing

The training and evaluation code can be accessed at <https://github.com/Guoxs/xTP-LLM>. A portion of the CATraffic dataset, covering traffic flow data for 100 traffic nodes over three months, is available at Google Cloud, and the download link can be found at <https://github.com/Guoxs/xTP-LLM>.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This study is supported by the National Natural Science Foundation of China (No. 52302379), Guangdong Provincial Natural Science Foundation-General Project (No. 2024A1515011790), Guangzhou Basic and Applied Basic Research Projects (Nos. 2023A03J0106 and 2024A04J4290), Guangdong Province General Universities Youth Innovative Talents Project (No. 2023KQNCX100), Guangzhou Municipal Science and Technology Project (No. 2023A03J0011), and Nansha District Key R&D Project (No. 2023ZD006).

Appendix A Prompt design

The prompt for xTP-LLM in traffic flow prediction is carefully designed, as shown in Table A1. The complete prompt contains both a system prompt and a user input prompt. The system prompt sets the role of the LLM, and it also contains a context knowledge part that provides additional background information in traffic flow prediction, as well as the CoT prompt that guides LLMs through the human reasoning process. The system prompt remains the same throughout the dataset, and the change parts are the user input prompt as well as the corresponding ground truth.

Table A1
Details of our designed prompt in xTP-LLM for traffic flow prediction task, including system prompt, input prompt, as well as the ground truth.

<p>System prompt:</p> <p>You are an expert traffic volume prediction model, that can predict the future volume values according to spatial temporal information. We want you to perform the traffic volume prediction task, considering the nearby environment and historical traffic volume data.</p> <p>Context knowledge:</p> <p>Context knowledge you could consider:</p> <ul style="list-style-type: none">• Traffic volume: the number of vehicles passing a specific region in an hour.• Traffic pattern characteristic: Traffic flow patterns in a city are influenced by various area attributes. Also, traffic volume has a periodic daily and weekly pattern.• Spatial-temporal factors correlation: Traffic flow in an area will be affected by its nearby infrastructures, during specific periods for different areas. You should think about how the volume will change in a specific area, during a specific time. For example,<ul style="list-style-type: none">- Airports, and train stations - increased volume on weekends and holidays.- Residential areas - more activities during morning and evening rush hours.- Commercial areas - busy during lunch hours and after-work periods.- Educational locations - high volume during peak hours near schools. <p>Chain of thought:</p> <p>Think carefully about the following questions about how spatial-temporal factors affect traffic flow.</p> <ul style="list-style-type: none">• What is the attribute of this area and what is the predicted time zone located in special periods (like rush hours, weekdays, weekends, and holidays)?• What are the traffic patterns of this area, and what is the change in different time slots?• What is the historical temporal trend according to temporal information, considering the weekdays, around holidays?
<p>Input prompt:</p> <p>Some important information is listed as follows:</p> <ul style="list-style-type: none">• Location: District 3 in Yolo, California, USA, along the US50-E freeway, lane 4, direction of eastbound.• Today’s weather: Sunny. The temperature is 6.0 °C and the visibility is 10.0 miles.• Region information: including transportation areas, commercial areas, and educational areas within a range of 5 km.• Current time: 3 p.m., 2018-2-19, Monday, Washington’s Birthday.• Traffic volume data in the past 12 h were 19, 44, 98, 150, 156, 178, 208, 246, 248, 257, 263 and 269, respectively. <p>According to the above information and careful reasoning, please predict traffic volumes in the next 12 h (from 4 p.m. to 3 a.m.). Format the final answer in a single line as a JSON dictionary like: Traffic volume data in the next 12 h: [V1, V2, V3, V4, V5, V6, V7, V8, V9, V10, V11, V12].</p>
<p>Ground truth:</p> <p>Traffic volume data in the next 12 h: [262, 229, 221, 214, 152, 127, 100, 58, 38, 25, 22, 18].</p>

Appendix B
Explanatory example

A complete example of interpretative output is given in Table B1. This example shows the traffic flow prediction in the next 12 h at 5 p.m. on Christmas Day 2019. By adding explanatory demand instructions to the original prompt (bold text in Input Prompt of Table B1), xTP-LLM can output both the prediction results as well as the explanatory text. It can be found that our model is capable of adjusting the traffic flow prediction strategy by taking into account holidays and the impact of weather on human activities, which shows that our model is highly reliable and accountable.

Table B1
Example of the complete response with explanatory output of our proposed xTP-LLM. This example corresponds to the Example (d) in Table 6.

System prompt: (This part is the same as that in Table A1) ...
Input prompt: Some important information is listed as follows: <ul style="list-style-type: none">• District 3 in Yolo, California, USA, along the US50-W freeway, lane 3, direction of westbound.• Today's weather: Rain. The temperature is 8.1 °C and the visibility is 9.7 miles.• Region information: including transportation areas, commercial areas and educational areas within a range of 5 km.• Current time: 5 p.m., 2019-12-25, Wednesday, Christmas Day.• Traffic volume data in the past 12 h were 119, 161, 218, 244, 260, 277, 276, 275, 282, 288, 274 and 240, respectively. According to the above information and careful reasoning, please predict traffic volumes in the next 12 h (from 2 a.m. to 1 p.m.) and explain it . Format the final answer in a single line as a JSON dictionary like: {Traffic volume data in the next 12 h: [V1, V2, V3, V4, V5, V6, V7, V8, V9, V10, V11, V12], Explanation: xxx }. Please think step by step.
xTP-LLM's response: { "Traffic volume data in the next 12 h": [214, 183, 158, 157, 119, 69, 47, 36, 31, 26, 27, 33], "Explanation": I will provide a step-by-step explanation for predicting traffic volumes in the next 12 h based on the given information. 1. Time and Date: It is 5 p.m. on a Wednesday, December 25, 2019, which is Christmas Day, with rainy weather and good visibility. 2. Historical Traffic Volume: In the past 12 h, the traffic volumes were: [119, 161, 218, 244, 260, 277, 276, 275, 282, 288, 274, 240]. 3. Current Situation: Given the holiday, Christmas Day, and the rainy weather, traffic patterns may differ from regular weekdays. 4. Christmas Day Factor: On holidays, traffic volume might be lower due to decreased commuter and commercial activities. People may stay indoors for celebrations. 5. Rain Impact: The rain could further reduce traffic volume as people may prefer to avoid driving in adverse weather conditions. 6. Traffic Trend Prediction: <ul style="list-style-type: none">- The evening hours (6 p.m. to 9 p.m.): Traffic volume is likely to be significantly lower than usual as people prioritize holiday celebrations.- Late-night hours (10 p.m. to 5 a.m.): Traffic volume may remain low throughout the night. }
ChatGPT's summary: Considering the current time at 5 p.m. on a rainy Christmas Day, traffic volume is expected to decrease significantly over the next 12 h. As it is a holiday, commuter and commercial activities are minimal, leading to reduced traffic. Additionally, the rainy weather may further deter people from traveling or driving, contributing to the decline in traffic volumes. The focus on indoor celebrations and festivities is likely to keep the roads quiet throughout the night.

References

- Asanjan, A., Yang, T., Hsu, K., Sorooshian, S., Lin, J., Peng, Q., 2018. Short-term precipitation forecast based on the PERSIANN system and LSTM recurrent neural networks. *J. Geophys. Res. Atmos.* 123, 12–543.
- Ates, E., Aksar, B., Leung, V.J., Coskun, A.K., 2021. Counterfactual explanations for multivariate time series. In: 2021 International Conference on Applied Artificial Intelligence, pp. 1–8.
- Bai, L., Yao, L., Li, C., Wang, X., Wang, C., 2020. Adaptive graph convolutional recurrent network for traffic forecasting. In: Proceedings of the 34th International Conference on Neural Information Processing Systems, pp. 17804–17815.
- Barredo-Arrieta, A., Laña, I., Del Ser, J., 2019. What lies beneath: a note on the explainability of black-box machine learning models for road traffic forecasting. In: 2019 IEEE Intelligent Transportation Systems Conference. IEEE, pp. 2232–2237.
- Cai, L., Janowicz, K., Mai, G., Yan, B., Zhu, R., 2020. Traffic transformer: capturing the continuity and periodicity of time series for traffic forecasting. *Trans. GIS* 24, 736–755.
- Cambria, E., Malandri, L., Mercorio, F., Nobani, N., Seveso, A., 2024. XAI meets LLMs: a survey of the relation between explainable AI and large language models. <https://doi.org/10.48550/arXiv.2407.15248>.
- Chang, E.Y., 2023. Examining GPT-4's capabilities and enhancement with socrasynth. In: The 10th International Conference on Computational Science and Computational Intelligence, pp. 7–14.
- Chen, Y., Chen, H., Ye, P., Lv, Y., Wang, F.Y., 2022. Acting as a decision maker: traffic-condition-aware ensemble learning for traffic flow prediction. *IEEE Trans. Intell. Transport. Syst.* 23, 3190–3200.
- Chen, Y., Wang, X., Xu, G., 2023. GATGPT: a pre-trained large language model with graph attention network for spatiotemporal imputation. <https://doi.org/10.48550/arXiv.2311.14332>.
- Dai, L., Wang, L., 2020. Nonlinear analysis of high accuracy and reliability in traffic flow prediction. *Nonlinear Eng.* 9, 290–298.
- Dimitrakopoulos, G., Demestichas, P., 2010. Intelligent transportation systems. *IEEE Veh. Technol. Mag.* 5, 77–84.
- Ding, N., Qin, Y., Yang, G., Wei, F., Yang, Z., Su, Y., et al., 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nat. Mach. Intell.* 5, 220–235.
- Du, S., Li, T., Gong, X., Hornig, S.J., 2020. A hybrid method for traffic flow forecasting using multimodal deep learning. *Int. J. Comput. Intell. Syst.* 13, 85–97.
- Fan, Z., Song, X., Xia, T., Jiang, R., Shibasaki, R., Sakuramachi, R., 2018. Online deep ensemble learning for predicting citywide human mobility. In: Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, pp. 1–21.
- Fang, Z., Long, Q., Song, G., Xie, K., 2021. Spatial-temporal graph ODE networks for traffic flow forecasting. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pp. 364–373.
- Gao, Q., Trajcevski, G., Zhou, F., Zhang, K., Zhong, T., Zhang, F., 2018. Trajectory-based social circle inference. In: Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 369–378.
- Gruver, N., Finzi, M., Qiu, S., Wilson, A.G., 2024. Large language models are zero-shot time series forecasters. In: Proceedings of the 37th International Conference on Neural Information Processing Systems, pp. 19622–19635.
- Guo, S., Lin, Y., Feng, N., Song, C., Wan, H., 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 922–929.
- Hochreiter, S., 1997. Long short-term memory. *Neural Computation MIT-Press* 9, 1735–1780.
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* 79, 2554–2558.
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., et al., 2021. Lora: low-rank adaptation of large language models. <https://doi.org/10.48550/arXiv.2106.09685>.
- Huang, S., Mamidanna, S., Jangam, S., Zhou, Y., Gilpin, L.H., 2023. Can large language models explain themselves? A study of LLM-generated self-explanations. <https://doi.org/10.48550/arXiv.2310.11207>.
- Jiang, J., Han, C., Zhao, W.X., Wang, J., 2023. PDFormer: Propagation delay-aware dynamic long-range transformer for traffic flow prediction, pp. 4365–4373.
- Jin, M., Wang, S., Ma, L., Chu, Z., Zhang, J.Y., Shi, X., et al., 2023. Time-LLM: time series forecasting by reprogramming large language models. <https://doi.org/10.48550/arXiv.2310.01728>.
- Jin, W., Lin, Y., Wu, Z., Wan, H., 2018. Spatio-temporal recurrent convolutional networks for citywide short-term crowd flows prediction. In: Proceedings of the 2nd International Conference on Compute and Data Analysis, pp. 28–35.
- Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y., 2022. Large language models are zero-shot reasoners. In: Proceedings of the 36th International Conference on Neural Information Processing Systems, pp. 22199–22213.
- Lan, S., Ma, Y., Huang, W., Wang, W., Yang, H., Li, P., 2022. Dstagnn: dynamic spatial-temporal aware graph neural network for traffic flow forecasting. In: International Conference on Machine Learning, pp. 11906–11917.
- Li, Y., Yu, R., Shahabi, C., Liu, Y., 2017. Diffusion convolutional recurrent neural network: data-driven traffic forecasting. <https://doi.org/10.48550/arXiv.1707.01926>.
- Liang, Y., Ke, S., Zhang, J., Yi, X., Zheng, Y., 2018. Geoman: multi-level attention networks for geo-sensory time series prediction. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence, pp. 3428–3434.
- Liu, S., Cheng, H., Liu, H., Zhang, H., Li, F., Ren, T., et al., 2023. Llava-plus: learning to use tools for creating multimodal agents. <https://doi.org/10.48550/arXiv.2311.05437>.
- Liu, X., Xia, Y., Liang, Y., Hu, J., Wang, Y., Bai, L., et al., 2024. LargeST: a benchmark dataset for large-scale traffic forecasting. In: Proceedings of the 37th International Conference on Neural Information Processing Systems, pp. 75354–75371.
- Marblestone, A.H., Wayne, G., Kording, K.P., 2016. Toward an integration of deep learning and neuroscience. *Front. Comput. Neurosci.* 10, 2159–42202.
- Peng, M., Guo, X., Chen, X., Zhu, M., Chen, K., Wang, X., et al., 2024. LC-LLM: explainable lane-change intention and trajectory predictions with large language models. <https://doi.org/10.48550/arXiv.2403.18344>.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al., 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 1–9.
- Ranjan, N., Bhandari, S., Zhao, H.P., Kim, H., Khan, P., 2020. City-wide traffic congestion prediction based on CNN, LSTM and transpose CNN. *IEEE Access* 8, 81606–81620.
- Ray, P.P., 2023. ChatGPT: a comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things and Cyber-Physical Systems* 3, 121–154.
- Shi, X., Xue, S., Wang, K., Zhou, F., Zhang, J., Zhou, J., et al., 2023. Language models can improve event prediction by few-shot abductive reasoning. In: Proceedings of the 37th International Conference on Neural Information Processing Systems, pp. 29532–29557.
- Thirunavukarasu, A.J., Ting, D.S.J., Elangovan, K., Gutierrez, L., Tan, T.F., Ting, D.S.W., 2023. Large language models in medicine. *Nat. Med.* 29, 1930–1940.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., et al., 2023. Llama 2: open foundation and fine-tuned chat models. <https://doi.org/10.48550/arXiv.2307.09288>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., et al., 2017. Attention is all you need. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 6000–6010.
- Wan, H., Feng, S., Tan, Z., Wang, H., Tsvetkov, Y., Luo, M., 2024. DELL: generating reactions and explanations for LLM-based misinformation detection. In: Findings of the Association for Computational Linguistics ACL 2024, pp. 2637–2667.
- Wang, S., Cao, J., Yu, P.S., 2022. Deep learning for spatio-temporal data mining: a survey. *IEEE Trans. Knowl. Data Eng.* 34, 3681–3700.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., et al., 2022. Chain-of-thought prompting elicits reasoning in large language models. In: Proceedings of the 36th International Conference on Neural Information Processing Systems, pp. 24824–24837.
- Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., et al., 2023. BloombergGPT: a large language model for finance. <https://doi.org/10.48550/arXiv.2303.17564>.
- Wu, Z., Pan, S., Long, G., Jiang, J., Zhang, C., 2019. Graph wavenet for deep spatial-temporal graph modeling. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence, pp. 1907–1913.
- Xu, M., Dai, W., Liu, C., Gao, X., Lin, W., Qi, G.J., et al., 2020. Spatial-temporal transformer networks for traffic flow forecasting. <https://doi.org/10.48550/arXiv.2001.02908>.
- Yan, J., Wang, H., 2023. Self-interpretable time series prediction with counterfactual explanations. In: Proceedings of the 40th International Conference on Machine Learning, pp. 39110–39125.
- Yang, H.F., Dillon, T.S., Chen, Y.P.P., 2016. Optimized structure of the traffic flow forecasting model with a deep learning approach. *IEEE Transact. Neural Networks Learn. Syst.* 28, 2371–2381.
- Yannis, G., Dragomanovits, A., Laiou, A., La Torre, F., Domenichini, L., Richter, T., et al., 2017. Road traffic accident prediction modelling: a literature review. *Proceedings of the Institution of Civil Engineers - Transport* 170, 245–254.
- Yu, B., Yin, H., Zhu, Z., 2018. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence, pp. 3634–3640.
- Yu, C., Ma, X., Ren, J., Zhao, H., Yi, S., 2020. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In: Proceedings of the 16th European Conference on Computer Vision, pp. 507–523.
- Zhang, J., Zheng, Y., Qi, D., 2017. Deep spatio-temporal residual networks for citywide crowd flows prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 1655–1661.
- Zhang, L., Fu, K., Ji, T., Lu, C.T., 2022. Granger causal inference for interpretable traffic prediction. In: 2022 IEEE 25th International Conference on Intelligent Transportation Systems, pp. 1645–1651.
- Zhou, Y., Li, Y., Zhu, Q., Chen, F., Shao, J., Luo, Y., et al., 2019. A reliable traffic prediction approach for bike-sharing system by exploiting rich information with temporal link prediction strategy. *Trans. GIS* 23, 1125–1151.
- Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M., 2023. MiniGPT-4: enhancing vision-language understanding with advanced large language models. <https://doi.org/10.48550/arXiv.2304.10592>.
- Ziegler, D.M., Stiennon, N., Wu, J., Brown, T.B., Radford, A., Amodei, D., et al., 2019. Fine-tuning language models from human preferences. <https://doi.org/10.48550/arXiv.1909.08593>.
- Zytek, A., Pidò, S., Veeramachaneni, K., 2024. LLMs for XAI: future directions for explaining explanations. <https://doi.org/10.48550/arXiv.2405.06064>.



Xusen Guo received the B.S. degree in material chemistry from Sun Yat-sen University, China in 2017, and received the M.S. degree in computer science and technology from Sun Yat-sen University in 2020. Currently, he is pursuing the Ph.D. degree in the Intelligent Transportation Thrust, Systems Hub, The Hong Kong University of Science and Technology (Guangzhou). His research interests include 3D perception in autonomous driving and the application of large language models in intelligent transportation.



Mingxing Peng received the B.S. degree in digital media technology from Huaqiao University, China in 2018, and received the M.S. degree in computer technology from Sun Yat-sen University in 2021. Currently, she is pursuing the Ph.D. degree in the Intelligent Transportation Thrust, the Systems Hub, The Hong Kong University of Science and Technology (Guangzhou). Her research interests include motion planning, trajectory prediction, and large language models for autonomous driving.



Qiming Zhang received the B.S. degree in automation, artificial intelligence, and automation academy from Huazhong University of Science and Technology, China in 2020, and received the M.S. degree in electrical engineering from Columbia University, USA in 2023. He is currently pursuing the Ph.D. degree in Intelligent Transportation Thrust, Systems Hub, The Hong Kong University of Science and Technology (Guangzhou). His research interests include the applications of foundation models in transportation and autonomous driving.



Meixin Zhu is a tenure-track Assistant Professor in the Thrust of Intelligent Transportation (INTR), the Systems Hub, The Hong Kong University of Science and Technology (Guangzhou), and with Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things. He obtained the Ph.D. degree in intelligent transportation from the University of Washington in 2022. His research interests include autonomous driving decision making and planning, driving behavior modeling, traffic-flow modeling and simulation, traffic signal control, and (multi-agent) reinforcement learning.



Junyue Jiang earned the B.S. degrees in electrical and computer engineering from Zhejiang University and the University of Illinois at Urbana-Champaign in 2024. He is presently pursuing the Ph.D. degree in the Department of Civil and Systems Engineering, Johns Hopkins University. His research primarily concentrates on the development of machine learning applications within the context of enhancing the sustainability of traffic systems.



Hao Frank Yang is currently a tenure-track Assistant Professor at the Department of Civil and Systems Engineering, Johns Hopkins University, a faculty member at the Johns Hopkins Data Science and AI Institute. He previously was a Research Scientist at the NSF AI Institute for Edge Computing Leveraging Next Generation Networks at Duke University. He completed the M.S. and Ph.D. degrees at the Department of Civil and Environmental Engineering, the University of Washington. His innovative research centers around developing trustworthy computing and representation learning methods to improve the safety, equity and sustainability in transportation. This includes developing hardware-software co-optimization solutions for traffic sensing and computing, proposing trustworthy machine-learning methods, and enhancing the human-machine collaborations in traffic systems.