PGET-NET: PHYSICS-GUIDED NEURAL NETWORK WITH EVENT-TEXT CONDITIONING FOR TRAFFIC FORECASTING

Anonymous authors
Paper under double-blind review

ABSTRACT

Accurate traffic forecasting during disruptive events is essential for dynamic traffic management and emergency decision-making, yet remains challenging. Disruptive events, such as extreme weather and traffic accidents, can disturb routine urban traffic patterns, thereby limiting the ability of methods that rely primarily on historical observations to accurately capture abrupt changes in traffic states. Although event-related text has been incorporated into forecasting models, it is often treated as auxiliary information rather than explicitly coupled with traffic dynamics. We propose PGET-Net, an event-text-conditioned physics-guided neural network that uses event semantics to modulate traffic-state evolution for multi-step forecasting. PGET-Net learns a nonlinear mapping from observed traffic states to a latent traffic potential field and models its evolution using diffusion-advection dynamics. The diffusion component describes routine propagation over the static road topology. The advection component maps event semantics to a graph representing event-impact propagation, which parameterizes event-induced perturbations to the latent dynamics. Both mechanisms are represented within a graph differential equation network, where a learnable gate adaptively combines their contributions to the evolution of the latent field. The resulting trajectory is decoded into future traffic prediction. Experiments on the Beijing Text-Traffic dataset (BjTT) show that PGET-Net outperforms the evaluated baselines, supporting the effectiveness of coupling event-conditioned perturbations with routine traffic propagation.

1 INTRODUCTION

Disruptive events, such as traffic accidents, extreme weather, and infrastructure failures, often severely disrupt the routine operation of urban road networks and may cause abrupt changes in traffic conditions within a short period. Under such emergencies, accurate traffic forecasting is crucial for supporting real-time decision-making, optimizing traffic signal timing, and planning efficient evacuation routes. It helps reduce traffic delays and improve road safety. Meanwhile, disruptive events involve considerable uncertainty, and their spatial extent and duration may change continuously. Traffic authorities need to promptly assess traffic conditions, implement response measures, and issue real-time alerts to drivers. This requires forecasting models to accurately predict traffic states under disruptive events. Such predictions provide a reliable basis for traffic situation assessment and emergency response.

Existing traffic forecasting methods can be broadly divided into physics-based and data-driven methods. Physics-based models typically rely on traffic flow theory Ni (2016) and represent traffic systems as coupled differential equations. Forecasts are then obtained by simulating the systems governed by these equations. However, such models are usually built on strong assumptions about traffic flow and a limited set of parameters Mo et al. (2021). They may therefore fail to fully capture complex travel behavior and uncertainty in real-world traffic. Data-driven methods typically use historical observations to train statistical learning models and then use the trained models to generate forecasts. In recent years, with the continuous accumulation of large-scale traffic data, deep learning-based data-driven methods have attracted widespread attention from both industry and academia. They have also achieved good performance in many practical applications. When

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

future traffic states continue to follow periodic patterns under routine operating conditions, these methods can generally achieve satisfactory forecasting performance. However, they rely heavily on historical traffic data and periodic traffic patterns, making it difficult to accurately characterize non-routine traffic changes caused by disruptive events.

Rapid advances in information technology have increased the availability of multimodal urban data, prompting researchers to incorporate social media content into traffic forecasting as auxiliary features. For example, Yao & Qian (2021) leveraged indicators derived from Twitter data, such as residents' late-night activities and emotional states, to improve morning peak congestion forecasting. Tsai et al. (2022) proposed a Traffic-Twitter Transformer that combines historical traffic data with tweet-based features, including keyword frequencies and counts of event-related terms. By capturing public discussions of disruptive events and cultural activities, the model improves forecasting accuracy under abnormal conditions. Melhem et al. (2024) extracted traffic events and location information from tweet streams, dynamically constructed an event knowledge graph, and used the graph to fine-tune a domain-specific large language model (LLM). This approach makes the forecasts more context-aware. Wang et al. (2024) iteratively used LLM-based agents to filter news and align relevant events with traffic time series, thereby improving traffic forecasting accuracy. Overall, these studies show that textual information from tweets and news reports complements traditional traffic sensor data and provides valuable event information for capturing traffic-state changes caused by social activities or unexpected accidents. Despite this progress, existing methods remain limited in combining event semantics with traffic dynamics. Textual inputs are typically treated as static auxiliary features and fused with spatio-temporal traffic representations at the feature level. Most existing studies have not explicitly coupled event semantics with traffic dynamics.

In summary, traffic forecasting still faces two key challenges:

- Unpredictability of disruptive events: Disruptive events, such as traffic accidents, concerts, and sporting events, occur frequently in transportation systems and cause irregular changes in traffic conditions. Traditional single-source data cannot accurately reflect the effects of these events on traffic conditions, resulting in reduced forecasting accuracy.
- Insufficient incorporation of physical knowledge: Although event-related text has been introduced into traffic forecasting models, it is usually treated as auxiliary information and fused with traffic data at the feature level rather than being explicitly coupled with traffic dynamics. These methods struggle to identify the perturbation effects of disruptive events on the spatio-temporal evolution of traffic states, thereby limiting forecasting accuracy.

To address these challenges, this paper proposes PGET-Net (Physics-Guided Neural Network with Event-Text Conditioning), an event-text-conditioned physics-guided neural network for multi-step traffic forecasting. PGET-Net no longer treats event text simply as an input parallel to traffic data features. Instead, it uses event semantics to condition the evolution of latent traffic states. Specifically, the model first nonlinearly maps observed traffic states into a latent traffic potential field that describes road-network traffic dynamics. It then employs diffusion–advection dynamics to model the evolution of this latent field. The diffusion component characterizes the routine propagation of traffic states over the static road topology. The advection component maps event semantics to an event-impact propagation graph and represents event impacts as an additional directional driving force acting on the evolution of the latent traffic potential field. The propagation graph further parameterizes the direction and strength of this driving force, thereby characterizing the directional perturbation of disruptive events on traffic-state evolution trajectories. The diffusion and advection mechanisms are unified within a graph differential equation network. A learnable gating mechanism adaptively coordinates their contributions to the evolution of the latent potential field. Finally, the model decodes the resulting latent evolution trajectory into future traffic states. In this way, PGET-Net explicitly couples event context with traffic dynamics while maintaining a structured distinction between routine propagation and event-conditioned perturbations. In summary, the main contributions of this study are threefold:

- This paper proposes PGET-Net, an event-text-conditioned physics-guided forecasting framework that introduces event semantics into the latent evolution of traffic states and enables multi-step traffic forecasting under disruptive events.
- This paper develops a graph differential equation model based on diffusion–advection mechanisms. It characterizes routine traffic propagation over the static road topology and

latent dynamical perturbations induced by disruptive events, respectively. A learnable gating mechanism is used to adaptively integrate their effects.

- Extensive experiments are conducted on the Beijing Text-Traffic Dataset (BjTT). Performance comparisons show that PGET-Net outperforms the baseline models, validating its effectiveness for traffic forecasting under disruptive events.

2 Related Work

Data-Driven Traffic Forecasting Methods. Traffic forecasting plays an important role in modern intelligent transportation systems. Forecasting methods have progressed from classical statistical models to deep spatio-temporal learning methods. Graph neural networks have become a dominant approach for modeling spatial dependencies over road networks, as exemplified by STGCN (Yu et al., 2017). Adaptive graph learning further allows latent spatial relations to be inferred directly from traffic observations (Wu et al., 2019). More recent studies extend this direction through trend-seasonality decomposition-based modeling (Cao et al., 2025), higher-order hypergraph structures (Li et al., 2025a), and Transformer-based representations (Li et al., 2025b). Overall, deep learning models based on graphs, hypergraphs, and self-attention have made considerable progress relative to traditional mathematical models and machine learning algorithms in addressing the need for accurate forecasting in intelligent urban transportation. However, most of these models rely on a single data source, making it difficult to meet the need to analyze the evolution of traffic conditions in nonroutine scenarios with sparse data. Moreover, most are data-driven methods that overlook the modeling of key physical dynamics in traffic systems, making it difficult to accurately analyze the spatio-temporal evolution of traffic data, particularly in complex nonroutine scenarios.

Physics-guided deep learning. Many recent studies have proposed combining physics-based modeling with state-of-the-art deep learning techniques, giving rise to a field known as "physics-guided deep learning" to address complex scientific and engineering problems (Karniadakis et al., 2021; Karpatne et al., 2024). The most direct approach incorporates physical knowledge as constraints in the loss function (Raissi et al., 2017; Shi et al., 2021), penalizing deviations from physical laws to encourage physically consistent patterns and improve model generalization (Greydanus et al., 2019; Choi et al., 2023b). For example, Jia et al. (2019) introduced a physics-based regularization loss for lake temperature prediction. However, constraint-based approaches may become restrictive when the prescribed physical knowledge is inaccurate or incomplete. An alternative is to embed physical mechanisms directly into neural architectures (Bao et al., 2021; Xu et al., 2024), allowing the learned dynamics to evolve under structured physical priors rather than relying solely on auxiliary losses. This idea has been increasingly explored in continuous spatio-temporal forecasting. Choi et al. (2023a) combined Neural ODEs with advection-diffusion dynamics for climate forecasting, while Hettige et al. (2024) incorporated physical knowledge into neural architectures for air-quality prediction. More recent studies further extend this direction by modeling open-system dynamics (Liang et al., 2025; Dey et al., 2026), delayed physical propagation (Wu et al., 2026), and Eulerian advection-diffusion-reaction processes (Zhang et al., 2026). In traffic forecasting, Ji et al. (2020) introduced latent spatio-temporal potential fields to characterize traffic-flow dynamics, which were further extended by Wang et al. (2022). However, these approaches rely on grid-based representations. Ji et al. (2022) subsequently extended physics-guided differential equation modeling to road networks by learning the evolution of latent traffic potential fields. Existing methods, however, have not sufficiently considered how disruptive events alter the propagation of traffic states, nor have they explicitly used event-text semantics to modulate the corresponding dynamical mechanisms. Therefore, explicitly coupling event-text semantics with traffic-state propagation dynamics while characterizing both routine propagation and perturbations caused by disruptive events is key to improving the accuracy of multi-step traffic forecasting under disruptive events.

3 METHODOLOGY

3.1 Problem Definition

Definition 1 (Traffic Forecasting.) Given a traffic network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ with N nodes and adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, let $\mathbf{x}_t \in \mathbb{R}^{N \times d}$ denote the traffic state at time step t, where d is the number of traffic features (e.g., traffic speed and flow). Given T historical observations $\mathbf{X}_{1:T} =$

$\{\mathbf{x}_1, \ldots, \mathbf{x}_T\} \in \mathbb{R}^{T \times N \times d}$, our goal is to predict the traffic states over the next τ time steps:

$$\hat{\mathbf{X}}_{T+1:T+\tau} = f\left(\mathbf{X}_{1:T}, \mathcal{G}\right), \tag{1}$$

where $\hat{\mathbf{X}}_{T+1:T+\tau} \in \mathbb{R}^{\tau \times N \times d}$.

Definition 2 (Disruptive-Event-Aware Traffic Forecasting.) This task extends conventional traffic forecasting by incorporating disruptive-event textual information. Let $\mathbf{s}_t$ denote a disruptive event occurring at time step t, such as a traffic accident, road construction, urban waterlogging, or flooding. The prediction function jointly processes the historical traffic sequence $\mathbf{X}_{1:T}$, historical event text $\mathbf{s}_t$, and graph structure $\mathcal{G}$:

$$\hat{\mathbf{X}}_{T+1:T+\tau} = f\left(\mathbf{X}_{1:T}, \mathbf{s}_T, \mathcal{G}\right). \tag{2}$$

3.2 Preliminaries

Diffusion-Advection Equation. The diffusion-advection equation describes the transport of a substance in a fluid under the combined effects of diffusion and advection Moreira et al. (1998). To characterize traffic state evolution in urban road networks during disruptive events, these two mechanisms are adopted within a latent traffic potential field. Diffusion describes propagation under normal conditions, whereas advection describes the propagation of event-driven perturbations. The continuous formulation is given by

$$\frac{\partial z}{\partial t} = k\nabla^2 z - \vec{v} \cdot \nabla z. \tag{3}$$

Here, z denotes the latent traffic potential; $\partial z/\partial t$ is its partial derivative with respect to time; k is the diffusion coefficient; ∇^2 is the Laplacian operator, and $k\nabla^2 z$ is the diffusion term; $\vec{v}$ is an implicit driving vector field represented by disruptive event information; ∇z is the potential gradient, and $-\vec{v} \cdot \nabla z$ is the advection term. The diffusion process represents the random movement of traffic potential from regions of higher potential to regions of lower potential under normal conditions. The advection process emphasizes directional transport induced by external forces, such as disruptive events, analogous to the rapid dispersion of pollutants driven by wind. To describe these transport processes on discrete graph structures, the diffusion-advection equation has recently been extended to graphs Hettige et al. (2024); Liang et al. (2025). Graph neural networks are used to approximate the corresponding spatial transport operators.

4 Physics-Guided Traffic Forecasting

4.1 Overview

Under abnormal traffic conditions, normal operations in urban road networks are often disturbed by disruptive events, causing traffic state evolution to deviate from normal patterns. To address this issue, PGET-NET is developed. Diffusion and advection mechanisms are adopted within a latent traffic potential field to describe propagation under normal conditions and event-driven perturbations, respectively. State evolution under their combined effects is then modeled.

The overall architecture is illustrated in Fig. 1. It consists of an encoder, a multi-graph differential equation network, and a decoder. Historical traffic states are first mapped by the encoder to the initial state of the latent traffic potential. A diffusion operator is then constructed from the static road network structure to provide a topological basis for modeling propagation under normal conditions. Meanwhile, the graph structure associated with advection and its corresponding operator are generated from disruptive event text. Event semantics are thereby incorporated into the modeling of transport relationships between nodes. The diffusion and advection processes are subsequently combined through a gating mechanism to construct a latent traffic state evolution function. Given the initial state and the evolution function, a sequence of future latent states is obtained through numerical integration using an ODE solver. Multi-step traffic predictions are then generated by the decoder.

4.2 Diffusion-Advection Graph Construction

Diffusion Graph Construction. Under the modeling assumption of a latent traffic potential field, a diffusion process is adopted to characterize latent state evolution under normal conditions. This

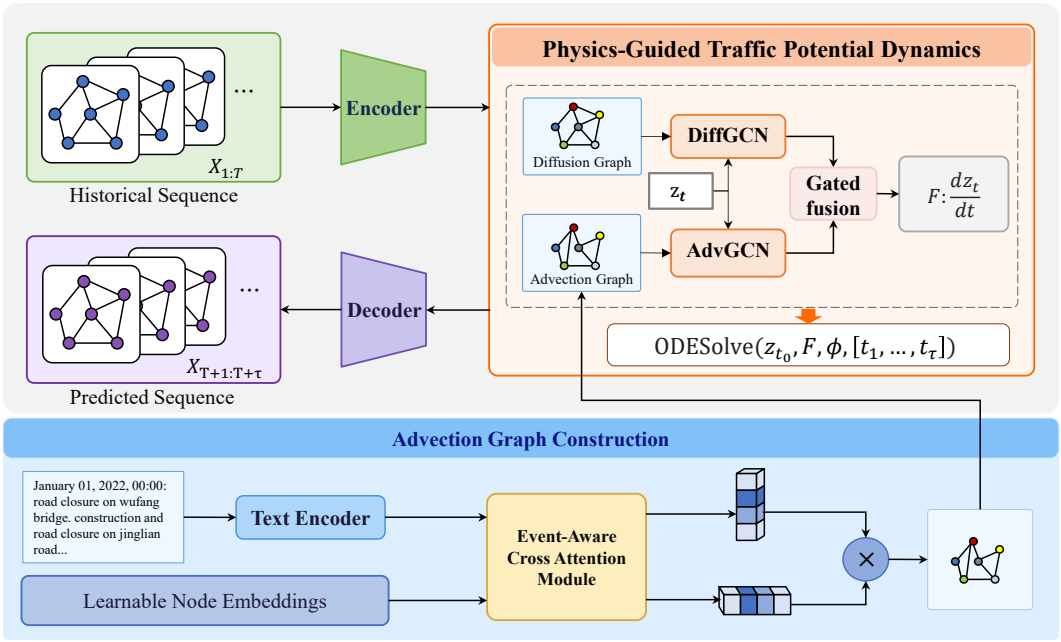

Figure 1: Overall architecture of PGET-Net.

process is guided by transport from higher-potential to lower-potential nodes, driven by potential differences between nodes. The diffusion graph is constructed from the road connections represented by the static road network topology, providing the connectivity for potential propagation between nodes. The graph is used to model latent traffic potential evolution under normal conditions.

Advection Graph Construction. Disruptive events may alter traffic propagation patterns in the road network. To characterize these changes, an advection mechanism is introduced to complement the propagation under normal conditions modeled by diffusion. This mechanism describes the directed transport of latent traffic potential under event-driven perturbations. Semantic information in event text is used to guide the learning of directed transport relationships between nodes. These relationships are used to construct an event-impact propagation graph, which serves as the advection graph for modeling event-induced perturbations to the latent dynamics. Specifically, event text is encoded into semantic embeddings $\mathbf{E}_T$ by a BERT-based (Devlin et al., 2019) text encoding layer. Meanwhile, learnable node embeddings $\mathbf{E}_S$ are introduced to represent differences between nodes. To incorporate event semantics into node representations, an event-aware cross-attention module based on multi-head attention, denoted as $\mathrm{ECA}(\cdot)$, is introduced to generate event-aware node embeddings $\mathbf{E}$:

$$\mathbf{E} = \mathrm{ECA}(\mathbf{E}_S, \mathbf{E}_T, \mathbf{E}_T), \tag{4}$$

where $\mathbf{E}_S$ is used as the query, $\mathbf{E}_T$ is used as the keys and values. Subsequently, the adjacency matrix $\mathbf{A}_{\mathrm{a}}$ is constructed from the embeddings $\mathbf{E}$ and normalized to represent spatial dependencies conditioned on events. The advection graph is then constructed from this matrix to model event-driven evolution of latent traffic potential.

4.3 PGET-Net

Traffic State Encoder. Given historical traffic states $\mathbf{X}_{1:T}$ as input, the initial latent traffic potential state $\mathbf{z}_{t_0}$ is generated by the encoder:

$$\mathbf{z}_{t_0} = \mathrm{Encoder}\left(\mathbf{X}_{1:T}\right), \tag{5}$$

Diffusion-Advection-Driven State Evolution. Guided by diffusion–advection mechanisms, a state evolution function is introduced to model the dynamics of the latent traffic potential field. During disruptive events, the latent traffic potential is affected by both diffusion and advection. A gated

fusion mechanism is adopted to account for both processes. Through adaptive fusion, the evolution function of the latent traffic potential is formulated as

$$\frac{d\mathbf{z}_t}{dt} = F(\mathbf{z}_t, t; \varphi) = \boldsymbol{\alpha}_t \odot k \cdot \mathrm{GNN}_{\mathrm{diff}}(\mathbf{z}_t) - (\mathbf{1} - \boldsymbol{\alpha}_t) \odot \mathrm{GNN}_{\mathrm{adv}}(\mathbf{z}_t), \tag{6}$$

where $\boldsymbol{\alpha}_t \in \mathbb{R}^{N \times L}$ denotes an adaptive gating tensor estimated from the current diffusion and advection components through a learnable linear layer followed by a Sigmoid activation, with L denoting the hidden dimension of the evolution function. It adaptively regulates the relative contributions of diffusion and advection across different nodes and latent dimensions. $\mathrm{GNN}_{\mathrm{diff}}$ and $\mathrm{GNN}_{\mathrm{adv}}$ perform K-hop message passing on the static diffusion graph and the event-conditioned advection graph, respectively. Through the latter graph, event semantics modulate the advection contribution to latent state evolution.

Given the initial state $\mathbf{z}_{t_0}$ and the evolution function F, the future potential evolution trajectory $\mathbf{z}_{t_1:t_\tau}$ is obtained using an ODE solver:

$$\mathbf{z}_{t_1:t_\tau} = \mathrm{ODESolve}(\mathbf{z}_{t_0}, F, \varphi, [t_0, \ldots, t_\tau]), \tag{7}$$

where $[t_0, \ldots, t_\tau]$ is the sequence of solution times, $\mathbf{z}_{t_1:t_\tau}$ denotes the future latent potential states, and φ represents all trainable parameters.

Decoder. Given the latent state sequence $\mathbf{z}_{t_1:t_\tau}$ obtained from the diffusion-advection evolution, future traffic predictions $\hat{\mathbf{X}}_{T+1:T+\tau}$ are generated by the decoder:

$$\hat{\mathbf{X}}_{T+1:T+\tau} = \mathrm{Decoder}(\mathbf{z}_{t_1:t_\tau}), \tag{8}$$

where $\hat{\mathbf{X}}_{T+1:T+\tau}$ denotes the predictions for the next τ time steps. The decoder learns the mapping from the evolved latent traffic potential states to the corresponding traffic states, transforming the latent trajectory into the predicted traffic sequence.

5 EXPERIMENT

5.1 EXPERIMENT SETTING

Dataset. The performance of PGET-Net is evaluated on the Beijing Text-Traffic dataset (BjTT) (Zhang et al., 2024). BjTT is a multimodal urban traffic dataset designed for traffic forecasting, with coverage of the major road network within Beijing's Fifth Ring Road. A total of $1{,}260$ road nodes and $32{,}400$ consecutive time steps are included in the dataset. The period from January to March 2022 is covered, and traffic observations are sampled at four-minute intervals. Unlike conventional datasets containing only traffic time series, both road traffic states and traffic event text are provided in BjTT. The textual records are temporally aligned with the traffic observations, and the date, time, road location, and anomalous event at the corresponding time are included in each record. The anomalous events are classified into 30 categories, including traffic accidents, road construction, road closures, and adverse weather. The dataset is divided into training, validation, and test sets at a ratio of 7:1:2. For multi-step forecasting, historical traffic states from the past hour (15 time steps) are used to predict traffic data for the next hour (15 time steps).

Baselines. To validate the effectiveness of PGET-Net, DCRNN (Li et al., 2017), STGCN (Yu et al., 2017), ASTGCN (Guo et al., 2019), GWN (Wu et al., 2019), STSGCN (Song et al., 2020), AGCRN (Bai et al., 2020), and STG-NCDE (Choi et al., 2022) are selected as comparison methods. These baselines predict future traffic states from historical traffic data, whereas PGET-Net uses both traffic events and historical traffic data as inputs.

Evaluation Metrics. To evaluate the performance of each method, three metrics commonly used in traffic forecasting tasks are employed, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). These metrics are formulated as

follows:

$$\text{MAE} = \frac{1}{N_s}\sum_{i=1}^{N_s}|\hat{y}_i - y_i|, \tag{9}$$

$$\text{RMSE} = \sqrt{\frac{1}{N_s}\sum_{i=1}^{N_s}(\hat{y}_i - y_i)^2}, \tag{10}$$

$$\text{MAPE} = \frac{100\%}{N_s}\sum_{i=1}^{N_s}\left|\frac{\hat{y}_i - y_i}{y_i}\right| \tag{11}$$

where the total number of samples is denoted by N_s, the ground-truth value by y_i, and the predicted value of the i-th sample by $\hat{y}_i$. Different aspects of model performance are measured by these metrics, thereby allowing the model to be comprehensively evaluated.

Implementation Details. The default settings provided in the publicly available implementations are adopted for all baseline models. For PGET-Net, the batch size is set to 16, the GRU hidden dimension to 64, the latent state dimension to 16, and the graph convolution order to 2. The neural ordinary differential equation is solved using `dopri5`, with the absolute and relative tolerances set to 10^{-4} and 10^{-3}, respectively. The model is optimized using Adam with an initial learning rate of 5×10^{-4} and is trained for 100 epochs. Early stopping is applied based on validation-set performance. The learning rate is reduced by a factor of 0.5 after the 25th, 45th, and 65th epochs, with the minimum learning rate set to 2×10^{-6}. The random seed is fixed at 2021. Our experiments are implemented using PyTorch 2.8.0 (CUDA 12.8) and conducted on NVIDIA GeForce RTX 4090 GPUs.

5.2 PERFORMANCE COMPARISON

Overall Performance. Table 1 presents the performance of PGET-Net and all baselines at prediction horizons of 20, 40, and 60 minutes, together with the average performance across all horizons. PGET-Net achieves the best performance, outperforming all baselines at every prediction horizon. Compared with the second-best method, PGET-Net reduces the average MAE by 0.84%, average RMSE by 4.66%, and average MAPE by 1.76%. These results support the benefit of incorporating diffusion–advection physical mechanisms into a deep learning framework and using event-text semantics to construct event-driven advection fields for modeling the dynamic evolution of latent traffic potential fields under sudden events. From Table 1 we can also observe that: 1) Earlier graph-based models such as STGCN and DCRNN mainly rely on predefined road topology to model spatial dependencies, while ASTGCN and STSGCN enhance spatio-temporal representation through attention or synchronous graph modeling. 2) In contrast, adaptive graph methods such as GWN and AGCRN can learn latent spatial dependencies directly from traffic observations, making them more flexible in modeling complex traffic variations. Nevertheless, such adaptive dependencies are primarily inferred from the historical traffic states themselves and lack explicit guidance from real external information. 3) STG-NCDE captures traffic evolution through continuous spatio-temporal dynamics, but the learned dynamics lack explicit physical interpretation. Overall, these results support the effectiveness of PGET-Net, achieving accurate traffic prediction by using event text to guide the physics-based modeling of traffic propagation.

Peak Time Performance. To further evaluate the effectiveness of PGET-Net, we focus on weekday peak hours (7:00–10:00 and 17:00–20:00) in the test set. These periods typically exhibit higher traffic volumes, and sudden events are more likely to cause rapid changes in the latent traffic potential field. As shown in Table 2, peak periods pose a greater challenge to all baseline models. Nevertheless, PGET-Net maintains the best average performance. During peak periods, event text provides complementary contextual information about disruptive events that is not explicitly captured by historical traffic observations. By incorporating this information into the event-conditioned advection process, PGET-Net can better capture traffic perturbations, while the diffusion component preserves routine propagation over the road topology, which may contribute to its superior forecasting performance during peak periods.

Table 1: Overall prediction performance comparison. The bold and underlined font show the best and the second best result respectively.

Model	20 min			40 min			60 min			Average		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
DCRNN	3.57	5.53	8.35%	3.81	6.06	9.40%	3.99	6.43	10.09%	3.68	5.77	8.84%
STGCN	3.67	5.65	8.70%	4.01	6.31	9.98%	4.27	6.78	10.88%	3.83	5.96	9.31%
ASTGCN	3.77	5.79	8.90%	4.08	6.38	10.16%	4.34	6.81	10.90%	3.92	6.06	9.49%
GWN	3.50	5.41	8.23%	3.71	5.83	9.08%	3.87	6.10	9.59%	3.59	5.58	8.56%
STSGCN	3.56	5.59	8.43%	3.82	6.14	9.36%	4.03	6.52	10.07%	3.69	5.85	8.87%
AGCRN	3.49	5.44	8.27%	3.66	5.83	8.95%	3.80	6.08	9.32%	3.56	5.59	8.52%
STG-NCDE	3.62	5.56	8.78%	3.82	6.05	9.70%	4.04	6.46	10.44%	3.72	5.81	9.19%
PGET-Net	**3.45**	**5.17**	**8.16%**	**3.62**	**5.51**	**8.76%**	**3.78**	**5.78**	**9.17%**	**3.53**	**5.32**	**8.37%**

Table 2: Performance comparison under peak-hour traffic conditions.

Model	20 min			40 min			60 min			Average		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
DCRNN	4.51	7.44	15.17%	5.39	9.13	20.34%	5.72	9.78	22.39%	4.88	8.13	17.48%
STGCN	4.78	7.77	16.96%	5.83	9.61	23.15%	6.23	10.30	25.49%	5.22	8.51	19.67%
ASTGCN	4.94	7.98	17.47%	5.95	9.67	23.71%	6.32	10.24	25.40%	5.36	8.64	20.14%
GWN	4.37	7.13	**14.48%**	5.13	8.48	18.70%	5.45	8.99	20.26%	4.68	7.64	16.22%
STSGCN	4.55	7.59	15.58%	5.37	9.19	20.02%	5.70	9.80	21.91%	4.90	8.23	17.53%
AGCRN	4.41	7.21	14.88%	5.09	8.54	18.41%	5.36	9.00	19.32%	4.67	7.70	16.17%
STG-NCDE	4.66	7.63	16.61%	5.49	9.20	21.83%	5.84	9.85	23.89%	5.03	8.42	18.97%
PGET-Net	**4.26**	**6.86**	**14.48%**	**4.78**	**7.90**	**17.03%**	**4.97**	**8.40**	**17.66%**	**4.43**	**7.24**	**15.20%**

5.3 ABLATION STUDY

Table 3: Ablation study of different physical dynamics modeling strategies on the BjTT dataset.

Model	20 min			40 min			60 min			Average		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
unkP	3.73	5.62	8.76%	4.10	6.21	10.20%	4.39	6.62	11.22%	3.91	5.88	9.46%
Diff.	3.66	5.50	8.59%	3.97	5.99	9.84%	4.22	6.35	10.69%	3.81	5.72	9.19%
Adv.	3.70	5.52	8.71%	4.02	6.04	9.86%	4.27	6.40	10.63%	3.85	5.75	9.22%
PGET-Net	**3.45**	**5.17**	**8.16%**	**3.62**	**5.51**	**8.76%**	**3.78**	**5.78**	**9.17%**	**3.53**	**5.32**	**8.37%**

To evaluate the roles of diffusion, advection, and their interaction in traffic forecasting, we compare PGET-Net with three variants: (1) **unkP**, which replaces the physics-guided ODE vector field with an MLP module, with unknown physical mechanisms; (2) **Diff.**, which retains only traffic diffusion over the road network; (3) **Adv.**, which retains only event-induced advection. Table 3 shows the comparison of the three variants. We have several findings: 1) PGET-Net substantially outperforms **unkP**, demonstrating the value of physical knowledge in traffic dynamics. 2) **Diff.** and **Adv.** also outperform **unkP**, indicating that each physical process provides useful information beyond unconstrained latent-state modeling. 3) Moreover, PGET-Net performs better than both single-process variants. Diffusion captures routine traffic propagation, whereas event-conditioned advection captures perturbations caused by disruptive events. Their combination therefore provides a more complete account of traffic evolution under disruptive events.

6 CONCLUSION AND FUTURE WORK

In conclusion, this paper addresses the critical challenge of multi-step traffic forecasting under disruptive-event conditions by introducing PGET-Net, an event-text-conditioned physics-guided

neural network. By nonlinearly mapping observed traffic states into a latent traffic potential field, PGET-Net employs diffusion-advection dynamics to characterize the evolution of traffic potential under disruptive events. The diffusion and advection mechanisms are jointly represented within a graph differential equation network, while their contributions to the evolution of the latent potential field are adaptively integrated through a learnable gating mechanism. The evolved latent states are subsequently decoded into future traffic states. Experiments on the Beijing Text-Traffic dataset (BjTT) demonstrate that PGET-Net outperforms the compared baseline models, supporting the effectiveness of combining event-conditioned perturbations with routine traffic propagation. The proposed framework provides a physics-guided solution for the unified modeling of disruptive-event semantics and traffic-state propagation mechanisms. Future research could focus on further characterizing the time delays, attenuation, and traffic recovery processes associated with event impacts over the forecasting horizon, thereby providing better support for traffic situation assessment and emergency decision-making under disruptive-event conditions.

AI USE STATEMENT

We primarily used AI for language polishing and code completion. All AI-assisted modifications to the text and code were manually reviewed by the authors. We take full responsibility for the final submission.

REFERENCES

Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems*, 33:17804–17815, 2020.

Tianshu Bao, Xiaowei Jia, Jacob Zwart, Jeffrey Sadler, Alison Appling, Samantha Oliver, and Taylor T Johnson. Partial differential equation driven dynamic graph networks for predicting stream water temperature. In *2021 IEEE International Conference on Data Mining (ICDM)*, pp. 11–20. IEEE, 2021.

Lingxiao Cao, Bin Wang, Guiyuan Jiang, Yanwei Yu, and Junyu Dong. Spatiotemporal-aware trend-seasonality decomposition network for traffic flow forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 11463–11471, 2025.

Hwangyong Choi, Jeongwhan Choi, Jeehyun Hwang, Kookjin Lee, Dongeun Lee, and Noseong Park. Climate modeling with neural advection–diffusion equation: H. choi et al. *Knowledge and Information Systems*, 65(6):2403–2427, 2023a.

Jeongwhan Choi, Hwangyong Choi, Jeehyun Hwang, and Noseong Park. Graph neural controlled differential equations for traffic forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 6367–6374, 2022.

Jeongwhan Choi, Seoyoung Hong, Noseong Park, and Sung-Bae Cho. Gread: Graph neural reaction-diffusion networks. In *International conference on machine learning*, pp. 5722–5747. PMLR, 2023b.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

Prasanjit Dey, Soumyabrata Dev, Angela Meyer, and Bianca Schoen-Phelan. Neuroddaf: Neural dynamic diffusion-advection fields with evidential fusion for air quality forecasting. *arXiv preprint arXiv:2604.01175*, 2026.

Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 32, 2019.

Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 922–929, 2019.

Kethmi Hirushini Hettige, Jiahao Ji, Shili Xiang, Cheng Long, Gao Cong, and Jingyuan Wang. Airphynet: Harnessing physics-guided neural networks for air quality prediction. In *International Conference on Learning Representations*, volume 2024, pp. 45307–45323, 2024.

Jiahao Ji, Jingyuan Wang, Zhe Jiang, Jingtian Ma, and Hu Zhang. Interpretable spatiotemporal deep learning model for traffic flow prediction based on potential energy fields. In *2020 IEEE international conference on data mining (ICDM)*, pp. 1076–1081. IEEE, 2020.

Jiahao Ji, Jingyuan Wang, Zhe Jiang, Jiawei Jiang, and Hu Zhang. Stden: Towards physics-guided neural networks for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 4048–4056, 2022.

Xiaowei Jia, Jared Willard, Anuj Karpatne, Jordan Read, Jacob Zwart, Michael Steinbach, and Vipin Kumar. Physics guided rnns for modeling dynamical systems: A case study in simulating lake temperature profiles. In *Proceedings of the 2019 SIAM international conference on data mining*, pp. 558–566. SIAM, 2019.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Anuj Karpatne, Xiaowei Jia, and Vipin Kumar. Knowledge-guided machine learning: Current trends and future prospects. *arXiv preprint arXiv:2403.15989*, 2024.

Qin Li, Zuocai Zheng, Chenyang Luo, Xuan Yang, Yong Wang, and Yuankai Wu. Joint prediction and understanding of multimodal traffic flow with a bidirectional temporal dynamic spatial hypergraph neural network model. *Transportation Research Part C: Emerging Technologies*, 180: 105358, 2025a.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.

Zetao Li, Zheng Hu, Peng Han, Yu Gu, and Shimin Cai. Ssl-stmformer self-supervised learning spatio-temporal entanglement transformer for traffic flow prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 12130–12138, 2025b.

Yuxuan Liang, Ronghui Xu, Peng Chen, Chenjuan Guo, Aoying Zhou, Lujia Pan, Zhongwen Rao, Bin Yang, et al. Air quality prediction with physics-guided dual neural odes in open systems. In *International Conference on Learning Representations*, volume 2025, pp. 55132–55154, 2025.

Wasen Melhem, Asad Abdi, and Farid Meziane. Traffic detection and forecasting from social media data using a deep learning-based model, linguistic knowledge, large language models, and knowledge graphs. In *KEOD*, pp. 235–242, 2024.

Zhaobin Mo, Rongye Shi, and Xuan Di. A physics-informed deep learning paradigm for car-following models. *Transportation research part C: emerging technologies*, 130:103240, 2021.

DM Moreira, U Rizza, GA Degrazia, C Mangia, T Tirabassi, et al. An analytical air pollution model: development and evaluation. *Contributions to Atmospheric Physics*, 71, 1998.

Daiheng Ni. *Traffic flow theory: Characteristics, experimental methods, and numerical techniques*. Elsevier, 2016.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.

Rongye Shi, Zhaobin Mo, and Xuan Di. Physics-informed deep learning for traffic state estimation: A hybrid paradigm informed by second-order traffic models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 540–547, 2021.

Chao Song, Youfang Lin, Shengnan Guo, and Huaiyu Wan. Spatial-temporal synchronous graph convolutional networks: A new framework for spatial-temporal network data forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 914–921, 2020.

Meng-Ju Tsai, Zhiyong Cui, Hao Yang, Cole Kopca, Sophie Tien, and Yinhai Wang. Traffic-twitter transformer: A nature language processing-joined framework for network-wide traffic forecasting. *arXiv preprint arXiv:2206.11078*, 2022.

Jingyuan Wang, Jiahao Ji, Zhe Jiang, and Leilei Sun. Traffic flow prediction based on spatiotemporal potential energy fields. *IEEE Transactions on Knowledge and Data Engineering*, 35(9):9073–9087, 2022.

Xinlei Wang, Maike Feng, Jing Qiu, Jinjin Gu, and Junhua Zhao. From news to forecast: Integrating event analysis in llm-based time series forecasting with reflection. *Advances in Neural Information Processing Systems*, 37:58118–58153, 2024.

Binqing Wu, Zongjiang Shang, Shiyu Liu, Jianlong Huang, Jiahui Xu, and Ling Chen. Airdde: Multifactor neural delay differential equations for air quality forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 15941–15949, 2026.

Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. *arXiv preprint arXiv:1906.00121*, 2019.

Wanghan Xu, Fenghua Ling, Wenlong Zhang, Tao Han, Hao Chen, Wanli Ouyang, and Lei Bai. Generalizing weather forecast to fine-grained temporal scales via physics-ai hybrid modeling. *Advances in Neural Information Processing Systems*, 37:23325–23351, 2024.

Weiran Yao and Sean Qian. From twitter to traffic predictor: Next-day morning traffic prediction using social media data. *Transportation research part C: emerging technologies*, 124:102938, 2021.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.

Chengyang Zhang, Yong Zhang, Qitan Shao, Jiangtao Feng, Bo Li, Yisheng Lv, Xinglin Piao, and Baocai Yin. Bjtt: A large-scale multimodal dataset for traffic prediction. *IEEE Transactions on Intelligent Transportation Systems*, 25(11):18992–19003, 2024.

Xukai Zhang, Shuliang Wang, Guangyin Jin, Ziqiang Yuan, Hanning Yuan, and Sijie Ruan. Eulerian neural network informed by chemical transport for air quality forecasting. *Advances in Neural Information Processing Systems*, 38:33576–33601, 2026.