



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Cheng, J., Homem de Almeida Correia, G., Cats, O., & Sharif Azadeh, S. (2026). T-STAR: A context-aware transformer framework for short-term probabilistic demand forecasting in dock-based shared micro-mobility. *Transportation Research Part C: Emerging Technologies*, 192, Article 105899. <https://doi.org/10.1016/j.trc.2026.105899>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.



T-STAR: A context-aware transformer framework for short-term probabilistic demand forecasting in dock-based shared micro-mobility

Jingyi Cheng ^{*}, Gonçalo Homem de Almeida Correia ¹, Oded Cats ¹,
Shadi Sharif Azadeh ¹

Transport and Planning, Delft University of Technology, Delft, Netherlands

ARTICLE INFO

Keywords:

Shared micro-mobility
Short-term demand forecasting
Demand uncertainty quantification

ABSTRACT

Reliable short-term demand forecasting is essential for managing shared micro-mobility services and ensuring responsive, user-centered operations. This study introduces T-STAR (Two-stage Spatial and Temporal Adaptive contextual Representation), a novel transformer-based probabilistic framework designed to forecast station-level bike-sharing demand at a 15-minute resolution. T-STAR addresses key challenges in high-resolution forecasting by disentangling consistent demand patterns from short-term fluctuations through a hierarchical two-stage structure. The first stage captures coarse-grained hourly demand patterns, while the second stage improves prediction accuracy by incorporating high-frequency, localized inputs, including recent fluctuations and real-time demand variations in connected metro services, to account for temporal shifts in short-term demand. Time series transformer models are employed in both stages to generate probabilistic predictions. Extensive experiments using Washington D.C.'s Capital Bikeshare data demonstrate that T-STAR outperforms existing methods in both deterministic and probabilistic accuracy. The model exhibits strong spatial and temporal robustness across stations and time periods. A zero-shot forecasting experiment further highlights T-STAR's ability to transfer to previously unseen service areas without retraining. These results underscore the framework's potential to deliver granular, reliable, and uncertainty-aware short-term demand forecasts, which enable seamless integration to support multimodal trip planning for travelers and enhance real-time operations in shared micro-mobility services.

1. Introduction

Shared micro-mobility services have become integral to urban transportation networks, offering flexible, low-emission options for short-distance and first- and last-mile travel (Yang et al., 2019). Among them, bike-sharing services have seen rapid global adoption, with large-scale networks deployed across major cities to enhance accessibility (Ma et al., 2020). As a key component of Mobility-as-a-

^{*} Corresponding author.

E-mail addresses: j.cheng-1@tudelft.nl (J. Cheng), g.correia@tudelft.nl (G. Homem de Almeida Correia), o.cats@tudelft.nl (O. Cats), s.sharifazadeh@tudelft.nl (S. Sharif Azadeh).

¹ Given his/her/their role as Editorial Board Member and Associate Editor, Gonçalo Correia, Oded Cats and Shadi Sharif Azadeh had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to a journal editor.

<https://doi.org/10.1016/j.trc.2026.105899>

Received 15 July 2025; Received in revised form 4 June 2026; Accepted 24 July 2026

Available online 30 July 2026

0968-090X/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Service (MaaS) ecosystems, bike-sharing services interact closely with public transport as well as other on-demand modes, facilitating seamless multimodal journeys within complex urban environments (Cleophas et al., 2019). However, their operational efficiency is often constrained by spatio-temporal imbalances between supply and demand, leading to an inefficient resource allocation and service disruptions. Accurate short-term demand forecasting is therefore essential for supporting real-time operational decisions, including fleet rebalancing and multimodal trip recommendations (Zhao et al., 2025; Chiariotti et al., 2020; Zhang et al., 2025). With shared mobility becoming increasingly integrated into multimodal transport systems, predictive tools that can handle these complexities are essential to support reliable, efficient, and user-centered operations.

Despite encouraging progress in short-term demand forecasting for bike-sharing systems, several important challenges remain that limit their operational effectiveness. Demand is shaped by complex, interrelated factors, ranging from exogenous influences like weather variability, public transport service disruptions, and other sudden local events to endogenous constraints such as fleet availability and inter-station demand interactions (Wilkesmann et al., 2023; Jia et al., 2021; Huang et al., 2021). Although some studies have started incorporating multi-source contextual data (Xu et al., 2023), integrating these dynamic signals in a structured and interpretable way remains technically demanding. Moreover, shared micro-mobility studies emphasize that high-resolution forecasts, such as at the station level and in short intervals, are important for effective operations (Chiariotti et al., 2020). However, such fine-grained predictions introduce modeling challenges due to data sparsity, intermittency, and noise, particularly if they are to be performed at intervals of 15 minutes or less (Kourentzes, 2013). Additionally, while probabilistic forecasting offers important advantages for uncertainty quantification and risk-aware planning (Zhao et al., 2025), related efforts in the forecasting literature remains limited. Despite growing demand for such capabilities, it remains largely unknown how to generate high-resolution, uncertainty-aware forecasts that effectively integrate heterogeneous contextual signals while maintaining robustness to sparsity and noise (Xu et al., 2018; Gammelli et al., 2022). These gaps underscore the need for context-aware probabilistic forecasting frameworks that can support effective real-time decision-making related to urban shared micro-mobility services.

To address the challenges of fine-grained, uncertainty-informed demand forecasting in dock-based bike-sharing systems, this study proposes the Two-stage Spatial and Temporal Adaptive contextual Representation (T-STAR) framework. While prior research has explored station-level and probabilistic forecasting, existing methods often struggle to balance the inherent noise and sparsity typical of ultra-short-term intervals (Liang et al., 2023; Xu et al., 2023; Gammelli et al., 2022). Unlike single-stage models that may conflate stochastic noise with recurring patterns, T-STAR distinguishes itself by employing a hierarchical architecture designed to explicitly navigate this resolution-uncertainty trade-off. Built on a Transformer backbone, the framework acts as a global predictor that generalizes across stations to support scalable, real-time forecasting. In the first stage, the model captures broad demand expectations using hourly aggregated network features; the second stage then refines these representations with fine-grained, station-specific contextual signals to model high-frequency fluctuations. This hierarchical design allows T-STAR to disentangle recurrent temporal patterns from localized disruptions, providing robust, uncertainty-aware predictions specifically optimized for the 15-minute resolution required for proactive rebalancing operations of urban shared micro-mobility services.

We demonstrate and validate the Transformer-based T-STAR framework through 15-minute, station-level demand forecasting experiments using real-world data from Washington D.C.'s Capital Bikeshare system. The evaluation covers both deterministic and probabilistic forecasting accuracy. We analyze spatial robustness by comparing model performance across stations with varying demand profiles, and assess temporal stability by examining prediction quality at each time step, including intervals under extreme conditions and abnormal observations. Additionally, a zero-shot forecasting experiment on an unseen service area is conducted to test model's adaptability and generalization capability. Collectively, the results demonstrate that T-STAR delivers accurate and robust predictions, making it a reliable and scalable tool for large-scale urban shared micro-mobility systems. A coding and application demo of the Transformer-based T-STAR model is provided through a public GitHub repository associated with this study¹.

The remainder of this study is organized as follows. Section 2 reviews the related literature and identifies the gaps. Section 3 introduces the formulation of probabilistic short-term demand forecasting problem for a dock-based bike-sharing system. We then explain the proposed T-STAR framework in Section 4, detailing its hierarchical architecture and contextual embedding strategy. Section 5 describes the case study, evaluation metrics, and design of benchmarking experiments. Section 6 analyzes the results of the forecasting experiments and discuss the key findings. Finally, Section 7 concludes this study and suggests future research directions.

2. Literature review

2.1. Demand forecasting approaches for one-way bike-sharing systems

State-space models (SSMs) have long served as the foundation for classical time series forecasting, offering a parametric framework that decomposes demand series into components such as trend, seasonality, and autoregressive error terms (Durbin and Koopman, 2012). Among these, Seasonal Auto-Regressive Integrated Moving Average (SARIMA) models have been frequently applied to forecast bike-sharing demand (Yoon et al., 2012; Xu et al., 2018). Despite their interpretability, SSMs are ill-suited for short-term bike-sharing scenarios, where demand is highly variable, zero-inflated, and non-negative (Lim et al., 2022). Furthermore, the need to model each station's dynamic demand time series individually poses scalability challenges for large-scale networks, making SSMs computationally expensive to be adapted for urban bike-sharing systems (Lim et al., 2022).

¹ <https://github.com/RinaPiggy/T-STAR>

To capitalize on the availability of large-scale bike-sharing demand data, recent studies have increasingly adopted machine learning (ML) approaches for forecasting. Compared to traditional state-space models, ML methods can efficiently process high-dimensional inputs and incorporate diverse contextual features to capture the complex, dynamic patterns of urban micro-mobility demand. Non-parametric tree-based models, such as Random Forest (RF) and eXtreme Gradient Boosting (XGBoost), have been widely applied in this context (Gao and Chen, 2022; Sathishkumar et al., 2020; Leem et al., 2024). Using data from the Seoul bike-sharing system, Gao and Chen (2022) and Sathishkumar et al. (2020) demonstrated that tree-based regressors can generate accurate city-level forecasts by leveraging multiple external inputs. Deep learning has shown promising performance in forecasting intermittent and highly variable demand series (Kourentzes, 2013). To better capture the complex temporal and spatial dynamics in bike-sharing networks, recent studies have increasingly focused on neural networks for demand forecasting (Zhang et al., 2018a; Xu et al., 2018; Lim et al., 2022; Gammelli et al., 2022; Lin et al., 2018; Li et al., 2023). Temporal dependencies have been modeled using recurrent architectures, such as Recurrent Neural Networks (RNNs) and long short-term memory networks (LSTMs). Zhang et al. (2018b) designed an LSTM-based predictor to learn long-term pattern while adaptively accounting for short-term fluctuations. Similarly, Xu et al. (2018) applied LSTMs for short-term inflow and outflow prediction in a large-scale free-floating system. To exploit shared spatial patterns across the network and model spatial dependencies among stations or service zones, studies have looked into convolutional neural networks (CNNs), graph neural networks (GNNs), and graph convolutional networks (GCNs). For example, Lin et al. (2018) proposed a GCN-based predictor that uncovers heterogeneous pairwise correlations between stations. Li et al. (2023) proposed a hybrid CNN-LSTM model to jointly capture spatiotemporal demand patterns.

Transformer models, originally developed for natural language processing, have demonstrated promising performance in time series forecasting, proving particularly effective at capturing the spatiotemporal complexity characteristic of large-scale shared mobility systems (Zerveas et al., 2021; Hu and Xiong, 2023; Li et al., 2024a). Unlike the sequential structure in RNNs, Transformer models leverage self-attention to capture both short- and long-range temporal dependencies, enabling more flexible and efficient learning of complex seasonality and non-linear demand patterns across stations or service zones (Hu and Xiong, 2023; Li et al., 2024a). Their parallel processing structure facilitates the seamless integration of diverse multivariate inputs, such as weather, time, and socioeconomic indicators. Models like the Interpretable Hierarchical Transformer (IHTF) (Hu and Xiong, 2023) and the Temporal Fusion Transformer (TFT) (Lim et al., 2021) enhance this capacity through variable selection networks and multi-head attention mechanisms that dynamically identify the most relevant features at each time step. Acting as global predictors, these models learn shared temporal-spatial representations across the network while forecasting adaptively to the unique dynamics of individual locations.

Transformers offer, next to their modeling capacity, several practical advantages over traditional sequential models. The self-attention architecture enables full parallelization of historical data processing, leading to significantly faster training on large datasets compared to RNNs or LSTMs, which rely on sequential computation (Vaswani, 2017). Transformers also demonstrate strong generalizability and have been successfully applied in cross-location and zero-shot forecasting tasks in mobility contexts, making them suitable for dynamic systems where new service areas are frequently introduced (Solatorio, 2023). Moreover, Transformer supports continuous fine-tuning from pre-trained models, enabling rapid adaptation to evolving urban patterns with limited new data (Solatorio, 2023). These scalability, adaptability, and transferability properties make Transformer a promising forecasting solution for data-rich, yet rapidly changing shared micro-mobility services.

While most past studies on shared micro-mobility have focused on point forecasting methods, recent research has increasingly emphasized the importance of probabilistic forecasting. Accurately probabilistic demand forecasting is critical for risk-aware management, as it quantifies the uncertainty inherent in fluctuating demand rather than providing single-point estimates. As reviewed by Cheng et al. (2024), these models equip managers to mitigate potential adverse conditions by evaluating the full distribution of possible outcomes (Zhao et al., 2025; Xu et al., 2023; Gammelli et al., 2022). Urban modes like ride-hailing, public transit, and bike-sharing share foundations in spatio-temporal dependence and heterogeneous context. Related work addressing demand uncertainty across various mobility sectors can be broadly categorized into three methodological frameworks: quantile-based estimation, generative and latent-variable modeling, and parametric density estimation.

Quantile-based methods bypass the need for explicit distributional assumptions by directly estimating empirical quantiles or prediction intervals. These approaches are particularly valued for their robustness against outliers and non-standard noise profiles. For instance, Quantile Regression Forests (QRF) and Temporal Fusion Transformers (TFT) have been successfully deployed to estimate demand quantiles for both on-demand ride-hailing and scheduled transit services (Lei and Wasserman, 2014; Cheng and Sharif Azadeh, 2025; Wang et al., 2025). While these models provide flexible interval estimates, they often lack a closed-form probability density function, which can complicate their integration into certain stochastic optimization pipelines.

To capture more complex, high-dimensional uncertainties, recent research has pivoted toward Generative Models and Neural Processes (NPs). These frameworks typically learn a mapping from input features to a latent representation z , which represents the underlying stochasticity of the demand process. An empirical distribution of future demand is performed by repeated sampling of z . Notable applications include VAE-based STGCNs for ride-hailing (Peng et al., 2025) and Bidirectional Spatial-Temporal Transformer NPs for ride-sourcing (Li et al., 2024b). Furthermore, Xu et al. (2023) extended this logic to dock-based bike-sharing by integrating Transformer encoders to refine spatiotemporal feature extraction, while Gammelli et al. (2022) introduced a Variational Poisson RNN to estimate posterior rates for bike-sharing pickup and drop-off activities. However, latent compression mechanisms of generative modeling often dismiss high-frequency volatility as noise through posterior collapse (Dai et al., 2020; Li et al., 2022). Consequently, prioritizing global distribution over sharpness produces conservative forecasts that cannot accurately capture the precise transitions required for high-resolution, sparse demand patterns (Choi and Min, 2022). Beyond the shared mobility domain, uncertainty quantification in spatio-temporal forecasting has seen rapid methodological advancement in the broader traffic forecasting literature (Qian et al., 2023a,b). Qian et al. (2023b) proposed a unified framework, DeepSTUQ, that jointly estimates both aleatoric and epistemic

uncertainty in traffic flow forecasting by combining Monte Carlo dropout, adaptive weight averaging, and conformal calibration, providing a comprehensive perspective on uncertainty decomposition strategies for spatio-temporal systems. While such approaches are developed primarily for dense, continuously observed traffic flow data under Gaussian distributional assumptions, they underscore the growing importance of reliable uncertainty estimation in spatio-temporal forecasting, and motivate the need for probabilistic frameworks specifically tailored to the sparse, zero-inflated demand characteristics of shared micro-mobility systems.

In contrast, Parametric Density Estimation methods assume that demand follows a predefined probability distribution and utilize deep learning to estimate its parameters. Although this approach carries the risk of distributional misspecification, it provides a mathematically grounded framework that is highly tractable for downstream stochastic optimization problems. Based on a Normal distribution assumption, Li et al. (2020) utilized GNNs for robust public transit forecasting. Prob-GNN was proposed by Wang et al. (2024) to model multi-modal demand via truncated Gaussian distributions. Lim et al. (2022) investigated the impact of different stochastic priors by fitting normal, truncated normal, and negative binomial distributions with DeepAR to generate probabilistic bike-sharing demand forecasts.

2.2. The granularity of demand forecasting problems in shared micro-mobility

Demand forecasting in shared micro-mobility systems can be categorized based on its spatial and temporal granularity, which are suitable to assist different operations. Coarse-grained forecasts, such as city-wide daily demand predictions, are suitable for long-term and strategic planning decisions. In contrast, spatially and temporally fine-grained forecasts, such as station-level per 15-minute forecasting, are essential for operational decisions that require both timely and accurate predictions.

The literature on demand forecasting for bike-sharing services can be categorized by the spatial granularity of the prediction task: city-level, zonal or cluster-level, and station-level forecasting. Early studies addressed this problem at the city level, predicting the total number of rental requests across the entire network for future time periods (Giot and Cherrier, 2014). To better capture spatial heterogeneity in usage patterns, subsequent work shifted toward regional-level forecasting. For example, some studies applied clustering algorithms such as K-means to group stations together based on similarity, then forecasts aggregate demand for each cluster (Li et al., 2015). Building on this, Chen et al. (2016) proposed dynamic clustering frameworks that incorporate real-time contextual signals, such as weather or events, to adaptively redefine spatial groupings over time, improving the model responsiveness to changes in demand distribution. These cluster-then-predict approaches aim to enhance forecast accuracy by aggregating sparse station-level signals within semantically or contextually coherent regions, thereby reducing noise and data sparsity. To overcome the loss of local detail inherent in cluster-based methods (Feng and Liu, 2024), recent studies have focused on station-wise demand forecasting, leveraging high-resolution data and advanced modeling techniques to provide more precise and actionable insights (Xu et al., 2023; Lim et al., 2022; Gammelli et al., 2022; Feng and Liu, 2024).

In short-term demand forecasting for dock-based bike-sharing services, temporal granularity plays a crucial role in determining the operational applicability of predictive insights. Most studies adopt an hourly aggregation level to balance modeling complexity and practical relevance (Xu et al., 2023; Li et al., 2023; Lin et al., 2018). Some researchers have explored coarser intervals, such as splitting the day into peak and off-peak periods or forecasting in 4-hour blocks, to capture broader demand trends (Chen et al., 2016; Liang et al., 2023). However, for operations with short lead times, such as real-time and continuous rebalancing (Chiariotti et al., 2020), higher temporal resolution is essential. To this end, several studies have investigated ultra-short-term forecasting at 15- or 20-minute intervals (Xu et al., 2018; Gammelli et al., 2022; Ma et al., 2022; Feng and Liu, 2024), enabling a more proactive and precise management of bike-sharing systems. Focusing on a zone-level forecasting problem, Xu et al. (2018) analyzed the tradeoff between temporal granularity and predictive accuracy by evaluating multiple interval lengths (10, 15, 20, and 30 minutes) across both state-space and neural network models. Their findings suggest that, while shorter intervals provide more timely information, longer intervals yield lower forecast error due to reduced noise and variability.

Studies have shown that spatially and temporally fine-grained predictions are crucial for enabling real-time operations in bike-sharing systems, such as dynamic fleet rebalancing, where vehicle routes and allocations are continuously updated (Chiariotti et al., 2020), and multimodal trip recommendations that account for anticipated bike availability (Zhang et al., 2025). However, generating accurate fine-grained forecasts is inherently challenging. Station-level demand is often sparse and highly volatile, especially during off-peak hours, and historical records may include irregular shocks caused by supply constraints, such as empty or full docking stations (Chen et al., 2016; Ashqar et al., 2022; Xu et al., 2018). Although high-resolution forecasts offer greater operational responsiveness, they tend to be more challenging due to noise and intermittency in demand data. In contrast, coarser temporal or spatial aggregation improves stability but sacrifices the detail needed for localized decision-making. To address this tradeoff, Chen and Jiang (2020) proposed a Bayesian hierarchical model that captures long-term demand shifts at the city level in the upper layer, while modeling daily station-level demand in the lower layer, incorporating dynamic features such as weather and holidays. Accounting for spatial dependencies across different levels of granularity, Khalesian et al. (2024) proposed a hierarchical reconciliation framework that refines zone-level traffic demand forecasts using aggregated regional predictions. Although the study of hierarchical modeling in bike-sharing demand forecasting remains relatively limited, this direction appears promising for integrating coarse- and fine-grained signals to enhance the accuracy and robustness of ultra-short-term, station-level predictions.

2.3. Determinants of bike-sharing demand

Accurate short-term demand forecasting in bike-sharing systems relies not only on historical usage patterns but also on a variety of contextual factors that influence user demand. These factors capture both recurring patterns and sudden fluctuations, enabling models

to effectively reflect the spatial, temporal, and environmental dynamics of the system. In this subsection, we review key contextual features discussed in the literature, including temporal patterns, fleet availability and station capacity, connectivity to public transit, weather conditions, and latent events.

Like other urban transit services, bike-sharing systems exhibit strong cyclic demand patterns that align with daily commuting routines, typically peaking during morning and evening rush hours (Zhao et al., 2019; Faghieh-Imani et al., 2014; Cats and Ferranti, 2022). To capture these recurring trends, seasonal features, such as hour of the day and day of the week, are commonly included as fundamental inputs in short-term demand forecasting models. However, these temporal patterns can vary significantly across stations due to local land use and neighborhood characteristics. For instance, stations at business and residential areas often experience demand peaks during commuting hours, while those located within pubs, restaurants, or entertainment districts tend to see higher demand in the evenings and on weekends.

Fleet availability and station capacity are key factors influencing short-term bike-sharing demand. High-capacity stations, typically located in areas with high passenger traffic, tend to attract greater usage due to the increased likelihood of bike availability and open docks for returns (Rixey, 2013). In contrast, low-capacity stations often experience suppressed rental activity, especially during peak periods, as limited bikes or docking spaces restrict access. These constraints give rise to censored demand, where true user intent is unobservable when stations are empty or full. As a result, discrepancies between actual and recorded demand can introduce bias and reduce forecasting accuracy. To mitigate this, recent studies have explicitly modeled the impact of censored demand, enhancing prediction accuracy and reliability in demand forecasting (Albiński et al., 2018; Huang et al., 2021).

The interaction between bike-sharing and public transit has been widely studied within the broader context of Mobility on Demand (MoD) and Mobility as a Service (MaaS). This relationship can be either complementary or competitive. For example, passengers may shift away from public transit when station-based bike-sharing offers a faster, more affordable, or more direct alternative (Shaheen and Martin, 2015). However, in lower-density or suburban areas, bike-sharing often plays a complementary role by bridging first- and last-mile gaps, thereby increasing the attractiveness of multimodal travel. Empirical studies have consistently shown that bike-sharing stations located near metro stations experience higher usage (Wang and Noland, 2021; Ma and Knaap, 2019; Tarpin-Pitre and Morency, 2020; Tushar and Buehler, 2024). There is a positive correlation between public transit passenger flows and nearby bike-sharing demand, especially during peak hours when commuters frequently combine multiple modes of transport. During major transit disruptions, bike-sharing systems also help sustain network resilience by serving as alternative transport modes (Cheng et al., 2021). For example, during the 2016 WMATA metro closure, bike-sharing demand surged by 191% in Washington D.C. (Jia et al., 2021). Reflecting this interdependence, recent forecasting models have begun incorporating public transit operations as exogenous features (Lv et al., 2021; Li et al., 2023). These findings underscore the operational and predictive value of integrating multimodal connectivity into short-term demand forecasting frameworks.

Prior studies have shown that the demand for urban bike-sharing services is greatly influenced by external conditions (Zhao et al., 2019; Wilkesmann et al., 2023). Among these, weather conditions play a particularly significant role in shaping user behavior. For example, Gebhart and Noland (2014) found that cold temperatures, rainfall, and high humidity are associated with a decline in Capital Bikeshare trips. Conversely, warm, dry, and sunny days tend to correspond with increased demand, reflecting users' greater willingness to travel by bike under favorable weather conditions (Gebhart and Noland, 2014; Wang and Noland, 2021).

Beyond commonly modeled factors, short-term bike-sharing demand can also be influenced by latent influences such as special events in the city, which are rare and has unclear impacts on the bike-sharing usage in the short run (Chen et al., 2016; Jia et al., 2021). These events often occur irregularly and may not be directly captured by input features, yet they can induce sudden, localized shocks in demand. To account for such latent influences, recent studies have explored the use of neural networks to learn latent spatiotemporal representations from historical data (Li et al., 2019; Ham et al., 2021). By leveraging deep learning models, these approaches infer the underlying patterns associated with unobserved factors, improving the ability of the predictors to generalize and respond to demand shifts that are not entirely explained by the observable variables.

2.4. Summary of literature and research gaps

Table 1 summarizes key studies closely aligned with this work on short-term bike-sharing demand forecasting, categorizing them by spatial and temporal granularity, forecasting type, the adopted predictive model, and the contextual inputs considered. While the field has seen significant progress, especially in moving from coarse, deterministic models to finer-grained and probabilistic approaches, several important research gaps still remain unaddressed:

1. *Lack of structured inclusion of contextual factors:* While weather conditions and temporal characteristic of demand patterns are often considered in demand forecasting for bike-sharing services, existing models often do not systematically incorporate dynamic contextual information, such as the station facility properties (i.e., its holding capacity and the multimodal connectivity to other mobility services (Liang et al., 2023) and the disruptions arising from censored demand (Xu et al., 2023; Gammelli et al., 2022) or external latent events which are not directly observable in available data.
2. *Challenges at high spatio-temporal resolution:* Existing studies vary widely in spatial and temporal resolution. While some models address station-level forecasting (Liang et al., 2023; Xu et al., 2023), few are equipped to handle ultra-short-term (e.g., 15-minute) predictions at such fine spatial resolution, which are more prone to noise, sparsity, and intermittency (Gammelli et al., 2022). It highlights the need for forecasting methods capable of generating accurate and robust predictions under high-resolution, volatile conditions. This is essential for supporting effective and proactive real-time operations in bike-sharing systems.

Table 1

Selected representative studies on short-term bike-sharing demand forecasting, categorized by spatial/temporal granularity, forecasting type, predictive model, and contextual inputs.

Reference	Spatial Granularity	Temporal Granularity	Forecasting Types	Predictive Model	Contextual Inputs
Lin et al. (2018)	station	hour	deterministic	GCN with Data-driven Graph Filter	Temporal
Liang et al. (2023)	station	4-hour	deterministic	Multi-Relational GNN	Temporal, Public Transport
Xu et al. (2018)	zone (TAZ)	10/15/20/30min	deterministic	LSTM	Temporal, Weather, Land Use
Feng and Liu (2024)	station	20/40/60min	deterministic	GCN with Group Multi-head Self-Attention	Temporal, Weather, Land Use
Xu et al. (2023)	station	hour	probabilistic	Transformer-Encoder-based Neural Process	Temporal, Censored Demand
Gammelli et al. (2022)	station	15/30/40min	probabilistic	Variational Poisson RNN	Temporal, Weather, Censored Demand
This study	station	15min	probabilistic	Transformer-based T-STAR	Temporal, Weather, Facility, Public Transport, Censored Demand

Table 2

A list of the key notations used throughout this study.

Notation	Description
$i \in I$	i denotes a dock-based bike-sharing station; I denotes the stations within the service network.
η	A fixed time interval used to discretize the timeline of the forecasting problem. It determines the temporal resolution at which demand is aggregated and predictions are made (e.g., $\eta = 15$ minutes or $\eta = 1$ hour).
t	The general index for time intervals, denoting the temporal order on a demand time series. There are in total T intervals in the time series generated from data.
y_t^i	The number of pickup or drop-off demand observed at station i during the t th time interval.
$\mathbf{y}^i$	The target demand time series for station i is $\mathbf{y}^i = (y_1^i, y_2^i, \dots, y_T^i)$, $y_t^i \in \mathbb{Z}$.
$\mathbf{X}_t^i$	The contextual vector associated with station i and time t . This includes: static station-level attributes $\mathbf{x}^i$, time-varying global features $\mathbf{x}_t$, shared across all stations, and dynamic, station-specific features $\mathbf{x}_t^i$ that vary over time.
h	Hour index on the hourly aggregated demand time series.
q	Quarter index on the per-15-minute aggregated demand time series.
H	The forecasting horizon of a model, i.e. the number of future time intervals to be predicted.
V	The historical look-back window of the forecasting model, which defines the number of past time steps considered as input for the predictions.

3. *Limited adoption of probabilistic forecasting methods:* Meanwhile, while probabilistic forecasts are essential for supporting risk-aware operations (Gast et al., 2015; Zhang et al., 2025), most existing studies have focused on point predictions (Liang et al., 2023; Feng and Liu, 2024), with limited exploration of uncertainty quantification (Xu et al., 2023; Gammelli et al., 2022). As a result, current models often fall short in answering operationally relevant questions under uncertain conditions, such as the likelihood of bike availability or expected waiting time. While probabilistic forecasting has advanced across the broader urban mobility landscape, dock-based systems remain distinct. Unlike the spatial fluidity of ride-hailing or the schedule-dependency of transit, dock-based micro-mobility is fundamentally constrained by fixed station capacities. This physical constraint induces censored demand, where true user intent is obscured by station stockouts or saturation. Consequently, there is a critical need for predictive methods designed to reconcile these localized constraints with the sparse, zero-inflated patterns inherent in high-granularity station data for shared micro-mobility. Furthermore, current research lacks efficient frameworks capable of maintaining accuracy amidst high volatility while simultaneously capturing inter-modal dependencies. Bridging this gap requires models that ensure long-term generalizability across evolving service landscapes and enable real-time adaptation for dynamic uncertainty reassessment (Cheng et al., 2024).

To address these gaps, we introduce transformer-based T-STAR, a two-stage framework for probabilistic, station-level demand forecasting in dock-based bike-sharing systems. This forecasting framework operates at a 15-minute temporal resolution, integrating a rich set of contextual features to capture the complex demand dynamics under real-time conditions.

3. Problem formulation

This section presents the probabilistic demand forecasting problem for dock-based bike-sharing services. We begin by outlining the problem context, followed by key definitions of this study. Table 2 provides an overview of the notations used in this study.

3.1. Dock-based bike-sharing services

A dock-based bike-sharing service operates through a network of stations I , each is identified by a unique ID. Each station $i \in I$ is equipped with a fixed number of parking docks C^i , defining its maximum bike-holding capacity. At any time t , the operational status of a station is characterized by two dynamic metrics, fleet availability and parking availability. Fleet availability represents

the number of bikes available for pick-up, bounded between 0 and the station's hold capacity C^i . Parking availability indicates the number of vacant docks at the station.

Dock-based bike-sharing services typically provide 24/7 on-demand access for users to pick up and drop off bikes. Users may approach a station from various origins, such as shopping centers, campuses, or by transferring from other transit services as part of a multimodal journey. A typical trip starts when a user unlocks a bike at one station and ends when the bike is returned at another. Historical records capture completed trips, including timestamps and the origin and destination stations. A pickup is not possible when no bikes are available, and a drop-off can be denied when the parking is full. In such situations, users may reroute to an alternative station or cancel their trip. These unfulfilled trips are not logged in the historical data, which only records completed journeys, limiting insights into unmet demand.

As reviewed in Section 2.3, bike-sharing demand is influenced by a combination of internal and external factors. These include recurrent usage patterns, weather conditions, and the operations of nearby public transit services. In this study, contextual features are integrated into the forecasting framework and categorized based on whether they are exogenous to the system and whether they are dynamic or static. Specifically, to account for multimodal connectivity, a bike-sharing station i is considered connected to a nearby public transit service m if its location L^i is within a threshold distance ν (Ma et al., 2018).

3.2. Probabilistic short-term demand forecasting for bike-sharing services

To ensure efficient service and customer satisfaction, bike-sharing providers require timely forecasts of both pickup and drop-off demand at each station. These activities directly impact station-level inflow and outflow, dynamically affecting fleet availability in the short term. Robust operational planning further benefits from probabilistic forecasts that quantify uncertainty in future demand. In this study, we propose a probabilistic forecasting approach that integrates historical trip records with station-level attributes and exogenous contextual features to predict short-term pickup and drop-off demand at individual stations.

We define a fixed time interval η (e.g., 15 minutes or 1 h) as the temporal resolution of the forecasting task. This interval determines both the aggregation of demand data and the frequency of prediction updates. For each station i , the historical trip records are processed to generate a demand time series $y_t^i = (y_1^i, y_2^i, \dots, y_T^i)$, where each $y_t^i \in \mathbb{Z}$ represents the number of demand (either pickups or drop-offs) that occurred during the t -th time interval. Depending on the specific forecasting objective, y_t^i may refer to the sequence of pickup counts or drop-off counts. Each target value y_t^i is associated with a multivariate feature vector X_t^i , representing contextual information available at time t to support forecasting.

The goal of the probabilistic (multi-step ahead) demand forecasting problem is to estimate the conditional distribution of future demand over the prediction horizon, given historical observations and available contextual information. Given a future horizon H , the task is to estimate,

$$p_D(y_{t+1:t+H} | X_{t-V:t}, y_{t-V:t}, X_{t+1:t+H}^*), \quad (1)$$

where,

- $y_{t+1:t+H}$ denotes the demand to be predicted over the next H time intervals;
- $X_{t-V:t}$ and $y_{t-V:t}$ represent the historical demand and contextual features observed during the past V intervals (i.e., the observable input window);
- $X_{t+1:t+H}^*$ denote the knowledge about the contextual information over the future horizon. In this study, we assume temporal information (like day of the week and hour of the day) and the estimated weather conditions are available for the future forecasting horizon as inputs.

Probabilistic forecasting models assume that future demand follows an underlying distribution D and aim to characterize this distribution rather than predict a single point estimate. By providing a distributional view of future outcomes, such as predicted quantiles or prediction intervals, probabilistic forecasting captures the inherent uncertainty in demand. These distributional insights enable uncertainty-aware decision-making, which is critical for robust planning and can be directly integrated into stochastic optimization frameworks.

4. Methodology

This section presents the Transformer-based Two-stage Spatial and Temporal Adaptive contextual Representation (T-STAR) developed for short-term demand probabilistic forecasting in dock-based shared micro-mobility systems. T-STAR builds on the hierarchical modeling of global and local random effects in multilevel time series forecasting (Wang et al., 2019), as well as recent advances in hierarchical prediction frameworks (Chen and Jiang, 2020; Khalesian et al., 2024). Its two-stage structure is designed to disentangle stable demand patterns from short-term fluctuations. At each stage of T-STAR, a multivariate time series Transformer models sequential demand in relation to the selected contextual features. Through this hierarchical integration of contextual information, T-STAR effectively captures both global structures and local dynamics, enabling accurate and uncertainty-aware forecasts.

Unlike multi-layer forecasting models such as the Temporal Fusion Transformer or ResNet-based architectures, which operate on time series data at a single temporal resolution (Lim et al., 2021; Shen et al., 2020), T-STAR introduces granularity-aware hierarchical supervision, where forecasts at each temporal level are trained independently before being fed to the downstream. This design reduces the risk of error propagation from noisy high-resolution time series and allows fine-grained predictions to benefit from the stability of aggregated representations learned at coarser scales, enhancing performance under sparse and volatile demand conditions.

Section 4.1 details the architecture and design rationale of the two-stage framework, explaining how contextual embeddings from Stage 1 are refined in Stage 2 to improve predictive accuracy. Section 4.2 introduces the Time Series Transformer, the selected base predictor in T-STAR, and explains how it generates probabilistic demand predictions while accommodating multivariate inputs.

4.1. A two-stage spatial and temporal adaptive contextual representation (T-STAR) framework

The T-STAR framework addresses the challenge of forecasting both stable demand patterns and short-term fluctuations by decomposing the task into two sequential stages. This two-stage design allows the model to align the temporal granularity of input features with their corresponding predictive targets, enabling the learning of more meaningful contextual representations while improving both predictive accuracy and computational efficiency. Specifically, Stage 1 models the underlying regular demand trends by leveraging coarse-grained contextual features, such as hourly weather conditions, time-of-day, and day-of-the-week features, to generate probabilistic estimate of hourly demand. Stage 2 then models short-term variations by focusing on localized deviations from this baseline, incorporating fine-grained temporal inputs and dynamic signals of recent supply-demand imbalances. This separation is designed to mitigate the impact of data sparsity in high-frequency time series and support more targeted use of context at each stage. Fig. 1 illustrates the hierarchical feature processing pipeline of the T-STAR framework, which generates probabilistic pickup and drop-off demand forecasts for the next 15-minute interval based on historical observations and multivariate contextual inputs for each station.

Stage 1: Modeling hourly demand patterns

In the first stage of the T-STAR framework, the model estimates the underlying regular demand trend by forecasting future hourly demand based on recent demand history and coarse-grained contextual information. Specifically, the model takes as input the past V_1 observations on the hourly demand sequence $y_{h-V_1:h}^i$ and the associated contextual features $X_{h-V_1:h}^i$ for station i , along with known or externally estimated future features $X_{h+1:h+H_1}^i$ over the forecasting horizon of length H_1 . The subscript in V_1 and H_1 indicates that these hyperparameters pertain to Stage 1 of the T-STAR framework. The output is an estimated conditional distribution over the future hourly demand for station i , denoted by $\hat{y}_{h+1:h+H_1}^i$, and is defined as:

$$\hat{y}_{h+1:h+H_1}^i \sim F_1\left(y_{h-V_1:h}^i, X_{h-V_1:h}^i, X_{h+1:h+H_1}^{*i}\right), \quad (2)$$

where F_1 is the Stage 1 forecasting model. Note that, T-STAR framework is model-agnostic and allows for different base predictors, such as Time Series Transformer and XGBoost. $y_{h-V_1:h}^i$ is the historical hourly demand series. $X_{h-V_1:h}^i$ includes temporal and environmental covariates aligned with those observations. And $X_{h+1:h+H_1}^{*i}$ contains forward-looking contextual features used to inform the forecast, which are either known in advance or estimated from external sources.

For the experiments in this study, we set $V_1 = 24$ to incorporate one full day of past hourly observations and $H_1 = 1$ to predict demand for the next hour. As the targeted forecasting problem in this study focuses on predictions for the next 15-minute, predicting the one-hour ahead demand from Stage 1 is sufficient. The value of H_1 can be adjusted as needed for longer-term forecasting problems. To capture regular demand patterns, the contextual factors, denoted by $X_{h-V_1:h}^i$ and $X_{h+1:h+H_1}^{*i}$, include temporal variables hour-of-the-day, day-of-the-week, and a binary indicator for national holidays. These features help describe the cyclical demand trends observed in historical data. In addition, hourly weather information (both past observations and one-hour-ahead forecasts) covering wind speed, precipitation, and temperature, are included to account for environmental influences on mobility behavior.

Intermediate demand variational signals generation

To generate dynamic demand variation signals, we compare fine-grained (15-minute) demand observations with coarse-grained (hourly) demand expectations derived from Stage 1 predictions. This process unfolds in three steps, as explained in the following.

Firstly, for each station i , the predicted mean $\hat{\mu}_{h+1}^i$ and standard deviation $\hat{\sigma}_{h+1}^i$ are extracted from the Stage 1 forecast of hourly demand for the next hour. The mean $\hat{\mu}_{h+1}^i$ represents the expected demand under the current temporal and contextual conditions (e.g., time of day and weather). As hourly forecasts are generated over time, these values are stored as sequences, forming a dynamic baseline of regular demand expectations and supporting the identification of short-term demand deviations.

Secondly, as hourly forecasts are continuously updated, we retrospectively compare them with the actual per 15-minute demand observations that occurred during the same hour. Say the current 15-minute interval is t , for a previous 15-minute interval indexed by $t - q'$, we identify its corresponding hour h' and compute the estimated 15-minute demand from Stage 1 as:

$$\hat{\mu}_{t-q'}^i = \frac{1}{4} \hat{\mu}_{h'}^i, \quad (3)$$

assuming demand is uniformly distributed across the four quarters of hour h' .

Thirdly, the dynamic demand variation signal is then calculated as the difference between the actual observation and the scaled hourly estimate:

$$\delta_{t-q'}^i = y_{t-q'}^i - \hat{\mu}_{t-q'}^i = y_{t-q'}^i - \frac{1}{4} \hat{\mu}_{h'}^i. \quad (4)$$

By repeating this process over previous time steps, we retain a sequence of demand variation signals $\{\delta_0^i, \dots, \delta_{t-1}^i\}$, which captures recent deviations from the expected demand. This sequence is a key dynamic input in Stage 2.




Fig. 1. Pickup and drop-off demand forecasting pipeline of Two-stage Spatial and Temporal Adaptive contextual Representation (T-STAR).

Fig. 2 illustrates the process of generating dynamic demand variation signals. The first panel from the top shows the predicted hourly demand $\hat{\mu}_h^i$ over the past 8 hours. The middle panel visualizes the actual per 15-minute demand values y_q^i alongside their corresponding hourly-derived estimates $\hat{\mu}_q^i$, for each short-term interval index by q . The bottom panel depicts the resulting demand variation signals δ_q^i , capturing localized fluctuations relative to regular hourly expectations.




Fig. 2. Illustration of the generation process of intermediate demand variational signals.

Stage 2: Capturing short-term demand variations

The second stage of the T-STAR framework focuses on capturing high-frequency fluctuations in short-term demand that deviate from the regular hourly patterns estimated in Stage 1. This refinement step enables adaptive 15-minute forecasts by incorporating additional signals not present in the coarse-grained predictions. Stage 2 models future short-term demand by integrating three primary sources of influence: (1) station-specific short-term demand patterns; (2) localized demand fluctuations driven by latent or unobserved disruptions, captured through recent demand variation signals; and (3) contextual factors arising from the surrounding mobility environment.

To establish a baseline understanding of current demand, the Stage 2 model first leverages recent station-level demand observations and temporal features. Specifically, the previous V_2 per 15-minute demand values $y_{t-V_2+1:t}^i$ are used to capture short-term autoregressive patterns that reflect local demand trends at station i . In addition, the predicted hourly demand distribution from Stage 1 provides a coarse-grained expectation of demand. The predicted mean $\hat{\mu}_t^i$ offers a reference level for typical demand under prevailing contextual conditions considered in Stage 1, while the standard deviation $\hat{\sigma}_t^i$ quantifies the associated uncertainty. These values are passed into Stage 2 as compact summaries of the expected demand, without the need to reintroduce the contextual inputs already processed in Stage 1. Finally, demand seasonality indicators, hour of the day and day of the week, are included to identify the seasonal short-term demand patterns.

In addition to recent demand observations, the Stage 2 model incorporates demand variation signals to account for sudden shocks or irregular changes in short-term demand. The signals, denoted by $\delta_{t-V_2:t}^i$, capture temporal clusters of deviations caused by latent events, such as concerts, local disruptions, or anomalies in the broader mobility system. They are often not directly observable by the bike-sharing systems but may significantly alter user behavior. Moreover, demand variation signals act as indirect indicators of operational constraints, including censored demand resulting from zero fleet availability or limited parking space at a station. For example, a surge in pickups unaccompanied by corresponding drop-offs may indicate a depletion of available vehicles, potentially suppressing subsequent pickup activity despite continued user demand. Since real-time fleet availability data is often unavailable, the model leverages recent variation patterns to infer such imbalances. To support this inference, the Stage 2 predictor incorporates the station's dock capacity C^i , along with both pickup and drop-off variation sequences $\delta_{t-V_2:t}^{i,p}$ and $\delta_{t-V_2:t}^{i,d}$, where superscripts p and d refer to pickup and drop-off demand.

In addition to temporal patterns and variation signals, the Stage 2 model incorporates contextual features that capture the station's surrounding mobility environment. A key component is the interaction with nearby public transit, as shared micro-mobility services often serve as first- and last-mile connectors. To capture this multimodal dependency, the model includes real-time deviations in metro activity. For each 15-minute interval, we compute the difference between the observed number of check-ins or check-outs at the nearby metro station (within predefined distance threshold ν) and their historical average for that interval, conditioned on the same hour of the day and day of the week. These deviations serve as proxies for irregular shifts in multimodal travel demand and

help the model adjust to sudden changes of usage from the connecting public transport services. Specifically, check-out deviations $\delta_t^{i,out}$ are used for pickup demand forecasting, while check-in deviations $\delta_t^{i,in}$ inform drop-off demand forecasting.

In summary, the Stage 2 predictor estimates the conditional distribution of the next 15-minute demand at station i and time t , incorporating recent demand trends, the predicted demand expectations from Stage 1, recent demand variation signals, and contextual features. Formally, the predicted demand distribution is given by:

$$\hat{y}_{t+1:t+H_2}^i \sim \mathcal{F}_2 \left(\underbrace{y_{t-V_2:t}^i}_{\text{recent demand}}, \underbrace{\delta_{t-V_2:t}^{i,p}, \delta_{t-V_2:t}^{i,d}}_{\text{past variation signals}}, \underbrace{\hat{\mu}_{t+1:t+H_2}^i, \hat{\sigma}_{t+1:t+H_2}^i}_{\text{Stage 1 estimations}}, \underbrace{X_{t-V_2:t}^i, X_{t+1:t+H_2}^{*i}}_{\text{contextual features}} \right), \quad (5)$$

where $\mathcal{F}_2$ is a predictor designed to process multivariate input streams and produce a probabilistic forecast of short-term demand.

In this study, which focuses on forecasting dock-based bike-sharing demand for the next 15-minute interval, the look-back window V_2 of the Stage 2 predictor is set to 24. This allows the model to access demand and contextual information from the past 6 hours (i.e., 24 intervals of 15 minutes). The prediction horizon H_2 is set to 1, as the task focuses on one-step-ahead forecasting of demand for the next 15-minute interval. As with the Stage 1 predictor, the values of V_2 and H_2 can be adjusted based on data availability and specific operational requirements.

Regarding external contextual inputs, the feature sets $X_{t-V_2:t}^i$ and $X_{t+1:t+H_2}^{*i}$ include the per 15-minute metro check-in or check-out flow variation signals ($\delta_t^{i,out}$ or $\delta_t^{i,in}$) from metro stations located near the target bike-sharing station. If a bike-sharing station is connected to multiple metro stations, the variation signal aggregates the passenger flow deviations from all nearby stations for the same interval. In addition to these multimodal signals, the contextual inputs also include temporal seasonality features such as hour of the day and day of the week, as well as the station's docking capacity C^i .

4.2. Multivariate time series transformer for probabilistic demand forecasting

Time series transformer (TST)

Originally introduced for natural language processing (Vaswani, 2017), Transformer architectures have recently shown strong performance in time series forecasting tasks, thanks to their ability to capture long-range temporal dependencies and process multivariate inputs efficiently (Zerveas et al., 2021). For the proposed T-STAR framework, TST is selected as the base predictor for both Stage 1 and Stage 2 models. TST is trained as a global model across all stations, learning station-specific representations while leveraging the shared patterns of demand, which makes it particularly suitable for large-scale bike-sharing systems.

Illustrated in Fig. 3, the architecture of TST consists of a contextual embedding layer, self-attention-based encoder and decoder blocks, and a probabilistic output head. The model takes as input a context window of past demand values and associated multivariate features and predicts a probabilistic forecast for the next time step(s). The input and output sequences are first transformed into fixed-dimensional embeddings using learnable projection layers and positional encodings. These embeddings are then processed by the encoder or decoder modules, which are composed of N_E and N_D stacked encoder and decoder blocks, respectively. Each block contains a multi-head self-attention mechanism followed by position-wise feed-forward sublayers (Vaswani, 2017). While the encoder summarizes historical patterns and contextual influences, The decoder enables multi-step forecasting ($H > 1$) by autoregressively attending to prior predictions at each step (Vaswani, 2017). Unlike RNNs, which process sequences sequentially due to their recursive structure, Transformers leverage self-attention mechanisms to process the entire input sequence in parallel. This eliminates sequential computation during training, enabling efficient mini-batching and accelerating training on large datasets.

While the T-STAR framework incorporates a decoder for multi-horizon flexibility, it is important to note that for all experiments in this study where $H = 1$, the decoder is bypassed. This reduces the computational overhead during training and inference, effectively leveraging an encoder-only configuration to capture the temporal dependencies on the time series.

The Transformer architecture offers several key advantages for demand forecasting, including the ability to model complex temporal dependencies without recurrence, enabling parallel training and scalability for large-scale systems (Vaswani, 2017; Zerveas et al., 2021). Its global learning framework enables both zero-shot generalization, where forecasts can be produced for new stations or service areas without additional training, as well as few-shot adaptation through lightweight fine-tuning over the pretrained Transformer. This flexibility is particularly valuable for fast-growing shared micro-mobility services, where new locations are frequently introduced. Moreover, the model can be incrementally updated using recent data, improving adaptability to evolving usage patterns and external conditions of the bike-sharing systems. When equipped with a distribution head (e.g., Negative Binomial), it enables uncertainty-aware probabilistic forecasting.

Contextual temporal embeddings

In addition to modeling temporal dependencies through past demand observations $y_{t-V_2:t}^i$ the contextual embedding layer of TST also learns contextual representations from associated historical covariates $X_{t-V_2:t}^i$ and future covariates $X_{t+1:t+H}^{*i}$ with known or externally estimated values. Here, the subscript t refers to the current time point at which the prediction is generated, while τ is used as a generic index that may refer to any point of time within the input window. To support generalization and structured modeling, contextual inputs are categorized into three types: the static station-level attributes x^i , the global time-varying features shared across all stations x_τ , and station-specific, time-dependent features x_τ^i . Together, at any given time point τ , the contextual input at time τ for station i is denoted as $X_\tau^i = \{x^i, x_\tau, x_\tau^i\}$. This structured input enables the model to represent both global and local influences on demand across stations.




Fig. 3. Architecture of a Transformer-based probabilistic forecasting model. Note: For one-step forecasting ($H = 1$), the decoder branch is inactive, and the encoder output is passed directly to the distribution head.

Our forecasting approach uses a global Transformer model trained on data from all stations. To capture station-specific characteristics, the model uses static features x^i , which include both the station ID (as a categorical feature) and operational attributes such as station capacity C^i from the Stage 2 model. The station ID is embedded through a learnable embedding layer, enabling the model to infer unique historical demand patterns associated with each station. These static inputs collectively produce a latent station representation $h^i \in \mathbb{R}^i$, reflecting the historical behavior and demand patterns of each station. If two stations exhibit similar demand dynamics and responses to contextual factors, their learned embeddings will also be similar, allowing the model to generalize effectively across the network while preserving station-specific nuances.

Dynamic global features x_τ , including weather conditions, hour of the day and day of the week, are encoded through a dense layer into a shared temporal embedding $h_\tau \in \mathbb{R}^s$. These inputs capture periodic and seasonal variations common across the network. Conversely, local dynamic features $x_\tau^i \in \mathbb{R}^p$, such as passenger flow variations and recent demand deviation signals, reflect the dynamic conditions specific to station i . To enable the Transformer to simultaneously reason across different data scales, we fuse heterogeneous inputs into a unified feature space. Specifically, we concatenate the static station latent representation h^i , the global temporal embedding h_τ , and station-specific dynamic features x_τ^i with recent demand observations y_τ^i . This fusion, formalized in Eq. (6), ensures that the model's view of a single time step is informed by both long-term station characteristics and dynamic contextual conditions:

$$h_\tau^i = [h^i, h_\tau, x_\tau^i, y_\tau^i]^T \in \mathbb{R}^d, \quad (6)$$

where $d = i + s + p + 1$, with i , s , and p representing the dimensions of the station embedding, global temporal embedding, and station-specific dynamic context, respectively. This concatenated vector is then passed through a shared projection layer:

$$z_\tau^i = h_\tau^i W_e + b_e, \quad (7)$$

where $W_e \in \mathbb{R}^{d \times e}$ and $b_e \in \mathbb{R}^e$ to map the input into a fixed-size embedding vector that can be uniformly processed by the Transformer encoder.

In addition to contextual embeddings, the model also incorporates positional information to preserve the temporal order of observations. Unlike recurrent models, Transformers process all input tokens in parallel and lack an inherent notion of sequence order. To address this, positional encodings $PE_\tau \in \mathbb{R}^e$ are incorporated to inject position-specific information into the input representations (Vaswani, 2017). These encodings are computed using the sinusoidal encoding scheme from Vaswani (2017), allowing the Transformer to infer token order through deterministic position-dependent signals. Adding the positional embedding element-wisely to contextual embeddings z_τ^i , the resulting embedding is,

$$e_\tau^i = z_\tau^i + PE_\tau. \quad (8)$$

Transformer requires a structured temporal sequence to perform self-attention. Eq. (9) formalizes the aggregation of these individual embeddings (e_τ^i) into a sequence matrix $E_t^i \in \mathbb{R}^{V \times e}$ over the look-back window V . This structured temporal sequence serves as the primary input for the Transformer Encoder, providing the necessary global context for the model to capture long-range dependencies within the bike-sharing network:

$$E_t^i = [e_{t-V}^i, \dots, e_t^i, \dots, e_t^i]^T \in \mathbb{R}^{V \times e}. \quad (9)$$

Probabilistic output

To support uncertainty-aware forecasting, the TST in T-STAR is trained to produce a predictive distribution rather than a single point estimate. Specifically, the model outputs the parameters of a Negative Binomial (NB) distribution at each prediction step. The NB distribution, also known as the Gamma-Poisson distribution, generalizes the Poisson model by introducing a dispersion parameter that




Fig. 4. Hourly average bike pickup and drop-off demand (Washington D.C., Oct–Dec 2022).

allows the variance to exceed the mean. This property makes it particularly suitable for modeling short-term bike-sharing demand, which is discrete, non-negative, and low in volume but highly variable over 15-minute intervals. The NB distribution has been widely adopted in probabilistic forecasting of discrete, intermittent demand. Prior studies highlight its suitability for highly variable customer demand (Salinas et al., 2020; Toubeau et al., 2018). A detailed formulation of the NB distribution and its likelihood function is provided in Appendix A.

The Transformer's probabilistic output layer estimates the parameters of the NB distribution, which are the mean and dispersion, for each future time step. To generate predictive uncertainty estimates, the model samples N values from the predicted distribution, forming an empirical forecast distribution from which statistics such as the median and standard deviation can be derived. Model training is performed by minimizing the Negative Log-Likelihood (NLL) of the observed demand under the predicted NB distribution. This learning objective encourages the model to correctly predict the expected value of demand and to capture the uncertainty of demand. In a general form, the NLL loss is defined as:

$$\mathcal{L}_{\text{NLL}} = -\log p_D(y_t | \hat{\theta}_t) \quad (10)$$

where y_t is the observed demand at time step t , D denotes the assumed predictive distribution (NB distribution in this study), and $\hat{\theta}_t$ represents the predicted distribution parameters.

5. Case study and experimental setup

In the following, we introduce the case study and describe how the proposed contextual two-stage demand forecasting framework is calibrated to this setting. We then present the performance evaluation scheme, which assesses both point and probabilistic prediction accuracy. To provide a comprehensive view of model performance, we evaluate not only overall accuracy but also robustness across stations and time. Finally, we outline the design of benchmarking experiments, including a summary of the selected baseline models used for comparison.

5.1. Case study description

To evaluate the performance of the proposed short-term demand forecasting methodology in a real-world setting, we apply it to a case study of a dock-based bike-sharing system. The objective is to predict pickup and drop-off demand for each bike-sharing station in downtown Washington D.C. for the future 15-minute interval. To balance computational efficiency with high predictive accuracy, we adopt a global prediction model which simultaneously forecasts future pickup and drop-off demand across all stations in the network. In our analysis, we utilize historical trip data from Washington D.C.'s Capital Bikeshare system, covering the period from October 1, to December 31, 2022². On average, the system processes over 6000 trips per day across 235 selected bike-sharing stations distributed throughout the city center of D.C.

The demand pattern of bike-sharing services is often characterized by strong seasonality (Eren and Uz, 2020). As shown in Fig. 4, the hourly demand exhibits significantly higher activity during daytime hours, from 8:00 a.m. to 8:00 p.m. The pattern clearly reflects commuter behavior, with pronounced peaks during the morning (8:00–10:00 a.m.) and early evening (4:00–7:00 p.m.) rush hours. Moreover, on weekdays, these commuting peaks appear consistently, indicating regular usage by urban commuters. In contrast, weekend demand is more evenly distributed throughout the day, suggesting a shift toward leisure-oriented usage. Furthermore, public holidays, while reducing commuting-related trips, often stimulate increased demand at bike-sharing stations near recreational areas and tourist attractions (Palaio et al., 2021). A similar pattern is observed in the demand of Washington D.C bike-sharing.

The demand time series for each BSS is constructed by aggregating the number of pickups and drop-offs into 15-minute intervals over the entire historical period. This results in 96 demand observations per day for each station. Temporal features are extracted from the timestamp of each interval to capture daily and weekly seasonal patterns. Additionally, a binary indicator is created to denote whether a given interval falls on a public holiday. To account for weather influences, historical hourly temperature, precipitation, and wind speed data are incorporated to support the predictive analysis (Open-Meteo, 2026). To capture the interaction between

² The dataset is publicly available via Capital Bikeshare's open data portal: <https://capitalbikeshare.com/system-data>




Fig. 5. Distribution of zero-demand percentages in the historical 15-minute intervals across 235 bike-sharing stations in Washington D.C case study.

public transit and bike-sharing demand, we incorporate metro passenger flow data, defined as the number of entries and exits at metro stations per 15-minute interval. This information is derived from boarding and alighting records provided by the Washington Metropolitan Area Transit Authority (WMATA) Metrorail system. To link metro activity to bike-sharing demand, we define a proximity threshold of 300 m, approximately a 5-minute walk, following the guideline in [Ma et al. \(2018\)](#). A bike-sharing station is considered connected to a metro station if it falls within this distance. Based on this criterion, a total of 63 BSS are identified as being connected to at least one metro station. Finally, the docking capacity of each BSS is retrieved from Capital Bikeshare dataset and imported as a static feature.

Demand data from the first 70 days (October 2, to December 10, 2022) are used for model training and validation, while the remaining 20 days are reserved for testing and evaluating predictive performance. In the analysis of the demand time series at the 15-minute resolution, we observe that a large proportion of data points are zero, particularly during early morning and late evening hours. The histogram in [Fig. 5](#) shows the distribution of zero-demand frequencies of 15-minute intervals across all stations. For most stations, zero demand occurs in over 70–90% of the 15-minute intervals, for both pickup and drop-off requests. It highlights a high degree of zero-inflation in the station-level demand time series for this case study.

The proposed T-STAR forecasting framework is applied to capture both multivariate long-term patterns and short-term fluctuations in future demand at a 15-minute resolution. It does so by integrating diverse factors from the external environment and the recent supply-demand dynamics of the bike-sharing system itself. These features either describe the long-term demand trends or drive localized short-term variations.

In the first stage, demand is aggregated into hourly intervals. An hourly prediction model is then trained using these aggregated time series, along with lower-frequency contextual features such as weather conditions, public holidays, and seasonal indicators (e.g., hour of the day and day of the week). This stage should produce reliable hourly demand estimates, which serve as coarse-grained inputs for the second stage.

In the second stage, a short-term forecasting model predicts demand at 15-minute intervals within each hour. This model incorporates multiple sources of information: (1) recently observed 15-minute demand, (2) seasonal indicators, (3) probabilistic hourly predictions from the first stage, and (4) high-frequency contextual features updated every 15 minutes.

In the Washington D.C. case study, the third source mentioned, i.e. the hourly predictions from the first-stage model, are represented by the mean and standard deviation of the predicted demand distribution. These serve as coarse-grained priors, establishing a baseline expectation of demand for short-term refinement.

The high-frequency features mentioned above as the fourth source include two types of demand deviation indicators, capturing fine-grained temporal dynamics. The first type refers to intermediate features from the two-stage framework, capturing recent deviations between observed 15-minute demand and hourly predictions from stage 1. These indicators, reflecting short-term rental and return activity, are jointly used to inform both pickup and drop-off predictions. The second type of deviation feature captures discrepancies in metro passenger flow. Specifically, it measures the difference between current check-ins or check-outs and their historical averages for the corresponding hour and weekday. For pickup demand forecasting, deviations in metro check-outs are included to reflect potential mode transfers from metro to bike-sharing. Conversely, deviations in metro check-ins are used for drop-off demand forecasting, capturing likely interchanges from bike-sharing to metro during commuting hours.

5.2. Predictive accuracy evaluation metrics and key performance measures

Predictive accuracy is evaluated using both deterministic and probabilistic metrics. These measures assess the model's ability to provide accurate point forecasts as well as to capture uncertainty in the predicted distribution.

Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) are used to evaluate the deterministic prediction accuracy. Given the group of actual values y_i and their corresponding predicted values $\hat{y}_i$, the Mean Absolute Error (MAE) is defined as,

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad (11)$$

where N is the total number of data points. And the Root Mean Squared Error (RMSE) is given by,

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}. \quad (12)$$

Probabilistic predictors typically return the estimated parameters of the predicted demand distribution, which can be used to generate predictive samples. In our experiments, we generate 100 sample points for each forecast, and the point forecast is taken as the median of these sampled values.

For probabilistic demand forecasting performance, metrics Continuous Ranked Probability Score (CRPS) and Interval Score (IS) are applied. CRPS is a comprehensive measure that quantifies forecasting uncertainties by taking the full probabilistic distributions into account. Based on Gneiting and Raftery's formulation (Gneiting and Raftery, 2007), when the actual value y is a scalar, the CRPS is defined as follows:

$$\text{CRPS}(F, y) = \int_{\mathbb{R}} [F(\hat{y}) - \mathbf{1}_{\{\hat{y} \geq y\}}]^2 d\hat{y}, \quad (13)$$

In this formulation, the empirical CDF of the $F(\hat{y})$ is estimated by a step function $\mathbf{1}(\hat{y} \geq y)$, returning 1 when the predicted value is higher or equal to the actual observation, otherwise 0. The values of CRPS are always non-negative, ranging from 0 to $+\infty$ by definition, with smaller values indicating better-calibrated probabilistic forecasts.

While CRPS measures the overall probabilistic difference between the actual and predicted distributions, the Interval Score (IS) focuses on the quality of specific prediction intervals. IS accounts for both the width of the interval and whether the actual observation falls within it. The calculation of IS follows:

$$\text{IS}_{\alpha} = (\hat{u} - \hat{l}) + \frac{2}{\alpha}(\hat{l} - y)\mathbf{1}_{\{y < \hat{l}\}} + \frac{2}{\alpha}(y - \hat{u})\mathbf{1}_{\{y > \hat{u}\}}, \quad (14)$$

where $[\hat{l}, \hat{u}]$ is the prediction interval associated with a nominal coverage probability denoted by $1 - \alpha$, and y is the actual value. In this study, we set α to be 0.1, meaning the lower bound $\hat{l}$ corresponds to the 5th percentile of the predicted distribution, while the upper bound $\hat{u}$ corresponds to the 95th percentile. A smaller IS value indicates a better quality of the generated prediction interval, which balances precision and reliability.

To summarize, both CRPS and IS evaluate how close the predicted distribution is to the actual value, while penalizing poorly calibrated probabilistic forecasts, such as under-confident predictions with wide confidence intervals. While both CRPS and IS assess the quality of probabilistic forecasts at the individual instance level, we compute the Mean CRPS (MCRPS) and Mean IS (MIS) to summarize the overall average performance across a group, such as per station or within the same time period.

The demand forecasting task for an urban bike-sharing service involves predicting multiple time intervals into the future across hundreds of locations within the service network. To provide a thorough and practical assessment of forecasting performance, we adopt a multi-dimensional evaluation approach. Rather than relying solely on aggregate accuracy, we examine model performance across spatial and temporal contexts to better understand its strengths and limitations in real-world operations. This includes evaluating overall performance to measure general accuracy, analyzing spatial variations to assess consistency across stations with differing demand profiles, and exploring temporal dynamics to detect fluctuations in predictive reliability over time. Such a comprehensive evaluation framework enables us to assess not only how well the model performs on average, but also its robustness across diverse scenarios (e.g., varying station characteristics, usage patterns, and operational conditions) that are critical for deploying predictive model in large-scale and dynamic bike-sharing systems.

5.3. Benchmarking design

To validate the performance of the Transformer-based T-STAR framework in the multivariate short-term demand forecasting task for the Washington D.C. case study, we conduct two series of benchmarking experiments. These aim to evaluate the impact of contextual input design and compare T-STAR's performance against both classic and state-of-the-art forecasting models, offering insights into the most effective configurations for accurate and robust bike-sharing demand prediction.

5.3.1. Series 1: Contextual input strategies

The first set of benchmarking experiments evaluates how contextual information contributes to demand forecasting performance. Prior work suggests that contextual features can improve both accuracy and reliability, particularly under changing operating conditions and external disruptions. Accordingly, we test progressively richer contextual feature sets to quantify their marginal impact on forecasting accuracy (see the overview in Table 3).

The features explored in this our experiments are categorized into four input combinations:

- **Baseline:** This basic category only considers temporal variables, hour of the day and day of the week, to capture typical cyclical patterns in demand. These features are derived directly from the datetime information of the demand time series and represent internal temporal structure of the past and future values on the time series.
- **Global Environment:** In addition to temporal features, this category includes contextual factors such as weather conditions and public holidays, which influence overall mobility patterns across the city. These features are shared across all stations and represent conditions that affect bike-sharing demand from the environment external to the bike-sharing operations.

Table 3

Summary of input combinations evaluated in the Washington D.C. case study. ✓ indicates that the input was included in the model.

Input Combination Types	Contextual Features				Time Series Info.	
	Temporal variables (hour of a day and day of the week)	External environment (weather & holiday)	environ- (weather & &	Bike-Sharing facility (dock capacity) and recent deviations from connected metro check- in/out flow	Past demand variations	Previous obser- vations from the time series
Baseline	✓					✓
Global Environment	✓	✓				✓
Environment, Facility and PT	✓	✓		✓		✓
T-STAR	✓	✓		✓	✓	✓

- **Environment, Facility and PT:** This advanced category extends the *Global Environment* setting by incorporating station-level dock capacity, which may influence future demand, and recent deviations in check-in and check-out passenger flows at nearby metro stations. Together, these features aim to capture the effects of local infrastructure and public transport dynamics on short-term bike-sharing usage.
- **T-STAR:** The proposed T-STAR framework adopts a comprehensive combination of inputs. It integrates all previous contextual features and additionally includes recent demand deviation features, matched precisely to hourly and short-term predictions based on these features' temporal granularity. These internal demand deviation features are calculated by comparing recent short-term observed demand against corresponding hourly demand forecasts generated by the first-stage long-term predictive model. The objective is to improve forecasting accuracy by explicitly capturing recent short-term demand fluctuations.

5.3.2. Series 2: Base predictors

The second series of our benchmarking exercise investigates the performance of different base predictive models, each capturing distinct forecasting assumptions and strengths. This exploration helps identify the effectiveness and appropriateness of various forecasting methods in the context of short-term bike-sharing demand prediction. The selected base predictors are introduced in the following.

Historical Average. This straightforward and deterministic baseline method forecasts future demand as the average of historical observations from the same day of the week and hour of the day, leveraging the recurring seasonal patterns of demand.

Myopic Approach. This approach utilizes the immediate past demand observation as the forecast for the subsequent time interval. This simplistic predictor serves as a practical baseline for short-term forecasting scenarios, under the assumption of minimal demand fluctuations between consecutive intervals.

Time Series Transformer (TST) - Hourly Model. This approach generates 15-minute demand forecasts by proportionally dividing hourly predictions produced by the TST model used in Stage 1 of the T-STAR framework. The hourly predictions leverage past hourly demand, temporal factors (day of the week and hour), and weather information. By operating at a coarser resolution, the model helps mitigate the sparsity and noise often present in high-frequency data, resulting in more stable predictions. However, unlike the T-STAR framework, this approach does not incorporate high-resolution temporal signals, limiting its responsiveness to short-term demand fluctuations.

XGBoost (Extreme Gradient Boosting). Introduced by [Chen and Guestrin \(2016\)](#), XGBoost is a machine learning technique renowned for its predictive performance and computational efficiency. Its effectiveness in short-term transportation demand forecasting arises from robustness against sparse data, efficient handling of nonlinear relationships, and resistance to overfitting. However, it is limited by its reliance on historical examples, leading to weaker performance when predicting scenarios unseen during training. Additionally, XGBoost's boosting structure is inherently designed for point prediction and does not support probabilistic forecasting by default.

In addition, we look into the application of six state-of-the-art neural-network predictors, each described below.

STGCN and STGCN-VAE. Originally introduced by [Yu et al. \(2017\)](#), the Spatio-Temporal Graph Convolutional Network (STGCN) addresses forecasting challenges in networked systems by integrating spatial graph convolutions with gated 1D temporal convolutions. This fully convolutional architecture enables the scalable learning of complex dynamics while explicitly leveraging the underlying network topology. Consistent with our Transformer-based predictors, we provide station-specific attributes as node features to learn representative station embeddings. To enable probabilistic forecasting for the standard STGCN, we extend the architecture with a distributional output head and train the model by minimizing the negative log-likelihood. In this study, we also include STGCN-VAE as a state-of-the-art graph-based benchmark for nonparametric probabilistic prediction ([Peng et al., 2025](#)). This model utilizes an STGCN backbone to extract spatio-temporal features, which are subsequently compressed into a latent space via a Variational Autoencoder (VAE) and processed through a Kernel Density Estimation (KDE) module to construct predictive distributions. For each graph-based predictor, we evaluated three commonly applied graph construction strategies to define the static adjacency matrix. To ensure an optimal baseline, we report the results corresponding to the best-performing configuration for each benchmark. Technical details regarding these construction methods and the final configurations selected for each model are provided in [Appendix B](#).

DeepAR. Introduced by Salinas et al. (2020), DeepAR is a probabilistic autoregressive forecasting model utilizing LSTM architecture. It excels at capturing temporal dependencies within time series data and serves as a global predictor for all bike-sharing stations. However, due to its sequential nature, it can struggle with sparse or intermittent demand series, resulting in potential instability and inaccuracies when forecasting infrequent or zero-inflated demand.

STAEformer. Spatio-Temporal Adaptive Embedding Transformer (STAEformer) was introduced by Liu et al. (2023). It augments a vanilla Transformer encoder with spatio-temporal adaptive embeddings that tie node representations to spatio-temporal attention patterns, thereby injecting spatial and temporal inductive bias and improving forecasting performance. Following the original design, STAEformer uses historical time-series from all locations as input, together with hour-of-day and day-of-week features as contextual attributes (Liu et al., 2023). For station-wise 15-minute bike-sharing demand forecasting, we adopt the same input structure and further extend STAEformer with a distributional output head. The model is trained by minimizing the negative log-likelihood to produce probabilistic demand forecasts.

Temporal Fusion Transformer (TFT). Recently proposed by Lim et al. (2021), TFT combines LSTM with a self-attention mechanism, effectively capturing joint temporal and contextual relationships. It incorporates a gating mechanism for adaptive feature selection, reducing the risk of overfitting despite numerous input features. Its complexity, however, can lead to extensive computational demands during training. Different from DeepAR and TST, TFT adopts pinball loss as training objective.

Simple Time Series Transformer (Simple TST). As introduced in Section 4.2, TST is used as the base predictor in the proposed T-STAR forecasting framework. To assess the added value of T-STAR's two-stage architecture, we implement a one-stage baseline model, referred to as *Simple TST*. This processes all available contextual information in a single stage to directly generate per-15-minute short-term demand forecasts.

We tuned model hyperparameters using a five-fold time-series cross-validation scheme that preserves temporal order and avoids information leakage. Specifically, we employed a rolling forecast origin design with an expanding training window and a fixed one-week validation window per fold. The five validation weeks were distributed across the three-month training period to ensure coverage of weekday/weekend demand patterns and to capture atypical demand intervals when present. Hyperparameter configurations were selected by minimizing the average validation RMSE across the five folds. For each model, 30 trials (i.e., 30 different hyperparameter combinations) were performed with Optuna package to efficiently explore the hyperparameter search space (Akiba et al., 2019). The types and search ranges of the tuned hyperparameters are detailed in the Appendix C, along with the selected values for each model.

6. Results and discussion

6.1. Short-term demand forecasting performance

6.1.1. Overall predictive accuracy and computational efficiency

This subsection reports the overall forecasting performance of selected predictors and corresponding analysis. To account for variability due to random initialization, we performed multi-seed training for each predictor. Specifically, we trained five independent model instances using five randomly selected seeds, while keeping the hyperparameter configuration and experimental protocol identical. The results reported in this section are aggregated across seeds by taking the average across runs.

We start with analyzing the overall predictive accuracy and computational efficiency of the proposed Transformer-based T-STAR framework for short-term bike-sharing demand forecasting, in comparison to selected benchmark methods. Each forecasting approach consists of two key components: the base predictive model and the contextual input features incorporated to support forecasting. We evaluate different families of forecasting approaches, including methods that produce only deterministic forecasts as well as those capable of generating both deterministic and probabilistic outputs. To provide a generalized assessment of model performance, we report the aggregated mean accuracy across all stations. In addition, we examine variations in predictive performance across different stations and over time to evaluate the robustness of each forecasting approach. This analysis allows us to assess whether a single global model can maintain consistent predictive accuracy under diverse spatial and temporal conditions.

Table 4 summarizes the overall predictive performance of the selected forecasting approaches for pickup demand, evaluated over a 20-day period across 235 Capital Bikeshare stations in Washington, D.C. The Transformer-based T-STAR framework consistently delivers accurate and robust performance across all key metrics, including deterministic accuracy (measured by MAE and RMSE) and probabilistic accuracy (measured by MCRPS and MIS). Among the selected SOTA probabilistic models, Transformer-based T-STAR achieves the lowest MAE, RMSE, and MCRPS. Compared to the historical average predictor, Transformer-based T-STAR reduces the MAE by 54.9% and the RMSE by 7.9%. Furthermore, relative to a Transformer predictor without contextual inputs, Transformer-based T-STAR improves performance by 19.7%, 7.4%, 27.7%, and 50.3% in MAE, RMSE, MCRPS, and MIS, respectively. These large improvements demonstrate the substantial benefits of two-stage adaptive contextual representation.

Although XGBoost and the hourly Transformer predictor achieve the lowest RMSE, this apparent advantage is partly driven by the demand series' pronounced zero inflation and heavy-tailed spikes, under which RMSE disproportionately rewards strategies that avoid large squared errors. In such cases, XGBoost tends to bias its predictions toward occasional high-demand instances in order to minimize the training loss, at the expense of accuracy on the more frequent low- or zero-demand observations. As a result, the forecasting experiment showed that XGBoost tends to over-predict during off-peak periods and underperforms at low-demand stations. The quarterly forecasts derived from the hourly Transformer benefit from temporal aggregation. The hourly forecasts tend to

Table 4

Predictive accuracy comparison of short-term pickup demand forecasting models across different approaches. The aggregated mean accuracy across all stations is reported before the parentheses. Metric values are averaged over five runs with different random seeds for each predictive model (except the Historical Average and Myopic approaches). The first value in parentheses indicates the standard deviation across stations, and the second value indicates the standard deviation across time intervals (aggregated across stations at each interval). Probabilistic metrics (MCRPS and MIS) are only reported for models supporting probabilistic forecasting.

Forecasting Type	Basic Predictor	Input Types	MAE	RMSE	MCRPS	MIS
Deterministic	<i>Historical Average</i>	<i>Baseline</i>	0.335 (0.150; 0.219)	0.504 (0.187; 0.260)	–	–
	<i>Myopic</i>	<i>Last demand observation of each station</i>	0.229 (0.115; 0.200)	0.578 (0.180; 0.312)	–	–
	<i>TST-Hourly</i>	<i>Global Environment</i>	0.204 (0.111; 0.165)	0.436 (0.138; 0.235)	–	–
	<i>XGBoost</i>	<i>Env., Facility and PT</i>	0.334 (0.068; 0.120)	0.459 (0.133; 0.205)	–	–
Probabilistic	<i>Simple TST</i>	<i>Baseline</i>	0.188 (0.123; 0.176)	0.501 (0.197; 0.310)	0.177 (0.109; 0.156)	3.753 (1.276; 2.750)
	<i>STAEformer</i>	<i>Baseline</i>	0.152 (0.096; 0.154)	0.475 (0.180; 0.289)	0.187 (0.054; 0.084)	3.628 (0.206; 0.286)
	<i>DeepAR</i>	<i>Env., Facility and PT</i>	0.159 (0.097; 0.156)	0.474 (0.169; 0.280)	0.133 (0.071; 0.119)	2.188 (0.756; 1.616)
	<i>STGCN</i>	<i>Env., Facility and PT</i>	0.422 (0.049; 0.332)	0.713 (0.071; 0.342)	0.399 (0.029; 0.200)	7.465 (0.446; 1.277)
	<i>STGCN-VAE</i>	<i>Env., Facility and PT</i>	0.404 (0.057; 0.087)	0.500 (0.124; 0.175)	0.398 (0.057; 0.087)	15.419 (2.286; 3.463)
	<i>TFT</i>	<i>Env., Facility and PT</i>	0.156 (0.095; 0.155)	0.473 (0.171; 0.281)	0.134 (0.067; 0.115)	2.123 (0.653; 1.142)
	<i>Transformer-Based T-STAR</i>	<i>T-STAR</i>	0.151 (0.089; 0.148)	0.464 (0.156; 0.271)	0.128 (0.062; 0.105)	1.864 (0.508; 0.925)

(1) Since the *Historical Average*, *Myopic*, *TST-Hourly*, and *XGBoost* models do not produce probabilistic forecasts in our experiments, the probabilistic evaluation metrics **MCRPS** and **MIS** are not reported for these models (denoted by – in the table).

(2) The best-performing values for each metric are highlighted in bold in the table.

be smoother and less volatile, so distributing them across four 15-minute intervals dampens extremes and improves RMSE. However, this uniform within-hour allocation cannot capture sub-hour demand bursts driven by short-term factors, yielding persistent interval-level deviations and therefore higher MAE than multivariate predictors that leverage short-term contextual inputs.

Both the historical average and myopic predictors exhibit significantly larger errors across all evaluation metrics. Moreover, STAEformer's spatio-temporal adaptive embeddings and attention improve point forecasting by learning flexible inter-station and temporal dependencies from historical demand, which aligns with its relatively low MAE/RMSE. However, with inputs largely limited to past demand and coarse calendar features, this method has less information to condition the uncertainties on exogenous drivers of dispersion and rare spikes. Context-rich probabilistic models (DeepAR, TFT, and Transformer-based T-STAR) therefore deliver better-calibrated and sharper predictive distributions, as evidenced by their lower MCRPS and MIS. This finding highlights the importance of incorporating contextual information and capturing dynamic temporal trends for adaptive forecasting and uncertainty quantification, compared to relying solely on the general seasonal demand patterns or the most recent observations.

Although finer temporal granularity increases challenges due to greater demand sparsity and variability, the short-term Transformer-based T-STAR forecasting approach maintains high predictive accuracy and performance robustness at this resolution. These improvements justify the efforts associated with T-STAR's two-stage training structure to enable more precise and timely operational decision-making in bike-sharing systems.

Table 5 presents the overall predictive performance of the selected forecasting approaches for short-term drop-off demand. Consistent with the results for pickup demand, the Transformer-based T-STAR framework demonstrates strong performance on the drop-off task, highlighting its ability to generalize predictive quality to a different but closely related demand series. Transformer-based T-STAR achieves the lowest MAE (0.151), MIS (1.895), and ties for the lowest MCRPS (0.126), indicating solid performance in both point and distributional predictions.

Interestingly, most models exhibit slightly better accuracy for drop-off than for pickup demand. This trend suggests that drop-off patterns may be inherently easier to forecast at a 15-minute resolution, potentially because drop-offs are less constrained by facility availability compared to pickups. For the Capital Bikeshare service, riders can opt to pay a small surcharge to park at public bike racks if docking stations are full. The flexibility in parking reduces the sensitivity of drop-off behavior to real-time bike-sharing station capacity (in parking docks). However, pickup is hard constrained by the fleet availability at the origin station.

As showed in **Table 5**, for drop-off demand, the vanilla Transformer without contextual inputs performs comparably to the T-STAR Transformer across MAE, RMSE, and MCRPS. This indicates that contextual information contributes less to forecasting short-term drop-offs than it does for pickups. However, T-STAR yields a markedly lower MIS. This indicates that T-STAR's two-stage design, together with short-term variation signals and station-specific attributes, improves uncertainty quantification by producing tighter and better-calibrated 95% prediction intervals, thereby reducing out-of-bound interval penalties on unseen data. We also observe the lowest RMSE for XGBoost and the hourly TST variant, likely due to temporal smoothing from their modeling and/or aggregation choices. Overall, results from the Capital Bikeshare case study indicate that drop-off dynamics are comparatively less sensitive to contextual variations, and that a context-free Transformer is sufficient for capturing short-term fluctuations, while T-STAR primarily adds value through improved interval reliability.

For both pickup and drop-off forecasting, graph-based predictors STGCN and STGCN-VAE are consistently outperformed by the T-STAR Transformer and other neural-network-based predictors. Their sub-optimal performance likely stems from a weak spatial signal relative to the dominant, station-specific temporal dynamics present at a 15-minute granularity.

Table 5

Predictive accuracy comparison of short-term drop-off demand forecasting models across approaches. Metric values are averaged over five runs with different random seeds for each predictive model (except the Historical Average and Myopic approaches). The aggregated mean accuracy across all stations is reported before the parentheses. The first value in parentheses indicates the standard deviation across stations, and the second value indicates the standard deviation across time intervals (aggregated across stations at each interval). Probabilistic metrics (MCRPS and MIS) are only reported for models supporting probabilistic forecasting.

Forecasting Type	Basic Predictor	Input Types	MAE	RMSE	MCRPS	MIS
Deterministic	<i>Historical Average</i>	<i>Baseline</i>	0.331 (0.153; 0.220)	0.505 (0.199; 0.270)	–	–
	<i>Myopic</i>	<i>Last demand observation of each station</i>	0.225 (0.114; 0.199)	0.572 (0.187; 0.316)	–	–
	<i>TST-Hourly</i>	<i>Global Environment</i>	0.199 (0.109; 0.165)	0.437 (0.146; 0.242)	–	–
	<i>XGBoost</i>	<i>Env., Facility and PT</i>	0.317 (0.071; 0.126)	0.452 (0.140; 0.217)	–	–
Probabilistic	<i>Simple TST</i>	<i>Baseline</i>	0.151 (0.091; 0.149)	0.462 (0.165; 0.276)	0.126 (0.065; 0.111)	2.202 (0.645; 1.446)
	<i>STAEformer</i>	<i>Baseline</i>	0.152 (0.097; 0.156)	0.475 (0.190; 0.297)	0.186 (0.055; 0.087)	3.655 (0.227; 0.345)
	<i>DeepAR</i>	<i>Env., Facility and PT</i>	0.155 (0.096; 0.155)	0.473 (0.181; 0.288)	0.133 (0.069; 0.116)	2.011 (0.743; 1.338)
	<i>STGCN</i>	<i>Env., Facility and PT</i>	0.602 (0.066; 0.479)	0.917 (0.069; 0.461)	0.535 (0.036; 0.283)	8.969 (0.201; 1.946)
	<i>STGCN-VAE</i>	<i>Env., Facility and PT</i>	0.395 (0.059; 0.092)	0.498 (0.132; 0.188)	0.390 (0.059; 0.092)	15.171 (2.369; 3.662)
	<i>TFT</i>	<i>Env., Facility and PT</i>	0.154 (0.095; 0.156)	0.472 (0.181; 0.290)	0.131 (0.068; 0.117)	2.083 (0.659; 1.223)
	<i>Transformer-Based T-STAR</i>	<i>T-STAR</i>	0.151 (0.109; 0.216)	0.462 (0.173; 0.284)	0.126 (0.078; 0.160)	1.895 (0.973; 2.334)

(1) Since the *Historical Average*, *Myopic*, *TST-Hourly*, and *XGBoost* models do not produce probabilistic forecasts in our experiments, the probabilistic evaluation metrics **MCRPS** and **MIS** are not reported for these models (denoted by – in the table).

(2) The best-performing values for each metric are highlighted in bold in the table.

To further investigate the performance gap of graph-based predictors, we conducted a seasonal detrended residual correlation analysis to isolate spatial dependencies from shared temporal patterns, for both pickup and drop-off demand time series. Specifically, we compared the station-wise pairwise correlations obtained for raw demand time series with those computed on demand residuals, the latter obtained by removing seasonal averages (grouped by day-of-week and hour) to isolate stochastic, high-frequency fluctuations. Our analysis show that, while raw data exhibits a modest average correlation (≈ 0.139) due to shared daily cycles, this correlation is reduced by over 68% (to ≈ 0.044) once seasonal effects are removed. Analysis of the demand residuals shows that the most correlated 5% of station pairs reach a maximum of only 0.076. These results indicate that at 15-minute intervals, station-wise demand fluctuations exhibit a low signal-to-noise spatial structure, with residual pairwise linear correlations remaining weak after accounting for shared seasonal patterns. While this does not preclude the existence of nonlinear or higher-order spatial dependencies that more expressive architectures might in principle exploit, it does suggest that static graph convolution mechanisms face critical challenges under these conditions, where stochastic noise tends to dominate over any coherent spatial signals.

Beyond the weakness of the spatial signal itself, the underperformance of STGCN may also reflect a structural limitation of static graph representations in highly dynamic systems such as bike-sharing, where station interactions evolve continuously due to commuting patterns, local demand shifts, weather conditions, and short-term operational dynamics. A fixed adjacency matrix cannot adapt to these evolving dependencies, potentially limiting the model's ability to extract meaningful spatial information even when it exists. Together, these two factors, low signal-to-noise spatial structure and the rigidity of static graph representations, place STGCN and STGCN-VAE at a structural disadvantage in this high-resolution setting. Temporal-focused predictors, which identify reliable station-specific contextual patterns, are better positioned to handle these conditions and consistently outperform static graph-based approaches in this high-resolution setting.

In addition, our experiments further reveal an informative contrast within the graph-based model family: while STGCN-VAE substantially improves point-prediction accuracy over the original STGCN, it yields the poorest predictive interval calibration among all benchmarks, with an MIS nearly double that of the next closest model. The point-prediction improvement suggests that some exploitable spatial structure may exist at this temporal resolution. The VAE's latent compression effectively filters out high-frequency noise, promoting smoother, more generalizable mean forecasts. However, this regularization mechanism undermines reliable uncertainty quantification. By forcing the latent space toward a smooth, low-dimensional Gaussian prior, the model inevitably discards the sharp, fine-grained stochastic variability required for well-calibrated prediction intervals. This trade-off is particularly pronounced in our setting, where 15-minute station-level demand is sparse and zero-inflated. Under these highly noisy and sparse conditions, generative models often suffer from a form of posterior collapse (Dai et al., 2020), in which the decoder learns to bypass the latent representation when input signals are too noisy or sparse, rendering the learned uncertainty estimates unreliable. This observation stands in contrast to the findings of Peng et al. (2025), where STGCN-VAE demonstrated strong calibration performance for bike-sharing demand aggregated at the postal-code and hourly level. At that coarser resolution, demand is denser and less zero-inflated, providing the VAE with sufficient signal to model latent demand variance accurately. The degradation observed in our setting highlights that the efficacy of generative graph-based models for probabilistic forecasting is highly contingent upon the spatio-temporal resolution of the use case.

In order to evaluate the statistical significance of the observed performance gains, we conducted a station-wise Wilcoxon signed-rank test comparing T-STAR transformer against all baseline models. For each station, error metrics were averaged across five independent experimental seeds to mitigate stochastic training effects ($N = 235$). This two-sided paired test was employed to determine

whether the distribution of station-wise metric values for T-STAR differed significantly from the benchmarks at a 95% confidence level. The test results for both pickup and drop-off demand consistently yielded p -values below 0.05 across all comparisons. These findings demonstrate that the improvements achieved by T-STAR are statistically significant and spatially consistent.

To assess robustness, we quantified run-to-run variability across five independent random seeds using the coefficient of variation (CV). Most architectures showed relative variability below 5%, whereas STGCN displayed markedly higher instability, with CV exceeding 10% on multiple metrics. Overall, this is consistent with greater sensitivity of fixed-adjacency graph-temporal convolutional models to sparse and intermittent demand, where the adjacency-induced inductive bias and convolutional parameterization may interact with stochastic optimization. In contrast, the attention-based baselines were comparatively more stable across seeds in our experiments.

Computational efficiency

From the perspective of computational efficiency, the Transformer-based T-STAR approach remains competitive relative to other baselines considered in this study. All selected predictors, except for Seasonal ARIMA, are able to generate forecasts for the next 15-minute interval within a few seconds.

In terms of training time, all deep learning models in our experiment were trained for 100 epochs using mini-batching with a batch size of 256. On a CPU-based computing environment, the average training times were approximately 5.5 minutes for Transformer, 6.5 minutes for DeepAR, and 11 minutes for TFT. Since the Transformer-based T-STAR framework consists of two sequential stages, each with its own Transformer model, the total training time is approximately 11 minutes. Training STGCN and STAEformer is more computationally demanding. On an NVIDIA A100 GPU, STGCN and STGCN-VAE requires around 6.5 minutes to train, while STAEformer requires 5.5 minutes.

For XGBoost, training a global model capable of making predictions across all stations increases data volume substantially. Although XGBoost does not require iterative training like deep learning models, it typically consumes a large amount of RAM. Nevertheless, model fitting was completed in approximately 2 minutes.

The historical average and myopic benchmark models are computationally lightweight, as they require minimal input features and no model training. In contrast, the standard Seasonal ARIMA model proved computationally infeasible in our experimentation. Since bike-sharing services operate continuously, modeling both hourly and weekly seasonal patterns requires a minimum seasonal period of 168. This long seasonal cycle significantly increases model complexity. Under these conditions, Seasonal ARIMA failed to converge within four hours of computation, and was therefore excluded from further analysis.

Predictive performance robustness of transformer-based T-STAR

To assess the predictive stability and robustness of the T-STAR Transformer against varying data availability and temporal demand shifts, we conducted two time-preserving evaluations: a rolling forecast origin (expanding-window) experiment and sliding window cross-validation. The former simulates a real-world deployment pipeline with periodic retraining, while the latter evaluates model stability under short-term distribution shifts across different historical contexts. Detailed experimental setups and result tables are provided in [Appendix D](#). The results indicate that T-STAR consistently outperforms the benchmark contextual TST (with Environment, Facility and PT inputs) across all experimental folds for both pickup and drop-off forecasting. Notably, T-STAR maintains superior performance even with limited training data in early rolling folds and exhibits lower station-wise variability when facing temporal shifts in demand patterns. These results confirm that T-STAR's spatiotemporal-aware architecture enhances generalizability and robustness against the inherent volatility of urban shared micro-mobility demand.

6.1.2. Station-level forecasting accuracy

We provide a focused analysis of station-level forecasting performance by presenting the deterministic and probabilistic accuracy of 15-minute pickup demand predictions across all 235 bike-sharing stations in [Figs. 6 and 7](#). The stations are ranked by their average daily demand, allowing us to examine how performance varies with demand intensity. A clear trend emerges, as both MAE and MIS generally increase with demand, indicating that short-term demand at high-activity stations is more challenging to predict accurately. This performance degradation is especially evident among the top five highest-demand stations. As shown in [Fig. 6](#), the Transformer-based T-STAR model consistently outperforms baseline models, with the most substantial accuracy gains observed at high-demand stations. Similarly, [Fig. 7](#) demonstrates significant improvements in probabilistic forecasting accuracy when comparing the T-STAR model to a basic Transformer, further confirming the benefit of the two-stage contextual approach.

These results suggest that as demand levels increase, short-term demand patterns become more complex, likely due to a broader set of interacting factors related to system dynamics and surrounding facilities. The multivariate contextual representations learned by the T-STAR model are critical for capturing these fluctuations and associated uncertainties. The model's forecasting advantage is especially pronounced at high demand nodes within the bike-sharing network, which typically play a central role in system operations. Improving accuracy at these high-impact locations can support more effective fleet management and service planning for micro-mobility services.

6.1.3. Forecasting stability across time

The temporal accuracy analysis is included to assess the consistency of model performance throughout the testing period, helping to detect potential drift, degradation, and identify challenging periods to predict. In this subsection, we first examine the predictive performance of selected models across each 15-minute interval during the testing period for the short-term pickup demand forecasting




Fig. 6. Station-level MAE by demand intensity in short-term pickup forecasting experiment. The left axis shows the MAE for each bike-sharing station, while the right axis indicates the corresponding daily average pickup demand. Stations along the horizontal axis are ordered by increasing average daily pickup demand. The left subplot presents MAE values for all 235 stations, and the right subplot provides a focused view of the top 30 highest-demand stations. The y-axis scales for both the left and right plots are matched to facilitate direct comparison.




Fig. 7. Station-level MCRPS and MIS by demand intensity in short-term pickup forecasting experiment. The left axis shows the MCRPS and MIS for each bike-sharing station, while the right axis indicates the corresponding daily average pickup demand. Stations along the horizontal axis are ordered by increasing average daily pickup demand. The subplots on the left column present MCRPS and MIS values for all 235 stations, and the subplots from the right column provide a focused view of the top 30 highest-demand stations. The y-axis scales for both the left and right plots are matched to facilitate direct comparison.

task. Next, to evaluate model performance under rare or extreme conditions, we conduct separate analyses for normal and abnormal demand observations at the 15-minute level.

The top panel of Fig. 8 shows the MAE values of different predictive models over time, where each point represents the forecasting error across all bike-sharing stations at a given 15-minute interval. The bottom panel displays the corresponding interval's total pickup demand across the city. As expected, models that incorporate recent observations (all except Historical Average) demonstrate greater adaptability, particularly during unexpected drops in demand throughout the testing period. Notably, although the multivariate XGBoost predictor can respond to sudden demand shifts (e.g., between time steps 960 and 1440), it still shows relatively larger errors during abrupt demand drops compared to parametric Transformer-based models. This is likely because such scenarios with




Fig. 8. Temporal dynamics of forecasting errors and actual pickup demand over the testing period. The top panel shows the per-timestep MAE of different forecasting models for short-term pickup demand. The bottom panel displays the corresponding total pickup demand at each 15-minute interval.




Fig. 9. Per-timestep comparison of probabilistic forecasting performance between the simple Transformer model (Simple TST), one-stage multi-variate Transformer (Multi. TST) model with contextual features (Env., Facility and PT) and Transformer-based T-STAR. The left panels show performance over a regular demand period (Dec.11–16, 2022), while the right panels correspond to an unexpectedly low demand period (Dec.22–27, 2022).

unexpected demand shifts were under-represented or absent in the training data, as XGBoost relies heavily on patterns observed during training and struggles to generalize beyond them. Among all evaluated models, the Transformer-based T-STAR consistently achieves the lowest MAE across time, with especially strong performance during peak demand intervals. These results highlight Transformer-based T-STAR's ability to capture regular demand patterns while also demonstrating robustness and responsiveness in forecasting short-term bike demand under varying and dynamic temporal conditions. This performance is attributed to its adaptive forecasting mechanism, which jointly leverages recent high-frequency demand variations and longer-term trends derived from hourly demand expectations.

Table 6

Forecasting accuracy of TST-based models with different inputs under normal and abnormal pickup demand scenarios. Metric values are averaged over five runs with different random seeds for each predictive model.

Predictive Model	Abnormal				Normal			
	MAE	RMSE	MCRPS	MIS	MAE	RMSE	MCRPS	MIS
<i>Simple TST</i>	2.210	2.497	1.997	56.662	0.170	0.440	0.160	3.264
<i>TST (Env., Facility and PT)</i>	2.149	2.424	1.879	46.784	0.138	0.412	0.128	1.758
<i>Transformer-Based T-STAR</i>	2.139	2.408	1.763	31.755	0.133	0.402	0.114	1.599

The best-performing values for each metric are highlighted in bold in the table.

We further examine the impact of contextual inputs by comparing a period with regular demand patterns to one with an unexpectedly sharp demand drop, focusing on the probabilistic performance of one-stage Transformer models versus the Transformer-based T-STAR. The regular demand period is selected to be the first 5-day period from Dec.11th to Dec.16th (time steps 0–480), and the unexpectedly low demand period to be 5-day period from Dec.22nd to Dec.27th (time steps 960–1440) on the testing period. This drop coincided with a major Arctic blast that brought record-low temperatures and strong winds to Washington D.C., creating unfavorable conditions for outdoor biking. Additionally, the period included the three-day Christmas holiday, further reducing commuting demand in the city.

Fig. 9 presents the per-timestep probabilistic forecasting performance of simple Transformer model, multivariate Transformer model with contextual features (Env., Facility and PT) and Transformer-based T-STAR, using MCRPS (top) and MIS (bottom) across two selected 5-day periods: a regular demand period (left) and an irregular low-demand period (right). Overall, T-STAR Transformer consistently outperforms the rest, achieving lower MCRPS and MIS values in both scenarios. The performance gain versus simple Transformer is especially evident during the regular demand period, suggesting the integration of short-term contextual embeddings effectively captures day-to-day demand fluctuations. During the irregular demand period with high uncertainty, all Transformer predictors recognize the dominant global influence (e.g., severe weather and holidays). However, the Transformer-based T-STAR variant consistently achieves better performance, as reflected by lower MCRPS and MIS. This pattern indicates that T-STAR's two-stage contextual design more effectively integrates the compound effects of extreme weather and holidays, yielding forecasts that are both more accurate and more stable under regime shifts.

Abnormal and normal demand prediction

To evaluate model robustness under unusual scenarios, we conduct a targeted analysis comparing forecasting performance specifically on rare abnormal observations. In this experiment, we define abnormal observations as time steps where a station's demand has a z-score of 3 or higher, indicating a significant deviation from its historical average. These abnormal intervals account for approximately 1% of the dataset. For this analysis, we focus exclusively on models using the TST as the base predictor, in order to isolate and highlight the contribution of spatio-temporal contextual embeddings generated by the T-STAR framework.

Table 6 presents the predictive accuracy of TST-based models using three different contextual input structures: *Baseline*, *Environment*, *Facility and PT*, and *T-STAR*, evaluated separately over normal and abnormal demand periods. The results reveal a clear difference in model performance between the two scenarios, where prediction errors are consistently higher across all metrics during abnormal events. This highlights the inherent difficulty in forecasting rare demand fluctuations, likely due to their unpredictable nature and limited representation in historical data. Among the models, the Transformer-based T-STAR consistently demonstrates superior predictive accuracy, particularly in uncertainty quantification. As expected, the TST model with *Environment*, *Facility*, and *PT* inputs outperforms the simpler baseline in accuracy. This improvement is largely driven by the inclusion of facility-related features that capture censored demand effects, and real-time metro passenger flow data that reflect delays or early arrivals. Together, these factors help account for anomalies in short-term bike-sharing demand caused by unfulfilled trips and intermodal connections between metro and bike usage. Its significantly lower MIS scores during abnormal periods suggest that the T-STAR framework enhances the model's ability to capture and represent uncertainty, making it more reliable under rare and extreme conditions. This is enabled by the inclusion of intermediate demand variation features, which reveal unexpected sequential fluctuations in recent demand and help reduce noise during training. These findings suggest that the Transformer-based T-STAR model is an advantageous option for supporting real-time operations in shared micro-mobility services, where anticipating and effectively managing extreme scenarios is critical in ensuring service reliability.

Table 7

Short-term pickup demand forecasting performance of time-series Transformer-based models using different contextual input configurations in the Washington D.C. case study. Metric values are averaged over five runs with different random seeds for each predictive model. The first value in parentheses indicates the standard deviation across stations, and the second value indicates the standard deviation across time intervals (aggregated across stations at each interval).

Basic Predictor	Input Types	MAE	RMSE	MCRPS	MIS
TST	<i>Baseline</i>	0.188 (0.123; 0.176)	0.501 (0.197; 0.310)	0.177 (0.109; 0.156)	3.753 (1.276; 2.750)
TST	<i>Global Environment</i>	0.152 (0.092; 0.151)	0.469 (0.163; 0.277)	0.131 (0.063; 0.108)	2.248 (0.536; 0.984)
TST	<i>Env., Facility and PT</i>	0.156 (0.091; 0.151)	0.472 (0.163; 0.276)	0.142 (0.060; 0.107)	2.009 (0.476; 0.910)
TST	<i>T-STAR</i>	0.151 (0.089; 0.148)	0.464 (0.156; 0.271)	0.128 (0.062; 0.105)	1.864 (0.508; 0.925)

The best-performing values for each metric are highlighted in bold in the table.

6.2. Discussion

Building on the experimental results, this section reflects on the advantages and practical implications of the proposed Transformer-based T-STAR framework for short-term demand forecasting in shared micro-mobility systems. We discuss three key aspects that contribute to model's effectiveness: the benefit of integrating diverse multivariate contextual inputs, the architectural advantage of T-STAR's two-stage design in infusing multivariate inputs, and the model's generalization capability for its zero-shot forecasting application on unseen stations. Together, these elements highlight the practical strengths of the proposed predictive framework and its potential for real-world deployment for shared micro-mobility services.

6.2.1. Analysis of contextual feature contributions

To quantitatively assess the benefits of progressively incorporating contextual information, we apply all contextual input strategies introduced in [Section 5.3.1](#), using the TST as the basic predictor. As shown in [Table 7](#), adding contextual features leads to steady improvements across all evaluation metrics (MAE, RMSE, MCRPS, MIS).

Notable gains are obtained when dynamic weather and holiday information is introduced. Moving from the simple baseline to the global environment Transformer predictor yields a consistent reduction across all error metrics. Incorporating facility and public transport signals (i.e., station capacity attributes and real-time metro flow deviations) further improves interval quality, reflected by a lower MIS, although MAE, RMSE, and MCRPS increase slightly. Across all variants, the T-STAR framework provides the best overall performance and the most stable results, as indicated by consistently low standard deviations across stations and time intervals. Compared to one-stage Transformer with the same set of information, the contextual T-STAR structure enhances predictive capability and model robustness, as evident by higher forecasting accuracy and smaller standard deviations. Results suggest that performance gains are driven not only by richer contextual inputs but also by T-STAR's temporally aware design. Its distinctive two-stage structure effectively captures demand variation signals and enhances model's adaptability to recent short-term fluctuations.

Our findings highlight the added value in integrating diverse multivariate contextual inputs in short-term demand forecasting for bike-sharing systems. As additional sources of contextual information, ranging from environmental, infrastructural, intermodal, to variational signals introduced by T-STAR, are integrated step by step, we observe consistent improvements in both predictive accuracy and uncertainty estimation. Our analysis underscores the importance of leveraging diverse signals to capture the complex demand dynamics inherent to shared micro-mobility systems.

6.2.2. Performance enhancement via contextual two-stage framework

Next, we assess the added value of T-STAR's two-stage architecture for processing multivariate contextual inputs across different predictive models. Specifically, we compare the short-term pickup demand forecasting performance of models using two contextual input strategies, one with standard *Env., Facility and PT* input strategy and the other with the T-STAR framework. Both strategies use the same set of external contextual features, where the only difference lies in the structural design, as T-STAR requires no additional contextual inputs. By evaluating multiple base predictors under these two configurations, we isolate the impact of the T-STAR architecture on predictive accuracy.

Results in [Table 8](#) show that the two-stage T-STAR framework consistently outperforms the standard one-stage contextual fusion approach across all four base models: XGBoost, STGCN, TFT, and TST. This improvement is consistent across all key metrics and is also reflected in lower spatial and temporal variability, indicating enhanced point prediction accuracy and more reliable probabilistic forecasts. In addition, the performance boost is not limited to Transformer-based models. The decision-tree-based XGBoost and the GNN-based STGCN also benefit from the T-STAR framework. This suggests that the gains are not model-specific, but stem from the architectural decomposition and the use of intermediate variational representations, which can enhance learning even for non-parametric predictors. In particular, because STGCN is sensitive to the high noise-to-signal ratio in 15-minute, station-level demand forecasting, the two-stage T-STAR design, by reducing input dimensionality and progressively revealing salient demand structure, can limit noise propagation and improve predictive performance. Among all configurations, the T-STAR-enhanced Transformer model delivers the best overall performance.

The findings from this analysis reinforce the value of combining temporal attention mechanisms with structured contextual processing. By comparing two contextual fusion strategies using the same input data, we show that performance gains can be achieved without adding new features. These results highlight that the strategic processing of contextual information is a crucial consideration,

Table 8

Comparison of short-term pickup forecasting performance between standard (Environment and Facility) and two-stage (T-STAR) contextual input strategies across different base predictors (XGBoost, TFT, TST). Metric values are averaged over five runs with different random seeds for each predictive model. All models use the same set of external contextual features. The key difference lies in T-STAR's unique two-stage architecture for processing these inputs. The first value in parentheses indicates the standard deviation across stations, and the second value indicates the standard deviation across time intervals (aggregated across stations at each interval).

Basic Predictor	Input Types	MAE	RMSE	MCRPS	MIS
XGBoost	Env., Facility and PT	0.334 (0.068; 0.120)	0.459 (0.133; 0.205)	–	–
XGBoost	T-STAR	0.320 [†] (0.073; 0.120)	0.450[†] (0.133; 0.206)	–	–
STGCN	Env., Facility and PT	0.422 (0.049; 0.332)	0.713 (0.071; 0.342)	0.399 (0.029; 0.200)	7.465 (0.446; 1.277)
STGCN	T-STAR	0.381 [†] (0.064; 0.363)	0.701 [†] (0.088; 0.377)	0.372 [†] (0.044; 0.221)	7.014 [†] (0.726; 1.416)
TFT	Env., Facility and PT	0.156 (0.095; 0.155)	0.473 (0.171; 0.281)	0.134 (0.067; 0.115)	2.123 (0.653; 1.142)
TFT	T-STAR	0.154 [†] (0.095; 0.154)	0.470 [†] (0.176; 0.286)	0.132 [†] (0.068; 0.114)	2.060 [†] (0.635; 1.157)
TST	Env., Facility and PT	0.156 (0.091; 0.151)	0.472 (0.163; 0.276)	0.142 (0.060; 0.107)	2.009 (0.476; 0.910)
TST	T-STAR	0.151[†] (0.089; 0.148)	0.464 [†] (0.156; 0.271)	0.128[†] (0.062; 0.105)	1.864[†] (0.508; 0.925)

(1) For each basic predictor, the best-performing value for each metric is marked with [†] to indicate the best model.

(2) Across all models, the best-performing values for each metric are highlighted in bold in the table.




Fig. 10. Spatial overview and demand comparison for zero-shot forecasting stations in the Arlington service area. Sub-figure (a) Locations of bike-sharing stations (blue circles) and metro stations (red circles) in Arlington, the unseen service area for the zero-shot forecasting experiment. Sub-figure (b) A daily average pickup demand heatmap that shows demand at the Arlington stations is generally lower than in most previously observed areas in Washington D.C.

next to the selection of contextual inputs. They also underscore the advantages of the two-stage T-STAR architecture, which effectively integrates diverse long-term and short-term signals to capture the complex and dynamic nature of shared micro-mobility demand.

6.2.3. Model generalizability and capability for zero-shot forecasting

In rapidly expanding micro-mobility systems, new service points are frequently introduced, both within existing cities as well as in entirely new markets. During these early launch phases, providers are under pressure to deliver reliable service from the very beginning. This is especially true in competitive environments or when the brand already carries strong market expectations. To meet these expectations, operators benefit from forecasting models that require no station-specific historical data yet still provide accurate predictions. A model with zero-shot forecasting capability offers this flexibility and can be deployed immediately at new locations. Moreover, it supports pre-launch scenario testing and simulation, helping providers make informed operational decisions even before a station becomes active. Generalizable, zero-shot models are therefore valuable tools for scalable and data-efficient service planning. Recent research has shown that Transformer-based models exhibit strong zero-shot capabilities across a range of forecasting applications (Dooley et al., 2023).

To evaluate the zero-shot forecasting capability of the proposed framework, we simulate the scenario of service expansion into a new geographic area. The original model was trained on 235 bike-sharing stations located in the Washington D.C. region, all situated to the northeast of the Potomac River. However, both the WMATA metro and Capital Bikeshare services also operate beyond the D.C. boundary, including the Arlington region to the southwest of the river which lies within the state of Virginia. In this experiment, we treat Arlington as a newly added service area and select 11 bike-sharing stations located there as the test set. As shown in Fig. 10, many of these stations are situated near WMATA metro stations, making the setting realistic and operationally relevant.

The experiment deploys the pre-trained predictors trained solely on data from the 235 D.C. stations. No retraining or fine-tuning is performed for the Arlington stations. Instead, these predictors are applied directly, relying only on the station information and real-time inputs specified in the predictive framework. In this zero-shot setting, we assume no historical demand data is available for the 11 new stations, so forecasts must be generated without station-specific demand history. For station-level embeddings, we assign

Table 9

Zero-shot per 15min pickup demand forecasting performance for 11 newly added Arlington stations using the pre-trained T-STAR model, compared with Historical Average and Myopic baselines. Metric values are averaged over five runs with different random seeds for each predictive model (except the Historical Average, Myopic approaches, and the pretrained zero-shot Lag-llama predictor).

The first value in parentheses indicates the standard deviation across stations, and the second value indicates the standard deviation across time intervals (aggregated across stations at each interval).

Predictive Model	Input Types	MAE	RMSE	MCRPS	MIS
Historical Average	<i>Baseline</i>	0.187 (0.032; 0.159)	0.320 (0.064; 0.215)	–	–
Myopic	<i>Last demand observation of each station</i>	0.105 (0.025; 0.150)	0.379 (0.066; 0.300)	–	–
Zero-Shot XGBoost	<i>Env., Facility and PT</i>	0.270 (0.014; 0.083)	0.332 (0.037; 0.148)	–	–
Zero-Shot Lag-Llama	<i>Previous demand observations of each station</i>	0.075 (0.020; 0.119)	0.307 (0.055; 0.255)	0.069 (0.018; 0.110)	2.357 (0.544; 4.038)
Zero-Shot DeepAR	<i>Env., Facility and PT</i>	0.063 (0.017; 0.108)	0.292 (0.055; 0.249)	0.059 (0.015; 0.094)	1.301 (0.213; 2.123)
Zero-Shot TFT	<i>Env., Facility and PT</i>	0.062 (0.017; 0.108)	0.291 (0.054; 0.200)	0.065 (0.015; 0.089)	1.487 (0.193; 1.374)
Zero-Shot Transformer-Based T-STAR	<i>T-STAR</i>	0.061 (0.016; 0.108)	0.290 (0.052; 0.249)	0.060 (0.012; 0.085)	1.367 (0.126; 1.268)

The best-performing values for each metric are highlighted in bold in the table.

each new station the mean embedding learned across the training stations, representing a typical demand characteristics in the D.C. network. Therefore, this setup emulates a practical deployment scenario in which an operator expands into a new area and needs immediate forecasting capability without historical data.

To assess model's generalization performance, we evaluate forecasts for the same testing period as previous experiments, focusing on the 11 Arlington stations. We benchmark the zero-shot Transformer-based T-STAR model against *Historical Average*, *Myopic*, *XG-Boost*, *DeepAR*, and *TFT*. Historical Average predictor estimates are based on the historical demand pattern of the new stations from the training period.

STGCN and STAEformer were excluded from zero-shot evaluations as they are inherently transductive. STGCN requires a pre-defined adjacency matrix tied to specific graph topologies, while STAEformer relies on node-specific learnable embeddings. These architectures lack the inductive capability to map static features of unseen stations to predictive representations without retraining. Unlike the other contextual predictors, they cannot effectively generalize to new locations where specific node IDs or spatial structures were not encountered during training.

Foundation time series models are large-scale predictors pretrained on massive, cross-domain datasets to capture universal temporal dependencies, allowing them to be transferred to downstream forecasting tasks without task-specific training (Bommasani et al., 2021). In this study, we also include *Lag-Llama*, a decoder-only transformer-based foundation model, as a competitive zero-shot benchmark (Rasul et al., 2023). For our application, Lag-Llama utilizes the eight most recent demand observations per station as its sole input, aligning with the historical input window used by our other zero-shot predictors.

The zero-shot forecasting results are summarized in Table 9. Overall, contextual models with an explicit temporal learning component (DeepAR, TFT, and T-STAR) transfer substantially better to unseen stations than naive or context-only baselines (Historical Average, Myopic, and zero-shot XGBoost). Interestingly, these pretrained temporal-contextual predictors also consistently outperform Lag-Llama across all four metrics. This suggests that while foundation models effectively capture universal temporal dependencies, the integration of domain-specific contextual knowledge remains vital for urban micro-mobility systems. Specifically, relying exclusively on historical lags, as Lag-Llama does, overlooks the distinct spatial and environmental profiles of new stations.

Among the temporal deep models, performance differences are small. T-STAR achieves the best point accuracy (MAE and RMSE). DeepAR attains the lowest MCRPS and MIS, while T-STAR delivers comparable probabilistic accuracy with similar variability across stations and time intervals. Conceptually, in zero-shot deployment, the single-stage forecaster can be biased because station embeddings for new stations are not learned (which is represented by a mean embedding from the known stations). T-STAR mitigates this via stage 2, which conditions on recent deviation signals to adjust the forecasts accordingly. The same mechanism can also stabilize forecasts under abrupt demand regime shifts of stations once post-change deviations are observed. Overall, these results support T-STAR's reliable generalization in realistic service expansion scenarios where forecasting is required immediately without station-specific historical demand.

While the zero-shot forecasting results highlight the strong generalization capability of the Transformer-based T-STAR model, its flexibility extends beyond initial deployment. As new service areas begin operating and collect data, the demand patterns of newly added stations become more distinguishable. As shown in Fig. 10, Arlington stations exhibit relatively low daily pickup demand compared to most of D.C., resembling patterns seen among the lower-demand D.C. stations. In such cases, the model can be further refined through few-shot learning with limited data to fine-tune station-level embeddings and improve forecast accuracy.

Beyond zero-shot forecasting, the generalizability of Transformer-based models offers key advantages for long-term operations. Urban shared mobility markets are highly dynamic, with demand patterns continuously shifting due to changes in public transit,

new points of interest, and the emergence of competing services. Transformer-based T-STAR architecture is deployment-ready to these evolving conditions thanks to their ability to support online learning and incremental model updates. These updates are both efficient, as they reuse prior training knowledge, and adaptive, since they require only a small amount of new data. This enables the model to respond not only to newly added service areas, but also to shifting demand patterns of current operational areas.

7. Conclusions

This study addresses the challenge of reliable and accurate short-term forecasting of station-level demand in dock-based bike-sharing systems, particularly under sparse, zero-inflated, and highly variable time series. To tackle this problem, we introduced a Two-stage Spatial and Temporal Adaptive contextual Representation framework that combines multivariate contextual inputs with a structured, sequential forecasting approach. By decomposing the task into two stages, T-STAR captures both long-term demand patterns and short-term fluctuations, through efficient process of coarse-grained and high-frequency contextual signals. Built on a Transformer architecture, the global T-STAR predictor enables real-time forecasting across large-scale bike-sharing networks, while adapting effectively to varying urban demand conditions. To evaluate the proposed framework, we conducted 15-minute demand forecasting experiments using real-world data from Washington D.C.'s Capital Bikeshare. We assessed both deterministic and probabilistic forecasting across different demand intensities, time periods, and station locations. The Transformer-based T-STAR model consistently delivered robust and accurate predictions, particularly at high-demand stations and under irregular demand conditions.

From a managerial perspective, this study highlights the practical value of adopting a hierarchical forecasting framework to effectively process the diverse contextual signals inherent in shared micro-mobility systems. Leveraging long-term demand trends, T-STAR enhances the robustness of short-term forecasts through reduced training noise. Incorporating rich external context—such as weather, infrastructure, and real-time intermodal dynamics—further improves forecast accuracy and resilience. These findings emphasize the importance of strategic data sharing and collaboration between public transport and shared mobility providers to strengthen multimodal integration and enhance the user experience within MaaS ecosystems. Transformer-based T-STAR's strong generalization capabilities also support zero-shot forecasting for new service areas and enable few-shot adaptation as data becomes available, offering crucial flexibility for fast-growing, dynamic urban mobility networks. Furthermore, integrating fine-grained, real-time forecasts into operational decision-making presents a valuable opportunity. Applications such as dynamic fleet rebalancing (Zhao et al., 2025) and multimodal trip coordination via MaaS platforms (Zhang et al., 2025) can greatly benefit from these predictive insights, improving the efficiency and reliability of shared mobility services in evolving urban environments.

Future research could build upon this study in several key dimensions, addressing both the current data constraints and potential architectural extensions. While the rolling-origin and sliding-window evaluations demonstrate the robustness of our results across different temporal splits, the three-month dataset limits the assessment of long-term predictive performance under diverse seasonal related impacts, such as weather changes, holiday periods, and tourism-driven demand changes. Therefore, a key future research direction involves extending the T-STAR framework to multi-year datasets to capture broader cyclical patterns and seasonal shifts. Beyond temporal coverage, future work could evolve T-STAR into a unified multi-task learning framework. Rather than modeling pickup and drop-off demand independently, a joint formulation could simultaneously forecast both demand streams and station-level fleet availability, once historical inventory observations are accessible. Such integration would better capture the coupled spatio-temporal dynamics of supply and demand. In particular, conditioning on real-time fleet distribution would enable direct estimation of station-wise net flow, providing an operationally meaningful signal for anticipating stock-outs and surpluses. A comprehensive predictor of this form could support earlier identification of rebalancing requirements and improve system reliability by enabling more informed, forward-looking control decisions. Future work can also strengthen T-STAR by coupling its two-stage temporal design with dynamic, attention-driven GNNs for high-granularity, intermittent demand. A learnable time-varying graph inferred from demand co-movement, operational interactions, and context (weather, events, commute regimes) would enable message passing only when informative. Attention-based neighbor selection and heterophily-robust layers can reduce noise, while interaction edges capture competition/substitution to improve fluctuation modeling and calibration. Finally, future research could explore the few-shot learning capabilities of the T-STAR framework and its generalizability across different service regions with limited historical data. While this study focuses on zero-shot forecasting, T-STAR's broader potential lies in its ability to adapt to new environments through lightweight fine-tuning. Investigating how the model can respond to evolving urban conditions would further demonstrate its value as a flexible and scalable tool for long-term shared micro-mobility management.

CRedit authorship contribution statement

Jingyi Cheng: Writing – original draft, Visualization, Software, Methodology, Investigation, Data curation, Conceptualization; **Gonçalo Homem de Almeida Correia:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization; **Oded Cats:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization; **Shadi Sharif Azadeh:** Writing – review & editing, Supervision, Project administration, Methodology, Funding acquisition, Conceptualization.

Acknowledgments

This research is supported by the project Seamless Shared Urban Mobility (SUM), which received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 101103646. The authors would also like to thank the Washington Metropolitan Area Transit Agency (WMATA) for their support in providing the data used in this study.

Data availability

A coding and application demo of the transformer-based T-STAR model is provided through a public GitHub repository associated with this study.

Appendix A. Negative binomial distribution

The Negative Binomial (NB) distribution is a discrete probability distribution used to model the number of event occurrences (e.g., bike pick-ups or drop-offs) within a fixed time interval. Unlike the Poisson distribution, which assumes equal mean and variance, the NB distribution includes a dispersion parameter that allows for greater variance. This flexibility makes it well-suited for over-dispersed count data and enables more accurate and robust probabilistic forecasting.

Assuming k is the number of observed demand, r is the shape parameter that reflecting the degree of demand variability, p is the probability parameter relating to the demand intensity level, the probability mass function of Negative Binomial distribution is given by,

$$P(X = k) = \binom{k+r-1}{k} (1-p)^r p^k.$$

The expected demand μ , which is the mean of a NB distribution, is expressed as:

$$\mu = \frac{r \cdot p}{1-p}.$$

Substituting p in terms of μ , we obtain:

$$p = \frac{\mu}{r + \mu}.$$

Thus, the probability mass function can be rewritten in terms of the expected demand μ and the shape parameter r , both of which are estimated by the Time Series Transformer model, as:

$$P(X = k) = \binom{k+r-1}{k} \left(\frac{r}{r+\mu} \right)^r \left(\frac{\mu}{r+\mu} \right)^k.$$

The variance of Negative Binomial distribution is given by $\mu + r^{-1}\mu^2$, where μ is the mean value.




Fig. A.1. Negative binomial distributions for different pairs of μ and r .

Appendix B. Graph construction strategies for STGCN and STGCN-VAE

Following the framework of Yu et al. (2017), we represent the bike-sharing network as a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where $\mathcal{V}$ denotes the set of N stations and the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ encodes the underlying spatial dependencies. To ensure architectural consistency with our Transformer-based predictors, station-specific attributes are incorporated as exogenous node features in both models, enabling the learning of high-dimensional latent station embeddings. In order to adapt STGCN for probabilistic demand forecasting, we appended a distributional output head to the final layer, allowing the model to parameterize a predictive distribution (e.g., Gaussian or Poisson). The parameters of this distribution are optimized by minimizing the NLL. In contrast, for STGCN-VAE,

probabilistic density is captured through the inherent stochasticity of the Variational Autoencoder framework, where the decoder reconstructs the future demand distribution based on the sampled latent variables.

To ensure a robust evaluation of our STGCN and STGCN-VAE benchmarks, we systematically investigated three distinct strategies for constructing the static adjacency matrix A : (i) Location Proximity, (ii) Trip-Flow Connectivity, and (iii) Demand Correlation. These strategies allow the models to capture dependencies driven by geographic distance, physical movement patterns, and functional usage similarity, respectively.

Location proximity graph

In the first graph construction strategy, the static adjacency matrix is calculated based on geographic proximity. Specifically, we employ an Inverse Distance Weighting (IDW) scheme with a local threshold (1000m) to define the edge weights A_{ij} between stations i and j :

$$A_{ij} = \begin{cases} \frac{1}{d_{ij}}, & \text{if } 0 < d_{ij} \leq 1000\text{m} \\ 0, & \text{otherwise} \end{cases} \quad (\text{B.1})$$

where d_{ij} denotes the pairwise geodesic distance in meters. This construction, adapted from [Liang et al. \(2022\)](#), ensures the model captures strong local spatial correlations while maintaining a sparse adjacency matrix to mitigate computational overhead.

Trip-flow based graph

The second strategy constructs the adjacency matrix based on the observed origin-destination (OD) trip flows, capturing the latent connectivity induced by the physical movement of the bike fleet. Let F_{ij} denote the total number of trips originating at station i and terminating at station j during the training period. Since STGCN typically operates on undirected graphs, we define a symmetric connectivity weight W_{ij} by aggregating bidirectional flows:

$$W_{ij} = F_{ij} + F_{ji} \quad (\text{B.2})$$

This matrix W reflects the global intensity of direct interactions between station pairs. To ensure a sparse and computationally efficient graph structure, we apply a top- K selection rule. The final adjacency matrix A is constructed by retaining only the strongest connections for each node:

$$A_{ij} = \begin{cases} W_{ij}, & \text{if } W_{ij} \in \text{top-}K \text{ connections for station } i \text{ or } j \\ 0, & \text{otherwise} \end{cases} \quad (\text{B.3})$$

where the A_{ij} is non-zero when the trip flow connection is among the top- K for either station i or j . Parameter k is selected to be five in our case study.

Demand correlation graph

In the final strategy, spatial dependencies are derived from the statistical correlations between station-level demand time series. This approach identifies functional similarities between stations that may be geographically distant but exhibit synchronized usage patterns. For each station pair (i, j) , we compute the Pearson correlation coefficient $\rho_{i,j}$ based on their historical demand profiles, following the design of [Ke et al. \(2021\)](#). To focus on meaningful positive dependencies and ensure network sparsity, we define the adjacency matrix A using a percentile-based threshold τ :

$$A_{i,j} = \begin{cases} \rho_{i,j} & \text{if } \rho_{i,j} \geq \tau \text{ and } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

where τ is set to the 95th percentile of all positive off-diagonal correlation values. By retaining only the top 5% of correlations, the resulting graph effectively isolates the most significant functional relationships within the network.

Comparative performance

The comparative performance of STGCN and STGCN-VAE across the three graph construction strategies is summarized in [Tables B.1](#) and [B.2](#). To ensure a fair and robust comparison, we conducted extensive hyperparameter tuning for each model-graph configuration, and all reported metrics represent the average of five independent runs with different random seeds. Our findings indicate that for the standard STGCN, the Location Proximity graph consistently yielded the most stable and accurate results, outperforming both trip-flow and demand-correlation configurations. Conversely, the probabilistic STGCN-VAE exhibited greater resilience to the underlying graph topology, with all three strategies achieving significantly improved and comparable accuracy. Notably, the combination with a Demand Correlation graph provided the best performance for STGCN-VAE in drop-off forecasting, while the one with Trip-Flow graph was slightly superior for pickups. For the final results reported in [Section 6](#), we have selected the best-performing graph construction choice for both STGCN and STGCN-VAE to reflect the best performance of these models.

Table B.1

Pickup forecasting performance of STGCN and STGCN-VAE across different graph construction strategies.

Model	Graph Type	Pickup			
		MAE	RMSE	CRPS	IS
STGCN	Location Proximity Graph	0.422 (0.049, 0.332)	0.713 (0.071, 0.342)	0.399 (0.029, 0.200)	7.465 (0.446, 1.142)
STGCN	Trip Flow Connection Graph	0.650 (0.066, 0.514)	0.986 (0.084, 0.502)	0.577 (0.052, 0.307)	9.889 (0.938, 2.249)
STGCN	Demand Correlation Graph	0.545 (0.117, 0.477)	0.897 (0.122, 0.475)	0.495 (0.075, 0.297)	8.416 (0.735, 2.139)
STGCN-VAE	Location Proximity Graph	0.405 (0.057, 0.087)	0.501 (0.124, 0.175)	0.402 (0.057, 0.087)	15.832 (2.282, 3.456)
STGCN-VAE	Trip Flow Correlation Graph	0.404 (0.057, 0.087)	0.500 (0.124, 0.175)	0.398 (0.057, 0.087)	15.419 (2.286, 3.463)
STGCN-VAE	Demand Correlation Graph	0.410 (0.056, 0.085)	0.503 (0.122, 0.173)	0.402 (0.056, 0.085)	15.398 (2.246, 3.394)

Table B.2

Drop-off forecasting performance of STGCN and STGCN-VAE across different graph construction strategies.

Model	Graph Type	Drop-off			
		MAE	RMSE	CRPS	IS
STGCN	Location Proximity Graph	0.602 (0.066, 0.479)	0.917 (0.069, 0.461)	0.535 (0.036, 0.283)	8.969 (0.201, 1.946)
STGCN	Trip Flow Connection Graph	0.792 (0.026, 0.521)	1.107 (0.039, 0.493)	0.695 (0.024, 0.307)	11.388 (0.753, 3.522)
STGCN	Demand Correlation Graph	0.881 (0.163, 0.780)	1.340 (0.181, 0.757)	0.725 (0.104, 0.491)	10.418 (0.918, 3.353)
STGCN-VAE	Location Proximity Graph	0.405 (0.058, 0.089)	0.503 (0.128, 0.184)	0.404 (0.058, 0.089)	16.046 (2.306, 3.557)
STGCN-VAE	Trip Flow Correlation Graph	0.411 (0.057, 0.088)	0.506 (0.127, 0.181)	0.405 (0.057, 0.088)	15.739 (2.272, 3.498)
STGCN-VAE	Demand Correlation Graph	0.395 (0.059, 0.092)	0.498 (0.132, 0.188)	0.390 (0.059, 0.092)	15.171 (2.369, 3.662)

Appendix C. Hyperparameter tuning for predictive models

This section presents the hyperparameter ranges explored during model tuning, along with the final selected values for each predictive model. All models use a consistent look-back window of 30 time steps, meaning the previous 30 observations from the target time series are used as input for forecasting. The hyperparameter tuning procedures follow the methods described in [Section 5.3.2](#). Below, we outline the specific hyperparameters tuned during the experiments and the corresponding values chosen based on validation performance.

Time series transformer

Both stages of the Transformer-based T-STAR model, as well as the benchmark Simple TST model, were implemented using the Hugging Face library. Three hyperparameters are selected for fine-tuning: *number of encoder layers* (N_E) determines the number of encoder blocks included in TST, *hidden size* sets the hidden layer dimension for the self-attention layer in a Transformer model, *dropout rate* is a regularization hyperparameter applied to avoid overfitting, which controls the ratio of neurons being randomly deactivated during a forward pass. The *learning rate* controls how quickly the model updates its weights during training process. The value ranges explored for key architectural hyperparameters are summarized in [Table C.1](#), along with the final selected values for each model.

Table C.1

Tuning ranges and selected hyperparameter values for time series Transformer-based models.

Hyperparameter	Value Range	Pickup			Dropoff		
		T-STAR: Stage 1	T-STAR: Stage 2	Simple TST	T-STAR: Stage 1	T-STAR: Stage 2	Simple TST
N_E	[1, 2, 3]	1	3	2	2	2	2
hidden size	[16, 32, 64]	64	16	64	64	64	64
dropout rate	[0.1, 0.2, 0.3]	0.1	0.1	0.3	0.2	0.2	0.2
learning rate	$[5 \times 10^{-5}, 0.01]$	6×10^{-4}	–	2.8×10^{-4}	1.5×10^{-4}		9.5×10^{-5}

N_E : number of encoder layers.

XGBoost

Five key hyperparameters were fine-tuned for the XGBoost model. The *maximum depth* parameter controls the maximum depth of each decision tree. While deeper trees can capture more complex patterns, they also increase the risk of overfitting. The *number of estimators* defines the total number of trees used in the boosting process. To introduce randomness and prevent overfitting, the *subsampling* parameter specifies the fraction of the training data used to grow each tree. Similarly, *colsample bytree* controls the proportion of input features randomly selected for each tree, promoting model diversity and robustness. The *learning rate* determines how much each tree contributes to the final prediction. A smaller learning rate results in more stable learning and better generalization but requires more trees, whereas a larger learning rate speeds up training at the cost of a higher risk of overfitting.

Table C.2

Hyperparameter tuning ranges and selected values for XGBoost model.

Hyperparameter	Value Range	Pickup	Dropoff
maximum depth	[10, 25, 30, 45, 60]	30	30
number of estimators	[50, 100, 150, 200]	150	150
subsample	[0.3, 0.5, 0.7, 1.0]	1.0	1.0
colsample bytree	[0.3, 0.5, 0.7, 1.0]	0.7	0.7
learning rate	$[5 \times 10^{-5}, 0.01]$	8×10^{-3}	9.6×10^{-3}

Table C.3

Hyperparameter tuning ranges and selected values for DeepAR model.

Hyperparameter	Value Range	Pickup	Dropoff
number of layers	[1, 2, 3]	3	3
hidden size	[32, 64, 128]	128	32
dropout rate	[0.1, 0.2, 0.3]	0.1	0.2
learning rate	$[5 \times 10^{-5}, 0.01]$	8×10^{-5}	1.3×10^{-4}

Table C.4

Hyperparameter tuning ranges and selected values for Temporal Fusion Transformer (TFT).

Hyperparameter	Value Range	Pickup	Dropoff
number of attention heads	{1, 2, 3}	1	3
hidden size	$[4, 8, 16, 32] \times \text{heads}$	32	16
embedding dimension	$[4, 16, 32] \in \mathbb{N}$	16	16
dropout rate	[0.1, 0.2, 0.3]	0.1	0.3
learning rate	$[5 \times 10^{-5}, 0.01]$	8×10^{-5}	1.3×10^{-4}

Table C.2 presents the tuning ranges of the hyperparameters for the XGBoost predictor, along with the values selected through the fine-tuning process.

DeepAR

The DeepAR model implemented via GluonTS was fine-tuned using three key hyperparameters. The *number of layers* parameter specifies the number of RNN layers in the model, determining the depth of temporal learning. The *hidden size* (also referred to as number of cells) defines the dimensionality of the hidden state in RNN layer. Table C.3 summarizes the value ranges used during tuning and the selected values of these hyperparameters for the DeepAR model.

Temporal fusion transformer (TFT)

For TFT predictors, we tuned the classic neural network hyperparameters *learning rate*, *dropout rate*, and three hyperparameters specified for the Transformer structure. The first structural hyperparameter determines the *number of attention heads* in self-attention layer in the decoder, allowing the model to focus on different aspects of the input sequence. The *hidden size* controls the size of hidden states for LSTM and Transformer layers. The *hidden size* should be divisible by number of attention heads. The *embedding dimension* sets the size of feature embedding tensors. Table C.4 presents the tuned hyperparameters for TFT, including their search ranges and selected values.

Spatio-temporal graph convolutional network (STGCN) and the variational autoencoder variant (STGCN-VAE)

For the STGCN and STGCN-VAE models, in addition to the *learning rate* and *dropout rate*, we tuned the *hidden channels* and the *number of STGCN blocks*. The *number of STGCN blocks* corresponds to the number of spatio-temporal convolution blocks stacked sequentially, and thus controls the depth of the network. Increasing this depth can improve the model's ability to capture higher-order and longer-range spatio-temporal dependencies, but it also increases computational cost and may raise the risk of overfitting. The *hidden channels* parameter specifies the dimensionality of the latent feature representation within the STConvBlock backbone, thereby controlling the capacity of the model's internal embeddings for each node and time step. Table C.5 shows the fine-tuned hyperparameters for STGCN models, and Table C.6 shows those for STGCN-VAE models. Note that the number of STGCN blocks in STGCN-VAE is fixed to be two according to Peng et al. (2025).

Table C.5

Hyperparameter tuning ranges and selected values for Spatio-Temporal Graph Convolutional Network (STGCN).

Hyperparameter	Value Range	Pickup	Dropoff
N_B	[1, 2, 3]	1	3
hidden channels	[16, 32, 64]	16	16
dropout rate	[0.1, 0.2, 0.3]	0.1	0.1
learning rate	$[5 \times 10^{-5}, 0.01]$	5.4×10^{-5}	5.6×10^{-5}

 N_B : number of STGCN blocks in the predictive model.**Table C.6**

Hyperparameter tuning ranges and selected values for Spatio-Temporal Graph Convolutional Network Variational Autoencoder (STGCN-VAE).

Hyperparameter	Value Range	Pickup	Dropoff
hidden channels	[16, 32, 64]	64	64
dropout rate	[0.1, 0.2, 0.3]	0.2	0.2
learning rate	$[5 \times 10^{-5}, 0.01]$	9.5×10^{-5}	2.5×10^{-4}

 N_B : number of STGCN blocks in the predictive model.**Table C.7**

Hyperparameter tuning ranges and selected values for Spatio-Temporal Adaptive Embedding Transformer (STAEformer).

Hyperparameter	Value Range	Pickup	Dropoff
N_S	[1, 2, 3]	2	1
spatial embedding size	[8, 16, 32]	16	32
adaptive embedding size	[8, 16, 32]	8	8
dropout rate	[0.1, 0.2, 0.3]	0.2	0.3
learning rate	$[5 \times 10^{-5}, 0.01]$	1.5×10^{-4}	5×10^{-4}

 N_S : number of self-attention layers in the predictive model.

Spatio-temporal adaptive embedding transformer (STAEformer)

In addition to the standard *learning rate* and *dropout rate*, we tuned three architecture-specific hyperparameters for STAEformer. The *number of self-attention layers* controls the encoder depth by specifying how many self-attention blocks are stacked in the temporal-attention module and the spatial-attention module. The *spatial embedding size* determines the dimensionality of the learned static node (station) embeddings, which provide a node-identity representation shared across time. The *adaptive embedding size* sets the dimensionality of the learned time-by-node adaptive embeddings, which supply an additional spatio-temporal latent representation concatenated to the input features. Table C.7 presents the tuning ranges and selected values for these hyperparameters.

Appendix D. Robustness analysis of transformer-based T-STAR

To evaluate the robustness of the Transformer-Based T-STAR predictor, we performed a rolling forecast origin (expanding-window) experiment and sliding window cross-validation using the Washington D.C. dataset. These evaluations are designed to test the model's sensitivity to training data volume and its stability across varying temporal contexts. To isolate the benefits of the T-STAR architecture, we benchmark its performance against a Contextual TST baseline, which incorporates the same exogenous features (Environment, Facility, and Public Transport inputs) but lacks the T-STAR structural design.

The rolling forecast origin experiment simulates a real-world deployment pipeline where the model is periodically retrained as new data is ingested. We initialized the training set using the first four weeks (w_1-w_4) and expanded the horizon incrementally to ten weeks (w_1-w_{10}). To prevent data leakage, each model was evaluated on the two-week period immediately following the training boundary. As detailed in Table D.1, MAE generally follows a downward trend for both models as the training history increases, confirming the expected benefit of larger data volumes. Notably, however, the T-STAR Transformer consistently outperforms the Contextual TST baseline across every fold. The performance margin is particularly significant in the initial folds (1st and 2nd), where training data is most constrained. This demonstrates that T-STAR Transformer is not only more accurate overall but also significantly more data-efficient, maintaining robust predictive performance even when historical observations are sparse.

To assess the model's sensitivity to short-term distribution shifts and varying historical contexts, we conducted a sliding window cross-validation. This setup utilizes a fixed 8-week training duration and a 2-week validation period, advancing the window forward by one week for each subsequent fold. This approach ensures that the model is tested against different temporal segments of the demand distribution. As shown in Table D.2, although forecast difficulty fluctuates across folds (likely reflecting localized events or inherent seasonality in urban mobility), T-STAR Transformer delivers uniformly lower MAE and reduced station-level standard

Table D.1

Rolling forecast origin robustness results. Each MAE entry reports the mean MAE across stations. Values in parentheses are the station-wise standard deviation. w_i denotes the i^{th} week in the Washington D.C. case study dataset.

Fold	Train weeks	Validation weeks	Pickup MAE		Drop-off MAE	
			Contextual TST	T-STAR Transformer	Contextual TST	T-STAR Transformer
1 th	w_1-w_4	w_5-w_6	0.365 (0.175)	0.326 (0.158)	0.361 (0.173)	0.317 (0.152)
2 th	w_1-w_6	w_7-w_8	0.237 (0.123)	0.210 (0.113)	0.469 (0.263)	0.207 (0.115)
3 th	w_1-w_8	w_9-w_{10}	0.253 (0.146)	0.238 (0.129)	0.253 (0.143)	0.239 (0.133)
4 th	w_1-w_{10}	$w_{10}-w_{11}$	0.172 (0.106)	0.163 (0.096)	0.183 (0.105)	0.163 (0.094)

Table D.2

Sliding forecast window cross validation results. Each MAE entry reports the mean MAE across stations. Values in parentheses are the station-wise standard deviation. w_i denotes the i^{th} week in the Washington D.C. case study dataset.

Fold	Train weeks	Validation weeks	Pickup MAE		Drop-off MAE	
			Contextual TST	T-STAR Transformer	Contextual TST	T-STAR Transformer
1st	w_1-w_8	w_9-w_{10}	0.253 (0.146)	0.238 (0.129)	0.253 (0.143)	0.239 (0.133)
2nd	w_2-w_9	$w_{10}-w_{11}$	0.273 (0.159)	0.228 (0.130)	0.256 (0.134)	0.225 (0.127)
3rd	w_3-w_{10}	$w_{11}-w_{12}$	0.174 (0.108)	0.166 (0.098)	0.175 (0.102)	0.169 (0.104)
4 th	w_4-w_{11}	$w_{12}-w_{13}$	0.134 (0.078)	0.121 (0.071)	0.124 (0.076)	0.119 (0.072)

deviation compared to the Contextual TST benchmark. This consistent performance across shifting temporal origins underscores T-STAR's superior ability to generalize across evolving demand patterns while maintaining high spatial precision. The reduction in standard deviation also indicates that T-STAR provides more reliable and equitable predictions across the entire station network.

Across both evaluations, T-STAR Transformer consistently maintains a decisive performance edge over the Contextual TST benchmark. The simultaneous reduction in mean absolute error and station-wise variance confirms that the T-STAR architecture is not only more accurate in its point estimates but also significantly more stable under diverse operational conditions.

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M., 2019. Optuna: a next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2623–2631.
- Albiński, S., Fontaine, P., Minner, S., 2018. Performance analysis of a hybrid bike sharing system: a service-level-based approach under censored demand observations. *Transp. Res. Part E: Logist. Transp. Rev.* 116, 59–69.
- Ashqar, H.I., Elhenawy, M., Rakha, H.A., Almanna, M., House, L., 2022. Network and station-level bike-sharing system prediction: a san francisco bay area case study. *J. Intell. Transp. Syst.* 26 (5), 602–612.
- Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al., 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Cats, O., Ferranti, F., 2022. Unravelling individual mobility temporal patterns using longitudinal smart card data. *Res. Transp. Bus. Manag.* 43, 100816.
- Chen, L., Zhang, D., Wang, L., Yang, D., Ma, X., Li, S., Wu, Z., Pan, G., Nguyen, T.-M.-T., Jakubowicz, J., 2016. Dynamic cluster-based over-demand prediction in bike sharing systems. In: Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing, pp. 841–852.
- Chen, T., Guestrin, C., 2016. Xgboost: a scalable tree boosting system. In: Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining, pp. 785–794.
- Chen, X., Jiang, H., 2020. Detecting the demand changes of bike sharing: a bayesian hierarchical approach. *IEEE Trans. Intell. Transp. Syst.* 23 (5), 3969–3984.
- Cheng, J., Sharif Azadeh, S., 2025. A short-term predict-then-cluster framework for meal delivery services. *Data Sci. Transp.* 7 (3), 26.
- Cheng, L., Lei, D., Tao, S., 2024. Recent advances in deep learning for traffic probabilistic prediction. *Transp. Rev.* 44 (6), 1129–1135.
- Cheng, L., Mi, Z., Coffman, D., Meng, J., Liu, D., Chang, D., 2021. The role of bike sharing in promoting transport resilience. *Netw. Spat. Econ.* 22 (3), 1–19.
- Chiariotti, F., Pielli, C., Zanella, A., Zorzi, M., 2020. A bike-sharing optimization framework combining dynamic rebalancing and user incentives. *ACM Trans. Auton. Adapt. Syst.* 14 (3), 1–30.
- Choi, D., Min, K., 2022. Hierarchical latent structure for multi-modal vehicle trajectory forecasting. In: European Conference on Computer Vision. Springer, pp. 129–145.
- Cleophas, C., Cottrill, C., Ehmke, J.F., Tierney, K., 2019. Collaborative urban transportation: recent advances in theory and practice. *Eur. J. Oper. Res.* 273 (3), 801–816.
- Dai, B., Wang, Z., Wipf, D., 2020. The usual suspects? reassessing blame for VAE posterior collapse. In: International Conference on Machine Learning. PMLR, pp. 2313–2322.
- Dooley, S., Khurana, G.S., Mohapatra, C., Naidu, S.V., White, C., 2023. Forecastpf: synthetically-trained zero-shot forecasting. *Adv. Neural Inf. Process. Syst.* 36, 2403–2426.
- Durbin, J., Koopman, S.J., 2012. Time series analysis by state space methods. Oxford University Press (UK).
- Eren, E., Uz, V.E., 2020. A review on bike-sharing: the factors affecting bike-sharing demand. *Sustain. Cities Soc.* 54, 101882.
- Faghih-Imani, A., Eluru, N., El-Geneidy, A.M., Rabbat, M., Haq, U., 2014. How land-use and urban form impact bicycle flows: evidence from the bicycle-sharing system (BIXI) in montreal. *J. Transp. Geogr.* 41, 306–314.
- Feng, J., Liu, H., 2024. An adaptive spatial-temporal method capturing for short-term bike-sharing prediction. *IEEE Trans. Intell. Transp. Syst.* 25 (11), 16761–16774.
- Gammelli, D., Wang, Y., Prak, D., Rodrigues, F., Minner, S., Pereira, F.C., 2022. Predictive and prescriptive performance of bike-sharing demand forecasts for inventory management. *Transp. Res. C: Emerg. Technol.* 138, 103571.
- Gao, C., Chen, Y., 2022. Using machine learning methods to predict demand for bike sharing. In: Information and Communication Technologies in Tourism 2022: Proceedings of the ENTER 2022 eTourism Conference, January 11–14, 2022. Springer, pp. 282–296.
- Gast, N., Massonnet, G., Reijbergen, D., Tribastone, M., 2015. Probabilistic forecasts of bike-sharing systems for journey planning. In: Proceedings of the 24th ACM International Conference on Information and Knowledge Management, pp. 703–712.
- Gebhart, K., Noland, R.B., 2014. The impact of weather conditions on bikeshare trips in washington, DC. *Transportation* 41, 1205–1225.

- Giot, R., Cherrier, R., 2014. Predicting bikeshare system usage up to one day ahead. In: 2014 IEEE Symposium on Computational Intelligence in Vehicles and Transportation Systems (CIVTS). IEEE, pp. 22–29.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* 102 (477), 359–378.
- Ham, S.W., Cho, J.-H., Park, S., Kim, D.-K., 2021. Spatiotemporal demand prediction model for e-scooter sharing services with latent feature and deep learning. *Transp. Res. Rec.* 2675 (11), 34–43.
- Hu, S., Xiong, C., 2023. High-dimensional population inflow time series forecasting via an interpretable hierarchical transformer. *Transp. Res. Part C: Emerg. Technol.* 146, 103962.
- Huang, Z., Wang, D., Yin, Y., Li, X., 2021. A spatiotemporal bidirectional attention-based ride-hailing demand prediction model: a case study in Beijing during COVID-19. *IEEE Trans. Intell. Transp. Syst.* 23 (12), 25115–25126.
- Jia, J., Zhang, H., Shi, B., 2021. Exploring bike-sharing behavior affected by public transportation disruption: case of Washington, DC, metro shutdown. *J. Transp. Eng. A: Syst.* 147 (3), 04020163.
- Ke, J., Qin, X., Yang, H., Zheng, Z., Zhu, Z., Ye, J., 2021. Predicting origin-destination ride-sourcing demand with a spatio-temporal encoder-decoder residual multi-graph convolutional network. *Transp. Res. C: Emerg. Technol.* 122, 102858.
- Khalesian, M., Furno, A., Leclercq, L., 2024. Improving deep-learning methods for area-based traffic demand prediction via hierarchical reconciliation. *Transp. Res. Part C: Emerg. Technol.* 159, 104410.
- Kourntzes, N., 2013. Intermittent demand forecasts with neural networks. *Int. J. Prod. Econ.* 143 (1), 198–206.
- Leem, S., Oh, J., Moon, J., Kim, M., Rho, S., 2024. Enhancing multistep-ahead bike-sharing demand prediction with a two-stage online learning-based time-series model: insight from Seoul. *J. Supercomput.* 80 (3), 4049–4082.
- Lei, J., Wasserman, L., 2014. Distribution-free prediction bands for non-parametric regression. *J. R. Stat. Soc. B: Stat. Methodol.* 76 (1), 71–96.
- Li, C., Bai, L., Liu, W., Yao, L., Waller, S.T., 2020. Graph neural network for robust public transit demand prediction. *IEEE Trans. Intell. Transp. Syst.* 23 (5), 4086–4098.
- Li, C., Geng, M., Chen, Y., Cai, Z., Zhu, Z., Chen, X.M., 2024a. Demand forecasting and predictability identification of ride-sourcing via bidirectional spatial-temporal transformer neural processes. *Transp. Res. C: Emerg. Technol.* 158, 104427.
- Li, C., Geng, M., Chen, Y., Cai, Z., Zhu, Z., Chen, X.M., 2024b. Demand forecasting and predictability identification of ride-sourcing via bidirectional spatial-temporal transformer neural processes. *Transp. Res. C: Emerg. Technol.* 158, 104427.
- Li, X., Xu, Y., Zhang, X., Shi, W., Yue, Y., Li, Q., 2023. Improving short-term bike sharing demand forecast through an irregular convolutional neural network. *Transp. Res. Part C: Emerg. Technol.* 147, 103984.
- Li, Y., Lu, X., Wang, Y., Dou, D., 2022. Generative time series forecasting with diffusion, denoise, and disentanglement. *Adv. Neural Inf. Process. Syst.* 35, 23009–23022.
- Li, Y., Zheng, Y., Zhang, H., Chen, L., 2015. Traffic prediction in a bike-sharing system. In: Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 1–10.
- Li, Y., Zhu, Z., Kong, D., Xu, M., Zhao, Y., 2019. Learning heterogeneous spatial-temporal representation for bike-sharing demand prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 33, pp. 1004–1011.
- Liang, Y., Huang, G., Zhao, Z., 2022. Bike sharing demand prediction based on knowledge sharing across modes: a graph-based deep learning approach. In: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC). IEEE, pp. 857–862.
- Liang, Y., Huang, G., Zhao, Z., 2023. Cross-mode knowledge adaptation for bike sharing demand prediction using domain-adversarial graph neural networks. *IEEE Trans. Intell. Transp. Syst.* 25 (5), 3642–3653.
- Lim, B., Arik, S.O., Loeff, N., Pfister, T., 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* 37 (4), 1748–1764.
- Lim, H., Chung, K., Lee, S., 2022. Probabilistic forecasting for demand of a bike-sharing service using a deep-learning approach. *Sustainability* 14 (23), 15889.
- Lin, L., He, Z., Peeta, S., 2018. Predicting station-level hourly demand in a large-scale bike-sharing network: a graph convolutional neural network approach. *Transp. Res. C: Emerg. Technol.* 97, 258–276.
- Liu, H., Dong, Z., Jiang, R., Deng, J., Chen, Q., Song, X., 2023. Staeformer: spatio-temporal adaptive embedding makes vanilla transformer SOTA for traffic forecasting. *arXiv preprint arXiv:2308.10425*.
- Lv, Y., Zhi, D., Sun, H., Qi, G., 2021. Mobility pattern recognition based prediction for the subway station related bike-sharing trips. *Transp. Res. C: Emerg. Technol.* 133, 103404.
- Ma, T., Knaap, G.-J., 2019. Estimating the impacts of capital bikeshare on metrorail ridership in the Washington Metropolitan area. *Transp. Res. Rec.* 2673 (7), 371–379.
- Ma, X., Ji, Y., Yang, M., Jin, Y., Tan, X., 2018. Understanding bikeshare mode as a feeder to metro by isolating metro-bikeshare transfers from smart card data. *Transp. Policy* 71, 57–69.
- Ma, X., Ji, Y., Yuan, Y., Van Oort, N., Jin, Y., Hoogendoorn, S., 2020. A comparison in travel patterns and determinants of user demand between docked and dockless bike-sharing systems using multi-sourced data. *Transp. Res. A: Policy Pract.* 139, 148–173.
- Ma, X., Yin, Y., Jin, Y., He, M., Zhu, M., 2022. Short-term prediction of bike-sharing demand using multi-source data: a spatial-temporal graph attentional LSTM approach. *Appl. Sci.* 12 (3), 1161.
- Open-Meteo, 2026. Open-meteo: free weather APIs for non-commercial use. <https://open-meteo.com/>. Accessed: 2025-03-24.
- Palacio, L., Vo, T., Maness, M., Bertini, R.L., Menon, N., 2021. Multicity investigation of the effect of holidays on bikeshare system ridership. *Transp. Res. Rec.* 2675 (7), 404–423.
- Peng, T., Gao, J., Cats, O., 2025. Uncertainty-aware probabilistic travel demand prediction for mobility-on-demand services. *Transp. Res. C: Emerg. Technol.* 181, 105383.
- Qian, W., Zhang, D., Zhao, Y., Zheng, K., Yu, J. Q., 2023a. Uncertainty quantification for traffic forecasting: a unified approach. In: 2023 IEEE 39th International Conference on Data Engineering (ICDE). IEEE, pp. 992–1004.
- Qian, W., Zhao, Y., Zhang, D., Chen, B., Zheng, K., Zhou, X., 2023b. Towards a unified understanding of uncertainty quantification in traffic flow forecasting. *IEEE Trans. Knowl. Data Eng.* 36 (5), 2239–2256.
- Rasul, K., Ashok, A., Williams, A.R., Ghonia, H., Bhagwatkar, R., Khorasani, A., Bayazi, M.J.D., Adamopoulos, G., Riachi, R., Hassen, N., et al., 2023. Lag-llama: towards foundation models for probabilistic time series forecasting. *arXiv preprint arXiv:2310.08278*.
- Rixey, R.A., 2013. Station-level forecasting of bikesharing ridership: station network effects in three US systems. *Transp. Res. Rec.* 2387 (1), 46–55.
- Salinas, D., Flunkert, V., Gasthaus, J., Januschowski, T., 2020. DeepAR: probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* 36 (3), 1181–1191.
- Sathishkumar, V.E., Park, J., Cho, Y., 2020. Using data mining techniques for bike sharing demand prediction in metropolitan city. *Comput. Commun.* 153, 353–366.
- Shaheen, S., Martin, E., 2015. Unraveling the modal impacts of bikesharing. *Access Mag.* 1 (47), 8–15.
- Shen, Z., Zhang, Y., Lu, J., Xu, J., Xiao, G., 2020. A novel time series forecasting model with deep learning. *Neurocomputing* 396, 302–313.
- Solatorio, A.V., 2023. Geoformer: predicting human mobility using generative pre-trained transformer (GPT). In: Proceedings of the 1st International Workshop on the Human Mobility Prediction Challenge, pp. 11–15.
- Tarpin-Pitre, L., Morency, C., 2020. Typology of bikeshare users combining bikeshare and transit. *Transp. Res. Rec.* 2674 (10), 475–483.
- Toubeau, J.-F., Bottieau, J., Vallée, F., De Grève, Z., 2018. Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets. *IEEE Trans. Power Syst.* 34 (2), 1203–1215.
- Tushar, M.S., Buehler, R., 2024. Bikeshare-metrorail integration in Washington, DC: what are the characteristics of neighborhoods that encourage capital bikeshare trips to and from the metrorail? *Transp. Res. Rec.* 2678 (10), 282–291.
- Vaswani, A., 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30, 5998–6008.
- Wang, H., Noland, R.B., 2021. Bikeshare and subway ridership changes during the COVID-19 pandemic in New York City. *Transp. Policy* 106, 262–270.
- Wang, Q., Wang, S., Zhuang, D., Koutsopoulos, H., Zhao, J., 2024. Uncertainty quantification of spatiotemporal travel demand with probabilistic graph neural networks. *IEEE Trans. Intell. Transp. Syst.* 25 (8), 8770–8781.

- Wang, Y., Liu, R.W., Liu, J., Yang, L., Liu, Y., He, P., 2025. Co-prediction of multiple transportation demands in ro-ro terminal based on dates fine-grained temporal fusion transformer. *Comput. Ind. Eng.*, 208, 111373.
- Wang, Y., Smola, A., Maddix, D., Gasthaus, J., Foster, D., Januschowski, T., 2019. Deep factors for forecasting. In: *International Conference on Machine Learning*. PMLR, pp. 6607–6617.
- Wilkesmann, F., Ton, D., Schakenbos, R., Cats, O., 2023. Determinants of station-based round-trip bikesharing demand. *J. Public Transp.* 25, 100048.
- Xu, C., Ji, J., Liu, P., 2018. The station-free sharing bike demand forecasting with a deep learning approach and large-scale datasets. *Transp. Res. Part C: Emerg. Technol.* 95, 47–60.
- Xu, M., Di, Y., Yang, H., Chen, X., Zhu, Z., 2023. Multi-task supply-demand prediction and reliability analysis for docked bike-sharing systems via transformer-encoder-based neural processes. *Transp. Res. C: Emerg. Technol.* 147, 104015.
- Yang, Y., Heppenstall, A., Turner, A., Comber, A., 2019. A spatiotemporal and graph-based analysis of dockless bike sharing patterns to understand urban flows over the last mile. *Comput. Environ. Urban Syst.* 77, 101361.
- Yoon, J.W., Pinelli, F., Calabrese, F., 2012. Cityride: a predictive bike sharing journey advisor. In: *2012 IEEE 13th International Conference on Mobile Data Management*. IEEE, pp. 306–311.
- Yu, B., Yin, H., Zhu, Z., 2017. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*.
- Zerveas, G., Jayaraman, S., Patel, D., Bhamidipaty, A., Eickhoff, C., 2021. A transformer-based framework for multivariate time series representation learning. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 2114–2124.
- Zhang, C., Zhang, L., Liu, Y., Yang, X., 2018a. Short-term prediction of bike-sharing usage considering public transport: a LSTM approach. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 1564–1571.
- Zhang, C., Zhang, L., Liu, Y., Yang, X., 2018b. Short-term prediction of bike-sharing usage considering public transport: a LSTM approach. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 1564–1571.
- Zhang, Y., Cheng, J., Cats, O., Azadeh, S.S., 2025. Stochastic optimization under supply uncertainty for multimodal trip planning based on demand prediction. In: *Proceedings of the 12th Triennial Symposium on Transportation Analysis (TRISTAN XII)*.
- Zhao, D., Ong, G.P., Wang, W., Hu, X.J., 2019. Effect of built environment on shared bicycle reallocation: a case study on nanjing, china. *Transp. Res. Part a: Policy Pract.* 128, 73–88.
- Zhao, R., Tian, Z., Tian, L., Liu, W., Wang, D. Z.W., 2025. Research on rebalancing of large-scale bike-sharing system driven by zonal heterogeneity and demand uncertainty. *Transp. Res. C: Emerg. Technol.* 170, 104933.