

Towards Generative Modeling of Urban Flow through Knowledge-enhanced Denoising Diffusion

Zhilun Zhou
Tsinghua University
Beijing, China
zzl22@mails.tsinghua.edu.cn

Jingtao Ding*
Tsinghua University
Beijing, China
dingjt15@tsinghua.org.cn

Yu Liu
Tsinghua University
Beijing, China
liuyu2419@126.com

Depeng Jin
Tsinghua University
Beijing, China
jindp@tsinghua.edu.cn

Yong Li
Tsinghua University
Beijing, China
liyong07@tsinghua.edu.cn

ABSTRACT

Although generative AI has been successful in many areas, its ability to model geospatial data is still underexplored. Urban flow, a typical kind of geospatial data, is critical for a wide range of applications from public safety and traffic management to urban planning. Existing studies mostly focus on predictive modeling of urban flow that predicts the future flow based on historical flow data, which may be unavailable in data-sparse areas or newly planned regions. Some other studies aim to predict OD flow among regions but they fail to model dynamic changes of urban flow over time. In this work, we study a new problem of urban flow generation that generates dynamic urban flow for regions without historical flow data. To capture the effect of multiple factors on urban flow, such as region features and urban environment, we employ diffusion model to generate urban flow for regions under different conditions. We first construct an urban knowledge graph (UKG) to model the urban environment and relationships between regions, based on which we design a knowledge-enhanced spatio-temporal diffusion model (KSTDiff) to generate urban flow for each region. Specifically, to accurately generate urban flow for regions with different flow volumes, we design a novel diffusion process guided by a volume estimator, which is learnable and customized for each region. Moreover, we propose a knowledge-enhanced denoising network to capture the spatio-temporal dependencies of urban flow as well as the impact of urban environment in the denoising process. Extensive experiments on four real-world datasets validate the superiority of our model over state-of-the-art baselines in urban flow generation. Further in-depth studies demonstrate the utility of generated urban flow data and the ability of our model for long-term flow generation and urban flow prediction. Our code

is released at: <https://github.com/tsinghua-fib-lab/KSTDiff-Urban-flow-generation>.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; • **Information systems** → **Spatial-temporal systems**.

KEYWORDS

Generative model, urban flow, knowledge graph, diffusion model

ACM Reference Format:

Zhilun Zhou, Jingtao Ding, Yu Liu, Depeng Jin, and Yong Li. 2023. Towards Generative Modeling of Urban Flow through Knowledge-enhanced Denoising Diffusion. In *The 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23)*, November 13–16, 2023, Hamburg, Germany. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3589132.3625641>

1 INTRODUCTION

Generative AI, especially large pre-trained models, have shown great success in many areas such as natural language processing and computer vision, with recently popular GPT-4 [26] and Stable Diffusion [29] as examples. However, their ability to model geospatial data is still underexplored [23, 30].

Urban flow is a typical kind of geospatial data that depicts the dynamic human mobility patterns in urban regions [38], playing an important role in public safety, traffic management, and urban planning. For example, the inflow and outflow refer to the number of people entering or leaving an urban region during a given time interval. With such data, the government can implement preventive measures in advance to avoid stampede or traffic congestion in regions experiencing a large inflow. Moreover, ride-sharing platforms can dispatch more taxis to regions with large outflows to meet the demand. In recent years, many studies have been conducted on *predictive modeling* of urban flow. As shown in Figure 1(a), they mostly focus on urban flow prediction, which trains a prediction model based on historical flow data to predict the future urban flow of regions [14, 17, 48]. However, these studies rely heavily on historical flow, and thus cannot be adapted to data-sparse areas such as newly planned regions or suburban regions. In addition, some studies have made efforts to predict static origin-destination flows from one region to another given their characteristics [21, 33, 51].

*Corresponding Author.

Nevertheless, they fail to capture temporal patterns since they cannot generate dynamic urban flow data, where the flow may vary significantly at different times of the day.

In this work, we take a step towards *generative modeling* of urban flow by studying a new problem of urban flow generation, which aims to train a generation model based on existing regions in order to estimate the urban flow of new regions without historical flow data, as shown in Figure 1(b). Urban flow generation can help estimate the urban flow for data-sparse regions or evaluate the flow patterns of newly planned regions in advance of region construction, which is essential for urban planning [49].

The urban flow of a region is affected by multiple factors including regional characteristics, urban environment, and urban flow in nearby regions. Therefore, urban flow generation can be formulated as a conditional generation problem, which aims to generate flow for different regions under different conditions. Recently, diffusion models have shown outstanding performance in a wide range of tasks like image synthesis [8], audio synthesis [12], time series modeling [27, 34] and spatio-temporal events prediction [47]. The key idea of diffusion models is to first gradually add noise to data until it becomes random noise, and then train a model to learn the reverse process, i.e., convert random noise into expected data distribution step by step. Moreover, additional information can be injected into each step of the reverse process, making it easy to generate data under given conditions using diffusion models. As a result, we propose to leverage conditional diffusion models to solve the urban flow generation problem so as to better control the generation process by multiple factors.

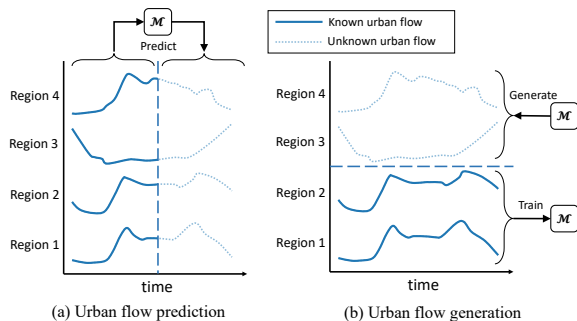


Figure 1: Comparison of urban flow prediction and urban flow generation.

However, accurately generating urban flow under diverse conditions still faces the following challenges: (1) **Significant variance in the spatial distribution of urban flow volume.** The volume of urban flow varies significantly across regions [10]. For example, newly planned regions located in suburban areas typically exhibit lower volumes than established regions in downtown areas due to underdeveloped infrastructure and transportation networks. We visualize the flow volume in different regions of Washington, D.C., and Baltimore, which is defined as the average flow in each hour of a region. As shown in Figure 2, the flow volumes of regions in downtown areas are generally much larger than those in suburban areas. Such spatial heterogeneity in flow volume distribution makes it a challenging task to generate urban flow accurately for different regions with different volumes without historical flow data. (2)

Highly complex and varying spatio-temporal dependencies of urban flow across different regions. There exist complex spatio-temporal dependencies of urban flow among regions, and such dependencies are affected by urban environment and vary across different regions. For example, in downtown business areas with dense road networks, the urban flow of a region may be largely influenced by nearby regions, resulting in strong spatial correlations. On the contrary, for residential regions in suburban areas, the flow transition among regions may be rather small, and urban flow there exhibits stronger temporal patterns, e.g., large outflow during weekday morning and large inflow at night. Most of the existing diffusion models can only model temporal correlation in time series [27, 34], while they fail to consider the spatio-temporal dependencies in urban flow and the effect of urban environment.

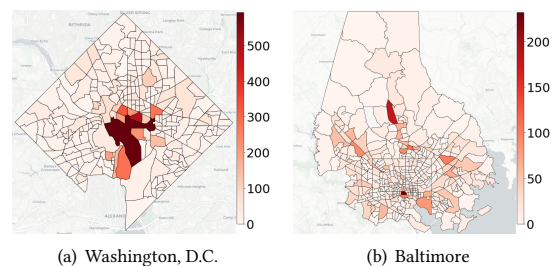


Figure 2: The average flow volume varies significantly in different regions.

To tackle the complexity of modeling urban flow and regions, we construct an urban knowledge graph (UKG) to model the urban environment as well as various relationships among regions. Based on UKG, we propose a **Knowledge-enhanced Spatio-Temporal Diffusion model (KSTDiff)** to generate dynamic urban flow for regions. Specifically, in the forward diffusion process, we gradually add Gaussian noise to the urban flow data of each region. To handle the variance of flow volume among regions, we design a volume estimator to directly estimate the flow volume for each region, which is then used to guide the diffusion process by adjusting the mean of added Gaussian noise accordingly. Moreover, the volume estimator is jointly trained with the diffusion model to make the diffusion process of each region learnable and customized, and thus overcome the first challenge. In the denoising process, we aim to train a denoising network to convert random noise to urban flow data step by step. To solve the second challenge, we design a knowledge graph enhanced spatio-temporal block (KGST block) to capture spatio-temporal dependencies of urban flow based on environmental information, where we utilize relation-aware graph convolution layers to capture multiple types of spatial dependencies based on UKG and use transformer layers to model temporal dependencies. Moreover, we leverage KG embeddings of regions with environmental information preserved to guide the spatio-temporal fusion. In this way, we manage to capture the spatio-temporal dependencies as well as the environmental factors that affect such dependencies.

Overall, our contributions can be summarized as follows:

- We study the problem of urban flow generation, which aims to generate urban flow for regions without historical flow data,

by proposing a knowledge-enhanced spatio-temporal diffusion model. To our best knowledge, we are the first to introduce diffusion model for generative modeling of geospatial data.

- We employ diffusion model to control the urban flow generation process based on multiple factors. Specifically, we design a volume estimator to guide the diffusion process with estimated flow volume for each region separately. Moreover, we construct an urban KG to model urban environment as well as complex relationships among regions, and devise an STKG block to capture spatio-temporal dependencies of urban flow.
- We conduct extensive experiments on four real-world urban flow datasets, and our model outperforms state-of-the-art models by over 13.3% in terms of MAE on all datasets, highlighting the superiority and robustness of our model. Several in-depth studies further demonstrate the effectiveness of our model design and its ability to generate long-term flow data. Moreover, experiments also demonstrate the utility of generated data for downstream applications and our model’s capability for predictive modeling of urban flow.

2 RELATED WORK

2.1 Urban Flow Modeling

Understanding urban flow is crucial for many areas like urban planning and transportation system. In recent years, plenty of works have studied the problem of urban flow prediction through deep learning methods [11, 13, 48], which leverage neural networks to predict the future inflow and outflow for citywide regions based on historical flow. The core of urban flow prediction is to model the spatio-temporal dependencies of urban flow. Existing works mostly use convolutional neural networks (CNN) or graph neural networks (GNN) to capture spatial dependencies, and use recurrent neural networks (RNN) or transformers to capture temporal dependencies. However, these works depend on historical flow and thus cannot be used for new regions without such data.

Another series of works aims to predict OD flow among regions. Deep Gravity [33] leverages neural networks to predict OD flow among each pair of regions based on their features, distance as well as outflows. GMEL [21] first learns an embedding for each region through graph attention network (GAT), and then use regression model to estimate the OD flow based on region embeddings. However, these works can only generate static urban flow, which is not enough for modeling urban flow with dynamic changes over time.

2.2 Diffusion Model for Time Series Modeling

Diffusion models have been widely used in time series forecasting, imputation, and generation owing to their ability to model high dimensional data distributions. For example, TimeGrad [27] is used for time series forecasting in an autoregressive way. It combines diffusion model with RNN and generates the value of the next time step conditioned on the hidden state of RNN. CSDI [34] uses diffusion models for probabilistic time series imputation by generating missing values conditioned on observed values. Diffusion models have also been used in time series generation applications like electronic health records (EHR) synthesis [1, 7, 46].

Most of these works use the denoising network architecture proposed in Diffwave [12], which leverages bidirectional dilated

convolution to capture the correlation between different time steps. However, these works fall short of modeling spatial dependencies among different regions in our urban flow generation task. Moreover, a recent work named DiffSTG [40] designs a novel UGnet for spatio-temporal graph forecasting, which can capture spatio-temporal dependencies of different nodes, while it depends on historical data to learn such dependencies and cannot manage to generate urban flow without historical flow data.

3 PRELIMINARIES

3.1 Problem Statement

In this section, we formally define our problem based on some concepts, and give a brief introduction of diffusion model. The notations used in the following sections are summarized in Table 1.

Definition 3.1 (Urban Region). Urban regions are defined as non-overlapping areas in a city partitioned by main road networks, or defined as administrative areas such as census tracts.

Definition 3.2 (Urban Flow). In this study, we focus on two kinds of urban flow, namely the inflow and outflow, which are defined as the number of people entering or leaving the region in a given time interval. They are usually calculated using user trajectory data or taxi trip data. In addition, the inflow of a region can also be calculated as the total number of visits to POIs in the region based on POI check-in data. The urban flow at time t can be represented as $F_t \in \mathbb{R}^{N_l \times d_f}$, where N_l is the number of target regions and d_f is the dimension of urban flow, e.g., $d_f = 2$ for inflow and outflow.

Based on these concepts, we define the research problem as follows.

Definition 3.3 (Urban Flow Generation). Given a set of regions $S_l = \{l_1, l_2, \dots, l_{N_l}\}$ with attributes like population and POI information, generate urban flow for these regions in different time intervals F_1, \dots, F_T without historical flow data. The distribution of generated urban flows should be similar to real flows.

The urban flow of each region is affected by many factors like region features, urban environment, and interactions among regions. To better utilize such factors for urban flow generation, we employ denoising diffusion probabilistic model (DDPM) as the backbone of our model. Moreover, we construct an urban knowledge graph (UKG) to provide a comprehensive description of urban environment as well as complex relationships among regions, and generate urban flow based on the UKG. We first give a brief introduction to DDPM and the construction of UKG.

3.2 Denoising Diffusion Probabilistic Model

Denoising diffusion probabilistic models (DDPM) [8] are deep generative models that learn a data distribution of variable x by gradually denoising a Gaussian noise. DDPM consists of a forward noising process and a reverse denoising process.

In the forward process, Gaussian noise is gradually added to the data $x_0 \sim q(x_0)$ to produce a Markov chain:

$$q(x_{1:N}|x_0) = \prod_{n=1}^N q(x_n|x_{n-1}), \quad (1)$$

Table 1: Notations.

Notations	Description
N_l	The number of regions
T	The length of generated flow (number of time intervals)
d_f, d_{fea}	Dimension of urban flow and region features
d_{KG}, d_h	Dimension of KG embedding and hidden size
ϵ	Gaussian noise
ϵ_θ	Denosing module with parameter θ
$\hat{\epsilon}$	Noise predicted by denoising module
N	The number of diffusion steps
$\mathbf{x}_0, \mathbf{x}_n$	Sample from real data, and the noised data at the n -th diffusion step
β_n	Noise level at the diffusion step n
$\alpha_n, \bar{\alpha}_n$	$\alpha_n = 1 - \beta_n, \bar{\alpha}_n = \alpha_1 \alpha_2 \dots \alpha_n$
f_ϕ	Volume estimator with parameter ϕ
c	Region features
\hat{s}, s	Predicted and real flow volume
E_{KG}	KG embedding of regions
\mathbf{h}_l	Input of region l to the spatial block
H_s, H_t	Output of the spatial temporal block and temporal block
M_1	Number of epochs for pretraining volume estimator
M_2	We update the volume estimator once every M_2 epochs

where $q(\mathbf{x}_n|\mathbf{x}_{n-1}) = \mathcal{N}(\sqrt{1-\beta_n}\mathbf{x}_{n-1}, \beta_n\mathbf{I})$ and $\beta_n \in (0, 1)$ represents the noise level. With the notation $\alpha_n = 1 - \beta_n$ and $\bar{\alpha}_n = \alpha_1 \alpha_2 \dots \alpha_n$, sampling \mathbf{x}_n at an arbitrary step n can be written in a close form:

$$q(\mathbf{x}_n|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_n; \sqrt{\bar{\alpha}_n}\mathbf{x}_0, (1 - \bar{\alpha}_n)\mathbf{I}), \quad (2)$$

which can also be reparameterized as

$$\mathbf{x}_n = \sqrt{\bar{\alpha}_n}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_n}\epsilon \quad (3)$$

with ϵ sampled from a standard gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

The reverse process recurrently denoises pure Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to recover the original data \mathbf{x}_0 :

$$p_\theta(\mathbf{x}_{0:N}) = p(\mathbf{x}_N) \prod_{n=1}^N p_\theta(\mathbf{x}_{n-1}|\mathbf{x}_n), \quad (4)$$

$$p_\theta(\mathbf{x}_{n-1}|\mathbf{x}_n) = \mathcal{N}(\mathbf{x}_{n-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_n, n), \sigma_\theta(\mathbf{x}_n, n)\mathbf{I}).$$

To obtain \mathbf{x}_{n-1} from \mathbf{x}_n , we train a denoising network ϵ_θ to predict the noise ϵ added to \mathbf{x}_0 using \mathbf{x}_n . Then, \mathbf{x}_0 can be calculated by Equation 3 as $\mathbf{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{1 - \bar{\alpha}_n}\epsilon_\theta(\mathbf{x}_n, n))$. On the other hand, we have the forward process posteriors:

$$q(\mathbf{x}_{n-1}|\mathbf{x}_n, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{n-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_n, \mathbf{x}_0), \tilde{\beta}_n\mathbf{I}), \quad (5)$$

where

$$\tilde{\boldsymbol{\mu}}(\mathbf{x}_n, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{n-1}}\beta_n}{1 - \bar{\alpha}_n}\mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_n}(1 - \bar{\alpha}_{n-1})}{1 - \bar{\alpha}_n}\mathbf{x}_n, \quad (6)$$

$$\tilde{\beta}_n = \frac{1 - \bar{\alpha}_{n-1}}{1 - \bar{\alpha}_n}\beta_n.$$

Consequently, we can obtain the parameterization of $p_\theta(\mathbf{x}_{n-1}|\mathbf{x}_n)$ in Equation 4:

$$\boldsymbol{\mu}_\theta(\mathbf{x}_n, n) = \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \frac{\beta_n}{\sqrt{1 - \bar{\alpha}_n}}\epsilon_\theta(\mathbf{x}_n, n)), \quad (7)$$

$$\sigma_\theta(\mathbf{x}_n, n) = \frac{1 - \bar{\alpha}_{n-1}}{1 - \bar{\alpha}_n}\beta_n.$$

The denoising network ϵ_θ can be trained by optimizing the following L1 loss function:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), n} \|\epsilon - \epsilon_\theta(\mathbf{x}_n, n)\|. \quad (8)$$

3.3 Urban Knowledge Graph Construction

A KG is a graph that consists of an entity set \mathcal{E} , a relation set \mathcal{R} and a fact set \mathcal{F} , where each fact is represented as a triplet $(h, r, t) \in \mathcal{F}$, denoting a directional edge from head entity h to tail entity t with relation type r . KG has been widely used in urban computing because of its ability to model urban environment as well as capture complex relationships between entities [16, 19, 37].

Inspired by previous works [18, 20, 50], we construct an urban KG (UKG) to model the environmental information in the city and interactions among regions. Specifically, we model the urban regions as entities in UKG, and use relations *BorderBy* and *NearBy* to describe the spatial adjacency relationship between regions because urban flows in nearby regions are easily affected by each other. Furthermore, the flow pattern of a region is also correlated with its function [14, 43], which is mainly reflected by the point of interests (POIs) in the region. As a result, we add POIs and POI categories as entities into UKG, and use relations *LocateAt* and *CateOf* to describe the location and category of POIs. In addition, we consider the geographical influence and competitive relationships between POIs by relation *CoCheckin* and *Competitive*. We also calculate the functional similarity for each pair of regions, i.e., the cosine similarity of POI category distribution, and use relation *SimilarFunc* to link regions with similar function. To further enhance the semantics of urban environment, we model business areas (BA) as entities and their relationships with regions and POIs by *ProvideService* and *BelongTo*. The details of relations in UKG are shown in Table 2.

Table 2: The details of relations in UKG. BA represents business area.

Relation	Head & Tail Entity Types	Semantic Information
<i>BorderBy</i>	(Region, Region)	Regions share part of the boundary
<i>NearBy</i>	(Region, Region)	Regions lie within a certain distance
<i>CoCheckin</i>	(POI, POI)	POIs visited by a user consecutively
<i>Competitive</i>	(POI, POI)	Nearby POIs with the same brand
<i>SimilarFunc</i>	(Region, Region)	Regions with similar POI distribution
<i>LocateAt</i>	(POI, Region)	POI locates at the region
<i>CateOf</i>	(POI, Category)	Category of POI
<i>ProvideService</i>	(BA, Region)	BA covers the region
<i>BelongTo</i>	(POI, BA)	POI locates at the BA

To better make use of the environmental knowledge in UKG, we leverage KG embedding model to learn an embedding vector for each region from UKG. Specifically, we choose a state-of-the-art model TuckER [3], which uses Tucker decomposition [35] as the scoring function to measure the plausibility of triplets in UKG:

$$\phi(h, r, t) = \mathcal{W} \times_1 \mathbf{e}_h \times_2 \mathbf{e}_r \times_3 \mathbf{e}_t, \quad (9)$$

where $\mathcal{W} \in \mathbb{R}^{d_{KG} \times d_{KG} \times d_{KG}}$ is a learnable core tensor, \times_n is the tensor product along the n -th dimension, and $\mathbf{e}_h, \mathbf{e}_r, \mathbf{e}_t \in \mathbb{R}^{d_{KG}}$ are the embeddings of head entity h , tail entity t and relation r , respectively. The goal of KG embedding model is to make the scoring function high for triplets that exist in the UKG, and thus the knowledge in UKG can be preserved in KG embeddings.

The benefits of UKG are two fold. First, it integrates various types of entities and relationships in the city comprehensively, and thus the learned embedding of a region can provide a description of the urban environment where it locates. Second, various types

of relationships between regions in UKG can help capture the spatial dependencies of urban flow between regions, which will be elaborated in the following section.

Definition 3.4 (Spatio-temporal Flow Generation on KG). The UKG organizes regions in a graph structure, and the urban flow F_t of different regions is not independent but connected by various types of relations in the UKG. Consequently, the *urban flow generation* problem can be reformulated as generating a dynamic graph signal for a subset of nodes corresponding to the target regions in the UKG, denoted as $\mathbf{x} = [F_1, \dots, F_T] \in \mathbb{R}^{N_I \times T \times d_f}$, given the node features and graph structure $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{F})$.

4 METHODS

Based on the aforementioned DDPM and UKG, we propose a knowledge-enhanced spatio-temporal diffusion model for urban flow generation. The framework of our model is shown in Figure 3(a). Specifically, we design an auxiliary module to estimate the flow volume of each region, and use the predicted volume to guide the diffusion process for each region separately, which enables learnable and more accurate flow generation process for regions with different volumes. Moreover, we design a novel knowledge-enhanced denoising network that leverages UKG to capture spatio-temporal dependencies of urban flow and fuse them adaptively based on urban environment.

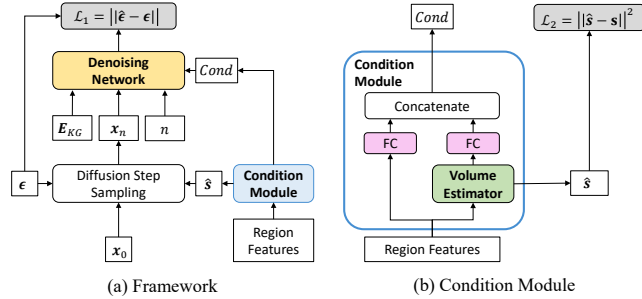


Figure 3: The (a) framework of our model and (b) design of condition module.

4.1 Region Customized Diffusion Process

The urban flow of a region is closely correlated with its features, such as socioeconomic indicators and demographics [14, 43]. Therefore, we leverage region features to guide the flow generation process through a condition module, as shown in Figure 3(b). Specifically, we design a volume estimator to estimate the flow volume for each region, which is trained together with the diffusion model. On one hand, the estimated flow volume is combined with region features as the condition of the diffusion model. On the other hand, we design a novel learnable and customized diffusion process for each region based on the estimated flow volume.

The volume of urban flow, which we define as the average flow per hour, may vary significantly across different regions. For example, regions in downtown areas usually have a much larger flow volume than suburban regions because of more developed transportation system. Such spatial variance may lead to the bias of

model towards popular regions with larger volume, and thus should be considered directly [10]. However, the vanilla diffusion models assume the same endpoint of diffusion process. In other words, the generation process for all regions starts from the same Gaussian noise $\mathcal{N}(0, I)$, which makes it difficult to distinguish regions with different flow volumes. Therefore, we employ the technique proposed in [6] and make the diffusion process learnable and customized for each region.

Specifically, We design a volume estimator f_ϕ with parameters ϕ to predict the flow volume based on region features (denoted as c), and use the predicted volume $f_\phi(c)$ to guide the diffusion process for each region. Here $c \in \mathbb{R}^{N_I \times d_{fea}}$ represents the features of all target regions, and the output of volume estimator is expanded to the same dimension as \mathbf{x} , i.e., $f_\phi(c) \in \mathbb{R}^{N_I \times T \times d_f}$. Then we change the endpoint of diffusion process to Gaussian noise with different mean values:

$$p(\mathbf{x}_N | c) = \mathcal{N}(f_\phi(c), I). \quad (10)$$

Accordingly, with the same notations as in Section 3.2, the forward diffusion process is modified as:

$$q(\mathbf{x}_n | \mathbf{x}_0, f_\phi(c)) = \mathcal{N}(\mathbf{x}_n; \sqrt{\alpha_n} \mathbf{x}_0 + (1 - \sqrt{\alpha_n}) f_\phi(c), (1 - \alpha_n) I). \quad (11)$$

Besides, in the backward denoising process, the posterior in Equation 5 should be changed to:

$$q(\mathbf{x}_{n-1} | \mathbf{x}_n, \mathbf{x}_0, c) = \mathcal{N}(\mathbf{x}_{n-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_n, \mathbf{x}_0, f_\phi(c)), \tilde{\boldsymbol{\beta}}_n I), \quad (12)$$

where

$$\begin{aligned} \tilde{\boldsymbol{\mu}}(\mathbf{x}_n, \mathbf{x}_0, f_\phi(c)) = & \frac{\sqrt{\alpha_{n-1}} \beta_n}{1 - \alpha_n} \mathbf{x}_0 + \frac{\sqrt{\alpha_n} (1 - \alpha_{n-1})}{1 - \alpha_n} \mathbf{x}_n \\ & + \left(1 + \frac{(\sqrt{\alpha_n} - 1)(\sqrt{\alpha_n} + \sqrt{\alpha_{n-1}})}{1 - \alpha_n}\right) f_\phi(c). \end{aligned} \quad (13)$$

For the architecture of f_ϕ , we simply adopt a two-layer feed-forward neural network with Leaky ReLU as the activation function [6]. In this way, the start point of flow generation process is approximately the mean value of the flow for each region, which enables more accurate flow generation for regions with different flow volumes.

We consider volume estimation as an auxiliary task and use MSE loss to optimize the volume estimator:

$$\mathcal{L}_2 = \|\hat{s} - s\|^2, \quad (14)$$

where $\hat{s} = f_\phi(c)$ is the predicted flow volume and $s \in \mathbb{R}^{N_I \times T \times d_f}$ is the real flow volume expanded to the same shape. In the training process, the volume estimator f_ϕ is trained with the denoising network ϵ_θ in an alternative manner. As shown in Algorithm 1, we first pretrain f_ϕ for M_1 epochs, and then alternatively update the whole diffusion model and f_ϕ . Specifically, we update the diffusion model for M_2 epochs by taking gradient decent step on \mathcal{L}_1 , and then update f_ϕ once with loss function \mathcal{L}_2 .

4.2 Knowledge-enhanced Denoising Network

The core of DDPM lies in the denoising network ϵ_θ , which aims to predict the noise ϵ from noised data \mathbf{x}_n . The architecture of denoising network is usually designed according to the specific task. Previous works about sequence generation typically use WaveNet-based structure [12, 25], which shows great performance in modeling time series while it cannot capture spatial dependencies in

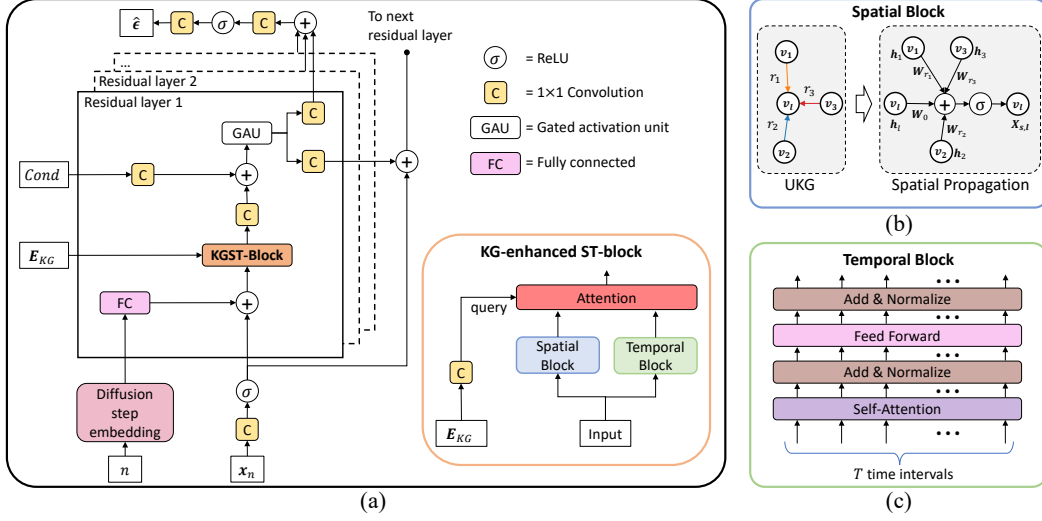


Figure 4: (a) The architecture of denoising network, (b) the detail of spatial block and (c) the detail of temporal block.

Algorithm 1 Training of our model

```

1: Step 1: PRETRAINING  $f_\phi$ 
2:   for  $epoch \in \{1, 2, \dots, M_1\}$  do
3:     Calculate predicted flow volume  $\hat{s} = f_\phi(c)$ 
4:     Update  $f_\phi$  using  $\mathcal{L}_2 = \|\hat{s} - s\|^2$ 
5: Step 2: JOINT TRAINING  $\epsilon_\theta$  AND  $f_\phi$ 
6:   repeat
7:     > Train diffusion model
8:     for  $epoch \in \{1, 2, \dots, M_2\}$  do
9:       Sample flow  $x_0 \sim q(x_0)$ , noise  $\epsilon \sim \mathcal{N}(0, I)$ 
10:      Sample diffusion step  $n \sim Uniform(\{1, 2, \dots, N\})$ 
11:      Calculate  $x_n$  with Equation 11
12:      Calculate predicted noise  $\hat{\epsilon}$  with denoising network
13:      Update  $\epsilon_\theta$  and  $f_\phi$  using  $\mathcal{L}_1 = \|\hat{\epsilon} - \epsilon\|$ 
14:     > Train volume estimator
15:     Calculate predicted flow volume  $\hat{s} = f_\phi(c)$ 
16:     Update  $f_\phi$  using  $\mathcal{L}_2 = \|\hat{s} - s\|^2$ 
17:   until Convergence

```

urban flow data. To bridge the gap, we propose a novel knowledge-enhanced network empowered by UKG to capture spatio-temporal dependencies of urban flow for better denoising ability.

4.2.1 Denoising Network Architecture. The architecture of our denoising network is shown in Figure 4. Specifically, we adopt the DiffWave architecture as the backbone, which consists of multiple residual layers. It takes noised data $x_n \in \mathbb{R}^{N_t \times T \times d_f}$, diffusion step n , KG embeddings of regions E_{KG} and conditions as input, and output the predicted noise: $\hat{\epsilon} = \epsilon_\theta(x_n, n, E_{KG}, Cond)$. The output $\hat{\epsilon} \in \mathbb{R}^{N_t \times T \times d_f}$ is of the same shape as x_0 and ϵ .

Specifically, we first use 128-dimensional diffusion step embedding for diffusion step n in the diffusion step embedding block:

$$n_{embedding} = [\sin(10^{0 \times 4/63} n), \dots, \sin(10^{63 \times 4/63} n), \cos(10^{0 \times 4/63} n), \dots, \cos(10^{63 \times 4/63} n)], \quad (15)$$

followed by two fully connected layers before being fed into residual layers. In the meanwhile, the input noised data x_n is mapped by a $\text{Conv}1 \times 1$ and sent to the first residual layer. In each residual layer, we first combine the input with projected diffusion step embedding to obtain flow representation $\mathbf{h} \in \mathbb{R}^{N_t \times T \times d_h}$, and then design a KG-enhanced ST-block (KGST-Block) to capture spatio-temporal dependencies of urban flow with the guidance of KG embedding, i.e., $H_{st} = KGST(\mathbf{h}, E_{KG})$, which will be elaborated in Section 4.2.2. The conditioner of each region $Cond$ is added to the output of KGST-Block H_{st} . After a gated activation unit, part of the output is connected to the next residual layer as input, while the rest is added to the final output through skip connection. The $\text{Conv}1 \times 1$ blocks in the network are used to map data to proper dimensions. Finally, the output $\hat{\epsilon}$ is the sum of data through skip-connections from each residual layer after projection by two $\text{Conv}1 \times 1$ blocks. We then introduce the details of KGST-Block.

4.2.2 KG-enhanced ST-Block Design. The KGST block consists of a spatial block and a temporal block for modeling spatial and temporal dependencies, after which an attention block is used for spatio-temporal fusion with the guidance of KG embeddings.

For the spatial block, we adopt 1-layer R-GCN [32] to capture the spatial dependency of urban flow among regions. R-GCN is a kind of graph convolutional network that aggregates information from the neighborhood of entities through different relations separately. As a result, the flow information of regions that are connected in UKG, i.e., nearby regions and functionally similar regions, can be propagated through the graph. Note that here we only use the subgraph of UKG that consists of training or testing regions and relations among them. Specifically, the output of R-GCN can be obtained as:

$$H_{s,l} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{N}_l^r} W_r h_j + W_0 h_l \right), \quad (16)$$

where $\mathbf{h}_l \in \mathbb{R}^{T \times d_h}$ is the input of region l , \mathcal{N}_l^r is the set of regions connected to region l via relation r and $W_r, W_0 \in \mathbb{R}^{d_h \times d_h}$ are learnable weight matrices. σ is a nonlinear activation function.

Table 3: The basic information of four real-world datasets and basic statistics of UKGs.

	Location	New York City	Beijing	Washington, D.C.	Baltimore
Basic information	Flow type	inflow & outflow	inflow & outflow	inflow	inflow
	Time span	2016.01.01-2016.06.30	2018.01.01-2018.01.31	2019.01.01-2019.05.31	2019.01.01-2019.05.31
	Time interval	1 hour	1 hour	1 hour	1 hour
	Flow length	24 hours	24 hours	24 hours	24 hours
	#Train regions	219	648	178	193
	#Test regions	61	362	59	210
UKG statistics	#Entities	26,604	36,752	11,988	16,579
	#Relations	5	9	5	5
	#Facts	67,418	149,350	33,112	47,884

In the temporal block, we use Transformer layer [36] to model the temporal dependency of urban flow, which consists of a multi-head attention layer, a fully-connected layer, and layer normalization.

For different regions, the impact of spatial and temporal flow dependencies may be different depending on the urban environment. For instance, regions in downtown areas usually have a large flow transition with nearby regions, and thus the spatial dependency should be emphasized. However, the flow patterns in suburban residential regions exhibit strong temporal patterns as people usually leave for work in the morning and go home in the evening, while the flow there is not influenced much by nearby regions.

Therefore, we utilize the KG embeddings of regions with environmental information preserved to guide the fusion of spatial and temporal dependencies. Specifically, we leverage the attention mechanism with KG embedding as the query. For each region l , let $H_{s,l}, H_{t,l} \in \mathbb{R}^{T \times d_h}$ denote the representation after spatial and temporal block respectively, and $E_{KG,l} \in \mathbb{R}^{d_{KG}}$ be its KG embedding. We first project the vectors to get the query vector $Q = E_{KG,l}W_Q$, the key vector $K_i = H_{i,l}W_K$ and the value vector $V_i = H_iW_V$ ($i \in \{s, t\}$). Then we calculate the importance of spatial and temporal representations as:

$$\alpha_i = \text{softmax}\left(\frac{QK_i^T}{\sqrt{d_h}}\right), i \in \{s, t\}. \quad (17)$$

Finally, the output is calculated as

$$H_{st,l} = (\alpha_s V_s + \alpha_t V_t)W_O, \quad (18)$$

where W_O is the output projection matrix.

5 EXPERIMENTS

5.1 Experiment Settings

5.1.1 Datasets. We conduct experiments on four real-world datasets to examine the effectiveness of our model, i.e., NYC dataset, Beijing dataset, Washington, D.C. dataset, and Baltimore dataset. We construct a UKG for each dataset according to methods in Section 3.3. The basic information of datasets and the statistics of corresponding UKGs are reported in Table 3. Due to the lack of related data in NYC, D.C., and Baltimore datasets, we do not incorporate business areas and several relations, i.e., *CoCheckin*, *Competitive*, *ProvideService* and *BelongTo*, in their UKGs.

Each dataset contains the hourly urban flow of all regions in multiple days, and we use each day (24 hours) as a training sample. Since the flow patterns are quite different on weekdays and weekends, we only use the data on weekdays in our experiments, and

our method can be directly applied to the weekends in the same way. We present the details of datasets in Appendix A.

5.1.2 Baselines. Since there are no other works that use the same setting as ours, we compare our model with conditional deep generative models, including **Conditional GAN** [24], **DGAN**, **RCGAN**, **CVAE**, and **Diffwave**. In addition, we also compare with several state-of-the-art urban flow prediction models, including **STGCN** [45], **CCRN** [44], **MSDR** [17], and **3DGCN** [41]. The urban flow prediction models above need historical flow as input to predict the future flow, which however is not available for test regions. Therefore, we sample the historical flow according to the distribution of train regions, and then generate the next steps recurrently [9, 15]. Furthermore, to evaluate the effectiveness of our denoising network design, we examine a variant of our model **KST-Diff(UGnet)** with a different denoising network [40]. The details of baselines and implementation can be found in Appendix B and Appendix C, respectively.

5.1.3 Evaluation Metrics. We adopt widely used Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Symmetric Mean Absolute Percentage Error (SMAPE) to measure the performance. We generate the same number of samples as real data for test regions, and use the mean of all generated samples to calculate MAE, RMSE, and SMAPE. In addition, we also use maximum mean discrepancy (MMD) to measure the distance between generated data distribution and real data distribution, the details of which are shown in Appendix D. Specifically, we follow previous work [5] and treat the flows for each region as vectors to calculate MMD, and report the average MMD of all test regions as the final metric.

5.2 Overall Performance

The comparison of our model with baselines on four datasets is shown in Table 4, from which we have the following observations.

First, our model outperforms all baselines on four datasets by a large margin. For example, our model outperforms existing methods by 14.6% to 35.2% in terms of MAE. Such great improvement demonstrates our model’s ability to accurately generate urban flow based on region characteristics.

Second, RCGAN and Diffwave achieve relatively better performance among baselines because they are able to capture temporal dependencies of urban flow. However, they still perform worse than our model mainly because they do not consider spatial dependencies among regions, while our model successfully captures spatio-temporal dependencies through UKG and KGST block. The

Table 4: Performance comparison with baselines on four datasets. Best results are presented in bold, and the second best results are underlined.

Model	NYC				Beijing				Washington, D.C.				Baltimore			
	MAE	RMSE	SMAPE	MMD	MAE	RMSE	SMAPE	MMD	MAE	RMSE	SMAPE	MMD	MAE	RMSE	SMAPE	MMD
STGCN	33.04	45.63	0.94	3.87	7.29	11.05	1.24	6.32	25.56	59.19	1.16	5.33	14.34	30.78	0.99	4.81
CCRNN	37.78	46.73	0.93	6.02	7.47	11.18	1.25	6.35	31.98	57.01	1.12	6.17	13.72	28.21	0.93	3.70
MSDR	32.98	43.31	0.81	3.36	6.87	11.27	1.22	5.06	32.45	58.30	1.12	5.38	15.28	29.32	1.04	4.16
3DGCN	27.82	41.38	0.80	3.94	4.37	10.58	1.08	3.51	22.44	56.68	0.91	2.14	13.89	29.92	0.92	2.86
Conditional GAN	26.75	41.73	0.71	4.88	3.93	9.98	1.04	4.46	19.51	54.43	0.86	4.31	35.90	87.03	1.02	4.80
DGAN	52.90	81.95	0.81	5.24	8.14	25.93	1.23	4.34	29.12	95.16	1.31	4.55	12.92	26.94	1.38	4.09
CVAE	36.51	58.48	0.70	4.19	2.67	6.61	<u>0.74</u>	4.17	28.76	64.30	1.82	4.17	82.97	254.28	0.97	4.50
RCGAN	<u>23.07</u>	<u>36.12</u>	<u>0.67</u>	5.09	4.11	10.47	1.21	5.45	<u>18.35</u>	48.28	<u>0.82</u>	2.40	40.54	128.67	1.08	5.68
Diffwave	24.53	36.12	0.68	3.20	2.30	6.37	0.78	1.85	36.82	51.07	1.19	5.70	12.88	21.26	1.05	3.71
KSTDiff(UGnet)	34.04	51.35	0.75	3.62	<u>2.26</u>	<u>5.31</u>	<u>0.74</u>	<u>1.82</u>	28.17	<u>46.48</u>	0.97	3.91	20.73	40.01	1.11	3.26
KSTDiff	19.46	32.07	0.52	2.58	1.96	5.15	0.65	1.43	15.67	41.83	0.65	1.79	8.35	15.56	0.68	2.00
Improvement	15.6%	11.2%	22.4%	19.4%	13.3%	3.0%	12.2%	21.4%	14.6%	10.0%	20.7%	16.4%	35.2%	26.8%	26.1%	30.1%

UGnet architecture is able to model spatio-temporal dependencies by GCN and TCN, while it does not show good performance, which further demonstrates the effectiveness of incorporating urban environmental information by UKG.


Third, the performance of many baselines is sensitive to the data. For example, Diffwave performs almost the second best on NYC, Beijing, and Baltimore datasets, while on D.C. dataset it fails to generate flow accurately. Similarly, the performance of RCGAN is considerable on NYC and D.C. datasets but relatively low on Beijing and Baltimore datasets. The volatile performance may result from that the four real-world datasets have different region features, time spans, and data sources of urban flow. However, our model consistently achieves the best performance on all these datasets, which demonstrates the robustness of our model.

At last, urban flow prediction models generally perform much worse than our model on all datasets. This is because they heavily rely on real historical flow, while the historical flow data on test regions is not available in our setting. The finding demonstrates that predictive modeling of urban flow, though extensively studied, cannot be adapted to data-sparse regions, which further shows the significance of generative modeling of urban flow.

5.3 Ablation Study


We conduct an ablation study to further evaluate the influence of region-customized diffusion process (RCDP) and KGST block in the denoising network. Specifically, to validate the flexibility of adapting the diffusion process to each region, we remove the volume prediction module in Figure 3 as well as its training process, and only keep region features in the condition module. In other words, the diffusion process is the same for all regions with $f_{\phi}(c) = 0$. As for KGST, we remove the spatial block as well as the attention module in the denoising network, and only preserve the temporal block in KGST block.

We visualize the SMAPE and MMD metrics on four real-world datasets in Figure 5. It can be observed that the performance drops in all cases after omitting RCDP or KGST, which demonstrates the effectiveness of such design. Besides, we find that the performance deteriorates significantly after removing KGST, which suggests that modeling spatial dependencies and urban environment plays an important role in urban flow generation.

**Figure 5: Performance comparison of models without KGST block or region-customized diffusion process (RCDP).**

5.4 Case Study

5.4.1 Long-term Flow Generation. In real-world scenarios, knowing long-term urban flow distribution in advance can better help urban planning [42]. To evaluate the performance of our model in generating longer flow sequences, we expand the generated flow length to one week, two weeks, and four weeks. For example, we generate the flow with 7×24 time steps for the one-week case. We conduct experiments on NYC dataset and D.C. dataset, and compare the performance of our model with the two best baselines, i.e. RCGAN and Diffwave. As shown in Figure 6, the error generally increases as the flow gets longer, while our model consistently achieves the best performance under various flow lengths, which demonstrates its robustness in modeling urban flow with various lengths. We further present the MMD metric in Appendix E.

**Figure 6: Performance comparison of models for long-term flow generation.**

5.4.2 Downstream Application. To evaluate the utility of generated urban flow in supporting other urban flow related tasks, we conduct experiments by using our generated urban flow for OD generation.

The OD generation task aims to generate origin-destination (OD) flow among regions, i.e., the number of people moving from one region to another, based on region features and total outflow of each region, which is important for a wide range of applications like transportation planning and epidemics spread modeling [2, 28, 39].

Specifically, we leverage the state-of-the-art OD generation model Deep Gravity [33], which uses feed-forward neural network with multiple layers to calculate the OD flow between each pair of regions based on their features and outflows. Note that OD generation task does not consider the temporal dynamics of flow, and only aims to generate the total OD flow among regions. Therefore, we use the sum of outflow in all time intervals as the input of Deep Gravity model. For evaluation, we adopt the same metrics as the original setting including Common Part of Commuters (CPC), Normalized Root Mean Squared Error (NRMSE), Pearson correlation coefficient, and Jenson-Shannon divergence (JSD). Here the benchmark is the performance of Deep Gravity model using real flow data, and a performance closer to the benchmark indicates the better utility of our generated flow.

We compare the performance of OD generation using real outflow data, data generated by our model or baseline models including RCGAN and Diffwave. It can be observed from Table 5 that real data achieves the best performance, while the data generated by our model perform closer to real data than other baselines. Furthermore, the performance of our generated data is not much worse than real data. For example, the CPC of our model is only 7.50% and 4.81% lower than the performance using real data on NYC and Beijing datasets, respectively, which demonstrates that our model can be used to help OD generation task in case of lacking real flow data without much performance loss.

Table 5: Performance comparison with baselines in OD generation. Corr. represents Pearson correlation coefficient.

Model	NYC				Beijing			
	CPC \uparrow	NRMSE \downarrow	Corr. \uparrow	JSD \downarrow	CPC \uparrow	NRMSE \downarrow	Corr. \uparrow	JSD \downarrow
RCGAN	0.656	7.87%	0.634	0.283	0.425	0.96%	0.503	0.468
Diffwave	0.671	8.19%	0.604	0.296	0.540	0.80%	0.638	0.410
KSTDiff	0.715	7.39%	0.683	0.258	0.554	0.79%	0.687	0.400
real data	0.773	5.96%	0.795	0.210	0.582	0.77%	0.713	0.377

5.4.3 Predictive Modeling of Urban Flow. The generative modeling of urban flow essentially aims to learn its conditional distribution. Therefore, it is possible to adapt our model for predictive modeling of urban flow by using historical flow as the condition. To validate this hypothesis, we slightly modify our model for urban flow prediction task. Specifically, in this task, we aim to predict the urban flow for all regions in the next T_2 time intervals in the future ($F_{t+1}, \dots, F_{t+T_2}$) based on urban flow in the past T_1 time intervals (F_{t-T_1+1}, \dots, F_t), where T_1 and T_2 are both 12 in our setting. Since historical flow is available in the flow prediction task, we first remove the volume estimator and adopt the vanilla diffusion process without guidance from the predicted volume. Following previous work [40], we consider the combination of historical and future flow $[F_{t-T_1+1:t}; F_{t+1:t+T_2}]$ as the data x_0 . Moreover, the data x_0 with future flow masked $[F_{t-T_1+1:t}; \mathbf{0}]$ is served as condition, with which we replace the condition module. The training process is similar to Algorithm 1 while we no longer need the pretraining and updating of f_ϕ .




Figure 7: The results of our model and baselines on urban flow prediction task.

We compare the performance of our model with the urban flow prediction baselines mentioned in Section 5.1.2, i.e., 3DGCN, CCRNN, STGCN, and MSDR, on NYC and D.C. datasets. We use four weeks' data as the training set, one week for validation, and one week for test in chronological order. We choose widely used MAE and RMSE to evaluate the performance and present the results in Figure 7. It can be observed that our model outperforms the baselines by 14.2% in terms of RMSE on D.C. dataset, and achieves comparable performance with the best baselines 3DGCN and CCRNN on NYC dataset (5.5% improvement in MAE and -7.6% in RMSE). The findings suggest that our model can be effectively adapted for urban flow prediction, which further demonstrates the power and versatility of generative modeling of urban flow.

6 CONCLUSION

In this paper, we investigate the generative modeling of a typical kind of geospatial data, i.e., urban flow, by studying the urban flow generation problem, which aims to generate urban flow in multiple time steps for regions without historical flow data. We leverage diffusion model to generate urban flow for different regions under different conditions. Specifically, we construct a UKG to model spatial relationships between regions as well as the urban environment, and further propose a knowledge-enhanced spatio-temporal diffusion model to generate urban flow based on UKG. Our model employs a region customized diffusion process to tackle the variance in flow volumes of different regions, and uses a knowledge-enhanced denoising network to capture the spatio-temporal dependencies of urban flow and the impact of environment. Extensive experiments show the superiority of our model design and the utility of generated flow data. Further studies demonstrate its ability for long-term flow generation and urban flow prediction. In the future, we aim to improve the efficiency of our model by accelerating the sampling process of diffusion model, e.g., adopting techniques like DPM-Solver [22]. In addition, a promising direction is to adapt our model to the cross-city scenario that generates urban flow for new cities [31].

ACKNOWLEDGMENTS

This research has been supported in part by the National Key Research and Development Program of China under Grant 2022YFF0606904; in part by the National Natural Science Foundation of China under Grant U22B2057, Grant U21B2036, and Grant U1936217.

REFERENCES

- [1] Juan Miguel Lopez Alcaraz and Nils Strodthoff. 2023. Diffusion-based Conditional ECG Generation with Structured State Space Models. *arXiv preprint arXiv:2301.08227* (2023).


- [2] Constantinos Antoniou, Jaume Barceló, Martijn Breen, Manuel Bulles, Jordi Casas, Ernesto Cipriani, Biagio Ciuffo, Tamara Djukic, Serge Hoogendoorn, Vittorio Marzano, et al. 2016. Towards a generic benchmarking platform for origin-destination flows estimation/updating algorithms: Design, demonstration and validation. *Transportation Research Part C: Emerging Technologies* 66 (2016), 79–98.
- [3] Ivana Balazević, Carl Allen, and Timothy M Hospedales. 2019. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590* (2019).
- [4] Hans Buehler, Blanka Horvath, Terry Lyons, Imanol Perez Arribas, and Ben Wood. 2020. A data-driven market simulator for small data environments. *arXiv preprint arXiv:2006.14498* (2020).
- [5] Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. 2017. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633* (2017).
- [6] Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. 2022. CARD: Classification and Regression Diffusion Models. In *Thirty-Sixth Conference on Neural Information Processing Systems*.
- [7] Huan He, Shifan Zhao, Yuanzhe Xi, and Joyce C Ho. 2023. MedDiff: Generating Electronic Health Records using Accelerated Denoising Diffusion Model. *arXiv preprint arXiv:2302.04355* (2023).
- [8] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.
- [9] Shuodi Hui, Huandong Wang, Zhenhua Wang, Xinghao Yang, Zhongjin Liu, Depeng Jin, and Yong Li. 2022. Knowledge Enhanced GAN for IoT Traffic Generation. In *Proceedings of the ACM Web Conference 2022*. 3336–3346.
- [10] Jiahao Ji, Jingyuan Wang, Chao Huang, Junjie Wu, Boren Xu, Zhenhe Wu, Junbo Zhang, and Yu Zheng. 2022. Spatio-Temporal Self-Supervised Learning for Traffic Flow Prediction. *arXiv preprint arXiv:2212.04475* (2022).
- [11] Guangyin Jin, Yuxuan Liang, Yuchen Fang, Jincui Huang, Junbo Zhang, and Yu Zheng. 2023. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. *arXiv preprint arXiv:2303.14483* (2023).
- [12] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. 2020. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761* (2020).
- [13] Guanyao Li, Shuhan Zhong, Xingdong Deng, Letian Xiang, S-H Gary Chan, Ruiyuan Li, Yang Liu, Ming Zhang, Chih-Chieh Hung, and Wen-Chih Peng. 2022. A Lightweight and Accurate Spatial-Temporal Transformer for Traffic Forecasting. *IEEE Transactions on Knowledge and Data Engineering* (2022).
- [14] Ziqian Lin, Jie Feng, Ziyang Lu, Yong Li, and Depeng Jin. 2019. Deepstn+: Context-aware spatial-temporal neural network for crowd flow prediction in metropolis. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 1020–1027.
- [15] Zinan Lin, Alankar Jain, Chen Wang, Giulia Fanti, and Vyas Sekar. 2020. Using gans for sharing networked time series data: Challenges, initial promise, and open questions. In *Proceedings of the ACM Internet Measurement Conference*. 464–483.
- [16] Chang Liu, Chen Gao, Depeng Jin, and Yong Li. 2021. Improving Location Recommendation with Urban Knowledge Graph. *arXiv preprint arXiv:2111.01013* (2021).
- [17] Dachuan Liu, Jin Wang, Shuo Shang, and Peng Han. 2022. Msdr: Multi-step dependency relation networks for spatial temporal forecasting. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1042–1050.
- [18] Yu Liu, Jingtao Ding, Yanjie Fu, and Yong Li. 2023. UrbanKG: An Urban Knowledge Graph System. *ACM Transactions on Intelligent Systems and Technology* (2023).
- [19] Yu Liu, Jingtao Ding, and Yong Li. 2021. Knowledge-driven site selection via urban knowledge graph. *arXiv preprint arXiv:2111.00787* (2021).
- [20] Yu Liu, Jingtao Ding, and Yong Li. 2022. Developing knowledge graph based system for urban computing. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Geospatial Knowledge Graphs*. 3–7.
- [21] Zhicheng Liu, Fabio Miranda, Weiting Xiong, Junyan Yang, Qiao Wang, and Claudio Silva. 2020. Learning geo-contextual embeddings for commuting flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 808–816.
- [22] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927* (2022).
- [23] Gengchen Mai, Chris Cundy, Kristy Choi, Yingjie Hu, Ni Lao, and Stefano Ermon. 2022. Towards a foundation model for geospatial artificial intelligence (vision paper). In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*. 1–4.
- [24] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [25] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* (2016).
- [26] OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774* [cs.CL]
- [27] Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. 2021. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*. PMLR, 8857–8868.
- [28] Jiangtao Ren and Qiwei Xie. 2017. Efficient od trip matrix prediction based on tensor decomposition. In *2017 18th IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 180–185.
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [30] Can Rong, Jingtao Ding, and Yong Li. 2023. An Interdisciplinary Survey on Origin-destination Flows Modeling: Theory and Techniques. *arXiv preprint arXiv:2306.10048* (2023).
- [31] Can Rong, Jie Feng, and Jingtao Ding. 2023. GODDAG: Generating Origin-destination Flow for New Cities via Domain Adversarial Training. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [32] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings 15*. Springer, 593–607.
- [33] Filippo Simini, Gianni Barlacchi, Massimiliano Luca, and Luca Pappalardo. 2021. A deep gravity model for mobility flows generation. *Nature communications* 12, 1 (2021), 6576.
- [34] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. 2021. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems* 34 (2021), 24804–24816.
- [35] Ledyard R Tucker. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika* 31, 3 (1966), 279–311.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [37] Huandong Wang, Qiaohong Yu, Yu Liu, Depeng Jin, and Yong Li. 2021. Spatio-Temporal Urban Knowledge Graph Enabled Mobility Prediction. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (2021), 1–24.
- [38] Senzhang Wang, Jiyue Li, Hao Miao, Junbo Zhang, Junxing Zhu, and Jianxin Wang. 2022. Generative-Free Urban Flow Imputation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2028–2037.
- [39] Sheng Wei and Lei Wang. 2020. Examining the population flow network in China and its implications for epidemic control based on Baidu migration data. *Humanities and Social Sciences Communications* 7, 1 (2020).
- [40] Haomin Wen, Youfang Lin, Yutong Xia, Huaiyu Wan, Roger Zimmermann, and Yuxuan Liang. 2023. DiffSTG: Probabilistic Spatio-Temporal Graph Forecasting with Denoising Diffusion Models. *arXiv preprint arXiv:2301.13629* (2023).
- [41] Tong Xia, Junjie Lin, Yong Li, Jie Feng, Pan Hui, Funing Sun, Diansheng Guo, and Depeng Jin. 2021. 3dgn: 3-dimensional dynamic graph convolutional network for citywide crowd flow prediction. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 6 (2021), 1–21.
- [42] Yulai Xie, Jingjing Niu, Yang Zhang, and Fang Ren. 2022. Multisize patched spatial-temporal transformer network for short-and long-term crowd flow prediction. *IEEE Transactions on Intelligent Transportation Systems* 23, 11 (2022), 21548–21568.
- [43] Fengli Xu, Pengyu Zhang, and Yong Li. 2016. Context-aware real-time population estimation for metropolis. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 1064–1075.
- [44] Junchen Ye, Leilei Sun, Bowen Du, Yanjie Fu, and Hui Xiong. 2021. Coupled layer-wise graph convolution for transportation demand prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 4617–4625.
- [45] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 3634–3640.
- [46] Hongyi Yuan, Songchi Zhou, and Sheng Yu. 2023. EHRDiff: Exploring Realistic EHR Synthesis with Diffusion Models. *arXiv preprint arXiv:2303.05656* (2023).
- [47] Yuan Yuan, Jingtao Ding, Chenyang Shao, Depeng Jin, and Yong Li. 2023. Spatio-temporal Diffusion Point Processes. *arXiv preprint arXiv:2305.12403* (2023).
- [48] Junbo Zhang, Yu Zheng, and Dekang Qi. 2017. Deep spatio-temporal residual networks for citywide crowd flows prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 31.
- [49] Yingxue Zhang, Yanhua Li, Xun Zhou, Xiangnan Kong, and Jun Luo. 2020. CurbGAN: Conditional urban traffic estimation through spatio-temporal generative adversarial networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 842–852.
- [50] Zhilun Zhou, Yu Liu, Jingtao Ding, Depeng Jin, and Yong Li. 2023. Hierarchical Knowledge Graph Learning Enabled Socioeconomic Indicator Prediction in Location-Based Social Network. In *Proceedings of the ACM Web Conference 2023*. 122–132.
- [51] George Kingsley Zipf. 1946. The P 1 P 2/D hypothesis: on the intercity movement of persons. *American sociological review* 11, 6 (1946), 677–686.

A DETAILS OF DATASETS

In this section, we present the details of four real-world datasets.

- **NYC Dataset.** In this dataset, we define regions as census tracts and aggregate taxi trips from NYC Open Data¹ to obtain the hourly inflow and outflow. The train and test regions are split based on community districts. Specifically, the Manhattan borough contains 12 community districts, and we choose 9 of them as the train regions and the rest 3 as test regions. The region features include number of POIs, area, population, number of takeaway orders, total price of takeaway orders, number of restaurants, and number of firms.
- **Beijing Dataset.** The regions in Beijing dataset are divided by main road networks. We obtain inflow and outflow for each region within Sixth Ring Road by aggregating user trajectories. As for the selection of train and test regions, we choose regions in downtown areas, i.e., Dongcheng, Xicheng, Chaoyang, Fengtai, Shijingshan, and Haidian districts, as the train regions, and the rest are test regions. The region features consist of number of POIs, population, education level, income level, unemployment rate, and number of crimes.
- **Washington, D.C. Dataset**². The regions in this dataset are defined as census tracts, and the inflows are calculated by POI-level hourly visits. The train and test regions are split by counties. Specifically, we choose regions in the District of Columbia as train regions, and regions in Arlington County as test regions. This dataset contains 23 region features including demographics, socioeconomic indicators and land use.
- **Baltimore Dataset**². The regions and inflows are obtained in the same way as D.C. dataset, and we choose regions in Baltimore City and Baltimore County as train regions and test regions, respectively. This dataset contains the same features as D.C. dataset.

The train regions and test regions of four datasets are visualized in Figure 8. Each region feature is normalized across all regions by z-score normalization.



(a) NYC Dataset (b) Beijing Dataset (c) D.C. Dataset (d) Baltimore Dataset

Figure 8: The division of train regions and test regions on four datasets. Blue regions denote train regions and orange regions denote test regions.

B DETAILS OF BASELINES AND IMPLEMENTATION

The details of baselines are as follows.

Conditional deep generative models.

¹<https://data.cityofnewyork.us/Transportation/2016-Yellow-Taxi-Trip-Data/k67s-dv2t>

²<https://github.com/SonghuaHu-UMD/MultiSTGraph>

- **Conditional GAN** [24]: It is the conditional version of vanilla GAN, whose generator and discriminator are both MLPs.
- **DGAN** [15]: It is a GAN-based model that first generates metadata and then generated time series conditioned on metadata. In our experiments, we omit the metadata generation process and use our condition as metadata.
- **RCGAN** [5]: It uses conditional GAN structure with LSTM as generator and discriminator to generate time series.
- **CVAE** [4]: It is a kind of conditional variational autoencoder (VAE) used for financial time series generation conditioned on market states. We adapt it for urban flow generation conditioned on region features.
- **Diffwave** [12]: It is a diffusion-based model for conditional waveform generation. The denoising network of it is similar to ours while it does not consider spatial dependency among regions.

Urban flow prediction models.

- **STGCN** [45]: It employs graph convolution layers and gated convolution layers to capture spatial and temporal correlations, respectively.
- **CCRNN** [44]: It proposes a novel graph convolution architecture for travel demand prediction where the adjacency matrices in different layers are self-learned during training, and it adopts GRU to capture temporal dynamics.
- **MSDR** [17]: It uses a new variant of RNN that considers multiple historical time step to capture temporal dependencies, and combine it with graph-based neural networks for spatio-temporal forecasting.
- **3DGCN** [41]: It uses an expanded GCN to capture spatio-temporal correlations of urban flow among regions at different time steps.

Variant of our model.

- **KSTDiff(UGnet)**: In this model, we change the denoising network in KSTDiff to UGnet proposed in [40]. It uses an Unet-like architecture and adopts Temporal Convolution Network (TCN) and graph convolution network (GCN) to capture temporal and spatial dependencies. Apart from the denoising network, the framework is the same as ours as shown in Figure 3.

C IMPLEMENTATION

In our experiments, we set the number of diffusion steps $N = 1000$, use 5 residual layers in the denoising module and the dimension of KG embedding is $d_{KG} = 32$. The learning rate is searched from $\{1e-4, 5e-4, 1e-3, 5e-3\}$ and the number of training epochs range from $\{200, 500, 1000, 1500, 2000\}$. The pretraining epoch number M_1 is set to 100 and M_2 is searched from $\{1, 10, 100, w/o\}$, where w/o means that we only pretrain the volume estimator and do not train it anymore in the following training process. For all models, we tuned the hyperparameters on each of the datasets and report the best results.

D DETAILS OF EVALUATION METRICS

The maximum mean discrepancy (MMD) is a metric that measures the similarity between two distributions by comparing statistics of the samples. Specifically, for each region, given the generated flow samples $\{f_g^i\}_{i=1}^n$ and real flow samples $\{f_r^i\}_{i=1}^m$, the MMD is

calculated as:

$$\begin{aligned}
 MMD = & \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n K(f_g^i, f_g^j) - \frac{2}{mn} \sum_{i=1}^n \sum_{j=1}^m K(f_g^i, f_r^j) \\
 & + \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq i}^m K(f_r^i, f_r^j),
 \end{aligned} \tag{19}$$

where $K(x, y) = \exp(-\|x - y\|^2 / (2\sigma^2))$ is the RBF kernel. In our study, we calculate MMD based on the code in this repository³.

E LONG-TERM FLOW GENERATION RESULTS

We present the MMD of long-term flow generation results in Figure 9, which shows a similar trend with MAE.




Figure 9: Performance comparison of models for long-term flow generation.

³https://github.com/easezyc/deep-transfer-learning/blob/master/MUDA/MFSAN/MFSAN_3src/mmd.py