
Satellites Reveal Mobility: A Commuting Origin-destination Flow Generator for Global Cities

Can Rong^{1,2} * Xin Zhang¹ Yanxin Xi³ Hongjie Sui¹ Jingtao Ding^{1†} Yong Li^{1‡}

¹Dept. Electronic Engineering, BNRist, Tsinghua University

²Singapore-MIT Alliance for Research and Technology

³Dept. Computer Science, University of Helsinki

Abstract

Commuting Origin-destination (OD) flows, capturing daily population mobility of citizens, are vital for sustainable development across cities around the world. However, it is challenging to obtain the data due to the high cost of travel surveys and privacy concerns. Surprisingly, we find that satellite imagery, publicly available across the globe, contains rich urban semantic signals to support high-quality OD flow generation, with over 98% expressiveness of traditional multisource hard-to-collect urban sociodemographic, economics, land use, and point of interest data. This inspires us to design a novel data generator, GIODGen (**G**lobal-scale **O**igin-**D**estination **F**low **G**enerator), which can generate OD flow data for any cities of interest around the world. Specifically, GIODGen first leverages Vision-Language Geo-Foundation Models to extract urban semantic signals related to human mobility from satellite imagery. These features are then combined with population data to form region-level representations, which are used to generate OD flows via graph diffusion models. Extensive experiments on 4 continents and 6 representative cities show that GIODGen has great generalizability across diverse urban environments on different continents and can generate OD flow data for global cities highly consistent with real-world mobility data. We implement GIODGen as an automated tool, seamlessly integrating data acquisition and curation, urban semantic feature extraction, and OD flow generation together. It has been released at <https://github.com/tsinghua-fib-lab/generate-od-pubtools>.

1 Introduction

Commuting origin-destination (OD) flows profile the regular population movement in daily life between every two urban regions within a given urban area [55], providing a critical foundation for various applications. For example, OD flows are essential inputs for simulations and analyses in traffic management and urban planning, serving as the travel demands of citizens for developing more informed policies [12, 66]. Recent studies have also explored the potential of OD flows in supporting research related to the United Nations Sustainable Development Goals (SDGs) on community detection [15], urban resilience [52, 28], public health [37] and environmental protection [81], further demonstrating their importance. Therefore, obtaining OD flows holds great significance for cities around the world, especially as they increasingly support interdisciplinary research efforts that span traditional and emerging urban challenges [55, 6].

However, it is very challenging to obtain this valuable data. Traditional methods for obtaining OD flows, such as door-to-door travel surveys [55] and the aggregation of large-scale individual mobility

*Work primarily done at Tsinghua University.

†Corresponding author: dingjt15@tsinghua.org.cn

‡Corresponding author: liyong07@tsinghua.edu.cn




Figure 1: An illustration of replacing the traditional hard-to-collect urban data by easily accessible global public data (e.g., satellite imagery and population data) in OD flow generation.

trajectories [11, 34, 22, 49], are often infeasible in target cities due to high costs and privacy concerns. Although recent studies have explored physical [85, 64, 6] and computational [51, 53, 63, 44, 58, 57, 56] models for OD flow generation, these approaches face key limitations. Physical models rely solely on population distribution and adopt overly simplistic assumptions, ignoring urban semantic differences and resulting in limited accuracy. Computational models perform better but still rely on hard-to-obtain, expensive, and fine-grained inputs such as sociodemographics, infrastructure data, land use, and points of interest (POIs), which limits their applicability to only one or a small number of data-rich, developed cities.

Recent studies [73, 82, 69, 25, 36, 50, 79, 30, 72, 24, 26, 61, 10, 54, 2, 1] have demonstrated that satellite images, publicly available around the world, can provide important urban semantics related to human activities in urban regions. For example, residential structures identifiable in satellite imagery often indicate regions with high outbound commuting demand, while the presence of commercial buildings typically corresponds to regions with concentrated inbound flows. In contrast, sparsely built-up areas tend to exhibit lower levels of human activity and mobility. Therefore, satellite imagery holds the potential to serve as an alternative to the hard-to-obtain inputs required by traditional computational OD generation models, and supports the generation of OD flows for cities with limited data, as shown in Figure 1.

Accordingly, we propose GIODGen (**G**lobal-scale **O**rigin-**D**estination Flow **G**enerator), which incorporates satellite imagery combined with population data as input for representing urban regions, and leverages recently developed vision-language geo-foundation models [84, 42] to extract high-quality semantic features for OD flow generation in any city worldwide. Specifically, satellite imagery is first preprocessed by cropping and stitching it according to the boundaries of each urban region. This step removes irrelevant pixels and avoids the noise outside the regions. Then, with the help of RemoteCLIP [43], the semantic features of urban regions are extracted. Finally, the extracted regional features are combined with population data to profile the urban space as input for the state-of-the-art OD flow generation model, WEDAN [56], and generate OD flows. To investigate the validity of adopting only the publicly available satellite imagery and population data as input, we conduct experiments on representative urban areas in the U.S., Europe, China, Brazil, and Africa. We surprisingly find that totally publicly available data can serve as a perfect alternative to the hard-to-obtain inputs mentioned above and generate OD flows with high consistency with real-world data. This may shed light on the potential for extending OD flow generation to any city worldwide with the help of ubiquitous satellite imagery and population data, which is of great significance for the sustainable development of cities around the world. To facilitate practical use, we release GIODGen as an open-source tool at <https://github.com/tsinghua-fib-lab/generate-od-pubtools>, which seamlessly automates the collection and preprocessing of satellite imagery and population data, urban semantic feature extraction, and OD flow generation together.

The contributions of this paper can be summarized as follows:




Figure 2: Cases of satellite imagery for urban regions.

- We find that publicly available satellite imagery and population data can replace hard-to-obtain urban data, such as sociodemographics, infrastructure data, land use, and POIs, to profile the urban space for inferring human mobility and generate OD flows.
- We propose GIODGen, a novel OD flow generator that leverages vision-language geo-foundation models to extract high-quality semantic features from satellite imagery and generate OD flows via graph diffusion models.
- We implement GIODGen as an efficient tool, which integrates data acquisition, curation and preprocessing, urban semantic feature extraction, and OD flow generation together. With only the boundaries of urban regions given, GIODGen can provide the OD flow data for cities of interest with *only a single line of code*.
- We conduct experiments on representative urban areas around the world to demonstrate the effectiveness of GIODGen, which may pave the way for research on generating OD flows for global cities and contribute to lowering the barrier of data access for sustainable urban development.

2 Preliminaries

2.1 Definitions and Problem Formulation

Definition 1. City. A city refers to a large, integrated urban area such as New York City, London, or Beijing. Our study focuses on commuting movements that occur within the spatial boundaries of a city.

Definition 2. Urban Region. Urban regions $\mathcal{R} = \{r_i | i = 1, 2, \dots, N\}$ are the basic spatial units within a city for the OD flow generation task. Population movements are modeled between these regions.

Every region corresponds to its unique satellite imagery, which consists of 2D, bird’s-eye views captured from satellites covering the region, as illustrated in Figure 2.

Definition 3. Commuting OD Flow. This refers to the number of people $\mathcal{F}_{r_{org}, r_{dst}}$ who live in one urban region r_{org} and work in another urban region r_{dst} . It captures the regular, static daily movement from home to the workplace and remains relatively stable for a long time.

PROBLEM 1. Commuting OD Flow Generation. Given a city and its region division, generate the commuting OD flows $\mathcal{F} = \{\mathcal{F}_{r_{org}, r_{dst}} | r_{org}, r_{dst} \in \mathcal{R}\}$ that represent the daily commuting patterns of the city. The generated OD flows should be as consistent as possible with real-world human mobility patterns.

2.2 Related Works

Urban Representation Learning via Satellite Imagery. Satellite imagery has been demonstrated to carry rich information about human activities, making it a feasible basis for profiling urban regions and generating OD flows. Specifically, satellite imagery offers a comprehensive overview of urban development from a macro perspective, while also enabling detailed analysis of regional contents at a micro level. At a coarse spatial level, satellite imagery has been leveraged to infer socio-economic indicators through tailored supervised and self-supervised learning methods, facilitating




Figure 3: The framework and data pipeline of GIODGen.

the extraction of high-level semantic representations of urban areas. These indicators include poverty levels [5, 4, 25, 36, 50, 79, 73, 27, 41, 82, 24, 61, 10, 2, 1, 72], crop yields [45, 46, 59, 71, 78, 80], land cover [31, 69, 62], commercial activeness [30, 40], and environmental metrics [72, 77, 76, 20, 60, 82, 54, 27]. At a finer spatial scale, satellite imagery has been extensively utilized to monitor urban geospatial features, enabling the identification of various socio-physical entities such as streets, airplanes, vehicles, and recreational facilities like baseball fields [74, 75, 39, 73, 83, 76, 20, 67, 76, 14]. These capabilities collectively highlight the strong potential of satellite imagery as a powerful and scalable modality for comprehensively representing urban regions, thereby enabling the generation of high-quality OD flow data around the world.

OD Flow Generation. OD flow obtaining, the task of obtaining OD flows for urban areas of interest lacking human mobility data, is an important yet traditionally costly and time-consuming process that relies on travel surveys [55, 3, 34]. Researchers have explored alternative data sources related to individual trajectories, such as call detail records (CDRs) [11, 34] and cellular network access (CNAs) [22, 49] to efficiently access OD flows. However, these methods face challenges regarding data accessibility and privacy concerns. To avoid these issues, researchers have introduced model-based methods for generating OD flows, which are often categorized into two distinct types. The first is physical models, e.g., the gravity model [85] and radiation model [64], which model the population movement by mimicking physical laws, such as Newton’s law of universal gravitation and the emission and absorption of the radiation process. These models are overly simplistic and are unable to effectively model the complexity inherent in human mobility. The second category is data-driven computational models that leverage machine learning and deep learning techniques to generate OD flows based on various urban factors, including sociodemographics, POI distributions, etc. [51, 53, 63, 44, 58, 57]. These methods achieve superior performance via sophisticated model structures but the high requirements for input data limits their applicability, especially in underdeveloped areas. Our framework leverages global publicly available data to replace the hard-to-collect features for profiling urban regions and generating OD flows. This breaks the barrier of data accessibility while reserving the superior performance advantages of computational models.

3 GIODGen: Commuting Origin-destination Flow Generator

In Figure 3, we illustrate the conceptual structure of the framework. It comprises three stages: *satellite imagery preprocessing*, *urban semantic feature extraction*, and *origin-destination flow generation*.

Satellite Imagery Preprocessing. For each region in the city, GIODGen first identifies specific tiles of satellite imagery corresponding to each region and downloads them from any online platform, such as Google Earth and Esri World Imagery. A tile refers to a smaller, square-shaped portion of a larger satellite image, which is created by dividing the satellite image of the whole world into

manageable pieces. This tiling process facilitates more efficient obtaining and processing, especially when working with large-scale satellite images. With all tiles obtained, the next phase involves concatenation and cropping, shaping the images to match the exact contours of each region. This procedure is critical for ensuring that the imagery faithfully represents the geographical and physical attributes of each region, creating an accurate visual portrayal. During cropping, the image content beyond region boundaries is removed and zero-padded to preserve spatial accuracy and exclude irrelevant information, as formalized below:

$$M_r(x, y) = \begin{cases} 1, & \text{if } (x, y) \in \text{Boundary}(r) \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where (x, y) denotes the pixel coordinates, M_r is the mask of region r . The image is then masked by M_r to remove the irrelevant parts, which is detailed in the equation below:

$$I_{\text{final}} = I_{\text{raw}} \odot M_r \quad (2)$$

where I_{raw} is the raw satellite image, I_{final} is the final satellite image of region r , and $\odot$ denotes the element-wise multiplication operation. Once processed, the satellite images for the regions appear as depicted in Figure 2.

Urban Semantic Feature Extraction. After completing the previous step, every region in the specified area is assigned a satellite image to represent its unique characteristics and functions, as illustrated in Figure 3. However, the raw pixel data from the images are unsuitable for direct use as urban regional features, due to the presence of excessive noise and irrelevant, redundant information unrelated to human mobility. As such, a feature extraction model is needed to process the satellite imagery, enabling the distillation of the most significant semantic information related to OD flows. This model could be any structure, such as VGG [65], ResNet [29], or ViT [18], and we choose the vision encoder from the large multimodal model, which also has the special design to handle the satellite images, RemoteCLIP [43]. Pre-trained with natural language supervision for the vision encoder, the large multimodal model exhibits zero-shot capabilities, allowing it to extract versatile semantic features from satellite imagery for various downstream tasks. We adopt RemoteCLIP as the semantic feature extractor in our framework. Specifically, we input the preprocessed satellite image of each region to the vision encoder of RemoteCLIP, whose output is a 1024-dimension high-level feature that captures the urban semantics of that urban region. During this stage, we only use the vision encoder of RemoteCLIP with the pre-trained weights frozen. Because recent works have demonstrated that finetuning the foundation model on the specific task will lead to forgetting the general knowledge learned from the pre-training and result in overfitting [17, 35, 32]. This process is formalized as follows:

$$E_r = \mathcal{V}(I_r), \quad (3)$$

where I_r is the satellite image of region r , $\mathcal{V}$ is the vision encoder of RemoteCLIP, and E_r is the 1024-dimension feature vector of region r . The final input to the OD flow generation model is built by merging the semantic features, extracted from satellite images, with the population of region r , as shown below:

$$X_r = [E_r, P_r], \quad (4)$$

where P_r means the population of r and $[]$ denotes a concatenation operation. X_r is the final features profiling region r , which will be input into the OD flow generation model before being processed by a multi-layer perceptron (MLP) whose parameters are trained while training the following OD flow generation model. In this way, the vision encoder can provide a general and versatile representation of the region, while the representation is further refined to fit the specific task of OD flow generation.

Origin-destination Flow Generation. We adopt WEDAN [56], a state-of-the-art OD flow generation model based on graph denoising diffusion [38, 47, 70, 23, 7, 21], to generate OD flows. In WEDAN, urban regions are represented as nodes and OD flows as weighted directed edges in a graph. The model follows a conditional generation paradigm, where node attributes serve as guidance for the denoising process and generating the directed edges and corresponding weights. In other words, the semantic characteristics of urban space are leveraged as conditioning inputs, guiding the denoising process to produce realistic and spatially consistent OD flows. Benefiting from its formulation of the city-wide OD network as a directed, weighted graph, WEDAN is capable of generating OD flows that closely align with real-world population mobility patterns. In our framework, the urban semantic feature X_r extracted for each region is used as the node attribute input, guiding the generation of

edge directions and weights during the denoising process. The procedure can be described as follows:

$$p_\theta(\mathbf{F}^{t-1}|\mathbf{F}^t, \mathcal{C}_\mathcal{R}) = \mathcal{N}(\mathbf{F}^{t-1}; \mu_\theta(\mathbf{F}^t, t, \mathcal{C}_\mathcal{R}), (1 - \bar{\alpha}^t)\mathbf{I}), \quad (5)$$

where

$$\mu_\theta(\mathbf{F}^t, t, \mathcal{C}_\mathcal{R}) = \frac{1}{\sqrt{\alpha_t}}(\mathbf{F}^t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(\mathbf{F}^t, t, \mathcal{C}_\mathcal{R})). \quad (6)$$

In the formulas, $\mathcal{C}_\mathcal{R} = \{X_r | r \in \mathcal{R}\}$ is the semantic features of urban regions within the city, t is the diffusion step, $\mathbf{I}$ is the identity matrix, ϵ_θ means denoising networks, μ_θ denotes μ of the Gaussian distribution, $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ refer to the noise scheduler. The denoising network utilized here is the graph transformer network [19], which predicts the noise needed to be removed from the current graph state $\mathbf{F}^t$ to reach the previous state $\mathbf{F}^{t-1}$.

4 Experiments

In this section, we conduct experiments to answer two key research questions:

- **RQ1:** For precision, can entirely public data, satellite imagery, and population data, sufficiently represent urban spatial characteristics and support the generation of high-quality OD flows?
- **RQ2:** For generalizability, does GIODGen demonstrate cross-continental transferability, enabling it to generate OD flows in diverse global cities with the help of global public input data?

All experiments were conducted on a single NVIDIA GeForce RTX 4090 GPU (24GB) and an Intel Xeon Platinum 8358 CPU @ 2.60GHz. For all trainable models, we performed the grid search to select optimal hyperparameters. The same training, validation, and testing splits were used when training all models and evaluating the performance.

4.1 Performance of Public Satellite Imagery and Population Data (RQ1)

We begin by evaluating the effectiveness of publicly available satellite imagery and population data in supporting OD flow generation. Specifically, we compare the performance of several existing OD flow generation models using two different types of input: (1) traditional fine-grained features that are often costly or difficult to obtain, and (2) publicly accessible features. In this experiment, we use detailed sociodemographic attributes, economic indicators, and POI distributions as the traditional inputs, and satellite imagery combined with population data as the public inputs. To evaluate representational capacity across input types, we apply two groups of existing models commonly used for OD flow generation: physical models, including the classical gravity and radiation models, which rely solely on population distribution act as lower-bound reference of performance; Computational models, including Random Forest [51], DeepGravity [63], GMEL [44], NetGAN [8], and WEDAN [56], which support richer semantic inputs and allow us to assess the performance impact of different input data sources. A detailed introduction of the models is provided in Appendix B.1.

Dataset. In this part, we use the dataset [56] from the United States to conduct experiments. This dataset totally consists of 3,333 urban areas, including 3,233 counties and 100 metropolitans. Each urban area is associated with its region division: census tracts within counties and census block groups within metropolitans. The dataset contains two parts: (1) city characteristics in urban regions, including population structure, education level, poverty, income, vehicle ownership, and other socioeconomic indicators from the American Community Survey (ACS) and (2) OD flow data between urban regions provided by the National Census Bureau through the Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES) [9].

Experimental Settings. To systematically evaluate the performance of different models and input types, we adopt root mean square error (RMSE), normalized RMSE (NRMSE), and common part of commuting (CPC) as metrics. We split the 3,333 urban areas into 8:1:1 proportions for training, validation, and testing. The evaluation metrics are averaged over all test urban areas. All experiments are repeated 5 times with different random seeds with average results and standard deviations reported.

Experimental Results. As shown in Table 1, public data demonstrates performance close to that of traditional hard-to-obtain urban data. We can see that the performances of global public data

Table 1: Performance of existing OD flow generation models on test urban areas in the United States with different input data. Traditional hard-to-obtain urban data and global public data are put together for better comparison. The former’s results are shown in the upper line, while the latter’s are shown in the lower line. Perc. in the table indicates percentage of the performance (CPC) achieved by the model trained with publicly available data relative to that of the model trained with traditional hard-to-collect fine-grained data. Gap means the remaining performance gap in RMSE between the model trained with publicly available data and the model trained with traditional hard-to-collect fine-grained data.

	CPC $\uparrow$	Perc.	RMSE $\downarrow$	Gap	NRMSE $\downarrow$
Gravity Model	0.321 ± 0.02	-	174.0 ± 10.4	-	2.222
Radiation Model	0.347 ± 0.04	-	196.9 ± 11.7	-	2.502
Random Forest	0.494 ± 0.02	-	100.4 ± 6.5	-	1.282
	0.480 ± 0.03	97.1%	114.0 ± 7.1	-13.5%	1.455
DeepGravity	0.449 ± 0.01	-	92.9 ± 9.9	-	1.186
	0.427 ± 0.05	95.1%	99.9 ± 12.0	-7.5%	1.275
GMEL	0.462 ± 0.01	-	94.3 ± 4.0	-	1.204
	0.451 ± 0.02	97.6%	105.4 ± 7.9	-11.8%	1.345
NetGAN	0.517 ± 0.05	-	89.1 ± 17.0	-	1.138
	0.468 ± 0.06	90.5%	98.0 ± 12.3	-10.0%	1.251
WEDAN	0.634 ± 0.01	-	64.06 ± 3.3	-	0.818
	0.623 ± 0.02	98.3%	67.88 ± 6.1	-5.9%	0.867

are slightly less accurate than those of traditional hard-to-collect information, but the deviations are almost indistinguishable. Across all models based on the data-driven schema, public data achieve over 90% of the performance of traditional features. This indicates that such data can further effectively profile and represent spatial characteristics in a manner comparable to traditional hard-to-obtain inputs. Moreover, computational models based on public data significantly outperform classic physical models, demonstrating the superior performance and widespread applicability of public data, including satellite imagery and population data.

4.2 Generalization Across Continents (RQ2)

In this part, we conduct experiments that transfer the OD flow generation model from one trained continent to another and investigate the generalization ability of GIODGen with global publicly available satellite imagery and population data. Specifically, we use the data from the United States to train the GIODGen model, and then transfer it to the United Kingdom for evaluation.

Dataset. The training data in the United States is the same as the dataset used in Section 4.1. The test data in the United Kingdom is from the Office for National Statistics (ONS) [48], which provides population-level commuting flows between all census units. Specifically, the data contains OD flows of 326 Local Authority Districts (LAD). In each district, the whole area is divided into urban regions by the boundaries of Middle layer Super Output Areas (MSOA).

Experimental Settings. We use 90% urban areas in the United States as the training data and the remaining 10% as the validation data to tune the hyperparameters. The trained GIODGen model is then transferred to the United Kingdom for generating OD flows. Generated OD flows will be compared with the ones from census data and evaluated by metrics introduced in Section 4.1. Metrics are the average of all LADs. We also introduce existing models to perform the same task and evaluate their performance as baselines. It is worth noting that existing models are only provided with population data and inter-region distances as input. This limitation is unavoidable, as fine-grained urban features that align across training and testing cities are typically unavailable in cross-continental experimental settings.

Experimental Results. As shown in Table 2, the proposed framework can be transferred between different continents with a robust performance. Compared with the baseline models, our framework

Table 2: Performance comparison of baseline methods with geographically distributed population and our proposed framework profiling urban regions based on global public data on the United Kingdom. All the methods are trained in the United States and evaluated in the United Kingdom. IMP. denotes improvement percentage relative to the baselines.

Model	CPC $\uparrow$	IMP.	RMSE $\downarrow$	IMP.	NRMSE $\downarrow$	IMP.
Gravity Model	0.240 ± 0.01	-27.5%	101.6 ± 3.42	+48.9%	1.752 ± 0.06	+48.9%
Radiation Model	0.323 ± 0.07	-2.4%	211.6 ± 4.27	-6.4%	3.647 ± 0.07	-6.4%
Random Forest	0.334 ± 0.03	+0.9%	223.2 ± 3.42	-12.2%	3.847 ± 0.06	-12.2%
DeepGravity	0.359 ± 0.04	+8.5%	157.0 ± 7.90	+21.1%	2.706 ± 0.13	+21.1%
GMEL	0.362 ± 0.01	+9.4%	149.1 ± 10.00	+25.0%	2.570 ± 0.17	+25.1%
NetGAN	0.331 ± 0.06	-	198.9 ± 15.21	-	3.429 ± 0.26	-
GIODGen	0.485 ± 0.03	+46.5%	72.68 ± 2.13	+63.5%	1.253 ± 0.04	+63.5%

can achieve a promising performance and outperform baselines on all metrics with a large margin, i.e., 34.0% improvement on CPC, and 28.5% improvement on RMSE and NRMSE. This indicates that 1) public data, i.e., satellite imagery and population data, can effectively be transferred between different continents; 2) the proposed framework can improve the transferability of the computational OD flow generation models; 3) computational models based on public data can beat physical models in cross-continental transfer experiments, supporting a great improvement on global scale OD flow generation. This strongly supports the conclusion that our framework is highly suitable for generating OD flows across diverse urban areas worldwide.

4.3 Case Studies on Typical Urban Areas Worldwide (RQ1 & RQ2)

To further validate the effectiveness of the proposed data generator, we collect human mobility-related data from all over the world and apply the OD flow generation to the corresponding urban areas. Specifically, we collect the data from the following typical urban areas and detailed data processing is provided in Appendix B.3.

- **Europe.** We select London and Paris as representative cities in Europe. The OD flow data for London is sourced from ONS [48], while the data for Paris is synthesized by Sebastian et al. [33] using the 2015 French population census provided by the National Institute of Statistics and Economic Studies (INSEE) [68]. Urban regions in Paris are defined at the commune level.
- **China.** We collect OD flow data for Beijing and Shanghai, two major metropolitan centers in China. The Beijing dataset is provided by a leading internet location service provider. The Shanghai data is extracted from Call Detail Records (CDRs) by China’s largest telecommunications company, following the method proposed by Iqbal et al. [34]. Urban regions in both cities are defined by subdistricts. All data was collected in 2016.
- **Brazil.** Julio et al. [13] extracted OD flows for the Rio de Janeiro Metropolitan Area (RJMA) in 2014 using CDRs. Urban regions are defined by Municípios.
- **Africa.** We use OD flow data extracted from CDRs collected in Senegal in 2013 as part of the D4D Challenge [16], processed following Iqbal et al. [34]. Urban regions are defined by Arrondissements.

It is worth noting that real-world OD flow data is inherently difficult to collect. Datasets differ in sampling methods, temporal coverage, and spatial granularity, and often contain noise or bias. As such, they do not serve as perfect ground truth, but only as approximate references. In our evaluation, we focus on comparing the spatial patterns and distribution characteristics between generated and observed data, aiming to assess the plausibility and consistency of the generated flows under real-world constraints.

Experimental Results. Based on the data collected from the representative urban areas above, we apply the proposed GIODGen framework to generate OD flows and compare the results with the corresponding real-world data. Given the heterogeneous data sources, varying sampling methods, and inherent noise across regions, a unified error-based evaluation metric is not applicable for direct comparison. To quantitatively assess the alignment between the generated and observed OD flows, we




Figure 4: Correlation analysis between generated OD flows and the flows extracted from diverse mobility-related data sources for typical urban areas around the world.

compute the Spearman rank correlation coefficients, visualized as rank-aligned normalized flow curve in Figure 4. Specifically, the data-oriented curve is constructed by the ascending ordered sort of the OD flow oriented from collected data, and the generated flow is also sorted in the same order. Then, the two curves are normalized to the same range and smoothed for better visualization. Although the distribution of the collected data is more skewed, the two curves are aligned with each other in a very similar manner, indicating that the generated OD flows are consistent with the real-world data in terms of data distribution. It is important to note that due to differences in data quality and sources, these correlation values are not directly comparable across cities. In addition, we qualitatively assess spatial patterns by visually comparing the generated and observed OD distributions on a case-by-case basis in Appendix B.4, under varying regional conditions, to provide a more intuitive understanding of the framework’s performance. The experimental results show that the OD flows generated by GIODGen closely align with the real-world data in terms of distributional structure, as evidenced by the high Spearman rank correlations observed in most cases. Although lower correlations may arise due to limitations in the reference data itself, high correlation values provide strong empirical support for the credibility and applicability of the proposed framework.

5 Discussion

Conclusion We find that the global publicly available data, i.e., satellite imagery and population data, has a strong representational power for profiling urban regions and supporting the generation of OD flows across diverse urban areas worldwide. Notably, public data can achieve 98% of the performance of the traditional hard-to-obtain urban data. Building on this insight, we propose a novel data generator, GIODGen, which can generate realistic OD flows for any urban area worldwide. GIODGen integrates vision-language geo-foundation models to extract urban semantic features from satellite imagery and region-level population statistics and generates OD flows through a graph denoising diffusion-based approach. Extensive experiments across all over the world validate both the representational power of public data and the generalizability of GIODGen. For the convenience of practical use, we release GIODGen as an open-source tool, which can automatically complete the data acquisition, curation, preprocessing, urban semantic feature extraction, and OD flow generation with only the boundaries of urban regions given.

Limitations and Boarder Impact GIODGen currently has two main limitations. First, the OD flow generation model is primarily trained on U.S. data. While we demonstrate strong cross-continental generalization with the aid of satellite imagery, broader geographic coverage would improve robustness. Second, the generator produces only static commuting OD flows, limiting its applicability to dynamic, fine-grained tasks such as time-dependent traffic analysis. Despite these limitations, GIODGen can benefit various domains, including urban planning, transportation design, energy consumption, carbon emissions, and public health research.

Ethical Claims All individual-level location data used in this study have been rigorously anonymized. Additionally, all computations are conducted locally, with no reliance on cloud-based infrastructure, to guarantee both privacy and data security.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant No. 62476152 and No. U24B20180; and research grants from the Beijing National Research Center for Information Science and Technology (BNRist).

References

- [1] Jacob Levy Abitbol and Marton Karsai. Interpretable socioeconomic status inference from aerial imagery through urban patterns. *Nature Machine Intelligence*, 2(11):684–692, 2020.
- [2] Donghyun Ahn, Jeasurk Yang, Meeyoung Cha, Hyunjoo Yang, Jihee Kim, Sangyoon Park, Sungwon Han, Eunji Lee, Susang Lee, and Sungwon Park. A human-machine collaborative approach measures economic development using satellite imagery. *Nature Communications*, 14(1):6811, 2023.
- [3] Kay W Axhausen, Andrea Zimmermann, Stefan Schönfelder, Guido Rindsfuser, and Thomas Haupt. Observing the rhythms of daily life: A six-week travel diary. *Transportation*, 29(2):95–124, 2002.
- [4] Kumar Ayush, Burak Uz Kent, Kumar Tanmay, Marshall Burke, David Lobell, and Stefano Ermon. Efficient poverty mapping from high resolution remote sensing images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12–20, 2021.
- [5] Kumar Ayush, Burak Uz Kent, Burak Uz Kent, Marshall Burke, David Lobell, and Stefano Ermon. Generating interpretable poverty maps using object detection in satellite images. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI’20*, 2021.
- [6] Hugo Barbosa, Marc Barthelemy, Gourab Ghoshal, Charlotte R James, Maxime Lenormand, Thomas Louail, Ronaldo Menezes, José J Ramasco, Filippo Simini, and Marcello Tomasini. Human mobility: Models and applications. *Physics Reports*, 734:1–74, 2018.
- [7] Andreas Bergmeister, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer. Efficient and scalable graph generation through iterative local expansion. *The Twelfth International Conference on Learning Representations*, 2024.
- [8] Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. Netgan: Generating graphs via random walks. In *International conference on machine learning*, pages 610–619. PMLR, 2018.
- [9] U.S. Census Bureau. Lehd origin-destination employment statistics data (2002-2021), 2024.
- [10] Marshall Burke, Anne Driscoll, David B Lobell, and Stefano Ermon. Using satellite imagery to understand and promote sustainable development. *Science*, 371(6535):eabe8628, 2021.
- [11] Francesco Calabrese, Giusy Di Lorenzo, Liang Liu, and Carlo Ratti. Estimating origin-destination flows using opportunistically collected mobile phone location data from one million users in boston metropolitan area. 2011.
- [12] Oded Cats and Erik Jenelius. Dynamic vulnerability analysis of public transport networks: mitigation effects of real-time information. *Networks and Spatial Economics*, 14:435–463, 2014.

- [13] Júlio César Chaves, Moacyr AHB da Silva, Ricardo de Souza Alencar, Alexandre G Evsukoff, and Vinícius da Fonseca Vieira. Human mobility and socioeconomic datasets of the rio de janeiro metropolitan area. *Data in brief*, 51:109695, 2023.
- [14] Jingbo Chen, Chengyi Wang, Dongxu He, Jiansheng Chen, and Anzhi Yue. Practical bottom-up golf course detection using multispectral remote sensing imagery. In *Image and Graphics: 9th International Conference, ICIG 2017, Shanghai, China, September 13-15, 2017, Revised Selected Papers, Part II* 9, pages 552–559. Springer, 2017.
- [15] Xiao-Jian Chen, Yuhui Zhao, Chaogui Kang, Xiaoyue Xing, Quanhua Dong, and Yu Liu. Characterizing the temporally stable structure of community evolution in intra-urban origin-destination networks. *Cities*, 150:105033, 2024.
- [16] Yves-Alexandre de Montjoye, Zbigniew Smoreda, Romain Trinquart, Cezary Ziemlicki, and Vincent D Blondel. D4d-senegal: the second mobile phone data for development challenge. *arXiv preprint arXiv:1407.4885*, 2014.
- [17] Yuxuan Ding, Lingqiao Liu, Chunna Tian, Jingyuan Yang, and Haoxuan Ding. Don’t stop learning: Towards continual learning for the clip model. *arXiv preprint arXiv:2207.09248*, 2022.
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *The Ninth International Conference on Learning Representations*, 2021.
- [19] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020.
- [20] Jie Feng, Jun Zhang, Tianhui Liu, Xin Zhang, Tianjian Ouyang, Junbo Yan, Yuwei Du, Siqi Guo, and Yong Li. Citybench: Evaluating the capabilities of large language models for urban tasks. page 5413–5424, 2025.
- [21] Xingcheng Fu, Yisen Gao, Yuecen Wei, Qingyun Sun, Hao Peng, Jianxin Li, and Xianxian Li. Hyperbolic geometric latent diffusion model for graph generation. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- [22] David Gundlegård, Clas Rydergren, Nils Breyer, and Botond Rajna. Travel demand estimation and network assignment based on cellular network data. *Computer Communications*, 95:29–42, 2016.
- [23] Kilian Konstantin Haefeli, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer. Diffusion models for graphs benefit from discrete state spaces. *arXiv preprint arXiv:2210.01549*, 2022.
- [24] Sungwon Han, Donghyun Ahn, Hyunji Cha, Jeasurk Yang, Sungwon Park, and Meeyoung Cha. Lightweight and robust representation of economic scales from satellite imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 428–436, 2020.
- [25] Sungwon Han, Donghyun Ahn, Sungwon Park, Jeasurk Yang, Susang Lee, Jihee Kim, Hyunjoon Yang, Sangyoon Park, and Meeyoung Cha. Learning to score economic development from satellite imagery. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2970–2979, 2020.
- [26] Xixuan Hao, Wei Chen, Yibo Yan, Siru Zhong, Kun Wang, Qingsong Wen, and Yuxuan Liang. Urbanvlp: Multi-granularity vision-language pretraining for urban socioeconomic indicator prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28061–28069, 2025.
- [27] Xixuan Hao, Wei Chen, Yibo Yan, Siru Zhong, Kun Wang, Qingsong Wen, and Yuxuan Liang. Urbanvlp: multi-granularity vision-language pretraining for urban socioeconomic indicator prediction. 2025.

- [28] Masahiko Haraguchi, Akihiko Nishino, Akira Kodaka, Maura Allaire, Upmanu Lall, Liao Kuei-Hsien, Kaya Onda, Kota Tsubouchi, and Naohiko Kohtake. Human mobility data and analysis for urban resilience: A systematic review. *Environment and Planning B: Urban Analytics and City Science*, 49(5):1507–1535, 2022.
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [30] Zhiyuan He, Su Yang, Weishan Zhang, and Jiulong Zhang. Perceiving commercial activeness over satellite images. In *Companion Proceedings of the The Web Conference 2018*, pages 387–394, 2018.
- [31] Danfeng Hong, Lianru Gao, Jing Yao, Bing Zhang, Antonio Plaza, and Jocelyn Chanussot. Graph convolutional networks for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 59(7):5966–5978, 2020.
- [32] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023.
- [33] Sebastian Hörl and Milos Balac. Open synthetic travel demand for paris and île-de-france: Inputs and output data. *Data in Brief*, 39:107622, 2021.
- [34] Md Shahadat Iqbal, Charisma F Choudhury, Pu Wang, and Marta C González. Development of origin–destination matrices using mobile phone call data. *Transportation Research Part C: Emerging Technologies*, 40:63–74, 2014.
- [35] Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd Omama, Tao Chen, Alaa Maalouf, Shuang Li, Ganesh Iyer, Soroush Saryazdi, Nikhil Keetha, et al. Conceptfusion: Open-set multimodal 3d mapping. *arXiv preprint arXiv:2302.07241*, 2023.
- [36] Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.
- [37] Jayson S Jia, Xin Lu, Yun Yuan, Ge Xu, Jianmin Jia, and Nicholas A Christakis. Population flow drives spatio-temporal distribution of covid-19 in china. *Nature*, 582(7812):389–394, 2020.
- [38] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In *International conference on machine learning*, pages 10362–10383. PMLR, 2022.
- [39] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS journal of photogrammetry and remote sensing*, 159:296–307, 2020.
- [40] Tong Li, Yanxin Xi, Huandong Wang, Yong Li, Sasu Tarkoma, and Pan Hui. Learning representations of satellite imagery by leveraging point-of-interests. *ACM Transactions on Intelligent Systems and Technology*, 14(4):1–32, 2023.
- [41] Tong Li, Shiduo Xin, Yanxin Xi, Sasu Tarkoma, Pan Hui, and Yong Li. Predicting multi-level socioeconomic indicators from structural urban imagery. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 3282–3291, 2022.
- [42] Xiang Li, Congcong Wen, Yuan Hu, Zhenghang Yuan, and Xiao Xiang Zhu. Vision-language models in remote sensing: Current progress and future trends. *IEEE Geoscience and Remote Sensing Magazine*, 2024.
- [43] Fan Liu, Delong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 2024.

- [44] Zhicheng Liu, Fabio Miranda, Weiting Xiong, Junyan Yang, Qiao Wang, and Claudio Silva. Learning geo-contextual embeddings for commuting flow prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 808–816, 2020.
- [45] Rose M Rustowicz, Robin Cheong, Lijing Wang, Stefano Ermon, Marshall Burke, and David Lobell. Semantic segmentation of crop type in africa: A novel dataset and analysis of deep learning methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 75–82, 2019.
- [46] Jorge Andres Chamorro Martinez, Laura Elena Cué La Rosa, Raul Queiroz Feitosa, Ieda Del’Arco Sanches, and Patrick Nigri Happ. Fully convolutional recurrent networks for multivariate crop recognition from multitemporal image sequences. *ISPRS Journal of Photogrammetry and Remote Sensing*, 171:188–201, 2021.
- [47] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *International conference on artificial intelligence and statistics*, pages 4474–4484. PMLR, 2020.
- [48] Office for National Statistics. User guide to origin-destination data for census 2021, england and wales. <https://www.ons.gov.uk>, October 2023. Accessed: 2023-10-26.
- [49] Changxuan Pan, Jiangang Lu, Shan Di, and Bin Ran. Cellular-based data-extracting method for trip distribution. *Transportation research record*, 1945(1):33–39, 2006.
- [50] Anthony Perez, Christopher Yeh, George Azzari, Marshall Burke, David Lobell, and Stefano Ermon. Poverty prediction with public landsat 7 satellite imagery and machine learning. *arXiv preprint arXiv:1711.03654*, 2017.
- [51] Nastaran Pourebrahim, Selima Sultana, Amirreza Niakanlahiji, and Jean-Claude Thill. Trip distribution modeling with twitter data. *Computers, Environment and Urban Systems*, 77:101354, 2019.
- [52] Paulo Jorge Gomes Ribeiro and Luís António Pena Jardim Gonçalves. Urban resilience: A conceptual framework. *Sustainable Cities and Society*, 50:101625, 2019.
- [53] Caleb Robinson and Bistra Dilkina. A machine learning approach to modeling human migration. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 1–8, 2018.
- [54] Esther Rolf, Jonathan Proctor, Tamma Carleton, Ian Bolliger, Vaishaal Shankar, Miyabi Ishihara, Benjamin Recht, and Solomon Hsiang. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature communications*, 12(1):4392, 2021.
- [55] Can Rong, Jingtao Ding, and Yong Li. An interdisciplinary survey on origin-destination flows modeling: Theory and techniques. *ACM Comput. Surv.*, 57(1), October 2024.
- [56] Can Rong, Jingtao Ding, Yan Liu, and Yong Li. A large-scale dataset and benchmark for commuting origin-destination flow generation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [57] Can Rong, Jingtao Ding, Zhicheng Liu, and Yong Li. City-wide origin-destination matrix generation via graph denoising diffusion. *arXiv preprint arXiv:2306.04873*, 2023.
- [58] Can Rong, Jie Feng, and Jingtao Ding. Goddag: Generating origin-destination flow for new cities via domain adversarial training. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [59] Marc Rußwurm and Marco Körner. Self-attention for raw optical satellite time series classification. *ISPRS journal of photogrammetry and remote sensing*, 169:421–435, 2020.
- [60] Erzhao Shao, Jie Feng, Yingheng Wang, Tong Xia, and Yong Li. One-shot transfer learning for population mapping. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1588–1597, 2021.

- [61] Makkunda Sharma, Fan Yang, Duy-Nhat Vo, Esra Suel, Swapnil Mishra, Samir Bhatt, Oliver Fiala, William Rudgard, and Seth Flaxman. Kidsat: satellite imagery to map childhood poverty dataset and benchmark. *arXiv preprint arXiv:2407.05986*, 2024.
- [62] Li Sibbo, Zhang Xin, Lin Yuming, and Li Yong. M3luc: Multi-modal model for urban land-use classification. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems*, pages 1–12, 2024.
- [63] Filippo Simini, Gianni Barlacchi, Massimiliano Luca, and Luca Pappalardo. A deep gravity model for mobility flows generation. *Nature communications*, 12(1):1–13, 2021.
- [64] Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. A universal model for mobility and migration patterns. *Nature*, 484(7392):96–100, 2012.
- [65] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *The Third International Conference on Learning Representations*, 2015.
- [66] Lijun Sun, Kay W Axhausen, Der-Hong Lee, and Xianfeng Huang. Understanding metropolitan patterns of daily encounters. *Proceedings of the National Academy of Sciences*, 110(34):13774–13779, 2013.
- [67] Xian Sun, Peijin Wang, Cheng Wang, Yingfei Liu, and Kun Fu. Pbnet: Part-based convolutional neural network for complex composite object detection in remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 173:50–65, 2021.
- [68] The National Institute of Statistics and Economic Studies (INSEE). Individus localisés au canton-ou-ville en 2015. <https://www.insee.fr/fr/statistiques/3625223>, 2015.
- [69] Burak Uzkent, Evan Sheehan, Chenlin Meng, Zhongyi Tang, Marshall Burke, David Lobell, and Stefano Ermon. Learning to interpret satellite images using wikipedia. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.
- [70] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *The Eleventh International Conference on Learning Representations*, 2023.
- [71] Anna X Wang, Caelin Tran, Nikhil Desai, David Lobell, and Stefano Ermon. Deep transfer learning for crop yield prediction with remote sensing data. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 1–5, 2018.
- [72] Yanxin Xi, Tong Li, Huandong Wang, Yong Li, Sasu Tarkoma, and Pan Hui. Beyond the first law of geography: Learning representations of satellite imagery by leveraging point-of-interests. In *Proceedings of the ACM Web Conference 2022*, pages 3308–3316, 2022.
- [73] Yanxin Xi, Yu Liu, Tong Li, Jintao Ding, Yunke Zhang, Sasu Tarkoma, Yong Li, and Pan Hui. A satellite imagery dataset for long-term sustainable development in united states cities. *Scientific data*, 10(1):866, 2023.
- [74] Yanxin Xi, Yu Liu, Zhicheng Liu, Sasu Tarkoma, Pan Hui, and Yong Li. From pixels to progress: generating road network from satellite imagery for socioeconomic insights in impoverished areas. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI '24*, 2024.
- [75] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018.
- [76] Congxi Xiao, Jingbo Zhou, Yixiong Xiao, Jizhou Huang, and Hui Xiong. Refound: Crafting a foundation model for urban region understanding upon language and visual foundations. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3527–3538, 2024.

- [77] Yibo Yan, Haomin Wen, Siru Zhong, Wei Chen, Haodong Chen, Qingsong Wen, Roger Zimmermann, and Yuxuan Liang. Urbanclip: Learning text-enhanced urban region profiling with contrastive language-image pretraining from the web. In *Proceedings of the ACM on Web Conference 2024*, pages 4006–4017, 2024.
- [78] Christopher Yeh, Chenlin Meng, Sherrie Wang, Anne Driscoll, Erik Rozi, Patrick Liu, Jihyeon Lee, Marshall Burke, David Lobell, and Stefano Ermon. Sustainbench: Benchmarks for monitoring the sustainable development goals with machine learning. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.
- [79] Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and deep learning to understand economic well-being in africa. *Nature communications*, 11(1):2583, 2020.
- [80] Jiaxuan You, Xiaocheng Li, Melvin Low, David Lobell, and Stefano Ermon. Deep gaussian process for crop yield prediction based on remote sensing data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [81] Jinwei Zeng, Yu Liu, Jingtao Ding, Jian Yuan, and Yong Li. Estimating on-road transportation carbon emissions from open data of road network and origin-destination flow data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14539–14547, 2024.
- [82] Xin Zhang, Yu Liu, Yuming Lin, Qingmin Liao, and Yong Li. Uv-sam: adapting segment anything model for urban village identification. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’24/IAAI’24/EAAI’24. AAAI Press, 2024.
- [83] Siru Zhong, Xixuan Hao, Yibo Yan, Ying Zhang, Yangqiu Song, and Yuxuan Liang. Urbancross: enhancing satellite image-text retrieval with cross-domain adaptation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6307–6315, 2024.
- [84] Yue Zhou, Litong Feng, Yiping Ke, Xue Jiang, Junchi Yan, Xue Yang, and Wayne Zhang. Towards vision-language geo-foundation model: A survey. *arXiv preprint arXiv:2406.09385*, 2024.
- [85] George Kingsley Zipf. The $p \propto 1/p^2$ hypothesis: on the intercity movement of persons. *American sociological review*, 11(6):677–686, 1946.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: See Abstract and Section 1.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: See Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[NA\]](#)

Justification: Our work is not about theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: See Section 4 and Github link provided in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: See Github link provided in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: See Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: See Section 5.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: See Section 5.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: See Section 5.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We have provided the license information in the Github repository.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: See Section 4.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: Our work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [\[NA\]](#)

Justification: Our work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [\[Yes\]](#)

Justification: We use LLMs for checking the writing errors in the paper.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Additional Details of GIODGen

A.1 Forward Diffusion Process

We give an introduction to the OD generation model based on graph diffusion in this part. Like the traditional diffusion model, graph diffusion-based models include two main processes: the forward diffusion and the reverse denoising process. During training, the forward diffusion process is utilized to create the training dataset for the denoising network. This process involves gradually adding small Gaussian noise to the original OD flow data, eventually transforming them into pure noise that adheres to a standard Gaussian distribution, as illustrated in the formula below:

$$\begin{aligned} q(F_{ij}^t | F_{ij}^{t-1}) &= \mathcal{N}(F_{ij}^t; \sqrt{1 - \beta_t} F_{ij}^{t-1}, \beta_t \mathbf{I}), \\ q(F_{ij}^1, \dots, F_{ij}^T | F_{ij}^0) &= \prod_{t=1}^T q(F_{ij}^t | F_{ij}^{t-1}), \end{aligned} \quad (7)$$

where F_{ij}^t denotes the flow starting at r_i flowing into r_j at t -th diffusion step, $\mathcal{N}$ is the Gaussian distribution, β_t is the noise level at time t , and $\mathbf{I}$ is the identity matrix.

The denoising process works in reverse of the forward diffusion process, which is already described in the main paper.

A.2 Training

The OD flow generation model requires training on diverse datasets to generalize effectively across global scales. We employ data from the whole United States to expose the model to varied mobility patterns, including those from developed and underdeveloped areas.

To generate training samples, the forward process is employed to produce OD flows with varying degrees of noise. Node inputs consist of semantic features extracted from satellite images and population data, while edge inputs are formed by noisy OD flows sampled from the forward process. The model is designed to predict the Gaussian noise disturbances to be removed at each noise level. The training employs the mean squared error (MSE) as the loss function, and optimization is conducted using the Adam optimizer, consistent with [57]. The loss function is formulated as follows:

$$\mathcal{L} = \mathbb{E}_{t, \epsilon \sim \mathcal{N}(0, \mathbf{I})} [\|\epsilon - \epsilon_\theta(\mathbf{F}^t, t, \mathcal{C}_R)\|_2^2] \quad (8)$$

where $\|\cdot\|$ denotes the L_2 norm.

B Experimental Details

B.1 Detailed Introduction of Existing OD Flow Generation Models

- **Random Forest.** [51] The Random Forest model is a type of ensemble learning method that builds multiple decision trees and combines their outputs to make a final prediction. It is a popular choice for its simplicity and effectiveness in handling complex relationships between various urban indicators profiling urban regions and OD flows.
- **DeepGravity.** [63] DeepGravity is a multi-layer perceptron (MLP)-based model inspired by the traditional gravity model. It models the process of decision making of destination choice as a classification problem and uses the softmax function to calculate the probability of each destination. With the outflow given, the model can generate the OD flows by multiplying the

probability of each destination and the outflow. In this paper, we directly generate the OD flows by using the model because there is no outflow for any urban areas around the world.

- **GMEL. [44]** GMEL is a graph-based model that uses the graph structure of the urban regions to generate the OD flows. It models the urban regions by aggregating the features of the regions and the features of the connections between the regions. After the geo-contextual embedding learned from the model, random forest is used to generate the OD flows based on the urban region embeddings.
- **NetGAN. [8]** NetGAN is a generative adversarial network (GAN)-based model that minimizes the Wasserstein distance between the random walk sequences of the generated and real OD flows, where the mobility flow networks are represented as graphs. The original NetGAN model is designed for unweighted networks. We adapt it to weighted networks by using the weighted adjacency matrix of the urban regions. The weighted matrix is the OD matrix of the city.
- **WEDAN. [56]** WEDAN is a graph denoising diffusion-based model that uses the graph structure of the urban regions to generate the OD flows. It models the urban regions as nodes and OD flows as the directed weighted edges. The model uses the conditional graph denoising diffusion process to generate the OD flows given the urban region features of the city.

B.2 Details of Evaluation Metrics

For evaluation, we adopt root mean square error (RMSE), normalized RMSE (NRMSE), and common part of commuting (CPC) as metrics, with the computations detailed below:

$$\begin{aligned}
 \text{RMSE} &= \sqrt{\frac{1}{|\mathbf{F}|} \sum_{r_i, r_j \in \mathcal{R}} \|\mathbf{F}_{ij} - \hat{\mathbf{F}}_{ij}\|_2^2}, \\
 \text{NRMSE} &= \frac{\text{RMSE}}{\sqrt{\frac{1}{N^2} \sum_{r_i, r_j \in \mathcal{R}} \|F_{ij} - \bar{F}_{ij}\|_2^2}}, \\
 \text{CPC} &= \frac{2 \sum_{r_i, r_j \in \mathcal{R}} \min(\mathbf{F}_{ij}, \hat{\mathbf{F}}_{ij})}{\sum_{r_i, r_j \in \mathcal{R}} \mathbf{F}_{ij} + \sum_{r_i, r_j \in \mathcal{R}} \hat{\mathbf{F}}_{ij}}
 \end{aligned} \tag{9}$$

where $\bar{\mathbf{F}}$ represents the expectation of the OD flows $\mathbf{F}$, which is extracted from collected data.

B.3 Data Processing for Typical Urban Areas Around the World

- **United States** OD flows of urban areas in the United States are collected and provided by the National Census Bureau through the Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES) [9]. This dataset provides commuting flows across all census blocks. In line with previous studies [51, 44, 53, 63], we aggregate this data to the census tract level, which defines the regions in our work, while counties represent areas. This dataset offers extensive population coverage and high accuracy, making it widely used in research related to human mobility. In this work, we utilized data collected in 2018.
- **China.** OD flow data are collected for Beijing and Shanghai, two key cities in China. The OD flow data for Beijing is provided by a major internet location service provider in China. Data from Shanghai are extracted from CNAs, provided by China’s largest telecommunications company, using the method proposed by Iqbal et al. [34].
- **Brazil.** Julio et al. [13] utilized CDRs to extract OD flows of the Rio de Janeiro Metropolitan Area (RJMA) in 2014. Urban regions in RJMA are defined by Municípios. This dataset provides total flows including repeated records of the same individuals. Therefore, there is a certain degree of bias.
- **Africa.** We identified the dataset of CDRs within Senegal in 2013 [16]. The data was processed to extract OD flows using the method proposed by Iqbal et al. [34]. Urban regions in Senegal are defined by Arrondissements. In Africa, the collection of mobility-related data

is much more challenging due to the lack of comprehensive and up-to-date infrastructure and limited technological resources, and funding of large-scale data gathering efforts. Because of the sparsity and scarcity, the spatial granularity of the data cannot be further refined.

B.4 Spatial Visualization of Comparison between Generated and Collected OD Flows

We would like to clarify that some visualizations for Beijing, Shanghai, Rio de Janeiro, and a region in Senegal also **appear in a separate submission currently under review at Scientific Data**. In that paper, we present a large-scale dataset of commuting OD flows for 1,625 global cities, generated using the data generator GLODGen introduced in this work, to support research on sustainable urban development.

The core contributions and goals of the two submissions are fundamentally different:

- The Scientific Data submission focuses on the global-scale dataset as a scientific contribution.
- The present paper focuses on the automated generator (GLODGen) and its methodology for world-wide OD flow generation, which highlights the potential of generalization for urban areas around the world.

Although these cities appear in both papers, the visualizations are styled differently and are used for distinct purposes: In Scientific Data, they are presented to validate the quality and scope of the dataset, whereas, in this work, they demonstrate the data generator’s generalizability and applicability across diverse global cities. We emphasize that the experimental results are not reused, but rather, the same generated data is used in different contexts to support complementary contributions in the two submissions.

The experimental results of spatial visualization are shown in the following.

To demonstrate the potential of our framework, which relies solely on public data, for generating OD flows globally, we conduct experiments in urban areas including the representative cities in the United States, Europe, China, Brazil, and Africa. Since OD flow data of these diverse areas are collected from different data sources with different sampling biases and noise, we cannot use a unified standard to evaluate the generation of these areas and compare generated OD flows with flows from data in terms of traditional error-based metrics. We visually and qualitatively compare the spatial distribution of the generation and data in a case-by-case manner under different scenarios and data conditions to provide an intuitive understanding of the performance of our framework.

B.4.1 China

We generate OD flows for Beijing and Shanghai, two representative cities in China. The visually qualitative comparisons of the generation and data are shown in Figure 5. Figures from above to below are Beijing and Shanghai respectively and the left column is the generation, while the right column is OD flows from data. The generated OD flows exhibit a notable spatial similarity to the data, with the city center and boundary shapes accurately reflected in their spatial distribution. The agreement between generated flows and observed data under the given sampling scheme highlights the framework’s validity. Nonetheless, slight discrepancies exist. For instance, flows from Beijing’s suburbs to the city center are underrepresented, and generated flows for Shanghai’s outskirts are somewhat denser than observed. These differences could stem from the unique urban characteristics of large cities, in which global public data may not fully follow a uniform regularity.

B.4.2 Europe

The Greater London and Greater Paris Metropolitan Areas are chosen as representative regions in Europe. The visually qualitative comparisons of the generation and data are shown in Figure 6. Figures from above to below are London and Paris respectively and the left column is the generation, while the right column is OD flows from data. The generated OD flows exhibit substantial similarity to the data. The city boundary is distinctly represented by blue OD flows, while the dense red and yellow lines highlight the city center. Unlike the results observed in China, the generated flows for London successfully capture long-distance OD flows between the city center and the outskirts. In Paris, the generation accurately reflects the limited long trip flows, suggesting that our framework effectively captures Paris’s centralized urban structure and efficient planning.




Figure 5: Visualization of the generated and data-oriented OD flows in Beijing and Shanghai. The generated OD flows show high consistency with the data-oriented data in both cities, demonstrating the effectiveness of GLODGen in capturing real-world mobility patterns.

B.4.3 Brazil

We select the Rio de Janeiro Metropolitan Area as the representative urban area for Brazil. The spatial visualization is presented in Figure 7, demonstrating strong consistency between the generated OD flows and the data. This result highlights the similarity in urban structure and human mobility patterns between Rio de Janeiro and the United States, suggesting that global public data and OD flows from the United States may effectively support OD flow generation for Brazil.

B.4.4 Africa

Considering the scarcity of data in Africa, we select a representative country in Africa, Senegal, to generate and evaluate our framework. OD flows of Senegal are aggregated from CDRs. The spatial visualization results are presented in Figure 8. As shown, the generated OD flows closely resemble the actual CDR data in terms of spatial distribution. Specifically, the generated OD flow data for Senegal correctly identifies the city center and the urban boundary, which highlights the capability of our framework in accurately modeling mobility patterns even in underdeveloped urban areas.




Figure 6: Visualization of the generated and data-oriented OD flows in London and Paris. The generated OD flows show high consistency with the data-oriented data in both cities, demonstrating the effectiveness of GIODGen in capturing real-world mobility patterns.




Figure 7: Visualization of the generated and data-oriented OD flows in Rio de Janeiro. The generated OD flows show high consistency with the data-oriented data, demonstrating the effectiveness of GIODGen in capturing real-world mobility patterns.

B.5

We conduct a feature importance analysis to investigate the importance of different input features for the OD flow generation. Specifically, we compare the performance degradation of the generation with only one input feature at a time. The results are shown in Table 3. The experimental settings are the same as Table 1. From a theoretical perspective, this is because satellite imagery inherently contains a degree of redundancy related to population information. Previous studies have shown that satellite imagery can be used to estimate population distribution. As a result, using satellite imagery alone can still achieve relatively good performance. However, using population data alone provides an insufficient characterization of urban regions and largely eliminates the representation of regional heterogeneity (e.g., regions with the same population but different functional roles contribute very differently to mobility patterns).




Figure 8: Visualization of the generated and data-oriented OD flows in Senegal. The generated OD flows show high consistency with the data-oriented data, demonstrating the effectiveness of GLODGen in capturing real-world mobility patterns.

Table 3: Comparison of the feature importance of population data and satellite imagery for the OD flow generation.

Input Feature	CPC	RMSE	NRMSE
Population + Satellite	0.623	67.88	0.867
Population Only	0.394	137.06	1.750
Satellite Only	0.500	92.42	1.180