Learning the complexity of urban mobility with deep generative collaboration network

Yong Li (liyong07@tsinghua.edu.cn)
Tsinghua University https://orcid.org/0000-0001-5617-1659
Yuan Yuan
Tsinghua University
Jingtao Ding
Tsinghua University
Depeng Jin
Tsinghua University

Article

Keywords:

Posted Date: December 12th, 2023

DOI: https://doi.org/10.21203/rs.3.rs-3666762/v1

Additional Declarations: There is **NO** Competing Interest.

Learning the complexity of urban mobility with deep generative collaboration network

Yuan Yuan[1+], Jingtao Ding[1+*], Depeng Jin[1], and Yong Li[1*]

[1]Beijing National Research Center for Information Science and Technology (BNRist), Department of Electronic Engineering, Tsinghua University, Beijing, P. R. China
[*]To whom correspondence should be addressed; Emails: liyong07@tsinghua.edu.cn, dingjt15@tsinghua.org.cn.
[+]Yuan Yuan and Jingtao Ding contribute equally to this work.

ABSTRACT

City-scale individual movements, resulting population flows, and urban morphology intricately intertwine, collectively contributing to the complexity of urban mobility, impacting critical aspects of a city, including socioeconomic exchanges and epidemic transmission. Existing models, derived from the fundamental laws governing human mobility, often capture only partial facets of this complexity. This paper introduces DeepMobility, a powerful deep generative collaboration network to bridge the heterogeneous behaviors of individuals and collective behaviors emerging from the entire population via constructing a unified model that encapsulates the multifaceted nature of complex urban mobility. Our experiments, conducted on mobility trajectories and flows in cities of China and Senegal, reveal that, in contrast to state-of-the-art deep learning models that simply "memorize" observed data, DeepMobility excels in learning the intricate data distribution and successfully reproduces the existing universal scaling laws that characterize human mobility behaviors at both the individual and population levels. DeepMobility also exhibits robust generalization capabilities, enabling it to generate realistic trajectories and flows for cities lacking corresponding training data. Our approach underscores the feasibility of employing generative deep learning to model the underlying mechanism of human mobility, and establishes an effective generative machine learning framework to capture the complexity of urban mobility comprehensively.

Human mobility, an indispensable component of urban functionality, serves as a linchpin for establishing vital connections across diverse city regions, thereby enabling residents to access and leverage urban services[1,2]. Beyond fostering commercial interactions and innovation diffusion, it concurrently engenders multifaceted challenges including traffic congestion[3,4] and epidemic transmission[5–7]. Consequently, human mobility plays a pivotal role in shaping urban dynamics across cultural, economic, and environmental dimensions[8–11]. City-scale individual movements, the resulting population flows, and the urban morphology are intricately intertwined, collectively contributing to the complexity of urban mobility. This complexity, with its extensive historical research context[12,13], further amplifies the dynamics governing urban economic and social systems, highlighting the enduring significance of modeling the complexities of urban mobility for understanding urban patterns and fostering sustainable development[14,15].

In the pursuit of understanding the intricate dynamics of urban mobility, statistical physicists have increasingly focused on the analysis of empirical mobility data to uncover universal patterns in human mobility since the turn of this century[16–18]. This leads to the discovery of scaling laws governing both individual movements[16,17,19–22] and population flows[18,23–27]. Individual human movements, unlike physical particles, can be approximated by a scale-free Lévy flight, with truncated power law distributed spatial distance[16,17], up to a distance characterized by the individual's radius of gyration, which also follows truncated power law distribution[17]. Conversely, temporal memory effects, representing the tendency to revisit particular locations, are characterized by the scaling laws including Zipf's law of visitation frequency[17], sublinear growth in the number of unique locations visited[20], and an ultraslow diffusion process[20]. In terms of collective behaviors, the flow of population mobility can be broadly characterized by the gravity law[23], which posits that the probability of movement between regions is proportionate to their respective populations. Furthermore, this flow can be more precisely predicted by the radiation model[18]. Temporal regularities emerge as well, notably the distance-frequency scaling law[27], revealing an inverse square relationship between the number of visitors to a location and their visit frequency. Additionally, power law distributions govern the number of trips between regions and trips originating or ending in specific regions[25,26]. Despite the success in identifying these fundamental laws, it is essential to note that the existing models, often developed through theoretical derivations, are limited in their capacity to fully encompass all facets of these laws. The significant disparities among various mobility laws, particularly those stemming from distinct levels of analysis of individual movements versus population flow, present a significant challenge in the pursuit of a unified theoretical model.

Recent advancements in Artificial Intelligence (AI), particularly in the domain of deep generative AI models, offer a promising alternative to mechanistic approaches in constructing high-capacity models capable of capturing various mobility laws. Deep learning models like generative adversarial networks (GANs)[28] or variational autoencoders (VAEs)[29] have demonstrated remarkable versatility in learning the distribution of real-world mobility data, and generating synthetic data with comparable statistical properties[30]. Previous research has successfully applied these models to specific modeling tasks, including the generation of human trajectories that mimic individual movements[31–36] and the prediction of population flows between pairs of regions[37–39]. Notably, these deep learning-based models have shown higher accuracy compared to traditional theoretical models[18,23]. However, despite the considerable realism achieved by these DL-based models in specific-level descriptions of urban mobility, they predominantly focus on either individual trajectories or aggregated flows, but cannot consider both simultaneously. The collective mobility patterns emerge from the bottom-up aggregation of individual movements, which in turn impose constraints that influence individual behaviors. This bi-directional influence between individual and population levels contributes to the complexity of urban mobility. Yet, effectively characterizing this intricate interplay in deep generative models remains an unresolved challenge.

In response to this challenge, we introduce DeepMobility, a novel generative deep-learning approach that captures the multifaceted nature of complex urban mobility. DeepMobility conceptualizes human movement as a sequential decision-making process and employs a GAN-based framework to train a deep generative collaboration network for simulating human mobility behaviors. This neural network comprises three components: a generator for producing individual trajectories and aggregated population flows, a discriminator for assessing the quality of these trajectories and flows against real data, and a critic for providing guidance from the discriminator so as to improve the generator. To characterize the dynamic interplay between individual behaviors and broader population trends, DeepMobility incorporates two innovative collaborative learning mechanisms: bottom-up interaction modeling and top-down feedback refinement. The bottom-up approach, implemented in the generator, effectively integrates social interactions into individual movement patterns. Concurrently, the top-down approach, functioning in the critic, allows for precise adjustments to individual behaviors based on aggregated population-level flow patterns. In this way, it successfully bridges the heterogeneous behaviors of individuals and collective behaviors emerging from the entire population to capture the multifaceted nature of complex urban mobility.

Utilizing data from three Chinese metropolises (Beijing, Shanghai, Shenzhen) and Senegal, we trained DeepMobility models to generate human mobility trajectories and the resulting flows across urban regions. Remarkably, despite being purely data-driven and free of predefined mechanisms, DeepMobility was able to simultaneously reproduce the existing multi-scale scaling laws previously discovered by physicists. The emergence of this capability to reproduce complex patterns indicates that DeepMobility goes beyond mere data memorization, capturing the underlying mechanisms of urban mobility in a way other deep learning models have not. In terms of the realistic generation of trajectories and flows, our results show the significant improvements over previous models. DeepMobility demonstrates advantages in five key statistical properties of trajectories and achieves substantial enhancements in flow generation: up to 120% in Beijing, 112% in Shanghai, 136% in Shenzhen, and 81% in Senegal. Moreover, we showcase DeepMobility's geographical transferability by verifying its effectiveness in cities with scarce mobility data, underscoring its broad applicability.

Results

DeepMobility framework

To fully model the complex urban mobility with both the individual movement laws and the emerging collective flow patterns, we propose a deep generative collaboration network for generating the multiscale realistic human trajectories and the resulting mobility flows in a city (Fig. 1). We aim to learn a mobility model that simulates an individual's mobility decision-making process based on observed data. Specifically, given an individual's travel history $x_{<t}$ at time t, it estimates the probability of visiting location l_t, i.e., $\pi(l_t|x_{<t})$, and generates a spatiotemporal trajectory by sequentially sampling $l_t \sim \pi(\cdot|x_{<t})$ to obtain a sequence of individual movements. Then, for the entire urban population, their movements are learned by following a joint policy, i.e., $\Pi(\mathbf{l}_t|\mathbf{x}_{<t})$, and aggregate into region-wide flows that reflect daily rhythms of urban activities. To capture human mobility patterns at both individual and population levels, we formulate the learning process of DeepMobility as the following multi-objective optimization problem with respect to π:

$$\min_{\pi} \left(L_{\text{dist}}\left(P_{\text{data}}(l_t|x_{<t}), \pi(l_t|x_{<t})\right), L_{\text{error}}\left(\mathbf{F}_{t,\text{data}}, \mathrm{F}(\Pi(\mathbf{l}_t|\mathbf{x}_{<t}))\right)\right) . \tag{1}$$

The first objective aims to minimize the distance between the statistical distribution of generated movements, i.e., $\pi(l_t|x_{<t})$, and that of observed data, i.e., $P_{\text{data}}(l_t|x_{<t})$, in terms of spatiotemporal regularity. The second objective aims to minimize the reconstruction error of generated flows, i.e., $\mathrm{F}(\Pi(\mathbf{l}_t|\mathbf{x}_{<t}))$, that are aggregated from population's movements. The complexity of this problem mainly lies in the bi-directional influence across individual and population levels of human mobility. First, the bottom-up aggregation of mobility flows, indicated by $\Pi(\mathbf{l}_t|\mathbf{x}_{<t})$, essentially incorporates the influence of social interactions

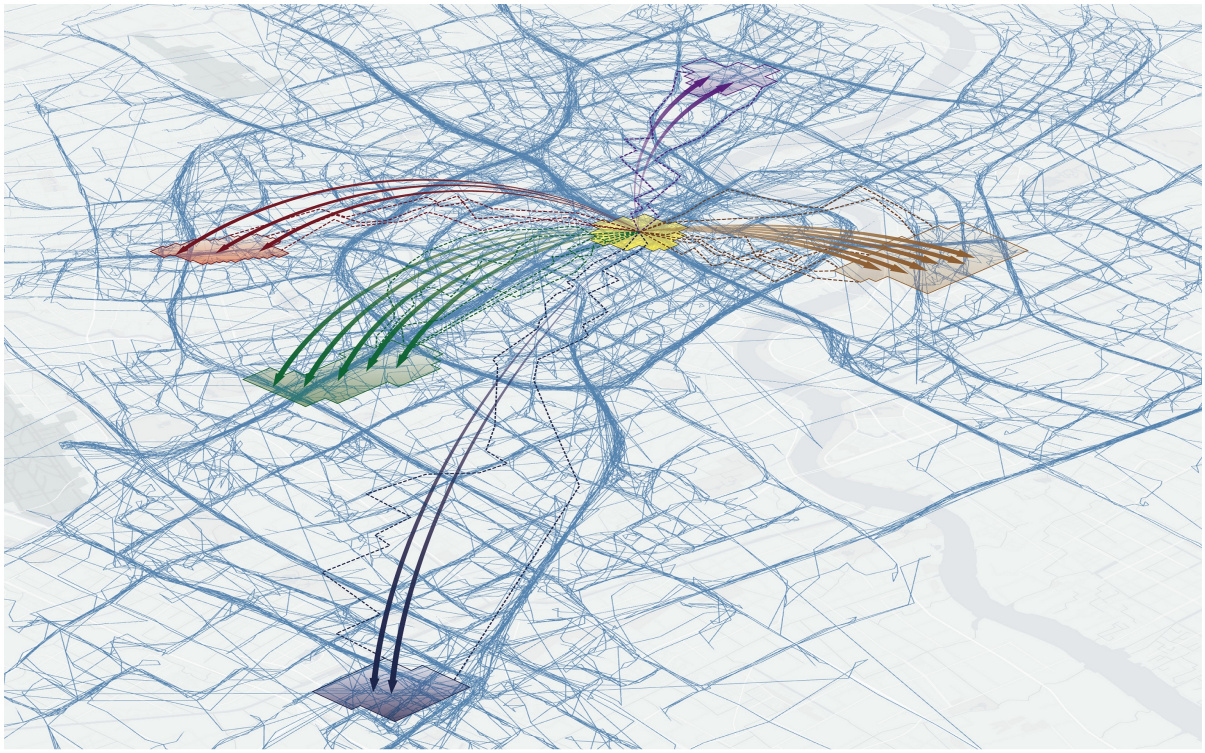

Fig. 1 Illustration of complex urban mobility from both the individual and population perspectives. The bottom layer represents individual movement trajectories between urban locations, and the top layer denotes population flows between urban regions, where more (fewer) arrow lines indicate larger (smaller) flows. Among trajectories, several samples belonging to each flow are highlighted with the same color. Regions are shown as geographical polygons.

from the population, which means $\Pi(\mathbf{l}_t|\mathbf{x}_{<t}) \neq \prod_{n=1}^{N} \pi(l_{n,t}|x_{n,<t})$ and prohibits the traditional independent modeling. Second, the individual mobility model π is simultaneously constrained by population-level mobility information, requiring a top-down learning process.

The modeling framework of DeepMobility is presented in Fig. 2. It consists of three components, i.e., a generator, a discriminator and a critic. The generator aims to generate individual trajectories with the resulting flows that are indistinguishable from empirical data by the discriminator. Fig. 2a illustrates the generation process at the generator, with a GRU-based state encoder that transforms location visit history into a fixed-length hidden vector and a hierarchical decoder that simulates the mobility decision process to first decide the next visit region and then choose a specific location in this region. To achieve collaborative learning from individual mobility to collective mobility and model the multiscale patterns and complexity, we design a bottom-up social interaction modeling mechanism at the trajectory decoder. Specifically, we use two modules to generate the next visit region based on individual preference and social interaction, respectively, and the final decision is made between these two according to a learnable probability score that measures an individual's uncertainty about following his/her preference. The preference-based module uses an MLP that takes state embedding as input and outputs a vector indicating the visitation probability of each region, while the interaction-based module also adopts an MLP-based structure[37] that takes regional attributes as input and predicts the visitation probability in terms of population movements. Fig. 2c illustrates the multilevel structure of the discriminator that evaluates the realism of generated trajectories and the aggregated flows. The individual-level discriminator also uses a GRU-based module as it needs to process a trajectory sequence and output the score indicating whether it is similar to actual data, while the population-level discriminator directly computes the relative error between the generated flow value and the ground truth. These feedbacks are sent back to improve the generator through another critic network using a well-established proximal policy optimization algorithm (PPO)[40], as illustrated in Fig. 2d-f. To achieve collaborative learning from collective to individual mobility data, we design a top-down feedback refinement mechanism at the critic. Specifically, the critic adopts a multilevel structure to approximate value functions for the generator output. Besides the individual-level critic, another population-level critic leverages a value decomposition technique[41] that transforms the overall assessment of population movements into individual-level feedback, which directly refines mobility behavior in a top-down manner (Method and Supplementary Section S1).

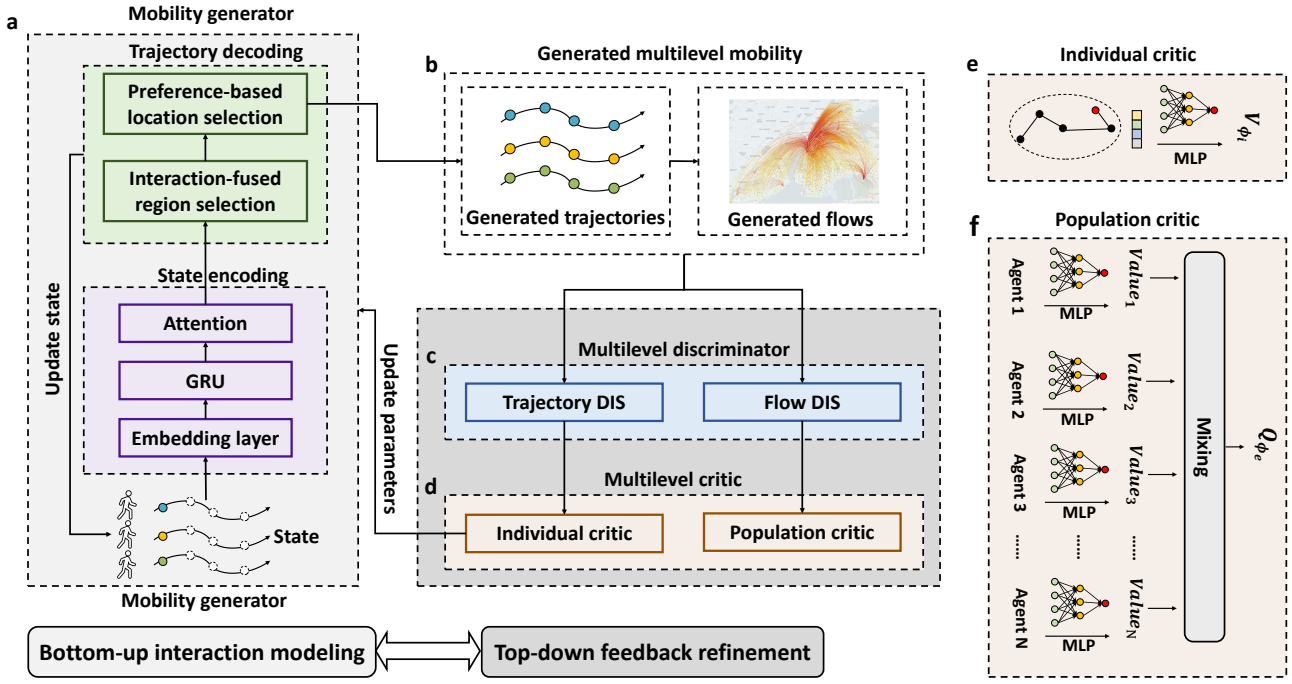

Fig. 2 Overview of the proposed deep generative collaboration network DeepMobility. It consists of three components to learn the complexity of urban mobility. The first component is a mobility generator, as shown in the left panel. This generator is implemented using a GRU-based state encoder and a trajectory decoder that produces the next visited location by employing a bottom-up social interaction modeling technique. Then a multilevel discriminator evaluates the utility of the generated movements from both individual and population perspectives. This feedback is used to train a multilevel critic (detailed in the right panel) that decomposes the overall guidance from population-level mobility and directly guides the optimization of the generator at the individual level (Top-down feedback refinement).

DeepMobility generates human mobility trajectories and the resulting flows at the urban scale

To assess the capability of the proposed DeepMobility modeling framework, we perform an experiment that utilizes it to generate synthetic data and evaluate whether they represent intricate mobility patterns at both individual and population levels (Experiment details are provided in Supplementary Section S2). First, for individual mobility patterns, we verify if the generated trajectories are statistically similar to the real data by quantify the distribution differences using the Jensen-Shannon divergence (JSD) and the Kolmogorov–Smirnov (KS) test[33,42], which are bound by $[0, 1]$, with 0 indicating a perfect match between two distributions. In particular, we focus on the following five fundamental metrics[13,30,43]: jump length Δr (distance of each travel), weekly trip distance r_w, radius of gyration r_g, waiting time Δt (time spent at the same location), and daily visited locations S_d of an individual. These metrics comprehensively cover empirically observed mobility patterns including spatial and temporal regularity (Δr, Δt)[17], population heterogeneity (r_g)[17] and ultraslow growth of travel distance and visited location number (r_w, S_d)[20]. Then, for population mobility patterns, we use the Common Part of Commuters (CPC)[27,37] and the Mean Absolute Error (MAE)[30,44] to calculate the distance between the real and the generated flows. The mobility flows are calculated as how many people move from one region to another within a period, and CPC is bounded by $[0, 1]$ with 1 indicating the two are identical and 0 suggesting no overlap. For performance comparison, we choose the four most widely used mechanistic approaches and deep-learning approaches, including the CTRW model[17], the EPR model[20], the TimeGeo model[42] and the GAN model[33,45].

To verify the modeling capability of DeepMobility at the individual mobility level, in Fig. 3a-d we calculate the JS divergence of five metrics (Δr, r_w, r_g, Δt and S_d) for generated trajectories of each model, finding that our model generates individual trajectories with the highest statistical similarity (See complete numerical results in Supplementary Table 3). Further results of the KS test (Supplementary Table 4) also verify such generation realism with 15 (out of 20) KS statistics of distributions no larger than 0.2. Unlike mechanistic models, DeepMobility leverages the power of deep learning to learn individual features from the mobility data. Thus it can capture diverse travel behaviors observed in the population, especially the $P(r_g)$ corresponding to the individuals' characteristic distance that has been empirically observed to have a high population heterogeneity. Unlike the one-shot generation of a complete trajectory as in the GAN model, DeepMobility simulates the mobility decision process of an individual and generates the next visit location in a sequential way, achieving significant realism improvement.

To test the capability of DeepMobility in generating complex mobility flows at the population level, in Fig. 3e-h we measure

the CPC between the generated flows and the real flows for each model, observing a remarkable performance improvement up to 120% (Beijing, 0.787 vs. 0.357), 112% (Shanghai, 0.686 vs. 0.324), 136% (Shenzhen, 0.560 vs. 0.216) and 81% (Senegal, 0.650 vs. 0.359). Our DeepMobility also achieves much less absolute error (MAE), equivalent to 17-30% of the best-performing baseline model in four different cities (Supplementary Table 3). To compare further the generated flows with the empirical data, we measure the number of travels between each pair of locations (Fig. 3i-l). Unlike our DeepMobility, both the mechanistic model TimeGeo and the deep-learning model GAN generate unrealistic flows that deviate from the empirical data. In particular, the GAN model tends to generate overestimated flows for those less-traveled pairs (number of travels < 1,000), and the TimeGeo model performs the opposite in the cases of Shanghai and Shenzhen. Correspondingly, in Supplementary Fig. 5-6 we plot the mobility networks describing the observed flows and the flows generated by three models (DeepMobility, GAN and TimeGeo), finding that DeepMobility captures the overall structure of the flow network while GAN and TimeGeo generate much denser connections and sparser connections, respectively.

To understand the origin of the aforementioned exceptional ability to capture intricate complexity within empirical urban mobility data, we remove two collaborative learning mechanisms in DeepMobility and retrain three model variants to evaluate their generation realism (Table 1). We find that, compared to a vanilla version without the designed collaborative learning mechanisms, DeepMobility achieves a significant improvement of the CPC up to 60.6% (Beijing), 29.2% (Shanghai), 20.7% (Shenzhen) and 8.5% (Senegal). It confirms the necessity of designing such a deep generative collaboration network to resolve the current discrepancy between individual mobility modeling and population mobility modeling. Furthermore, bottom-up interaction modeling and top-down feedback refinement are incorporated in different modules, allowing us to remove either of them to compare performance, finding that incorporating population influences into the decision process of individual-level movements contributes more than refining the learning of individual mobility behavior. Note that neither of the two designed mechanisms hurts generation realism at the individual mobility level, showing their strong compatibility.

By generating realistic individual trajectories and the resulting population flows at the urban scale, DeepMobility successfully preserves the organic nature of the urban population, as individuals' daily activities and lives are closely related to their mobility. To demonstrate its capability of reconstructing individuals' activities at various urban locations, we apply a location type inference method[42] to identify a collection of home locations from the generated trajectories, finding that their spatial distribution is in good agreement with the empirical data (Supplementary Fig. 7 and Section S3.2). Another important empirical observation of urban life through a mobility lens is the distribution of the most frequent daily mobility networks, i.e., daily motifs[46]. As we show in Supplementary Fig. 8, the distribution of the identified nine distinct motifs is again consistent with the empirical data, with JS divergence statistics less than 0.084 in four cities (Supplementary Section S3.2).

Reproduction of scaling laws governing human mobility

In order to validate whether DeepMobility is capable of reproducing scaling laws that have been empirically observed, we subsequently investigate whether the generated trajectories and flows exhibit corresponding statistical patterns. As a first step, in Fig. 4**a**-**d** we calculate the jump length Δr at hourly intervals in the DeepMobility-generated data, finding that $p(\Delta r) \sim (\Delta r + \Delta r_0)^{-\beta_1} \exp(-\Delta r/\kappa_1)$, the same form of truncated power law as the previous finding[17]. The exponent β_1 ($\sim 1.2 - 1.3$) remains similar in different cities. However, we observe that the scaling exponent β_1, though in excellent agreement with the empirical value β_1 ($\sim 1.1 - 1.3$), is relatively smaller than that in the literature (~ 1.75), suggesting a possible discrepancy between their underlying mechanisms. To further understand the origin of this discrepancy, in Fig. 4**e-h** we measure the degree of spatial regularity in the generated population by calculating the radius of gyration r_g for all individual trajectories, finding that their distribution $P(r_g)$ is again consistent with the previous finding[17,20], well-approximated with a truncated power-law. Indeed, the measured value of the scaling exponent $\beta_2 \sim (1.10, 1.17)$ is also smaller (vs. 1.65[17]), suggesting that this difference on population heterogeneity (r_g) may account for above observed discrepancy (Δr). Most important, however, is the fact that an increased heterogeneity of travel patterns indicated by smaller β_2 can be a major challenge for not only mechanistic models but also deep-learning models. As we show in these figures, the GAN model generates trajectories directly and fails to reproduce empirically observed scaling laws, confirming the importance of learning to simulate individual travel decisions as in DeepMobility.

Another key scaling property regarding individual mobility is the frequency f of the kth most visited location follows Zipf's law[17] $f_k \sim k^{-\zeta}$ ($\zeta \sim 1.2$), arising from the memory effect that individuals tend to return to previously visited locations preferentially. By encoding individuals' states using their visitation history $x_{<t}$, DeepMobility can capture the aforementioned preferential return mechanism[20] in the learned policy $\pi_\theta(l_t|x_{<t})$. In Fig. 4**i-l** we measure f_k for mobility data generated by DeepMobility, finding that the observed scaling behavior is in excellent agreement with the previous finding, with scaling exponent $\zeta \sim (1.0, 1.1)$ not far from that in the literature. To further understand how DeepMobility accounts for this visitation pattern of humans, we remove the memory effect by not encoding the visitation history, finding that the newly generated data exhibits a much more even distribution of f_k as in scale-free random walk models[17] (a smaller ζ, in Supplementary Fig. 9). On the other hand, we show that DeepMobility-generated individuals also have a tendency to decrease the exploration of other

locations over time[20], as manifested by a sublinear growth in the number of distinct locations (Supplementary Fig. 10a-d and Section S3.3), and their long-term spatial movements up to one year follow an ultraslow growth (< logarithmic)[17] with time evolution (Supplementary Fig. 10e-h and Section S3.3). The above reproduced mobility patterns associated with the memory effect contribute to the high predictability of our generated mobility trajectories (Supplementary Fig. 11 and Section S3.3), which is consistent with earlier findings[19].

Finally, to test whether DeepMobility can reproduce the population-level scaling patterns, in Fig. 4**d**-**p** we show the 'spectral' flow $\rho_i(r, f)$, i.e., the number of visitors with a given visitation frequency f to each location i from a given distance r (to their home location), finding that $\rho_i(r, f)$ decreases as the inverse square of the product of f and r. This power-law scaling relation reproduced by DeepMobility is in excellent agreement with the discovered distance–frequency scaling[27], largely owing to our designed collaborative learning mechanisms that bridge the heterogeneous movements of individuals and collective patterns emerging from the entire population (See deviated results of GAN in Fig. 4**m**-**p**, as well as corresponding results of DeepMobility removing collaborative learning in Supplementary Fig. 12). Furthermore, we find that DeepMobility successfully reproduces other empirically observed scaling patterns, including power law distributions in the number of trips between regions[25] and similar statistical properties of origin-destination demand networks across different cities[26] (Supplementary Fig. 13-14 and Section S3.3). The remarkable consistency between mobility patterns reproduced by DeepMobility and fundamental laws established by physicists[17,19,20,25–27] suggests that DeepMobility successfully captures the intricate interplay between individual heterogeneous movements and collective behaviors in a manner unmatched by previous models.

DeepMobility captures the inherent mechanism of complex urban mobility that is geographically transferable

For generative models of urban mobility, the capability for geographic transferability reflects their ability to generalize in capturing consistent mobility patterns across different cities. This is also essential for DeepMobility, as it means the model captures the inherent mechanisms of complex urban mobility, rather than mere data memorization. Moreover, due to the increasing cost of data acquisition, high-quality human mobility data is often scarce or even absent in some underdeveloped urban regions. In real-world applications of urban mobility generation, practitioners may have to develop generative models using available mobility data collected from some cities and then use these models to generate urban trajectories in other targeted cities without any mobility data.

To accommodate transferable mobility generation between different cities, we improve the model design of DeepMobility (Method, Supplementary Section S1.5), train this improved model on mobility data of one Chinese city and test its generation realism in two other Chinese cities. To verify the geographic transferability of DeepMobility, in Fig. 5 **a-d** we calculate evaluation metrics of generation realism at both individual and population levels and compare with three mechanistic models (Complete results are shown in Supplementary Table 5). Note that, unlike deep-learning models such as GANs, mechanistic models are intrinsically transferable. We find that DeepMobility-generated trajectories have the highest statistical similarity to real data in terms of five individual-mobility metrics, i.e., $(\Delta r,\ r_w,\ r_g,\ \Delta t,\ S_d)$, and the resulting flows reconstruct realistic population-level mobility with the highest accuracy (CPC and MAE). In particular, DeepMobility trained in one city reproduces the spatio-temporal mobility patterns ($P(\Delta r)$, $P(\Delta t)$) in another city, with JS divergence less than 0.1, and finally captures the population heterogeneity ($P(r_g)$). The improvement in flow generation is significant, over 60% (CPC) in 3×2 source-target pairs. In Fig. 5 **e-g** we show complete results of transferability evaluation for $P(\Delta r)$, $P(\Delta t)$ and flow similarity, finding that the transferred DeepMobility is on par with its counterpart trained on the target city. Individual-level trajectories are equally realistic (JS divergence < 0.1) and population-level flows yield CPCs that are remarkably close to non-transferred ones (56-95%). These results indicate the potential capability of extending DeepMobility to generate realistic mobility data with high utility for any given cities around the world.

Discussion

Recent advancements in generative AI technologies have markedly enhanced content generation capabilities, spanning text, images, and videos. However, generating human behavior, in contrast to these forms of content, presents a more formidable challenge due to the complexity of intricate linkages between individual actions and collective population dynamics. Our research, focusing on human mobility behavior as an initial endeavor, demonstrates that a novel generative deep-learning approach, enriched with effective collaborative learning mechanisms, can successfully bridge this gap and enable the generation of complex urban mobility data across various cities. This development further implies, with minimal urban context information on demographics and geography, how generative AI helps generate a vibrant, 'organic' urban population with intricate dynamics by modeling the way individuals interact, engage, and utilize services in the course of their daily movements.

This new deep-learning approach is designed to generate both individual movements and the resulting population flows in a city, making it possible to answer a long-standing question of whether deep-learning models can capture underlying mechanisms driving complex urban mobility across both individual and population levels. Our results of reproducing mobility scaling laws give a "yes" answer to this question (Fig. 4), confirming the importance of achieving collaborative learning between individual

and population mobility. In this regard, machine intelligence can further augment understanding and learning of complex mechanisms behind individuals' mobility decisions[47], and finally enable the data-driven discovery of new theories regarding human mobility[14,48,49].

An interesting question here is how trustworthy the designed deep model is in generating realistic synthetic data without compromising privacy since empirical studies have found that individuals can be identified from mobility data due to the uniqueness of their trajectories[50,51]. To that end, we examine the overlap ratio between real and generated trajectories, and also the identifiability of a real individual from generated trajectories (Supplementary Section 3.5). We find that, for most individuals in real data, their trajectories share only a small portion with the most similar generated trajectories, and it is mathematically infeasible to identify a real individual when mixed with generated ones (Supplementary Fig. 15). Therefore, the proposed deep generative model does not simply "memorize" the real mobility trajectories, but instead learns the underlying mechanisms driving human mobility patterns.

From a practical standpoint, our developed DeepMobility framework has demonstrated the potential to generate realistic and complex urban mobility data. This is particularly significant in cities lacking available mobility data, providing immediate and valuable applications in the fields of epidemic disease containment, traffic engineering, and urban planning[52–54]. Looking ahead, the framework can be enhanced by integrating the impact of urban road networks on human mobility, thus offering a more comprehensive modeling approach[55]. In terms of future applications, DeepMobility has the potential to evolve into a transparent tool for constructing open urban mobility data that could offer detailed insights into population movements within cities globally. This advancement would serve as a valuable complement to the current static and coarse-grained mapping of world populations[56,57]. By providing a more nuanced understanding of urban mobility patterns, DeepMobility will be instrumental in supporting the development of sustainable and livable cities worldwide[58–60].

Methods

M1. Datasets

We use four datasets to demonstrate the DeepMobility generation framework. The first two datasets (DS1, DS2) use the anonymized location record of about 1.8 million users in Beijing, China and 0.32 million users in Shenzhen, China, respectively. These users have signed up for a location-based service and their locations are recorded every hour for a one-month period. The third dataset (DS3) covering Shanghai (China) consists of around 1.9 million anonymized users of China's major telecom company. The data are collected during a one-week period for billing purposes, recording the location at the beginning and the end of each service (a call, an SMS, or an Internet connection). The fourth dataset (DS4) of Senegal is based on anonymized call detail records (CDRs) from about 0.3 million users during a two-week period with a temporal resolution of 10 min. These datasets capture daily human movements at both individual and population levels, i.e., trajectories and the resulting flows. The details of data processing and feature extraction are provided in Supplementary Section S1.1.

M2. DeepMobility

GAIL based framework

To solve the generative learning problem of urban mobility in Equation (1), we resort to generative adversarial imitation learning (GAIL)[61] due to the analogy between mobility modeling and decision policy learning. Specifically, we define the set of locations L as the action space A, and the set of visitation history $X_{<t}$ as the state space S. Then learning $\pi(l_t|x_{<t})$ equals to finding optimal $\pi(s,a)$ in GAIL. Moreover, to describe the complex decision process of a group of individuals, we further extend MDP into Decentralized Partially Observable MDP (Dec-POMDP)[62] that can be represented by a 7-tuple $< S, \{A_n\}, P, R, \{\Omega_n\}, O, \gamma >$:

(1) S is a set of states and each state $\mathbf{s}$ consists of all agents. In the partially observable setting, agents have no access to the overall state.

(2) A_n is a set of actions for agent n, and $A = \times_n A_n$ is the set of joint actions. Specifically, an action $a_{n,t}$ indicates the next place to visit for individual n at time t.

(3) P is a set of conditional transition probabilities between states, with $P(\mathbf{s}'|\mathbf{s},\mathbf{a})$ denoting transition probability from $\mathbf{s}$ to $\mathbf{s}'$ given a joint action $\mathbf{a}$. The transition is deterministic in this problem.

(4) $R : S \times A \mapsto \mathbb{R}$ denotes the reward function $r(\mathbf{s},\mathbf{a})$.

(5) Ω_n is a set of observations for agent n, and $\Omega = \times_n \Omega_n$ is the set of joint observations.

(6) O is a set of conditional observation probabilities, i.e., $\{O(\mathbf{o},\mathbf{s},\mathbf{a})\}$. The observation is also deterministic in this problem. Specifically, for an individual n at time t and location l_{t-1}, his/her observation $o_{n,t}$ combines both historical movements

$X_{<t} = [x_1, x_2, ..., x_{t-1}]$ and the distribution of population movements from l_{t-1} to other locations, denoted as $\tilde{\mathbf{F}}_{l_{t-1}}$. The latter represents a limited ability of individuals to observe other people's travel decisions.

(7) $\gamma \in [0,1]$ is the discount factor.

Based on the above formulation, we propose a GAIL-based approach for learning $\pi(l_t|x_{<t})$ that aims to capture human mobility patterns at both individual and population levels. The preliminary on GAIL is detailed in Supplementary Section S1.2.

Collaborative learning mechanisms

In order to establish a bidirectional link between individual and population levels of mobility modeling in DeepMobility, we design the bottom-up and top-down collaborative learning processes, respectively.

(1) Bottom-up interaction modeling. We begin by developing a generator module capable of creating interconnected movements from a big urban population. Aside from personal taste, an individual's travel decisions are heavily influenced by social interactions[63,64]. For example, trajectories traveled by a person show his/her daily pattern, implying a memory effect in which historical movements influence future mobility behavior. Meanwhile, an individual may visit some unexpected areas advised by his or her friends on occasion, suggesting social interactions among a population of individuals. To directly characterize these interactions for N individuals, substantial pairwise linkages ($\sim N^2$) would have to be computed, which is not feasible. To address this issue, we propose the formulation of the mobility model $\pi_\theta(a_t|o_t)$ as a composite of two parameterized decision processes:

$$\pi(a_t|o_t) = \texttt{Collab}\big(\pi_{\theta_I}(a_t|X_{<t}), \pi_{\theta_P}(a_t|\tilde{\mathbf{F}}_{l_{t-1}})\big), \tag{2}$$

where $\pi_{\theta_I}(a_t|X_{<t})$ and $\pi_{\theta_P}(a_t|\tilde{\mathbf{F}}_{l_{t-1}})$ represent two distinct decision policies considering individual preference and population influence, respectively. The first part $\pi_{\theta_I}(a_t|X_{<t})$ captures individual preference by learning movement regularities from the historical trajectory $X_{<t}$, while the second part $\pi_{\theta_P}(a_t|\tilde{\mathbf{F}}_{l_{t-1}})$ characterizes the social interaction influence from population movements $\tilde{\mathbf{F}}_{l_{t-1}}$ that are shaped by the urban environment. The collaboration between these two parts is designed as making a discrete choice between them according to a parameterized Bernoulli distribution $Bernoulli(u_t)$, where $u_t = U_{\theta_u}(X_{<t}) \in [0,1]$ characterizes a learned probability that measures individual's uncertainty on following his/her preference. Formally,

$$\pi(a_t|o_t) = \begin{cases} \pi_{\theta_I}(a_t|X_{<t}), \ \texttt{Prob} = 1 - u_t \\ \pi_{\theta_P}(a_t|\tilde{\mathbf{F}}_{l_{t-1}}), \ \texttt{Prob} = u_t \end{cases}. \tag{3}$$

If the individual uncertainty on historical visitation is high, this individual is more likely to follow a population-level decision (π_{θ_P}) instead of his/her preference (π_{θ_I}). The designed collaboration between two decision processes successfully incorporates population influences at the individual level, achieving bottom-up social interaction modeling of complex urban mobility.

(2) Top-down feedback refinement. Next, we design a critic module that improves the generator in terms of capturing bidirectional influence between individual level and population level of urban mobility. The critic V_ϕ approximates the value function of mobility decisions at the generator π_θ and guides the optimization direction of π_θ accordingly using policy learning algorithms like PPO[40] (Detailed in Supplementary S1.2). As π_θ generates individual movements that are then aggregated into population flows, V_ϕ should evaluate the value of each movement based on not only its fitness to individual-level patterns but also its contribution to aggregated flows. The former is completed by an individual critic $V_{\phi_i}(o_t)$ that learns to predict the correct value function of individual movements based on whether $\pi_\theta(a_t|o_t)$ matches patterns in empirical data. However, the latter is far more challenging due to the high-dimensional space of joint actions $\mathbf{a}_t \in \times_n A_n$ and observations $\mathbf{o}_t \in \times_n \Omega_n$, arising from the entire population. Therefore, to solve this issue, we propose to decouple the joint space optimization by decomposing the population critic $Q_{\phi_p}(\mathbf{o}_t, \mathbf{a}_t)$ from the entire population ($\sim$N) into a sum of $Q_{\phi_p}(o_{n,t}, a_{n,t})$ from each individual. Formally,

$$Q_{\phi_p}(\mathbf{o}_t, \mathbf{a}_t) = \sum_{n=1}^{N} Q_{\phi_p}(o_{n,t}, a_{n,t}). \tag{4}$$

Then we optimize Q_{ϕ_p} using Monte-Carlo policy evaluation as follows,

$$\begin{aligned} &\min_{\phi_e} \mathbb{E}_\Pi \Big[\big(\hat{Q}_\Pi(\mathbf{o}, \mathbf{a}) - Q_{\phi_p}(\mathbf{o}, \mathbf{a})\big)^2 \Big], \\ &\text{where } \hat{Q}_\Pi(\mathbf{o}, \mathbf{a}) = \hat{\mathbb{E}}_\tau \Big[\sum_{k=1}^{T-t} \gamma^{k-1} \mathbf{r}_{t+k} \mid \mathbf{o}_t = \mathbf{o}, \mathbf{a}_t = \mathbf{a} \Big], \\ &\text{and } \mathbf{r}_{t+k} = \mathbf{r}(\mathbf{o}_{t+k}, \mathbf{a}_{t+k}) = \frac{\mathbf{F}_{t+k,\text{data}} - \mathbf{F}_{t+k,\text{model}}}{\mathbf{F}_{t+k,\text{data}}}, \ \mathbf{a}_{t+k} \sim \Pi_\theta(\cdot|\mathbf{o}_{t+k}). \end{aligned} \tag{5}$$

The global reward $\mathbf{r}_t$ measures the relative error between the real flow matrix $\mathbf{F}_{t,\text{data}}$ and $\mathbf{F}_{t,\text{model}}$ generated by following joint policy $\Pi(\mathbf{a}_t|\mathbf{o}_t)$. Note that $\frac{\partial Q_{\phi_p}(\mathbf{o}_t,\mathbf{a}_t)}{\partial Q_{\phi_p}(o_{n,t},a_{n,t})} \geq 0, \forall i$, thus the optimization of Q_{ϕ_p} yields the same optimization direction on each $Q_{\phi_p,n}$. In this way, we manage to make the optimization of the population-level critic $Q_{\phi_e}(\mathbf{o}_t,\mathbf{a}_t)$ feasible by turning global optimization into a series of coordinated optimization processes at the individual level, i.e., with respect to $Q_{\phi_p}(o_{n,t},a_{n,t})$. Finally, we combine value estimations of $\pi_\theta(a_t|o_t)$ from two critics, i.e., $V_{\phi_i}(o_t)$ and $Q_{\phi_p}(o_t,a_t)$, and update the generator parameters using PPO (Detailed in Supplementary Section S1.3). The above design of value decomposition from population level to individual level achieves the top-down refinement of $\pi_\theta(a_t|o_t)$ based on feedback from the quality of generated flows.

Generator architecture

Supplementary Fig. 1 shows the network architecture of the mobility generator.

(1) State encoder. We first utilize a Gated Recurrent Unit (GRU) to learn the state representation of the historical trajectory. Specifically, we transform previously visited locations represented by one-hot encoding, along with temporal information, into vectors $\mathbf{e}_t$ by an embedding layer; then we learn a representation of the individual's historical trajectory h_t using a GRU network, which captures non-Markov, memory effects of individual mobility.

(2) Hierarchical decoder. Next, we adopt a hierarchical structure at the trajectory decoder that first outputs the next region r_j to visit and then selects a specific location l_j that belongs to r_j.

In the first stage, we design an interaction-fused region selection process that characterizes the influences from both individual preference and social interaction. As in Equation (2), we select r_j according to an uncertainty-based probability score u_t, which is obtained based on the following multi-head uncertainty estimation module:

$$
\begin{aligned}
&u_t = \texttt{Sigmoid}(\texttt{var}(\mathbf{v}_m)),\\
&\text{where } \mathbf{v}_m = \texttt{MLP}(h_t) = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_H \end{pmatrix}
\end{aligned}
\tag{6}
$$

is a H-dimensional vector whose variance across different heads $\{1,2,...,H\}$ is used to estimate the state uncertainty. Note that we use the Sigmoid function to transform it into a probability score between 0 and 1. If the estimated uncertainty u_t is high, the individual will be more likely to consider following collective behaviors. Otherwise, the individual prefers to follow her own preference based on historical memory.

As for the specific network architecture that simulates the region selection process based on individual preference, we utilize an MLP to transform h_t into an n-dimensional vector and normalize it by a softmax function as follows:

$$
\pi_{\theta_l}(r_j|X_{<t}) = \frac{\exp(\text{MLP}(h_t))_j}{\sum\exp(\text{MLP}(h_t))}\ ,\ j \in \{1,2,...,N_r\}\ , \tag{7}
$$

where N_r is the number of regions, and $\pi_{\theta_l}(r_j|X_{<t})$ denotes the probability of visiting region r_j.

As for simulating the region selection process based on collective behaviors, we consider the spatial movement distribution of populations originating from the current region to other regions. We utilize the follow network architecture to learn a N_r-dimensional probability $\pi_{\theta_p}(r_j|\tilde{\mathbf{F}}_{l_{t-1}}), j \in \{1,2,...,N_r\}$ as in Equation (2). The vector is calculated in the following steps. First, an embedding layer transforms region attributes, including the population $x_{r,pop}$ and POI distribution $x_{r,poi}$, into region representation e_r, then the input vector e_{ij} is obtained by concatenating the representation of the origin region e_{r_i}, the representation of the destination region e_{r_j}, the distance representation $e(d_{ij})$ between two regions, and the time representation $e(t)$. Second, the input vectors e_{ij} are all fed into the same network, which is an MLP with ten 64-dimensional hidden layers and has the LeakyReLu as the activation function. The last layer outputs the probability to observe a trip from the origin region r_i to destination regions $r_j, j \in \{1,2,...,N_r\}$. The above process is formulated as follows:

$$
\begin{aligned}
&e_r = \texttt{Concat}(W_{pop}\texttt{one-hot}(x_{r,pop}) + W_{poi}\texttt{one-hot}(x_{r,poi}))\ ,\\
&e_{ij} = \texttt{Concat}(e_{r_i}, e_{r_j}, e(d_{ij}), e(t))\ ,\\
&\pi_{\theta_p}(r_j|\tilde{\mathbf{F}}_{l_{t-1}}) = \frac{\exp(\text{MLP}(e_{ij}))_j}{\sum\exp(\text{MLP}(e_{ij}))}\ ,\ j \in \{1,2,...,N_r\}\ .
\end{aligned}
\tag{8}
$$

In the second stage, based on selected region r_j, we design another location-selection process that chooses the next location belonging to r_j. To cope with a varied number of locations across different regions, we utilize an attention-based network to

calculate the probability of visiting a specific location l_j within r_j as follows:

$$\begin{aligned}
&\pi(l_j) = \frac{\exp(u_j^T u_w)}{\sum_k \exp(u_k^T u_w)} \ , j \in \{1,2,...,N_{l,r}\} \ , \\
&\text{where } u_j = \texttt{tanh}(\texttt{MLP}(e_j)) \\
&\text{and } e_j = \texttt{Concat}(h_t, e_{l,j}, e(d_{ij}), e(t)) \ .
\end{aligned} \tag{9}$$

N_{l,r_j} is the number of locations in region r_j, h_t is the history representation obtained from the state encoder, $e_{l,j} = \texttt{one-hot}(l_j)$ is the embedding of the location, $e(d_{ij})$ is the representations of distance between current location l_i and target location l_j, and $e(t)$ is the representation of the current time. u_w queries the location characteristics associated with the current state, which is a randomly initialized vector and is updated in the training procedure.

Discriminator architecture

We define multilevel reward functions that give a comprehensive evaluation of the generation outcome with respect to individual and population levels, respectively. We characterize the two reward functions separately, where the first individual-level reward function is characterized by a neural network-based discriminator D_ϕ with parameters ϕ, and the second population-level reward function is directly available. Supplementary Fig. 2 illustrates the discriminator architecture. Specifically, the multilevel rewards are calculated as follows:

$$r_{I,t} = \log(D_\phi(o_t, a_t)) \ , \tag{10}$$

$$r_{P,t} = \frac{\mathbf{F}_{t,\text{data}} - \mathbf{F}_{t,\text{model}}}{\mathbf{F}_{t,\text{data}}} \ , \tag{11}$$

where o_t, a_t denotes the state and action at time t, $\mathbf{F} \in \mathbb{R}^{N_t \times N_r \times N_r}$ is the flow matrix, N_t is the number of time periods, $r_{I,t}$ denotes the individual-level reward, and $r_{P,t}$ denotes the population-level reward. D_ϕ is trained by a binary classification task that distinguishes between real and generated state-action pairs. We adopt a non-parametric method for $r_{P,t}$, which calculates the relative distance between the generated flows and real-world cases.

As for the network architecture of D_ϕ shown in Supplementary Fig. 2, we also design a hierarchical structure that evaluates the decisions at two stages, i.e., region and location, respectively. At the region level, the state o_t is the historical sequence of visited regions, and the action a_t is the selected region based on the state. Correspondingly, the state and action at the location level are the historical location sequence and the selected locations, respectively. For both levels, the network consists of two components: 1) an embedding layer to transform the historical sequence into vector representations, 2) a GRU to obtain the sequence representation, 3) an output layer with a Sigmoid activation function to produce the classification result based on the sequence representation. The discriminator's output denotes the probability that the state-action pair comes from the real data. The non-parametric discriminator calculates the distance between real and generated population flows.

Critic architecture

Supplementary Fig. 3 shows the network architecture of the individual-level critic and population-level critic. The individual-level critic shares the same state encoder with the generator, including the Embedding layer and GRU layer, to obtain state representations. Then we adopt an MLP to predict the individual state value $V_{\phi_i}(o_t)$. The population-level critic is modeled by another MLP network, which takes in the concatenation of the state embedding and action embedding. Specifically, the state embedding is obtained by a similar network as the individual-level critic, and the action embedding is obtained by an Embedding layer. The joint value function $Q_{\phi_p}(\mathbf{o}_t, \mathbf{a}_t)$ is obtained by aggregating the value $Q_{\phi_p}(o_{n,t}, a_{n,t})$ from each individual n. Mathematically, the above process is formulated as follows:

$$\begin{aligned}
&V_{\phi_i}(o_t) = \texttt{MLP}(\texttt{GRU}(\texttt{Emb}(o_t))) \ , \\
&Q_{\phi_p}(o_{n,t}, a_{n,t}) = \texttt{MLP}(\texttt{Concat}(\texttt{GRU}(\texttt{Emb}(o_{n,t})), \texttt{Emb}(a_{n,t}))) \ , \\
&Q_{\phi_p}(\mathbf{o}_t, \mathbf{a}_t) = \sum_{n=1}^{N} Q_{\phi_p}(o_{n,t}, a_{n,t}) \ ,
\end{aligned} \tag{12}$$

where $o_{n,t}$ and ,$a_{n,t}$ are the state and action of the n_{th} individual, $\mathbf{o}_t$ and $\mathbf{a}_t$ are the joint state and action, and $\texttt{Emb}$ denotes Embedding layers.

Transferable mobility generation

To accommodate transferable mobility generation between different cities, we refine the design of DeepMobility as follows. First, we enhance the transferability of the location representation used in the generator. The widely used embedding

technique[30,33,65] is no longer applicable due to a lack of transferability. Instead, we encode all locations in different cities using two characteristics, visitation popularity and POI number grouped by category. These two features reflect important properties about location attractiveness and land use profile, respectively, and should have a general impact on individuals' travel decisions, independent of the city in which they are located. Moreover, they are both location-based aggregation metrics that can be readily collected at a low cost from location data providers (such as Safegraph) and crowd-sourcing platforms (such as OpenStreetMap). The detailed design of transferable location representation is presented in Supplementary Fig. 4 and Section S1.5. Second, to further guarantee the generalization capability of DeepMobility, we remove the design of top-down feedback refinement, as it provides accurate but prone-to-overfitting supervision on the generated flows during training. Then we train this improved DeepMobility on mobility data of one Chinese city and test its generation realism in the other two Chinese cities. Note that we choose a generator with good transferability rather than one that produces statistically similar data.

Data availability

The data supporting the results of this study is available on GitHub (https://github.com/tsinghua-fib-lab/DeepMobility).

Code availability

The code used in this research is released on Github (https://github.com/tsinghua-fib-lab/DeepMobility).

Author contributions

Y.Y., J.D., D.J. and Y.L. jointly launched this research, and Y.Y., J.D., and Y.L. contributed ideas. Y.Y., J.D., and Y.L. designed the research methods and provided the research outline. Y.Y. developed the DeepMobility framework and performed the experiments. J.D., D.J., and Y.L. provided critical revisions. D.J. and Y.L. managed the project. All authors jointly analyzed the results and participated in the writing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this manuscript.

	Individual-level metrics					Population-level metrics	
Model	dailyloc	jump length	trip distance	duration	radius	MAE	CPC
Beijing							
w/o M1 and M2	0.011	0.066	0.023	**0.013**	0.044	62.54	0.490
w/o M1	**0.008**	0.067	0.021	0.018	0.058	44.34	0.587
w/o M2	0.010	0.063	**0.019**	0.015	0.046	61.09	0.507
DeepMobility	0.021	**0.058**	0.023	0.026	**0.044**	**22.68**	**0.787**
Shanghai							
w/o M1 and M2	0.004	0.033	0.041	0.004	0.070	100.7	0.531
w/o M1	**0.002**	**0.028**	0.011	0.003	0.025	78.84	0.635
w/o M2	0.003	0.042	0.030	**0.003**	0.043	100.0	0.545
DeepMobility	0.003	0.029	**0.008**	0.005	**0.025**	**63.12**	**0.686**
Shenzhen							
w/o M1 and M2	0.100	0.141	0.042	0.013	0.148	346.5	0.464
w/o M1	0.077	0.150	0.063	0.022	0.170	266.4	0.548
w/o M2	0.073	0.158	0.053	**0.011**	0.156	287.4	0.490
DeepMobility	**0.070**	**0.046**	**0.019**	0.015	**0.073**	**253.0**	**0.560**
Senegal							
w/o M1 and M2	0.028	0.095	0.035	0.006	0.056	324.6	0.599
w/o M1	0.053	**0.035**	**0.004**	0.006	**0.032**	255.1	0.620
w/o M2	**0.009**	0.044	0.035	0.007	0.036	277.9	0.610
DeepMobility	0.010	0.037	0.019	**0.006**	0.034	**223.0**	**0.650**

Table. 1 Ablation study on the two collaborative learning mechanisms. M1 stands for the bottom-up interaction modeling at the generator. M2 stands for the top-down feedback refinement at the critic. The results show that the two designs have major contributions to the population-level performance, and at the same time, capture the mobility patterns at the individual level.

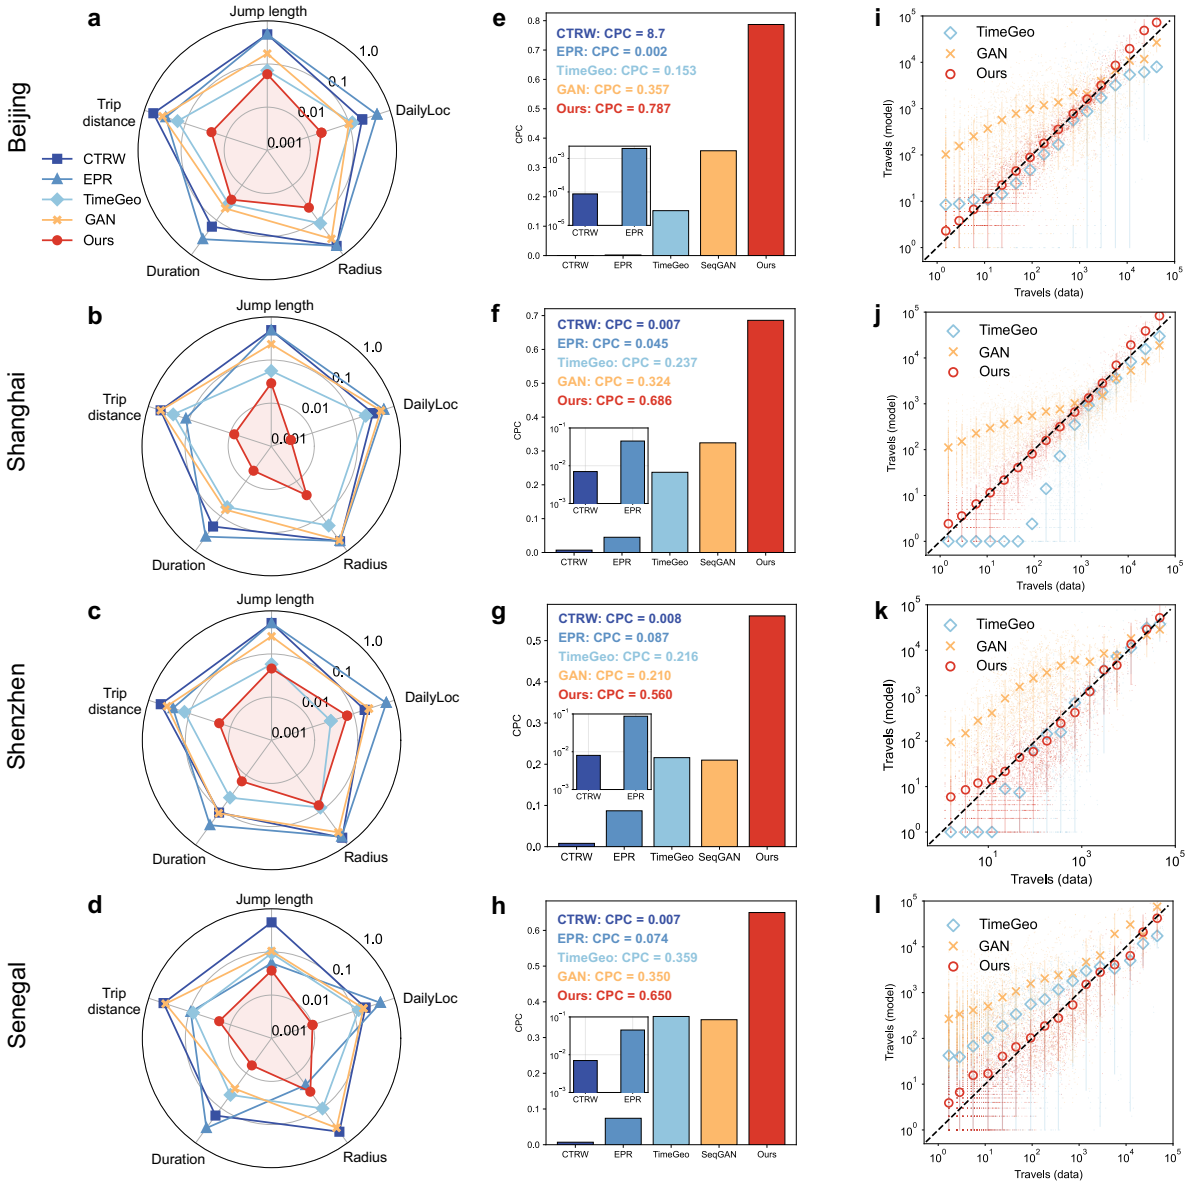

Fig. 3 Comparison of generation realism, in terms of individual and population scales. For the four empirical datasets, comparison of the performance in terms of JSD-based metrics (**a**-**d**), including *Jump length*, *DailyLoc*, *Radius*, *Duration*, and *Trip distance*, of the Continuous Time Random Walk (CTRW), exploration and preferential return model (EPR), TimeGeo, GAN, and DeepMobility (Ours), for Beijing (**a**), Shangahi (**b**), Shenzhen (**c**), and Senegal (**d**). Lower value of JSD-based metrics denotes a closer distribution with real data and thus represents better performance, and our framework clearly achieves the best performance. **e**-**h** Comparison of the performance in terms of Common Part of Commuters (CPC) of the CTRW, EPR, TimeGeo, GAN, and Ours. **i**–**l** Statistical illustration of the model-predicted and real values of population flows for the four datasets. Symbols denote the average number of generated flows for each bin and lines represent the 10%-90% percentiles. The dashed line is a perfect agreement between the observed flows and the generations. The points below symbols are scatter plot for each flow between a region pair. Our framework systematically outperforms other models.

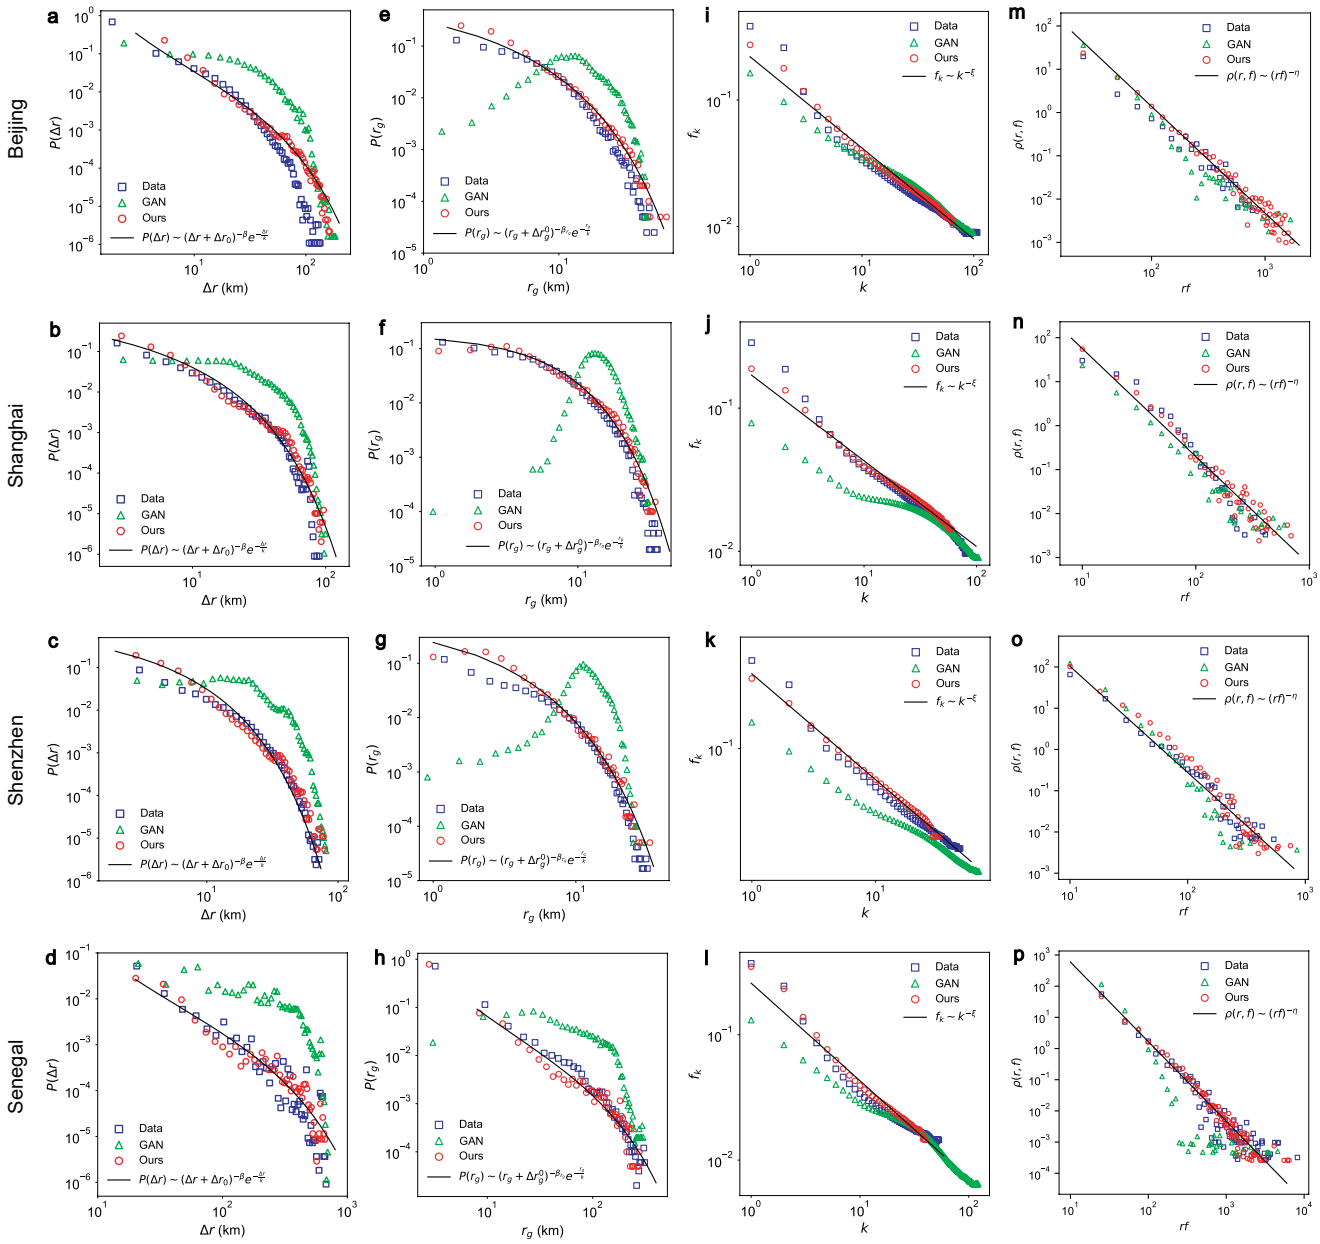

Fig. 4 Generated mobility data versus empirical data regarding different mobility laws at both individual and population levels. a-d Our generated mobility data reproduces the truncated power law of jump length (Δr), with the distribution approximated by $p(x) \sim (x+x_0)^{-\beta}\exp(-x/x^{\text{cut}})$. The solid black line represents the fitting result of Δr, with $\beta = \{1.28, 1.25, 1.34, 1.28\}$ and $R^2 = \{0.976, 0.942, 0.971, 0.914\}$ for the four datasets, respectively. **e-h** Similar results of the reproduced truncated power law of radius of gyration (r_g), with $\beta = \{1.20, 1.28, 1.22, 1.32\}$ and $R^2 = \{0.991, 0.981, 0.976, 0.954\}$. **i-l** The generated mobility data reproduces Zipf's law of visitation frequency, where the visitation frequency f_k to the k_{th} most visited location is well approximated by a power law $P(k) \sim k^{-\zeta}$, with $\zeta = \{1.15, 1.05, 1.01, 1.10\}$ and $R^2 = \{0.973, 0.981, 0.966, 0.963\}$. **m-p** The aggregated location visitation at the population level also reproduces the distance-frequency scaling law of mobility flow, where the number of visitors to a location with a specific frequency $\rho(r, f)$ is well-described by a power law fitting $\rho(r, f) = \mu/(rf)^{\eta}$, with $\eta = \{2.01, 2.04, 2.09, 2.08\}$ and $R^2 = \{0.926, 0.927, 0.940, 0.964\}$.

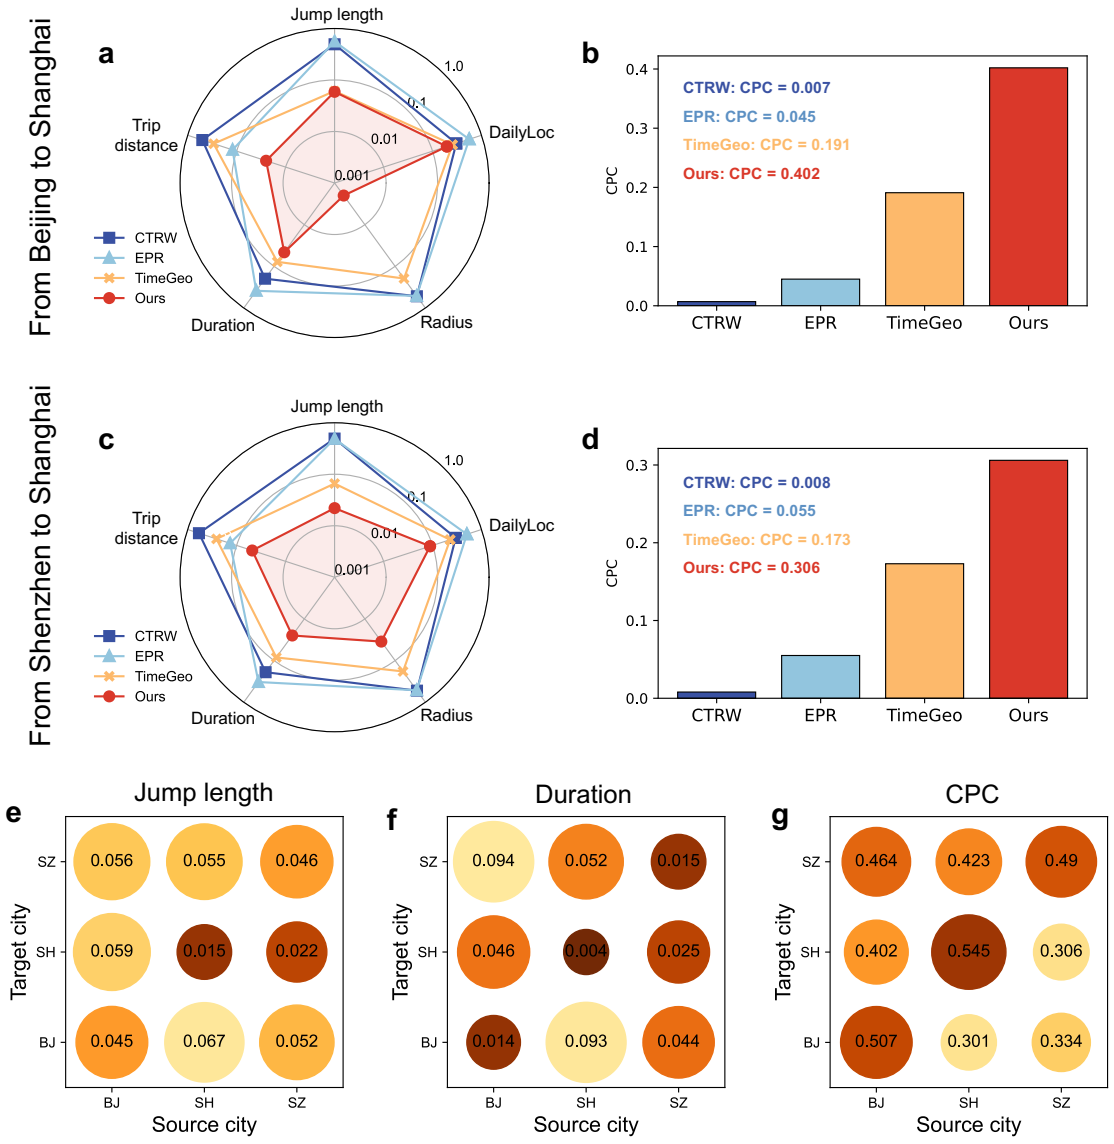

Fig. 5 Geographic transferability of DeepMobility. The models are trained on one source city and then evaluated on other target cities without finetuning. **a-d** The generation performance on the target city Shanghai with different source cities (**a-b** Beijing and **c-d** Shenzhen) regarding both individual and population scales. **e-f** The performance of DeepMobility on all source-target city pairs in terms of *jump length* (**e**), *duration* (**f**) and *CPC* (**g**). The size of the circle denotes the metric value, and the deeper color denotes better performance.

References

[1] Jing Yuan, Yu Zheng, and Xing Xie. "Discovering regions of different functions in a city using human mobility and POIs". In: *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2012, pp. 186–194.

[2] Veronique Van Acker, Phil Goodwin, and Frank Witlox. "Key research themes on travel behavior, lifestyle, and sustainable urban mobility". In: *International journal of sustainable transportation* 10.1 (2016), pp. 25–32.

[3] Hugo Barbosa, Surendra Hazarie, Brian Dickinson, Aleix Bassolas, Adam Frank, Henry Kautz, Adam Sadilek, José J Ramasco, and Gourab Ghoshal. "Uncovering the socioeconomic facets of human mobility". In: *Scientific reports* 11.1 (2021), p. 8616.

[4] R'emi Louf and Marc Barthelemy. "How congestion shapes cities: from mobility patterns to scaling". In: *Scientific reports* 4.1 (2014), pp. 1–9.

[5] Luıs MA Bettencourt, José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B West. "Growth, innovation, scaling, and the pace of life in cities". In: *Proceedings of the national academy of sciences* 104.17 (2007), pp. 7301–7306.

[6] Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J Ramasco, and Alessandro Vespignani. "Multiscale mobility networks and the spatial spreading of infectious diseases". In: *Proceedings of the National Academy of Sciences* 106.51 (2009), pp. 21484–21489.

[7] Srinivasan Venkatramanan, Adam Sadilek, Arindam Fadikar, Christopher L Barrett, Matthew Biggerstaff, Jiangzhuo Chen, Xerxes Dotiwalla, Paul Eastham, Bryant Gipson, Dave Higdon, et al. "Forecasting influenza activity using machine-learned mobility map". In: *Nature communications* 12.1 (2021), pp. 1–12.

[8] Kathleen A Cagney, Erin York Cornwell, Alyssa W Goldman, and Liang Cai. "Urban mobility and activity space". In: *Annual Review of Sociology* 46 (2020), pp. 623–648.

[9] Esteban Moro, Dan Calacci, Xiaowen Dong, and Alex Pentland. "Mobility patterns are associated with experienced income segregation in large US cities". In: *Nature communications* 12.1 (2021), pp. 1–10.

[10] Noli Brazil. "Environmental inequality in the neighborhood networks of urban mobility in US cities". In: *Proceedings of the National Academy of Sciences* 119.17 (2022), e2117776119.

[11] Yanyan Xu, Luis E Olmos, David Mateo, Alberto Hernando, Xiaokang Yang, and Marta C Gonzalez. "Urban Dynamics Through the Lens of Human Mobility". In: *Nature computational science* 3.7 (2023), pp. 611–620.

[12] Michael Batty. "Cities as Complex Systems: Scaling, Interaction, Networks, Dynamics and Urban Morphologies." In: (2009).

[13] 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". In: *Physics Reports* 734 (2018), pp. 1–74.

[14] Luca Pappalardo, Ed Manley, Vedran Sekara, and Laura Alessandretti. "Future directions in human mobility science". In: *Nature Computational Science* 3.7 (2023), pp. 588–600.

[15] Fernando Chirigati. "Gauging urban development with neural networks". In: *Nature Computational Science* 2.4 (2022), pp. 216–216.

[16] Dirk Brockmann, Lars Hufnagel, and Theo Geisel. "The scaling laws of human travel". In: *Nature* 439.7075 (2006), pp. 462–465.

[17] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. "Understanding individual human mobility patterns". In: *nature* 453.7196 (2008), pp. 779–782.

[18] Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. "A universal model for mobility and migration patterns". In: *Nature* 484.7392 (2012), pp. 96–100.

[19] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. "Limits of predictability in human mobility". In: *Science* 327.5968 (2010), pp. 1018–1021.

[20] Chaoming Song, Tal Koren, Pu Wang, and Albert-László Barabási. "Modelling the scaling properties of human mobility". In: *Nature physics* 6.10 (2010), pp. 818–823.

[21] Laura Alessandretti, Piotr Sapiezynski, Vedran Sekara, Sune Lehmann, and Andrea Baronchelli. "Evidence for a conserved quantity in human mobility". In: *Nature human behaviour* 2.7 (2018), pp. 485–491.

[22] Laura Alessandretti, Ulf Aslak, and Sune Lehmann. “The scales of human mobility”. In: *Nature* 587.7834 (2020), pp. 402–407.

[23] George Kingsley Zipf. “The P 1 P 2/D hypothesis: on the intercity movement of persons”. In: *American sociological review* 11.6 (1946), pp. 677–686.

[24] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. “A tale of many cities: universal patterns in human urban mobility”. In: *PloS one* 7.5 (2012), e37027.

[25] Xiao-Yong Yan, Wen-Xu Wang, Zi-You Gao, and Ying-Cheng Lai. “Universal model of individual and population mobility on diverse spatial scales”. In: *Nature communications* 8.1 (2017), pp. 1–9.

[26] Meead Saberi, Hani S Mahmassani, Dirk Brockmann, and Amir Hosseini. “A complex network perspective for characterizing urban travel demand patterns: graph theoretical analysis of large-scale origin-destination demand networks”. In: *Transportation* 44 (2017), pp. 1383–1402.

[27] Markus Schläpfer, Lei Dong, Kevin O’Keeffe, Paolo Santi, Michael Szell, Hadrien Salat, Samuel Anklesaria, Mohammad Vazifeh, Carlo Ratti, and Geoffrey B West. “The universal visitation law of human mobility”. In: *Nature* 593.7860 (2021), pp. 522–527.

[28] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. “Generative Adversarial Networks”. Preprint at https://arxiv.org/abs/1406.2661. 2014.

[29] Diederik P Kingma and Max Welling. “Auto-encoding variational bayes”. Preprint at https://arxiv.org/abs/1312.6114. 2013.

[30] Massimiliano Luca, Gianni Barlacchi, Bruno Lepri, and Luca Pappalardo. “A survey on deep learning for human mobility”. In: *ACM Computing Surveys (CSUR)* 55.1 (2021), pp. 1–44.

[31] Kun Ouyang, Reza Shokri, David S Rosenblum, and Wenzhuo Yang. “A non-parametric generative model for human trajectories.” In: *IJCAI*. Vol. 18. 2018, pp. 3812–3817.

[32] Dou Huang, Xuan Song, Zipei Fan, Renhe Jiang, Ryosuke Shibasaki, Yu Zhang, Haizhong Wang, and Yugo Kato. “A variational autoencoder based generative model of urban human mobility”. In: *2019 IEEE conference on multimedia information processing and retrieval (MIPR)*. IEEE. 2019, pp. 425–430.

[33] Jie Feng, Zeyu Yang, Fengli Xu, Haisu Yu, Mudan Wang, and Yong Li. “Learning to simulate human mobility”. In: *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 2020, pp. 3426–3433.

[34] Xin Zhang, Yanhua Li, Xun Zhou, Ziming Zhang, and Jun Luo. “Trajgail: Trajectory generative adversarial imitation learning for long-term decision analysis”. In: *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE. 2020, pp. 801–810.

[35] Xinyu Chen, Jiajie Xu, Rui Zhou, Wei Chen, Junhua Fang, and Chengfei Liu. “Trajvae: A variational autoencoder model for trajectory generation”. In: *Neurocomputing* 428 (2021), pp. 332–339.

[36] Huandong Wang, Changzheng Gao, Yuchen Wu, Depeng Jin, Lina Yao, and Yong Li. “PateGail: a privacy-preserving mobility trajectory generator with imitation learning”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 12. 2023, pp. 14539–14547.

[37] Filippo Simini, Gianni Barlacchi, Massimilano Luca, and Luca Pappalardo. “A Deep Gravity model for mobility flows generation”. In: *Nature communications* 12.1 (2021), pp. 1–13.

[38] Giovanni Mauro, Massimiliano Luca, Antonio Longa, Bruno Lepri, and Luca Pappalardo. “Generating mobility networks with generative adversarial networks”. In: *EPJ Data Science* 11.1 (2022), p. 58.

[39] Can Rong, Jie Feng, and Jingtao Ding. “GODDAG: Generating Origin-destination Flow for New Cities via Domain Adversarial Training”. In: *IEEE Transactions on Knowledge and Data Engineering* (2023).

[40] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. “Proximal policy optimization algorithms”. Preprint at https://arxiv.org/abs/1707.06347. 2017.

[41] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. “Value-decomposition networks for cooperative multi-agent learning”. Preprint at https://arxiv.org/abs/1706.05296. 2017.

[42] Shan Jiang, Yingxiang Yang, Siddharth Gupta, Daniele Veneziano, Shounak Athavale, and Marta C González. “The TimeGeo modeling framework for urban mobility without travel surveys”. In: *Proceedings of the National Academy of Sciences* 113.37 (2016), E5370–E5378.

[43] Jinzhong Wang, Xiangjie Kong, Feng Xia, and Lijun Sun. “Urban human mobility: Data-driven modeling and prediction”. In: *Acm Sigkdd Explorations Newsletter* 21.1 (2019), pp. 1–19.

[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*. Vol. 34. 01. 2020, pp. 808–816.

[45] Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. “Seqgan: Sequence generative adversarial nets with policy gradient”. In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 31. 1. 2017.

[46] Christian M Schneider, Vitaly Belik, Thomas Couronné, Zbigniew Smoreda, and Marta C González. “Unravelling daily human mobility motifs”. In: *Journal of The Royal Society Interface* 10.84 (2013), p. 20130246.

[47] Joshua C Peterson, David D Bourgin, Mayank Agrawal, Daniel Reichman, and Thomas L Griffiths. “Using large-scale experiments and machine learning to discover theories of human decision-making”. In: *Science* 372.6547 (2021), pp. 1209–1214.

[48] Christian Bongiorno, Yulun Zhou, Marta Kryven, David Theurel, Alessandro Rizzo, Paolo Santi, Joshua Tenenbaum, and Carlo Ratti. “Vector-based pedestrian navigation in cities”. In: *Nature Computational Science* 1.10 (2021), pp. 678–685.

[49] Fernando Chirigati. “From data to conservation laws”. In: *Nature Computational Science* 3.9 (2023), pp. 733–733.

[50] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. “Unique in the crowd: The privacy bounds of human mobility”. In: *Scientific reports* 3.1 (2013), pp. 1–5.

[51] Huandong Wang, Chen Gao, Yong Li, Gang Wang, Depeng Jin, and Jingbo Sun. “De-anonymization of mobility trajectories: Dissecting the gaps between theory and practice”. In: *The 25th Annual Network & Distributed System Security Symposium (NDSS’18)*. 2018.

[52] Yanyan Xu, Serdar Çolak, Emre C Kara, Scott J Moura, and Marta C González. “Planning for electric vehicle needs by coupling charging profiles with urban mobility”. In: *Nature Energy* 3.6 (2018), pp. 484–493.

[53] Edward Barbour, Carlos Cerezo Davila, Siddharth Gupta, Christoph Reinhart, Jasleen Kaur, and Marta C González. “Planning for sustainable cities by estimating building occupancy with mobile phones”. In: *Nature communications* 10.1 (2019), pp. 1–10.

[54] Lais MA Rocha, Aline Bessa, Fernando Chirigati, Eugene OFriel, Mirella M Moro, and Juliana Freire. “Understanding spatio-temporal urban processes”. In: *2019 IEEE International Conference on Big Data (Big Data)*. IEEE. 2019, pp. 563–572.

[55] Jiawei Xue, Nan Jiang, Senwei Liang, Qiyuan Pang, Takahiro Yabe, Satish V Ukkusuri, and Jianzhu Ma. “Quantifying the spatial homogeneity of urban road networks via graph neural networks”. In: *Nature Machine Intelligence* 4.3 (2022), pp. 246–257.

[56] Andrew J Tatem, Abdisalan M Noor, Craig Von Hagen, Antonio Di Gregorio, and Simon I Hay. “High resolution population maps for low income nations: combining land cover and census in East Africa”. In: *PloS one* 2.12 (2007), e1298.

[57] Pierre Deville, Catherine Linard, Samuel Martin, Marius Gilbert, Forrest R Stevens, Andrea E Gaughan, Vincent D Blondel, and Andrew J Tatem. “Dynamic population mapping using mobile phone data”. In: *Proceedings of the National Academy of Sciences* 111.45 (2014), pp. 15888–15893.

[58] Christoph Steinacker, David-Maximilian Storch, Marc Timme, and Malte Schröder. “Demand-driven design of bicycle infrastructure networks for improved urban bikeability”. In: *Nature Computational Science* 2.10 (2022), pp. 655–664.

[59] Yu Zheng, Yuming Lin, Liang Zhao, Tinghai Wu, Depeng Jin, and Yong Li. “Spatial planning of urban communities via deep reinforcement learning”. In: *Nature Computational Science* (2023), pp. 1–15.

[60] Fernando Chirigati. “Greener urban communities”. In: *Nature Computational Science* (2023), pp. 1–2.

[61] Jonathan Ho and Stefano Ermon. “Generative adversarial imitation learning”. In: *Advances in neural information processing systems* 29 (2016).

[62] Frans A Oliehoek and Christopher Amato. *A concise introduction to decentralized POMDPs*. Springer, 2016.

[63] Jameson L Toole, Carlos Herrera-Yaqüe, Christian M Schneider, and Marta C González. “Coupling human mobility and social ties”. In: *Journal of The Royal Society Interface* 12.105 (2015), p. 20141128.

[64] Fengli Xu, Yong Li, Depeng Jin, Jianhua Lu, and Chaoming Song. “Emergence of urban growth patterns from human mobility behavior”. In: *Nature Computational Science* 1.12 (2021), pp. 791–800.

[65] Jie Feng, Yong Li, Chao Zhang, Funing Sun, Fanchao Meng, Ang Guo, and Depeng Jin. “Deepmove: Predicting human mobility with attentional recurrent networks”. In: *Proceedings of the 2018 world wide web conference*. 2018, pp. 1459–1468.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- NMISI.pdf