


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

# Transportation Research Part C

journal homepage: [www.elsevier.com/locate/trc](http://www.elsevier.com/locate/trc)

## Learning universal human mobility patterns with a foundation model for cross-domain data fusion

Haoxuan Ma <sup>a</sup>,<sup>1</sup>, Xishun Liao <sup>a</sup>,<sup>\*</sup>,<sup>1</sup>, Yifan Liu <sup>a</sup>, Qinhua Jiang <sup>a</sup>, Chris Stanford <sup>b</sup>, Shangqing Cao <sup>c</sup>, Jiaqi Ma <sup>a</sup>

<sup>a</sup> Department of Civil and Environmental Engineering, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, 90095, CA, USA

<sup>b</sup> Novateur Research Solutions, 20110 Ashbrook Place, STE 170, Ashburn, 20147, VA, USA

<sup>c</sup> Department of Civil and Environmental Engineering, University of California, Berkeley, 760 Davis Hall, Berkeley, 94720, CA, USA

### ARTICLE INFO

#### Keywords:

Foundation model for mobility  
Multi-modal data fusion  
Cross-domain transfer learning  
Mobility pattern synthesis  
LLM-informed model  
Deep learning

### ABSTRACT

Human mobility modeling is critical for urban planning and transportation management, yet existing approaches often lack the integration capabilities needed to handle diverse data sources. We present a foundation model framework for universal human mobility patterns that leverages cross-domain data fusion and large language models to address these limitations. Our approach integrates multi-modal data of distinct nature and spatio-temporal resolution, including geographical, mobility, socio-demographic, and traffic information, to construct a privacy-preserving and semantically enriched human travel trajectory dataset. Our framework demonstrates adaptability through domain transfer techniques that ensure transferability across diverse urban contexts, as evidenced in case studies of Los Angeles (LA) and Egypt. The framework employs LLMs for semantic enrichment of trajectory data, enabling comprehensive understanding of mobility patterns. Quantitative evaluation shows that our generated synthetic dataset accurately reproduces mobility patterns observed in empirical data. The practical utility of this foundation model approach is demonstrated through large-scale traffic simulations for LA County, where results align well with observed traffic data. On California's I-405 corridor, the simulation yields a Mean Absolute Percentage Error of 5.85% for traffic volume and 4.36% for speed compared to Caltrans PeMS observations, illustrating the framework's potential for intelligent transportation systems and urban mobility applications.

## 1. Introduction

### 1.1. Motivation

Human mobility modeling has emerged as a critical component in urban planning, transportation management, business, and public policy development (González et al., 2008; Liao et al., 2024a; Stanford et al., 2024). Understanding how people move through cities and interact with the surrounding environment is fundamental to creating efficient, sustainable, and livable urban spaces. This importance has grown with increasing urbanization and the complexities of modern transportation systems (Hasan et al., 2013; Ma et al., 2024b).

\* Corresponding author.

E-mail addresses: [haoxuanma@ucla.edu](mailto:haoxuanma@ucla.edu) (H. Ma), [xishunliao@ucla.edu](mailto:xishunliao@ucla.edu) (X. Liao), [bmmliu@ucla.edu](mailto:bmmliu@ucla.edu) (Y. Liu), [qhjiang@ucla.edu](mailto:qhjiang@ucla.edu) (Q. Jiang), [cstanford@novateur.ai](mailto:cstanford@novateur.ai) (C. Stanford), [caoalbert@berkeley.edu](mailto:caoalbert@berkeley.edu) (S. Cao), [jiaqima@ucla.edu](mailto:jiaqima@ucla.edu) (J. Ma).

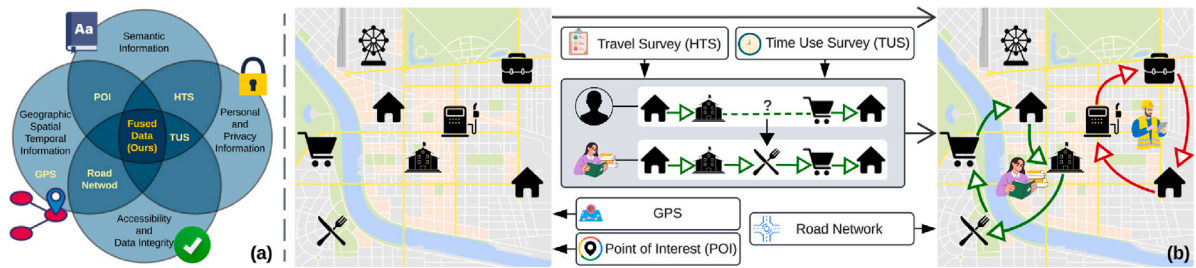
<sup>1</sup> Equal Contribution

<https://doi.org/10.1016/j.trc.2025.105311>

Received 14 April 2025; Received in revised form 8 July 2025; Accepted 14 August 2025

Available online 28 August 2025

0968-090X/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.



**Fig. 1.** Integration of cross-domain multi-modal mobility data with human mobility modeling. (a) The need for data fusion and (b) the fused data.

Foundation models and large language models (LLMs) offer transformative capabilities for mobility modeling by enabling multi-modal data integration, semantic enrichment, and cross-context generalization. These capabilities directly address the data fragmentation challenges that have historically limited mobility analysis, allowing for more comprehensive insights than previously possible with domain-specific approaches. Recent advances in multi-modal transportation systems have demonstrated the potential of sophisticated machine learning approaches, including graph neural networks for transportation recommendation systems (Yan et al., 2024) and ensemble learning methods for analyzing transportation behavior patterns (Yan et al., 2025).

One of the key research challenges in mobility modeling lies in integrating heterogeneous data sources to address incomplete or sparse datasets. Data fusion has emerged as a critical approach for combining diverse data types and enhancing modeling accuracy and robustness. Existing methods have demonstrated the potential of integrating sources such as GPS, social media, and environmental data to provide richer insights into mobility patterns (Ounoughi and Yahia, 2023; Afyouni et al., 2022). Researchers have also explored techniques for cross-domain integration and spatial-temporal modeling to improve movement predictions and scalability (Zheng, 2015; Wei et al., 2024). Although significant progress has been made, existing approaches still encounter challenges such as data sparsity, heterogeneity, and limited adaptability across diverse geographic contexts. To address these issues, this paper proposes a foundation model for mobility through cross-domain data fusion that enhances the applicability, precision, and generalizability of mobility modeling across diverse geographic and socio-demographic settings.

### 1.2. Challenges

Comprehensive datasets addressing the diverse requirements of mobility research remain scarce. An ideal mobility dataset would integrate rich semantic information with high-resolution trajectory data and demographic attributes while ensuring privacy protection. However, as shown in Fig. 1(a), existing datasets often focus on only a subset of these dimensions (Kapp and Mihaljevic, 2023). The complexity of combining these diverse data sources stems from two key challenges: cross-domain data harmonization and multi-modal information integration, as shown in Fig. 2.

Foundation models for mobility must address unique challenges in data integration that differ from those in natural language processing or computer vision. The multi-modal nature of mobility data – spanning continuous trajectories, discrete locations, semantic activities, and demographic attributes – requires sophisticated harmonization techniques to create coherent representations.

The fundamental challenge of cross-domain data fusion lies in combining data sources that complement each other but often contain conflicting information. Our foundation model addresses this by fusing data listed in Fig. 2, where each source has distinct limitations: trajectories lack semantic context and are usually low quality, surveys are biased and expensive, and infrastructure data requires synthesis with human mobility patterns to be meaningful.

Moreover, the integration process must also address disparities across data modalities. Spatio-temporal data requires bridging continuous trajectories with discrete locations while managing different temporal resolutions. Semantic information spans multiple representations requiring consistent interpretation, while statistical data must reconcile individual observations with aggregate patterns across appropriate spatial and temporal scales.

To address these challenges, we propose a foundation model through cross-domain data fusion that synthesizes mobility data with integrated semantic, trajectory, demographic, and location-based information, as shown in Fig. 1(b). By leveraging the semantic understanding capabilities of LLMs and the generalization power of transfer learning, our approach enables a more comprehensive understanding of mobility patterns while preserving privacy.

### 1.3. Contribution

This research establishes a foundation model for human mobility that bridges multiple data gaps: between high-resolution but semantically sparse trajectories and context-rich but limited surveys, between population-level statistics and individual-level behaviors, and between static infrastructure data and dynamic human movements. Our approach enables new insights into human mobility and its implications for societal infrastructure, which hold promise for numerous applications, including traffic optimization, urban planning, and anomaly detection (He et al., 2024; Ma et al., 2024; Stanford et al., 2024).

Our research makes several significant contributions:

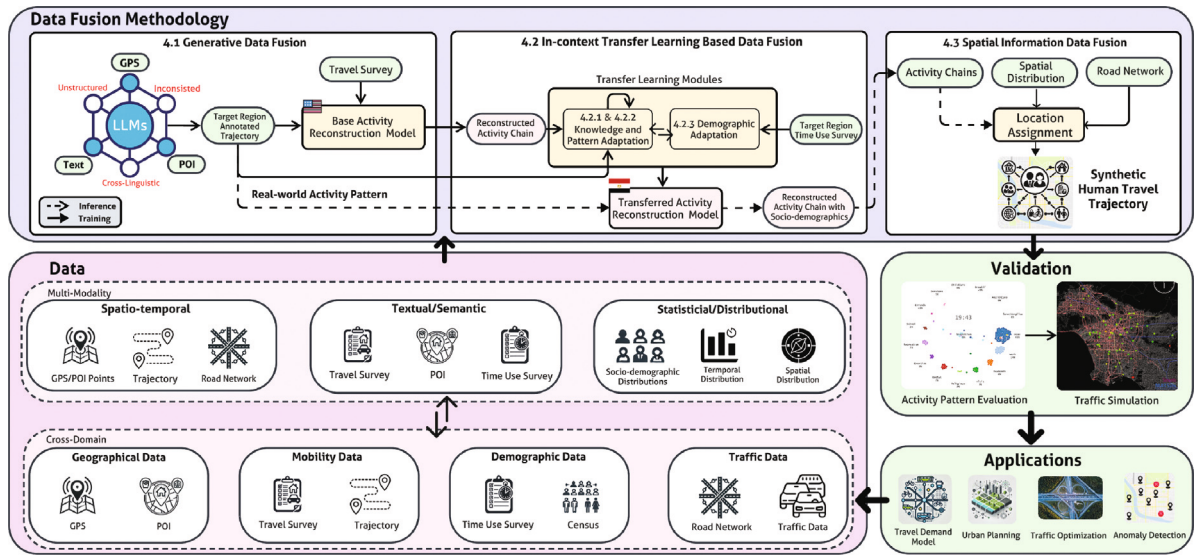


Fig. 2. An overview of the proposed data fusion framework. The captions within the methodology indicate their corresponding sections in the paper for reference.

- We developed a novel foundation model that processes and integrates cross-domain multi-modal data sources to create comprehensive synthetic datasets while balancing usability and privacy considerations. The framework harmonizes heterogeneous data types (spatial-temporal data: GPS (Global Positioning System) (Veraset, 2024) and POI (Point of Interest) (OpenStreetMap contributors, 2017), mobility data: Household Travel Survey (HTS) (Federal Highway Administration, 2017), demographic data: Time Use Survey (TUS) (U.S. BLS, 2024), and traffic data: road network data (OSM) (OpenStreetMap contributors, 2017)) and demonstrates generalization capabilities through validation with large-scale traffic simulation.
- We incorporated large language models (LLMs) as a core component for semantic enrichment of trajectory data, fusing the realism of GPS data with the knowledge of survey data and reconstructing missing activities with a survey-trained model. This LLM-powered approach significantly improves the completeness and interpretability of mobility trajectories.
- We implemented transfer learning mechanisms within our foundation model to address two key data challenges: (1) enabling knowledge and pattern adaptation in data-sparse regions through semi-supervised learning, leveraging limited GPS trajectories to adapt source domain patterns without requiring complete ground truth data, and (2) incorporating demographic variations into anonymized trajectories through unsupervised distribution adaptation, aligning mobility patterns with demographic characteristics without pairwise individual labels.

## 2. Related work

### 2.1. Human travel trajectory reconstruction

Trajectory reconstruction has become crucial in understanding human mobility patterns, especially when dealing with incomplete datasets. Recent advancements have addressed the challenges of data sparsity and irregularity through innovative techniques. Chen et al. (2019) introduced the Context-enhanced Trajectory Reconstruction (CTR) method, using tensor factorization to reconstruct complete individual trajectories from sparse Call Detail Records. Li et al. (2019) proposed the Multi-criteria Data Partitioning Trajectory Reconstruction (MDP-TR) method for large-scale, low-frequency mobile phone datasets, enhancing reconstruction performance by considering spatiotemporal patterns of missing data and individual similarities. For GPS data, Zheng et al. (2010) developed a collaborative system for location and activity recommendations, demonstrating significant improvements in inferring activity types. Alexander et al. (2015) emphasized the importance of comprehensive temporal and spatial analysis in trajectory reconstruction using mobile phone data.

Despite these advancements, challenges persist, including the diverse characteristics of data sources, limited data accessibility, and restricted model adaptability across different geographic regions. These challenges underscore the need for continued research to develop more robust and widely applicable reconstruction techniques.

### 2.2. Large language models for mobility modeling

Large language models (LLMs) have demonstrated significant potential in mobility modeling by leveraging their ability to process contextual and semantic information from diverse data sources. Recent advancements, such as those by Liu et al. (2024b) and Cai

**Table 1**  
Activity category codes with descriptions.

1. Home	6. Shop services	11. Social
2. Work	7. Meals out	12. Healthcare
3. School	8. Errands	13. Worship
4. Caregiving	9. Leisure	14. Other
5. Shop goods	10. Exercise	15. Pickup/Drop

et al. (2024), highlight the versatility of LLMs in tasks like trajectory reconstruction and decision-making. Studies on foundational models like LLaMA (Touvron et al., 2023), PaLM (Anil et al., 2023), and TEG-DB (Li et al., 2024) further underline their power in multi-modal data fusion and semantic interpretation. These capabilities make LLMs valuable complements to traditional methods, particularly in addressing challenges of data sparsity and complex semantic understanding in mobility applications.

### 2.3. Transfer learning

Recent advancements in transfer learning have shown significant potential for enhancing human mobility analysis, particularly in data-limited scenarios. Techniques originally developed for natural language processing offer valuable insights for trajectory reconstruction and mobility modeling.

Howard and Ruder (2018) gradual unfreezing technique could be adapted to preserve general mobility patterns while adjusting to specific regional characteristics. Peters et al. (2019) and Merchant et al. (2020) suggest that middle layers are often the most transferable, which could be useful when adapting mobility models from data-rich to data-scarce regions. Liu et al. (2019) explored freezing bottom layers while fine-tuning top layers, a strategy that could preserve fundamental human movement patterns while adapting to unique characteristics of specific urban environments.

These studies suggest enhancing trajectory reconstruction through gradual model adaptation and selective layer fine-tuning. Future work should adapt these methods to mobility data's unique characteristics, considering spatiotemporal dependencies and movement complexity, to improve mobility modeling in data-sparse regions.

## 3. Problem formulation

To address the multi-modal and cross-domain challenges, we propose a three-stage fusion framework that advances beyond existing approaches typically limited to combining similar data sources. As shown in Fig. 2, our framework integrates diverse data types to generate realistic human mobility patterns. The generative module (4.1) first employs LLMs to fuse unstructured POI data with GPS stay points, creating semantically annotated but temporally incomplete trajectories due to GPS fragmentation. These trajectories are then completed through an activity reconstruction model incorporating mobility pattern learned from travel survey.

The transfer learning stage (4.2) is designed to capture real-world mobility dynamics while enabling deployment in data-scarce regions. Through knowledge (4.2.1) and pattern (4.2.2) adaptation, the model can be adapted to different regions and cultures and learns to balance the generic knowledge learned from surveys with the realism in GPS trajectories, including region-specific patterns. Our demographic adaptation module further enhances realism by incorporating sociodemographic variations using time use survey (TUS) data, which provides rich activity patterns across different population groups, ensuring our synthetic trajectories reflect both demographic-specific behaviors and real-world mobility dynamics while maintaining privacy.

Finally, the third stage (4.3) grounds these behaviorally realistic patterns in physical space by fusing the generated activity chains with spatial distributions and traffic data to assign feasible locations, enabling validation through large-scale traffic simulation. This complete integration of dynamic human behavior with static infrastructure data creates a comprehensive framework that generates mobility patterns reflecting both the complexity of human decision-making and the constraints of physical infrastructure.

Starting with LLM-based trajectory annotation, we denote  $i$  for an agent. The  $j$ -th trajectory collected for the agent contains  $N$  stay points,  $Traj_j^i = \{P_1^{i,j}, P_2^{i,j}, \dots, P_N^{i,j}\}$ . Each stay point is linked with an activity, where stay point  $P_n^{i,j} = [T_n^{i,j}, (x, y)^{i,j}, S_n^{i,j}, E_n^{i,j}]$  consists of activity type  $T_n^{i,j}$  (as shown in Table 1), GPS location  $(x, y)^{i,j}$ , start time  $S_n^{i,j}$  and end time  $E_n^{i,j}$ . Then activity chains capturing the activity pattern in this region can be extracted from the annotated trajectory, as  $A_j^i = \{A_1^{i,j}, A_2^{i,j}, \dots, A_N^{i,j}\}$ , where  $A_n^{i,j} = [T_n^{i,j}, S_n^{i,j}, E_n^{i,j}]$ .

With our semantically annotated trajectories represented as  $A_j$ , we propose an activity reconstruction model  $M^{R1}$  tailored to a specific region  $R1$ . For an incomplete  $A_j$ , the model synthesizes a new activity chain  $A_j' = M^{R1}(A_j)$ , by filling the gaps based on the common activity patterns learned from the region  $R1$ .

The activity reconstruction adapts cultural effects to activity patterns through transfer learning, but incorporating sociodemographic correlation patterns poses a challenge. To address this, we introduce a demographic adaptation module that modifies reconstructed activity chains based on sociodemographic characteristics. Let us define the sociodemographic space  $S$  as:  $S = \{s_g | g \in G\}$ , where  $G$  represents the set of demographic groups. For a demographic space with  $N$  dimensions (e.g., age, gender, employment status), where each dimension  $i$  has  $C_i$  categories, the total number of demographic groups is:  $|G| = \prod_{i=1}^N C_i$ .

The adapter functions as an additional layer after the reconstruction model  $M^{R1}$ . Given a batch of reconstructed activity chains probability matrix  $Y \in \mathbb{R}^{B \times K \times |T|}$ , where  $B$  is the batch size,  $K = 96$  represents 96 time steps with 15 min time intervals for a day, and

$|T|$  represents the number of activity types defined in Table 1, we introduce an adapter matrix  $\mathbf{M} \in \mathbb{R}^{|\mathcal{G}| \times |T|}$ . For each demographic group  $g \in \mathcal{G}$ , its adapter vector  $\mathbf{a}_g \in \mathbb{R}^{|T|}$  modifies the reconstructed activity probabilities. The demographically-adapted activity probabilities for group  $g$  across  $K$  time intervals are computed as:

$$\mathbf{P}_g = \text{softmax}(\mathbf{Y} + \mathbf{a}_g, \text{axis} = |T|) \quad (1)$$

where  $\mathbf{a}_g$  is broadcast across the  $K$  time intervals.

Finally, we assign the activity chains to physical zonal locations while preserving spatial and temporal feasibility. Given a set of TAZs  $\mathcal{Z}$  and road network  $\mathcal{R}$ , we assign locations to activities by optimizing:

$$L^* = \arg \min_{L \in \mathcal{Z}} \sum_{i=1}^n \sum_{j=1}^{m_i} d(L_{i,j}, L_{i,j+1}) \quad (2)$$

subject to constraints on travel time feasibility via  $\mathcal{R}$  and consistency with observed spatial distributions. This enables validation against real-world traffic patterns.

## 4. Methodology

### 4.1. Generative data fusion

#### 4.1.1. LLM-informed trajectory semantic annotation

Human trajectory data collected from GPS devices include geo-location and time-stamped stay points that elucidate individual mobility patterns. However, these trajectories suffer from both semantic sparsity and temporal incompleteness. We develop the generation-based data fusion module, as Fig. 2 (4.1), to address these two challenges.

Adopting the methodology from our previous work (Liu et al., 2024), we use LLMs to annotate stay points using the POI dataset. This process utilizes LLM to assign semantic labels to POIs based on attributes such as “name”, “amenity”, and “construction type”. Each POI is then categorized into one of 15 activity codes as described in Table 1. The LLM’s prompt configuration includes a “Task Description”, “Activity Code Illustration”, and a description of each POI in natural language format, ensuring comprehensive understanding and accurate classification. The output from the LLM provides the top three most relevant activity codes for each POI, each with a corresponding probability.

Subsequently, a combination of rules and a Bayesian-based algorithm is used to associate stay points with the most appropriate POI and activity type, as  $T = \arg \max_{T_i} \{P(T_i | \text{POI}, S)\}$ , where  $T$  is the activity at the stay point, given POIs and start time  $S$ . This assignment takes into account both the inferred activity type and spatio-temporal information, thereby enhancing the accuracy of the trajectory’s semantic enrichment. Ultimately, this process results in human trajectories that are richly annotated with semantic information.

#### 4.1.2. Data augmentation by activity reconstruction

While this LLM-based annotation provides semantic context to the trajectories, their inherent incompleteness limits their direct use for mobility analysis. Therefore, our framework continues with an activity reconstruction process that fuses these annotated but incomplete trajectories with comprehensive mobility data from travel surveys. A base Activity Reconstruction Model  $M^0$  is trained on survey data to learn generic activity patterns and generate complete activity chains based on limited observed GPS stay points. It can be adapted to other regions with a small dataset, supported by the strong similarity in human activities across regions (Schneider et al., 2013; Cao et al., 2019). The adaptation of this base model to realistic geographical and cultural contexts is detailed in Section 4.2.

Encoding one-day activity chains  $A_j^i$  is the first step in preparing training data. This process converts a single day into a time vector with 96 time slots, each representing a 15 min interval, where activity types occupy their corresponding time slots (see Fig. 3). The training data is then processed using three masking techniques to simulate incomplete real-world data, enrich training examples, and improve the model’s robustness in representing activity patterns and interdependencies. The first technique, **activity-based masking**, randomly removes entire activities from the chain, simulating GPS signal loss in specific locations. The second, **period masking**, removes continuous periods of the day, mimicking situations where data is unavailable for extended periods, like when users are offline. The third, **time slot-based masking**, randomly removes individual time slots, representing sporadic or intermittent data loss.

The backbone of the model adopts Transformer-based structures (Vaswani et al., 2017) for sequence completion, enhanced with temporal features (time segments and day-of-week labels). We formulate the reconstruction as a time slot classification task with a tailored loss function that integrates three components to capture both activity accuracy and temporal dependencies:

**Cross-Entropy Loss** is commonly used to measure the dissimilarity between predicted probability  $\hat{y}_c$  and true activity labels  $y_c$  for given class  $c$  over the 96 time slots, as  $L_{CE} = -\sum_c w_c \cdot y_c \log(\hat{y}_c)$ . The use of class weights  $w_c$  balances the loss across potentially imbalanced activity classes.

**Transition Loss** uses binary cross-entropy to predict activity transitions by comparing predicted changes with true changes. It captures temporal dynamics by penalizing incorrect transitions, and encouraging accurate predictions at the correct time points. This approach helps generate realistic and coherent activity sequences, avoiding erratic jumps between activities.

$$L_{TR} = -\frac{1}{N} \sum_{i=1}^N [t_i \log(\hat{t}_i) + (1 - t_i) \log(1 - \hat{t}_i)] \quad (3)$$

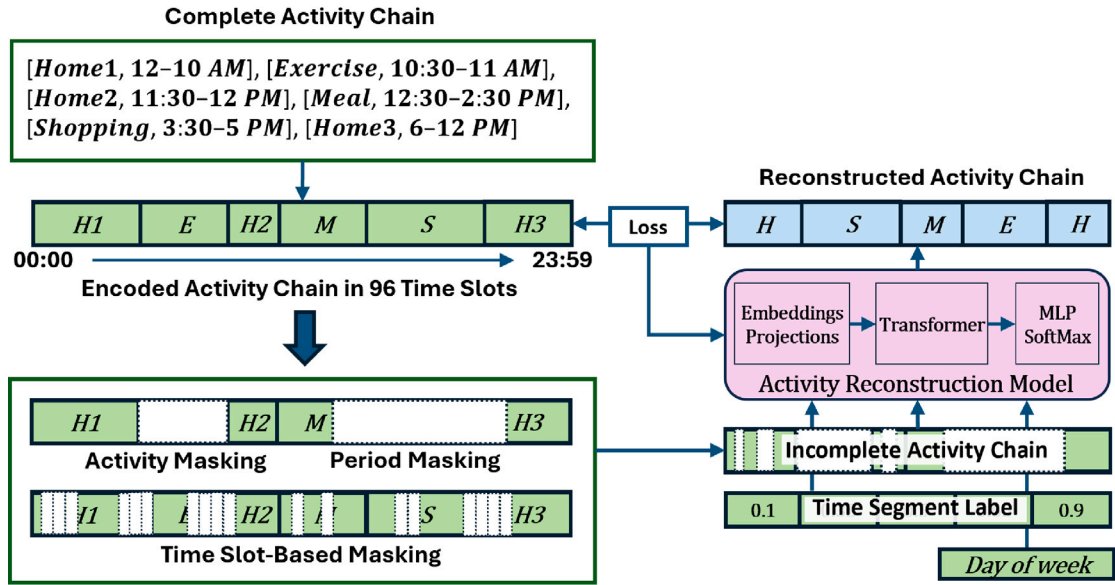


Fig. 3. Network architecture of the activity reconstruction model.

where  $N$  is the number of time steps,  $t_i$  and  $\hat{t}_i$  are true and predicted transitions at time step  $i$  (1 if activity changed, 0 otherwise).

**Dynamic Time Warping Loss (DTW)**,  $L_{DTW} = DTW(\hat{y}, y)$ , is adopted to evaluate similarity between predicted sequence  $\hat{y}$  and true sequence  $y$  using DTW distance (Müller, 2007).

Finally, the loss functions  $L$  combines three loss terms as  $L = w_1 \cdot L_{CE} + w_2 \cdot L_{TR} + w_3 \cdot L_{DTW}$ .

#### 4.2. In-context transfer learning

While our generation module effectively reconstructs activity chains, adapting to diverse regional contexts and demographic characteristics is crucial for realistic mobility patterns. As shown in Fig. 2 (4.2), our transfer learning approach addresses this through three adaptations: (1) **knowledge adaptation** for region-specific knowledge when regional survey data is available, (2) **mobility pattern adaptation** integrating survey data with GPS data while retaining universal activity motifs, and (3) **demographic-aware adaptation** for demographic-specific patterns. Based on research showing that over 90% of human activities share a common structure between regions (Schneider et al., 2013; Cao et al., 2019), this underlying similarity enables efficient adaptation to new regions with limited data.

##### 4.2.1. Supervised transfer learning for knowledge adaptation

When transferring to a region with a complete activity dataset (e.g., regional survey data), we employ progressive unfreezing as the transfer learning strategy. This approach balances knowledge preservation in the source domain with adaptation to the target domain, tuning three model components: embedding layers, Transformer layers, and an MLP, as illustrated in Fig. 3. Following Yosinski et al. (2014), input-near layers capture general features while output-near layers learn task-specific ones. By unfreezing the middle layers last (Peters et al., 2019), we mitigate catastrophic forgetting and facilitate fine-grained domain adaptation.

In particular, the unfreezing process proceeds in three phases. In the **Initial Phase** (first quarter of epochs), only the MLP and embedding layer are unfrozen. This allows the MLP, closest to the output, to adapt activity probability to the target domain, while the embedding layer refines semantic representations by incorporating new features. During the **Intermediate Phase** (second quarter of epochs), the Transformer's encoder and decoder layers closest to the input are unfrozen, enabling adjustments to foundational sequence data patterns while preserving core source-domain knowledge. Finally, in the **Final Phase**, the middle layers of the Transformer's encoder and decoder are unfrozen to refine intermediate-level representations, capturing more complex activity relationships and temporal dependencies.

##### 4.2.2. Semi-supervised transfer learning for mobility pattern adaptation

A supervised learning strategy is impractical for transferring the model to regions with only incomplete GPS data. To address this data sparsity issue, we develop a semi-supervised learning strategy as a key module to fuse survey data knowledge learned by the base model with the realism of GPS data in the target region.

As outlined in Algorithm 1, the base model  $M^0$  first synthesizes the initial dataset  $A_0$ , initiating an iterative adaptation process. At each step, synthetic datasets  $A_n$  are generated by  $M^n$  from raw activity chains  $A_{raw}$  to train subsequent models  $M^{n+1}$ . The process

**Algorithm 1** Semi-Supervised Transfer Learning

---

**Input:** Raw activity chain  $A_{raw}$ , initial model  $M^0$   
**Output:** Synthetic activity chain  $A_{syn}$ , transferred model  $M^n$   
Generate  $A_0 = M^0(A_{raw})$  and initialize  $n = 1$   
**repeat**  
  Train  $M^n$  using  $A_{n-1}$  with masking  
  Generate  $A_n = M^n(A_{raw})$   
  Compute JSD similarity  $J = \text{JSD}(A_{raw}, A_n)$   
  **if**  $J$  has converged **then**  
     $A_{syn} = A_n$   
    **break**  
  **end if**  
   $n = n + 1$   
**until** convergence or maximum iterations reached

---

continues until the JSD between synthetic and real-world data converges, ensuring the model's readiness for accurate activity chain reconstruction and application in new regions.

$$L_{real} = \frac{1}{N_r} \sum_{i=1}^{N_r} m_i \cdot L_{CE_i} \quad (4)$$

$$L_{synthetic} = \frac{1}{N_s} \sum_{i=1}^{N_s} (1 - m_i) \cdot L_{CE_i} \quad (5)$$

where  $m_i$  is the mask (1 for real, 0 for synthetic),  $N_r$  and  $N_s$  are the numbers of real and synthetic points, respectively.

The total loss function then becomes:

$$L = w_1 \cdot (w_l \cdot L_{real} + w_s \cdot L_{synthetic}) + w_2 \cdot L_{TR} + w_3 \cdot L_{DTW} \quad (6)$$

To avoid overfitting to the new synthetic data and ensure robustness, 20% of the previous iteration's dataset is retained in each new training cycle. This retention strategy allows the model to refine its understanding by integrating information from both new synthetic data and a portion of the previously learned data, promoting stability and preventing drastic shifts in learned patterns.

Furthermore, we differentiate real and synthetic data during training by applying a mask-based loss weighting technique. This method prioritizes real data in the loss computation, weighting it more heavily in the cross-entropy loss term as:

where  $w_1$ ,  $w_2$ , and  $w_3$  are weights balancing the different loss components, while  $w_l$  and  $w_s$  control the relative importance of real versus synthetic data within the cross-entropy loss.

#### 4.2.3. Optimization-based transfer learning for demographic adaptation

Incorporating sociodemographic patterns into mobility modeling is challenging due to scarce paired GPS-demographic data. While HTS provides such paired data, its limited regional availability constrains model generalizability. Our transfer learning approach addresses this by incorporating demographic dynamics through targeted adaptation of activity patterns. We leverage TUS as a privacy-preserving source of aggregated activity distributions across demographic groups. Harmonizing activity categories to classification space  $\mathcal{T}$  yields target distributions  $\mathbf{D}_g \in \mathbb{R}^{|\mathcal{T}|}$  for each demographic group  $g \in \mathcal{G}$ .

We propose a demographic adaptation mechanism through a learnable adapter matrix that modulates the activity reconstruction module's output distribution. The adapter optimization is formulated as an optimization problem with reward-guided updates. We formulate a dual-criteria objective function that captures both global and local distribution alignment:

$$\mathcal{L}(\mathbf{P}_g, \mathbf{D}_g) = e^{-\alpha \text{RMSE}(\mathbf{P}_g, \mathbf{D}_g)} + \beta \cdot \text{KL}(\mathbf{P}_g \parallel \mathbf{D}_g) \quad (7)$$

where the exponential RMSE term ensures coarse distribution matching while the KL divergence term enforces fine-grained consistency between generated and target distributions.

The adaptation weights are optimized through reward-weighted gradient descent, where  $\mathcal{L}(\mathbf{P}_g, \mathbf{D}_g)$  serves as an adaptive scaling factor:

$$\mathbf{a}_g^{t+1} = \mathbf{a}_g^t + \eta \cdot \mathcal{L}(\mathbf{P}_g, \mathbf{D}_g) \cdot \nabla_{\mathbf{a}_g} \mathcal{L} \quad (8)$$

with  $\nabla_{\mathbf{a}_g} \mathcal{L} = (\mathbf{D}_g - \mathbf{P}_g)$  providing the descent direction toward distribution matching.

To maintain numerical stability and interpretability, we impose bounded constraints on the adapter weights:

$$\mathbf{a}_g^{t+1} = \text{clip}(\mathbf{a}_g^{t+1}, \epsilon, M) \quad (9)$$

where  $\eta$  denotes the learning rate, and  $\epsilon, M$  define the feasible adaptation range. The optimization terminates upon reaching distribution convergence:

$$\|\mathbf{P}_g^t - \mathbf{D}_g\|_2 \leq \delta \quad \forall g \in \mathcal{G} \quad (10)$$

where  $\delta$  controls the distribution matching precision. This formulation ensures stable convergence while maintaining interpretable demographic adaptation weights, as shown in Algorithm 2.

**Algorithm 2** Gradient-Based Distribution Optimization for Demographic Adapters

---

**Input:** Reconstruction Model  $f$ , Time Use Survey Data  $\mathcal{T}$ , Learning rate  $\eta$ , Convergence threshold  $\delta$   
**Output:** Optimized adapter collection  $\mathcal{A}$   
 Process  $\mathcal{T}$  to get target distributions  $\{\mathbf{D}_g\}$   
 Initialize  $\mathcal{A} = \{\}$   
**for**  $g \in \mathcal{G}$  **do**  
   Initialize adapter weights  $\mathbf{a}_g$   
   **repeat**  
      $\mathbf{Y} = f(\text{input})$   
      $\mathbf{P}_g^t = \text{softmax}(\mathbf{Y} + \mathbf{a}_g^t)$   
      $\mathcal{L} = e^{-\alpha \cdot \text{RMSE}(\mathbf{P}_g^t, \mathbf{D}_g)} + \beta \cdot \text{KL}(\mathbf{P}_g^t | \mathbf{D}_g)$   
      $\nabla_{\mathbf{a}_g} \mathcal{L} = \mathbf{D}_g - \mathbf{P}_g^t$   
      $\mathbf{a}_g^{t+1} = \mathbf{a}_g^t + \eta \cdot \mathcal{L} \cdot \nabla_{\mathbf{a}_g} \mathcal{L}$   
      $\mathbf{a}_g^{t+1} = \text{clip}(\mathbf{a}_g^{t+1}, \epsilon, M)$   
     **if**  $\|\mathbf{P}_g^t - \mathbf{D}_g\|_2 \leq \delta$  **then**  
       **break**  
     **end if**  
   **until**  $\text{iteration} = \text{max\_iterations}$   
    $\mathcal{A} = \mathcal{A} \cup \{\mathbf{a}_g\}$   
**end for**  
**return**  $\mathcal{A}$

---

### 4.3. Activity location assignment

To translate our generated activity chains into real-world use, we developed a location assignment module that merges spatial limits and traffic data, as in Fig. 2 (4.3), which assigns each activity to specific TAZs through a layered approach: first fixing mandatory activities (home, work, school) based on commute patterns, then placing other activities while considering both distance spread and flow between fixed points. The process refines assignments until matching real spatial patterns that reflect regional traits.

The activity location assignment process transforms abstract activity chains into spatially grounded trajectories by assigning each activity to specific TAZs. This process must balance multiple objectives: maintaining realistic commuting patterns, preserving feasible trip chains, and matching observed spatial distributions of activities across the region, as detailed in Algorithm 3.

#### 4.3.1. Mandatory activity assignment

The first stage focuses on assigning locations for mandatory activities (work and school), using home locations as fixed anchor points. This stage is crucial as these locations strongly influence the spatial patterns of other activities throughout the day. The process includes:

**Distance Sampling:** For each individual, we sample commute distances from demographic-specific distributions ( $D_{md}$ ) derived from regional data. These distributions vary by sub-region to capture local commuting patterns, such as longer commutes in suburban areas versus shorter ones in urban cores.

**Land Use Compatibility:** Candidate TAZs are filtered based on land use compatibility. For work activities, we consider employment density and workplace location patterns; for school activities, we account for educational facility distribution.

**Spatial Matching:** The final TAZ selection optimizes for both the sampled distance and land-use suitability while ensuring the spatial distribution of mandatory activities matches regional patterns. This is formulated as:  $Z_{md} = \arg \min_{z \in \mathcal{Z}} \{d(h, z) - \hat{d}\}$  where  $h$  is the home TAZ,  $z$  is a candidate TAZ, and  $\hat{d}$  is the sampled target distance.

#### 4.3.2. Non-mandatory activity assignment

The second stage assigns locations for non-mandatory activities (shopping, leisure, etc.) between mandatory activity anchor points. This stage employs a more complex optimization approach:

**Distance and Angular Constraints:** Two key parameters guide the assignment: - Distance to next activity ( $D_{nmd}$ ): sampled from activity-specific distributions - Angular deviation ( $\theta$ ): measures the difference between the direct path to the next activity and the path to the next anchor point

**Trip Chaining Optimization:** Location selection considers both previous and next activities through a weighted objective function:  $Z_{nmd} = \arg \min_{z \in \mathcal{Z}} \{\alpha \cdot |d(z_{prev}, z) - \hat{d}| + \beta \cdot |\theta(z_{prev}, z, z_{next}) - \hat{\theta}|\}$  where  $\alpha$  and  $\beta$  balance the importance of distance and directional constraints.

**Temporal Feasibility:** Assignments must ensure travel times between activities are realistic given the transportation network. This adds a constraint:  $t(z_i, z_{i+1}) \leq T_{max}$  where  $t(\cdot)$  is the travel time function and  $T_{max}$  is the available time window.

**Algorithm 3** Location Assignment for Activity Chains

---

```

1: Input: Activity chains  $A$ , Home TAZ assignments  $H$ , Reference patterns  $R$ 
2: Output: Spatially-assigned activity chains  $A'$ 
3: for each activity chain  $a_i \in A$  do
4:    $h_i \leftarrow H[i]$  {Get home TAZ}
5:   // Assign mandatory activities
6:   for activity  $k \in a_i$  where  $\text{type}(k) \in \{\text{work, school}\}$  do
7:      $d \sim P(D|\text{type}(k))$  {Sample from distance distribution}
8:      $k.\text{taz} \leftarrow \text{argmin}_z |d(h_i, z) - d|$ 
9:   end for
10:  // Assign non-mandatory activities
11:  for activity  $k \in a_i$  where  $\text{type}(k)$  is non-mandatory do
12:     $p \leftarrow \text{prev\_anchor}(k)$ 
13:     $n \leftarrow \text{next\_anchor}(k)$ 
14:     $d \sim P(D|\text{type}(k), \text{region}(h_i))$ 
15:     $\theta \sim P(\Theta|\text{type}(k))$ 
16:     $k.\text{taz} \leftarrow \text{argmin}_z \{\alpha|d(p, z) - d| + \beta|\theta(p, z, n) - \theta|\}$ 
17:  end for
18: end for
19: return  $A'$ 

```

---

**4.3.3. Distribution refinement**

The final stage ensures that the aggregate spatial distribution of activities matches observed patterns through iterative refinement:

**Distribution Comparison:** Activity frequencies across sub-regions are compared with ground truth data using cosine similarity metrics.

**Parameter Adjustment:** Reference distributions ( $D_{md}$ ,  $D_{nmd}$ ,  $\theta$ ) are iteratively adjusted using:  $D^{t+1} = D^t + \eta \cdot (F_{target} - F_{current})$  where  $F$  represents activity frequency distributions and  $\eta$  is a learning rate.

**Convergence Criteria:** The refinement process continues until:  $|F_{target} - F_{current}| < \epsilon$  where  $\epsilon$  is a predefined threshold.

The complete algorithmic implementation is presented in Algorithm 3, which formalizes this three-stage process into a computational framework. The algorithm employs efficient data structures and optimization techniques to handle large-scale assignment tasks, making it suitable for regional-level applications.

This location assignment methodology ensures that the generated activity patterns are not only temporally and demographically realistic but also spatially feasible within the context of real urban environments. The multi-stage approach allows for fine-grained control over spatial distributions while maintaining the integrity of individual activity chains.

**5. Data source****5.1. Household travel surveys**

Household Travel Surveys have long been a fundamental data source in transportation planning and modeling. These surveys provide comprehensive insights into daily travel patterns by collecting detailed information about individual and household travel behaviors, including complete door-to-door trip chains, activity purposes, timing, and mode choices. This detailed trip chain information makes HTS particularly valuable for understanding the sequential nature of human mobility decisions and underlying patterns.

A key contribution of HTS is its standardized categorization of travel activities, which forms the basis for our model's activity type inference and pattern generation. This standardization provides a consistent vocabulary for describing different types of activities while capturing temporal relationships between activities and their durations. Furthermore, HTS data includes rich sociodemographic information, enabling the understanding of how different population groups engage in various travel behaviors.

In our framework, HTS data serves as the foundation for training our base model, enabling it to learn complex sequential patterns in daily activities and their temporal dependencies. The comprehensive nature of HTS allows our model to capture not only individual activity choices but also the logical connections between successive activities. This learning enables the generation of synthetic activity chains that maintain realistic temporal patterns while respecting logical constraints in human mobility behavior.

However, the limited availability of HTS data poses a significant challenge to mobility modeling, particularly in developing regions. Traditional HTS collection is resource-intensive and costly, making it impractical for many countries to conduct regularly. To address this limitation, our framework implements several innovative approaches. First, through domain adaptation, our model can transfer knowledge from regions with available HTS data to those without, while preserving cultural and regional characteristics. This adaptation process considers regional activity preferences and local transportation contexts.

We further enhance our framework's capability through a data augmentation approach that generates synthetic travel patterns while maintaining statistical consistency with observed behaviors. This is complemented by our AI network's novel self-completion

mechanism, which can intelligently fill gaps in incomplete travel patterns while maintaining consistency with learned behavioral patterns. The framework also integrates street network information to enhance the realism of generated patterns, ensuring that synthetic travel behaviors reflect realistic constraints of the physical transportation infrastructure.

### 5.2. GPS and POI data integration

GPS trajectory data provides rich spatiotemporal information about human movement patterns. In our framework, we utilize the Veraset dataset, which contains anonymized cellular points and their trajectories. While this data offers invaluable insights into real-world movement patterns, it comes with several inherent limitations. First, GPS data often contains measurement errors and noise, leading to inconsistent or inaccurate location recordings. More significantly, due to privacy concerns, the data is anonymized, stripping away crucial information about trip purposes and individual characteristics. This anonymization, while essential for protecting personal privacy, creates a significant gap in understanding the underlying motivations for observed movements.

These limitations of raw GPS data can be largely addressed through strategic integration with Point of Interest (POI) data. POI data provides the semantic context necessary to transform anonymous coordinate points into meaningful activity locations. By matching GPS stay points with nearby POIs, we can infer likely activity types and purposes, adding crucial semantic meaning to the raw movement data. Our framework employs Large Language Models (LLMs) for this matching process, enabling more nuanced and context-aware association between locations and activities.

However, the integration of GPS and POI data presents its own challenges. POI databases, particularly from open-source platforms, often suffer from completeness and accuracy issues. Some locations may be missing, incorrectly categorized, or outdated. Additionally, in dense urban areas, multiple POIs might exist within the error radius of a GPS stay point, making it difficult to determine the actual visited location. Our framework addresses these challenges through a probabilistic matching approach that considers both spatial proximity and temporal patterns, supplemented by the semantic understanding capabilities of LLMs.

The complementary nature of these data sources enables our framework to generate more realistic and contextually aware mobility patterns while maintaining individual privacy and data security. This integration forms a crucial foundation for subsequent stages of our mobility modeling pipeline, particularly in activity reconstruction and location assignment.

### 5.3. Time use survey data

Time Use Surveys (TUS) are systematic studies that collect detailed information about how individuals spend their time throughout the day. These surveys typically require participants to maintain detailed diaries, recording their activities in fixed time intervals (often 15–30 min), including information about what they do, when they do it, where they are, and who they are with. Unlike transportation-focused surveys, TUS captures the full spectrum of human activities, from personal care and work to leisure and social interactions.

A key advantage of TUS is its widespread availability across different countries and regions. National statistical offices regularly conduct these surveys as part of their standard data collection efforts, making them more accessible than specialized transportation surveys. This broad availability is particularly valuable for regions where detailed travel surveys are scarce or unavailable.

TUS data is structurally organized by sociodemographic characteristics, providing detailed insights into how different population groups allocate their time. The surveys typically segment data by age, gender, employment status, and other demographic factors, enabling fine-grained analysis of behavioral patterns across different population segments. This demographic sensitivity makes TUS particularly valuable for our framework, as it provides target distributions for activity patterns that vary by population group.

In our framework, these demographic-specific activity distributions serve as the foundation for our demographic adaptation approach. When generating synthetic activity patterns, our model uses TUS-derived distributions as reward signals, encouraging the generation of activities that statistically match observed patterns for each demographic group. This mechanism ensures that generated activities not only follow plausible sequences but also reflect realistic time-use patterns for different population segments.

However, integrating TUS data presents unique challenges. The primary difficulty lies in reconciling the different activity classification systems between TUS and transportation surveys. While HTS focuses on travel-related activities, TUS employs a broader, more general activity classification system. This difference necessitates careful mapping and harmonization of activity categories to maintain consistency in our modeling framework.

The cultural sensitivity of TUS data adds another valuable dimension to our framework. These surveys capture fundamental differences in how societies organize their daily activities — from variations in working hours and meal times to distinct patterns of social interaction and leisure activities. This cultural awareness is crucial for our framework's ability to adapt to different regional contexts, supporting our goal of generating culturally appropriate mobility patterns across diverse urban environments.

## 6. Los Angeles (LA) County case study

We evaluate our framework through two case studies: Los Angeles (LA) County and Egypt. The LA study demonstrates our complete pipeline capabilities, while Egypt showcases transfer learning in data-scarce regions. Each component in our pipeline – semantic annotation, activity reconstruction, transfer learning (including both activity pattern and demographic adaptation) – is essential in the data fusion process, as removing any component would prevent the generation of complete, demographically representative mobility patterns. We compare the output synthetic data against multiple sources, such as HTS, TUS, and raw GPS points, as our baseline, which helps reflect the disagreement on the same measurements across different data sources.

**Algorithm 4** Large-Scale Stay Point Extraction Pipeline

---

```

1: Input: Raw GPS trajectories  $T$ , Temporal threshold  $\delta_t = 300s$ , Spatial threshold  $\delta_s = 300m$ 
2: Output: Processed stay points  $S$  with H3 regional mapping
3: // Stage 1: Temporal Stay Detection
4: for each trajectory  $t_i \in T$  do
5:    $temporal\_stays \leftarrow detect\_temporal\_gaps(t_i, \delta_t)$ 
6:    $spatial\_stays \leftarrow sequential\_spatial\_clustering(temporal\_stays, \delta_s)$ 
7: end for
8: // Stage 2: Regional Filtering and H3 Mapping
9:  $regional\_users \leftarrow filter\_by\_boundary(spatial\_stays, LA\_bounds)$ 
10:  $h3\_stays \leftarrow map\_to\_h3(regional\_users, resolution=10)$ 
11: // Stage 3: Speed-based Trajectory Filtering
12: for each stay sequence  $s_i \in h3\_stays$  do
13:   if  $travel\_speed(s_i) > 30$  km/h then
14:      $classify\_as\_passing(s_i)$ 
15:   else
16:      $classify\_as\_stay(s_i)$ 
17:   end if
18: end for
19: // Stage 4: Quality Assurance
20:  $filtered\_stays \leftarrow apply\_coverage\_filter(h3\_stays, min\_coverage=0.33)$ 
21:  $S \leftarrow aggregate\_daily\_patterns(filtered\_stays)$ 
22: return  $S$ 

```

---

## 6.1. Experimental setup

### 6.1.1. Data preprocessing

For our Los Angeles County case study, we gathered half-year GPS points from Veraset (Veraset, 2024), a data vendor. We developed a robust stay point extraction methodology using PySpark-based distributed processing to handle million-scale GPS trajectory analysis efficiently.

**Stay Point Detection Algorithm and Validation:** We implemented a dual-threshold approach combining temporal (300-second) and spatial (300-meter) criteria with dynamic centroid updating. The algorithm employs haversine distance calculations to maintain spatial accuracy, incrementally updating centroid positions as new points are added to each stay cluster. Sequential processing ensures that genuine activity locations are distinguished from transient GPS noise. Speed-based trajectory filtering using a 30 km/h threshold distinguishes genuine stays from passing trajectories, preventing misclassification of highway travel as activity locations. We implemented H3 hexagonal indexing (resolution 10) for efficient spatial clustering and applied regional boundary validation (33.5–34.5° N, –118.8 to –118.0° W) to ensure LA County geographic focus (Uber Engineering, 2018). Additional quality filters retained only trajectories with minimum 33% time-of-day coverage and more than one activity per day, ensuring adequate daily activity representation.

Our large-scale stay point extraction employs a multi-stage processing pipeline detailed in Algorithm 4. The algorithm processes raw GPS trajectories through temporal clustering, spatial validation, and quality assurance to produce high-fidelity activity locations.

Following the extraction pipeline, we conducted comprehensive statistical validation of the processed stay points to ensure data quality and behavioral plausibility. The analysis reveals robust spatial and temporal patterns consistent with urban mobility characteristics.

The extracted stay points demonstrate realistic spatial clustering patterns across LA County, with higher density concentrations in urban centers and employment hubs. Geographic density analysis, as shown in Fig. 4, exhibits expected urban mobility patterns with concentrated activity in downtown LA, Hollywood, Santa Monica, and major employment corridors, which corresponds with higher population density associated with these areas. In the northern part of LA County, such as Lancaster, the activity density is relatively lower than the population density, because these places are typically underdeveloped residential areas that provide less employment and attractions.

Statistical analysis of stay point characteristics reveals well-distributed activity patterns throughout the week. As shown in Fig. 5, the average daily activity frequency shows optimal coverage (mean: 64.0%, median: 61.5%) exceeding our minimum 33% threshold, ensuring adequate temporal representation for reconstruction modeling. Activity duration analysis shows a mean stay duration of 3.6 h, indicating substantial dwell times consistent with meaningful activity engagement rather than transient GPS noise. Activity type duration distributions follow expected patterns with shorter durations for transitional activities (shopping, meals) and longer durations for anchoring activities (home, work).

The H3-based spatial clustering effectively captured regional mobility patterns while maintaining computational efficiency. Speed-based filtering successfully distinguished 89% of stay points from passing trajectories, preventing misclassification of

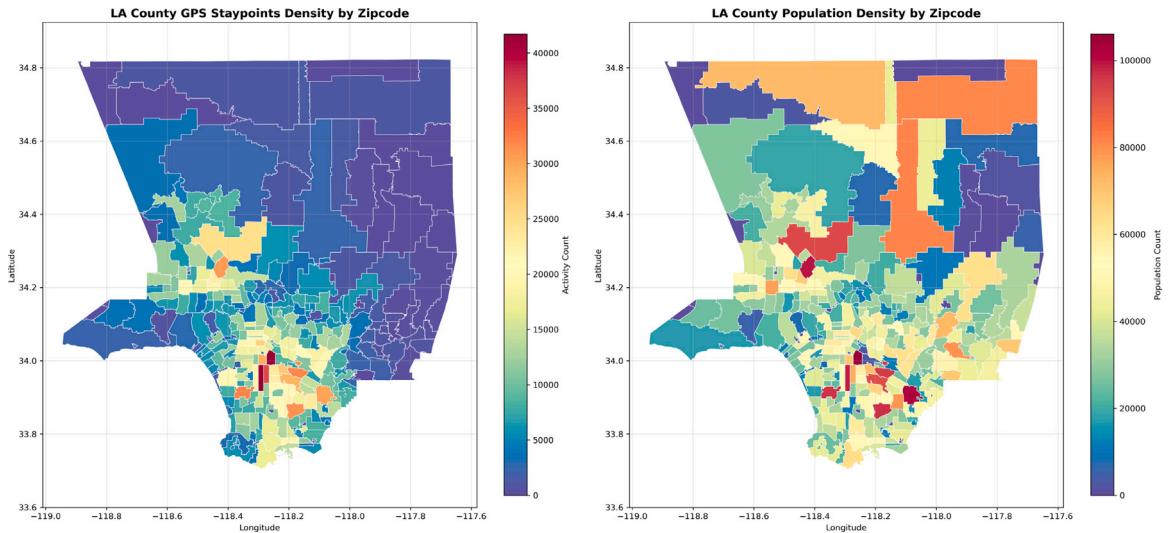


Fig. 4. GPS staypoints and population density of LA county.

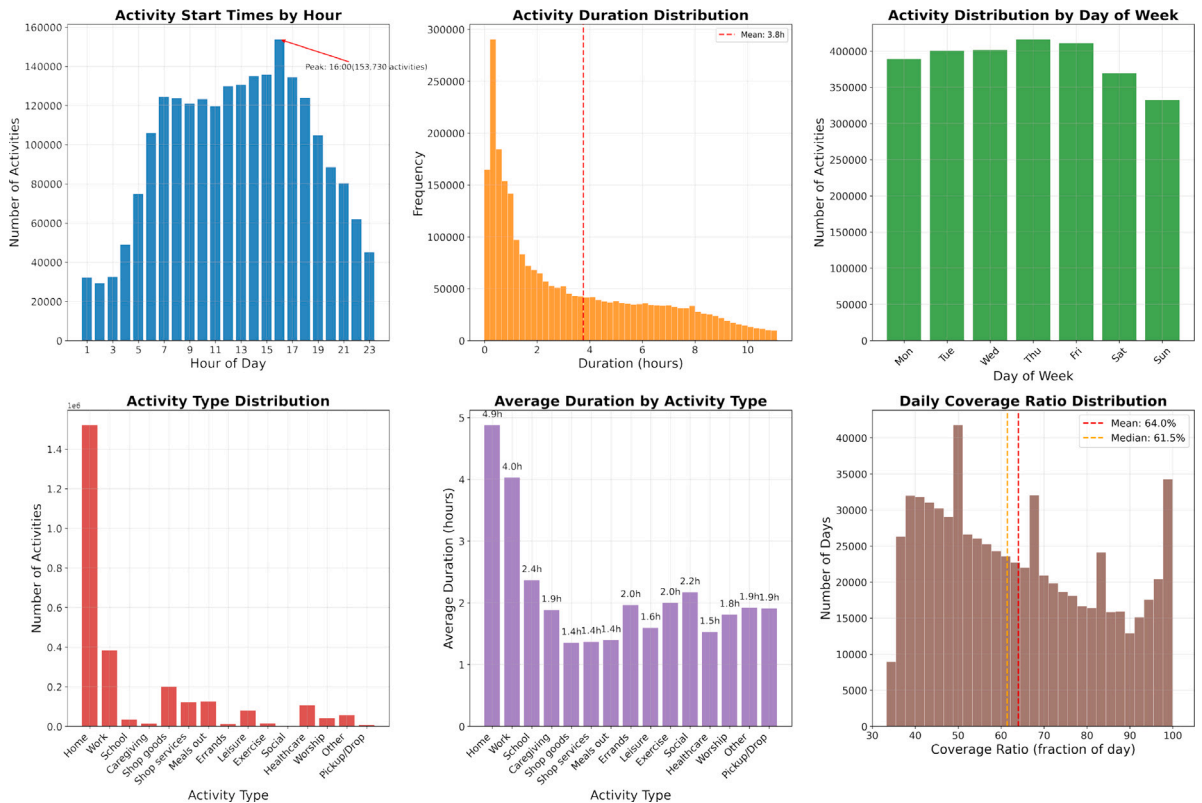


Fig. 5. Comprehensive statistical analysis of processed GPS trajectories.

highway travel as activity locations. Regional boundary validation ensured 100% geographic consistency within LA County bounds, supporting accurate regional modeling requirements for downstream transportation analysis.

Following preprocessing, we acquired approximately 1 million high-quality daily stay point trajectories, aggregated by individual agents to form complete daily activity patterns suitable for downstream reconstruction modeling.

We fused the daily stay points and POI data following our methodology (Section 4.1.1), which combines POI classification using Large Language Models with a Bayesian-based activity inference algorithm to identify the most probable activity type for each stay point. Examples of POI labeling and annotated trajectories are shown in Figs. 6 and 7.

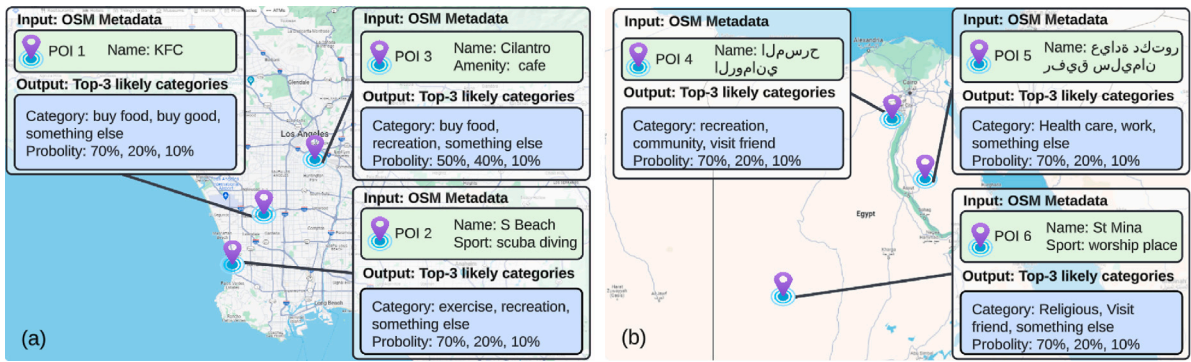


Fig. 6. POI labeling using LLM by associating three activity types with likelihood estimation (a) Los Angeles, USA (b) Cairo, Egypt (Liu et al., 2024).

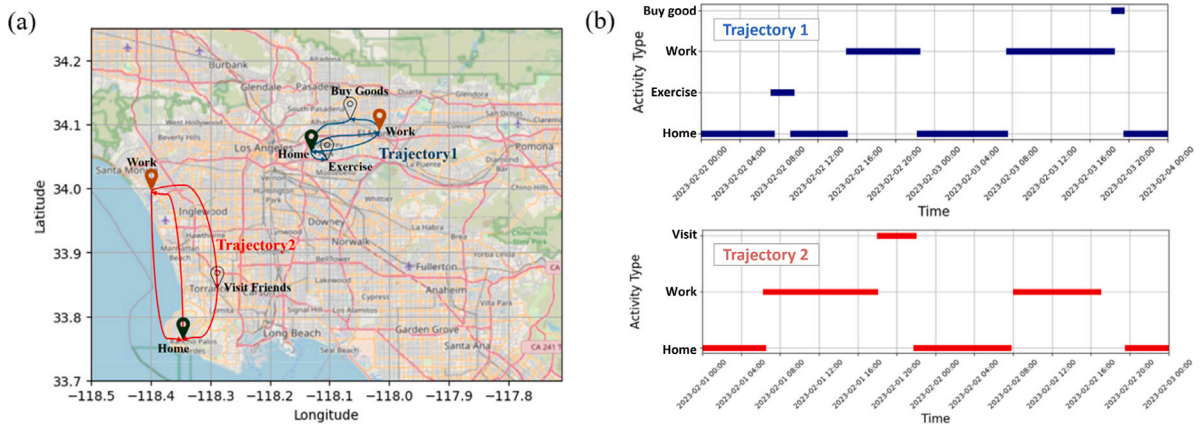


Fig. 7. Semantic annotated trajectories with activity type (Liu et al., 2024).

We constructed a representative synthetic population for Los Angeles County using a rule-based classification approach. Through analysis of our annotated trajectories, approximately 70% contain work-related activities, aligning with LA County’s employment rate (US Bureau of Labor Statistics, 2023). We leveraged distinctive activity patterns as classification rules — trajectories containing work activities were classified as belonging to workers, while those with school activities were strongly indicative of young demographics. This rule-based approach, rather than learning-based methods, was deliberately chosen to minimize error propagation in upstream processing stages that could compound into larger biases in downstream tasks. For trajectories where simple rules could not provide confident classification, we maintained them as unclassified rather than risk misattribution.

For base model training, we utilized the Southern California Association of Governments (SCAG) Activity-Based Model (ABM) dataset, which provides comprehensive synthetic activity chains for the Los Angeles region. The SCAG ABM integrates various activity-related choice models spanning both long-term and short-term decisions, generating detailed information about activity types, sequences, and temporal patterns (Liao et al., 2024b). We selected this dataset over alternatives like NHTS (National Household Survey Data) (Section 5.1) due to its substantial size – 700,000 samples for training, 200,000 for validation, and 100,000 for testing – which is crucial for effectively training transformer-based architectures.

For demographic adapter training, to establish target distributions  $D_g$  for each demographic group as defined in Section 4.2.3, we preprocessed Time Use Survey data at the individual survey level (Section 5.3). For LA County, each TUS entry contains detailed time allocation across daily activities for individual respondents. We mapped activity types from the original TUS classification system to our standardized 15-activity taxonomy used throughout the framework. Time use data was then aggregated to the demographic group level, computing the percentage of time allocated to each activity category. This preprocessing transforms raw survey responses into probability distributions  $D_g \in \mathbb{R}^{17}$  representing typical activity patterns for each of the 12 demographic groups (3 age groups  $\times$  2 gender categories  $\times$  2 employment statuses). These target distributions serve as targets for the demographic adapter optimization process described in Algorithm 2.

6.1.2. Evaluation metrics

We employ Jensen–Shannon Divergence (JSD) (Luca et al., 2021) as our primary metric to evaluate distributional similarities across multiple dimensions: activity chain length, duration, type, and temporal patterns (start/end times), as presented in Eq. (11).

Lower JSD values indicate better alignment. For spatial validation, we use cosine similarity for Origin–Destination matrices and Mean Absolute Percentage Error (MAPE) for traffic metrics.

$$JSD(P \parallel Q) = \frac{1}{2} \sum_{x \in X} \left[ P(x) \log \left( \frac{P(x)}{M(x)} \right) \right] + \frac{1}{2} \sum_{x \in X} \left[ Q(x) \log \left( \frac{Q(x)}{M(x)} \right) \right] \quad (11)$$

### 6.1.3. Training procedure

The training procedure follows a carefully structured three-phase protocol:

1. Initial Warm-up Phase: The model trains on unmasked data for 10% of total epochs, establishing baseline understanding of complete activity chains.
2. Intermediate Phase: Training continues with 40% masked data for the next 30% of total epochs, introducing the model to partial data scenarios.
3. Final Phase: The remaining epochs use 70% masked activity chains, preparing the model for heavily incomplete real-world data.

The training for activity reconstruction (Section 4.1.2), knowledge adapter (Section 4.2.1), and pattern adapter (Section 4.2.2) was conducted on an NVIDIA L40S GPU. As all three tasks share a similar network structure, the batch size was uniformly set to 512, with a learning rate of 0.001. To ensure stable training, we applied regularization techniques including dropout, L2 regularization, and early stopping. The model was trained up to 120 epochs using a combined loss function comprising cross-entropy, transition, and dynamic time warping losses, weighted at 0.7, 0.15, and 0.15, respectively.

The demographic adapter (Section 4.2.3) training was implemented on an NVIDIA RTX 4090 GPU, requiring approximately 10 h of computation time, averaging 40 min per demographic group. Detailed training specifications include: learning rate of 0.01, exponential reward function, and adapter initialization with all weights set to 1.0. Weight constraints maintain a minimum threshold of 0.01 to prevent negative values and apply GradNorm to dynamically handle the loss term weighting, optimizing across 15 activity categories. The training configuration employed a maximum of 1000 epochs for convergence, a batch size of 1500 samples, and a convergence threshold of 0.01 for early stopping. Data processing utilized 16,000 preprocessed activity sequences with 96 time slots (15 min intervals over 24 h). The masking strategy applied continuous segment masking with 70% coverage, 8–20 time steps per masked segment (maximum 8 segments), and 60% random activity masking during training. Best model selection was based on the lowest RMSE versus the target distribution from TUS across all training epochs, with automatic early stopping upon convergence detection. This relatively efficient training time, despite the complexity of demographic-specific pattern learning, demonstrates the practical feasibility of our approach for real-world applications.

After training the demographic-aware adapters, we integrated them with the transferred reconstruction model to create a comprehensive system capable of generating both demographically appropriate and temporally coherent activity patterns. This integrated model preserves the base model’s capabilities in temporal pattern recognition and activity dependencies while incorporating demographic sensitivity through the trained adapters.

Using the trained demographic-aware model, we conducted large-scale inference to generate synthetic activity patterns using NVIDIA L40S GPUs. We handled different trajectory data using different masking methods. For trajectories with relatively low loss or high completeness, we could often infer their socio-demographic information. In such cases, the problem became one of completion: we simply matched the corresponding socio-demographic info to the trajectory and used the model to fill in the missing data. However, for data with low completeness, it was not possible to infer any socio-demographic information. We treated this as a generation task, masking as much data as possible – up to 80% of the time in a day – and assigning it a socio-demographic profile. The model then generated a trajectory consistent with this socio-demographic information.

The model generated activity chains through multiple passes – 5 passes for weekdays and 2 passes for weekends – with each pass using different random masks to ensure diversity in the generated patterns. This process allowed the model to explore various possible activity combinations while maintaining demographic consistency. The generation process incorporated both demographic factors and variations in the day of the week to ensure realistic activity patterns between different segments of the population and throughout the week.

## 6.2. Model performance analysis

### 6.2.1. LLM-based POI classification and assignment validation:

We conducted comprehensive validation using ground truth data from both Los Angeles County and Egypt. Our evaluation reveals the critical advantages of LLM approaches in handling real-world data challenges.

We randomly sampled 500 POI observations from each region and created manual ground truth labels through five independent human annotators following uniform standards. This process enables quantitative assessment of our LLM-based classification accuracy against human judgment.

Using GPT-3.5 as our primary LLM, we achieved 93.4% accuracy and 96.1% F-1 score for Egypt dataset, and 91.4% accuracy with 92.9% F-1 score for LA County dataset. The model’s Hit@1 rates (where the top prediction matches ground truth) reached 74.7% for Egypt and 75.7% for LA, demonstrating consistent performance across diverse geographic and linguistic contexts.

We also validated the POI assignment procedure using our Bayesian-based approach that integrates LLM-derived POI classifications with temporal activity patterns. Using the NHTS California Add-on dataset as ground truth, we created a test dataset of 362 individuals with 2007 activities, consisting of 1724 mandatory activities and 283 non-mandatory activities.

**Table 2**  
Performance metrics for POI classification with different approaches.

Metric	Egypt (gpt3.5)	Egypt (gpt4)	Egypt (Heuristic)	LA (gpt3.5)	LA (gpt4)	LA (Heuristic)
Accuracy@1	90.3%	93.4%	64.3%	82.5%	91.4%	49.4%
Accuracy@2	22.6%	9.7%	4.4%	15.4%	8.7%	0.4%
Accuracy@3	6.1%	9.0%	0.0%	1.9%	7.0%	0.0%
F-1 Score	91.5%	96.1%	65.7%	82.1%	92.9%	51.7%

Our activity inference validation demonstrates strong overall performance with 91.7% accuracy and 92.3% F-1 score. The model shows particularly robust performance for mandatory activities (home, work, school), achieving 98.3%, 87.2%, and 90.0% accuracy, respectively, validating the LLM’s ability to classify POIs associated with these fundamental daily activities correctly. Non-mandatory activities show more varied performance, with categories like healthcare (100% accuracy) and religious activities (100% accuracy) performing excellently, while more ambiguous activities like “visit friends” (52.2% accuracy) prove more challenging due to their context-dependent nature.

To assess real-world robustness, we conducted comprehensive noise tolerance testing by introducing Gaussian noise at three levels (5 m, 10 m, and 20 m standard deviation) to simulate GPS inaccuracies common in urban environments. The model maintains strong stability across noise levels, with Accuracy@3 remaining above 84% even at 20 m noise levels (88.4% at 5 m, 85.3% at 10 m, 84.2% at 20 m), demonstrating the framework’s resilience to typical GPS positioning errors encountered in real-world deployment scenarios.

The Bayesian inference component effectively combines LLM-derived POI activity probabilities with temporal likelihood distributions derived from survey data, enabling context-aware activity classification. This dual validation approach – both direct POI classification and end-to-end activity inference – provides comprehensive evidence of the LLM framework’s effectiveness in handling the complete semantic trajectory annotation pipeline from raw GPS coordinates to meaningful activity sequences.

For baseline comparison, we implemented a rule-based heuristic POI classification method that maps OSM tags such as amenity, shop, tourism, leisure, etc., to the 15 activity type codes using a fixed priority and mapping scheme. This approach assigns each POI up to three most likely activity types based on the presence and priority of OSM tags, and evaluates performance using top-1, top-2, and top-3 accuracy as well as the weighted F1 score. As shown in Table 2, the heuristic method achieves an Accuracy@1 of 64.3% for Egypt and 49.4% for LA County, which is substantially lower than the LLM-based approaches (gpt-3.5 and gpt-4). The heuristic’s Accuracy@2 and @3 are also much lower, indicating limited ability to capture ambiguous or multi-purpose POIs. The F1 scores further highlight the gap, with the heuristic baseline reaching only 65.7% (Egypt) and 51.7% (LA), compared to over 90% for LLMs. These results demonstrate that while the heuristic method provides a transparent and interpretable baseline, it is significantly outperformed by LLM-based classification, especially in handling complex or ambiguous POI categories.

### 6.2.2. Base model evaluation on survey data

The performance of the base model  $M^N$  in reconstructing 70% masked NHTS data and its transferability to the SCAG datasets is quantified in Table 3. The model achieves strong performance on NHTS data, with JSD values ranging from 0.001 to 0.017 across all metrics. However, directly applying the model to SCAG data results in performance degradation indicated by increasing JSD value, revealing significant regional differences in mobility patterns. The JSD values between NHTS and SCAG, which range from 0.009 to 0.059 across various activity attributes, quantify these differences and are visualized in Fig. 8. Supervised transfer learning is applied to  $M^N$  to obtain  $M^S$ , yielding low JSD values when reconstructing the SCAG dataset. This result shows the algorithm effectively adapts to different geographical contexts and data formats, despite initial disparities between SCAG and NHTS datasets.

Beyond its commendable performance in system-level JSD evaluations, the proposed model exhibits proficiency in capturing patterns of common activities, with the transfer learning successfully adapting the model from regional to national contexts. As illustrated in Fig. 9, the model demonstrates remarkable accuracy in reconstructing the temporal patterns of activity start times across both datasets.

Essential activities such as home, work, and school show closely aligned patterns between the two datasets, evidenced by low JSD values in Table 3 (0.023, 0.021, and 0.024, respectively), highlighting the universality of these routines. Conversely, regional variations are apparent in other activities. Shopping patterns, for example, differ significantly, with SCAG data showing a midday peak (JSD 0.054) versus NHTS’s more even distribution, suggesting distinct regional shopping behaviors. Exercise activities also vary, with NHTS indicating a preference for early morning sessions (JSD 0.036). These differences likely mirror varying lifestyle choices between urban Los Angeles and the broader national context. The model’s consistently low  $M^S$  and  $M^N$  values, typically below 0.01, confirm its efficacy in capturing these nuances, underscoring its utility for urban planning and policy-making by identifying both universal and regional behavioral patterns.

### 6.2.3. Transfer learning evaluation on GPS and survey data

JSD values across iterations for the evaluation metrics are shown in Table 4, showing how the iterative transfer learning process improves alignment with the target region. In this table, only three metrics are included because the lack of ground truth in incomplete data renders activity duration and chain length unreliable for comparison, whereas the selected metrics are more accurate at the system level and can better reflect regional mobility patterns. All values reported for Iteration 1 through Iteration 6 are the mean  $\pm$  standard deviation over three random seeds. Activity type JSD drops from 0.0508 at baseline to 0.0028 by Iteration 3.

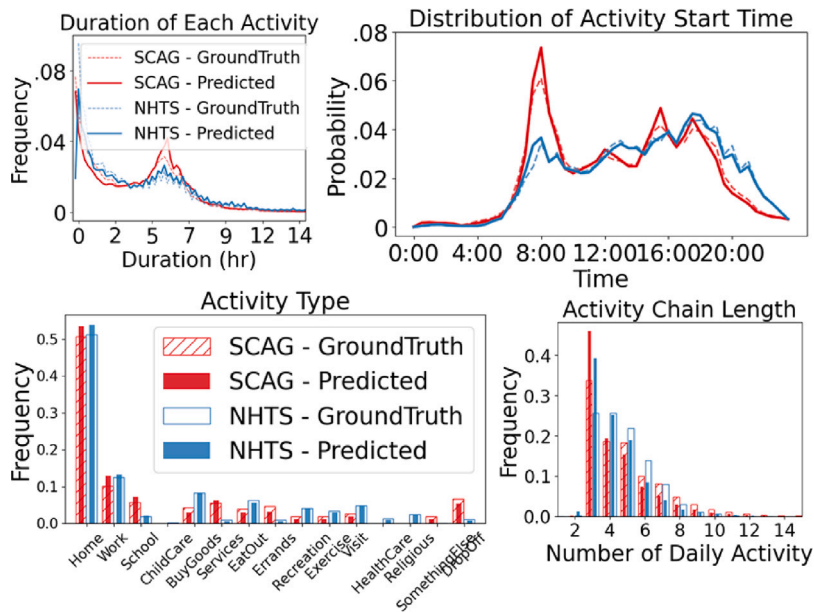


Fig. 8. Performance evaluation on NHTS and SCAG with 70% masked activity chain.

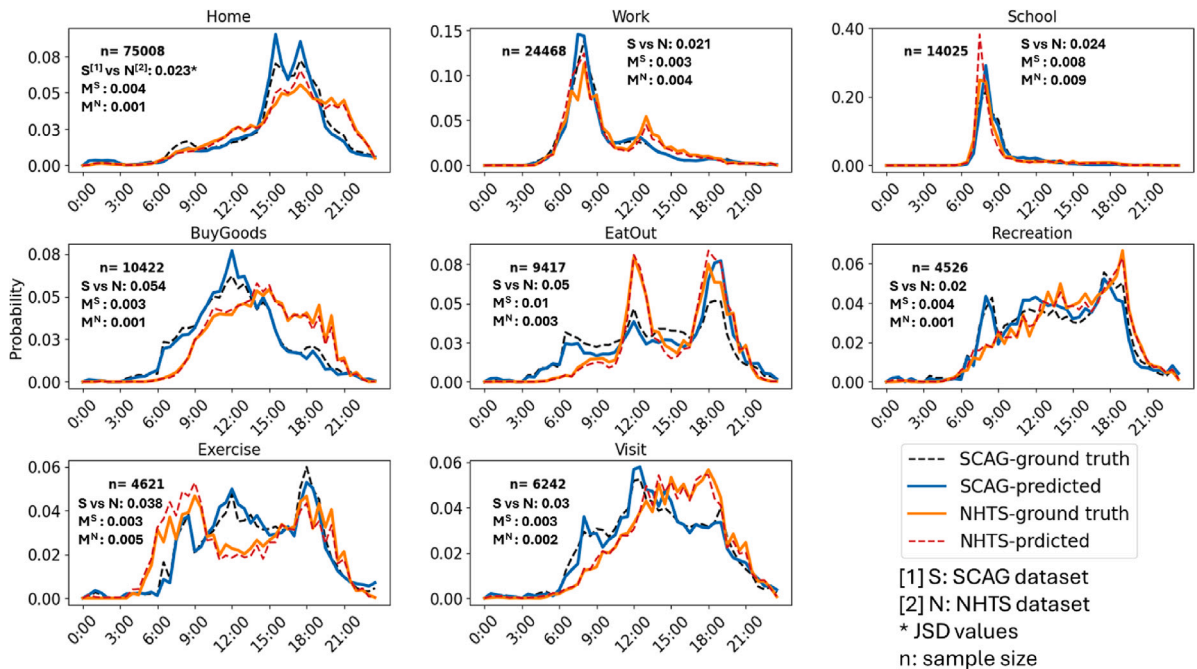


Fig. 9. The start time of common activities reconstructed by base and its transferred model using NHTS and SCAG.

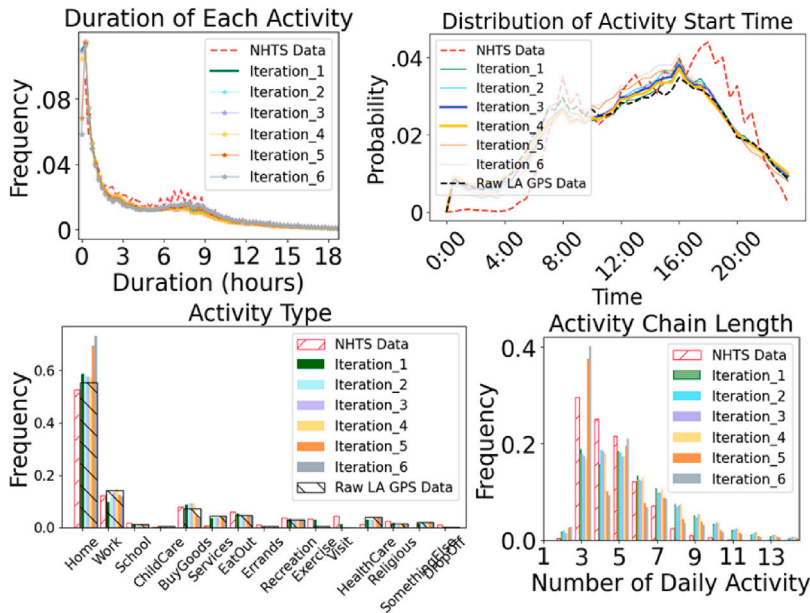
Table 3

JSD value for similarity measurement across datasets and model performance.

JSD values	Length	Duration	Type	Start time	End time
NHTS vs SCAG	0.017	0.009	0.059	0.018	0.013
$M^N$ on NHTS	0.017	0.002	0.001	0.001	0.005
$M^N$ on SCAG	0.021	0.003	0.017	0.004	0.007
$M^S$ on SCAG	0.011	0.004	0.004	0.001	0.001

**Table 4**  
Model evolution during iterative transfer learning for LA case.

JSD values	Activity type	Start time	End time
NHTS vs LA GPS	0.0508	0.0289	0.0367
Iteration_1	0.0131 ± 0.0033	0.0018 ± 0.0006	0.0044 ± 0.0007
Iteration_2	0.0041 ± 0.0011	0.0015 ± 0.0004	0.0042 ± 0.0006
<b>Iteration_3</b>	<b>0.0028 ± 0.0004</b>	0.0010 ± 0.0002	0.0030 ± 0.0003
<b>Iteration_4</b>	0.0030 ± 0.0003	<b>0.0008 ± 0.0001</b>	<b>0.0022 ± 0.0003</b>
Iteration_5	0.0052 ± 0.0008	0.0027 ± 0.0005	0.0054 ± 0.0007
Iteration_6	0.0081 ± 0.0012	0.0050 ± 0.0009	0.0085 ± 0.0011



**Fig. 10.** Iterative adaptation from NHTS to LA via semi-supervised transfer learning with LA GPS data.

Similarly, start time and end time JSDs reach their lowest values, 0.0008 and 0.0022, respectively, at Iteration 4. Model performance deteriorates beyond Iterations 3 and 4, which appear to strike the best balance for reducing dissimilarities and adapting the model. This decline aligns with findings that AI models collapse when trained on recursively generated data (Shumailov et al., 2024), as over-reliance on synthetic data erodes nuanced or rare activities in the original data. Fig. 10 demonstrates how the progressive adaptation of the NHTS-based model to the LA context evolves over multiple iterations via the proposed Semi-Supervised Transfer Learning using LA GPS data.

The second step of transfer learning is performing the demographic adaptation by fusing TUS data. We trained 12 demographic adapters in accordance with sex (male/female), age (young, mid-aged, elderly), and employment status (worker, non-worker). Analysis of the aggregated Time Use Survey distributions (Table 6) reveals distinct activity patterns across demographic groups with aspects of sex, age, and employment status.

The training patterns reveal interesting demographic-based variations in adaptation difficulty. Worker groups generally achieve higher reward values (0.95–0.98) and converge more quickly, typically within 400–500 epochs for Mid-aged and elderly workers. This faster convergence likely reflects more structured and consistent daily patterns among employed individuals. Young workers require more training epochs (> 1000) but ultimately achieve similarly high reward values, suggesting more complex activity patterns that require additional learning time. Non-worker groups present greater challenges in pattern adaptation, reflected in slightly lower reward values (0.82–0.93) and longer training times. This increased difficulty in modeling non-worker patterns aligns with the greater variability in their daily activities observed in the target distributions.

We evaluate our demographic transfer learning module by first comparing generated activity patterns across demographic groups with both TUS and HTS data (Table 5). Note that these JSDs measure alignment at the demographic group level, where matching patterns becomes inherently more challenging than population-level alignment due to the need to capture group-specific behavioral variations. The results demonstrate effective demographic adaptation, with output patterns showing strong alignment with TUS data (JSD: 0.006–0.018) across all demographic groups. This close alignment with TUS validates our adapter’s ability to learn demographic-specific patterns. The larger differences between output and HTS patterns (JSD: 0.048–0.094) are expected and reasonable, as they reflect the inherent differences between TUS and HTS data (TUS vs HTS JSD: 0.011–0.063). This suggests our framework successfully adapts to TUS patterns while maintaining reasonable consistency with HTS characteristics.

**Table 5**  
JSD value for similarity measurement of activity duration distribution across demographic group.

Group	HTS vs Output	TUS vs Output	TUS vs HTS
Young M	.083/.086	.008/.009	.014/.014
Young F	.094/.085	.015/.013	.016/.011
Mid-aged M	.048/.067	.018/.006	.016/.054
Mid-aged F	.055/.066	.018/.007	.022/.059
Elder M	.070/.066	.015/.009	.020/.055
Elder F	.077/.071	.013/.009	.029/.063

Note: Values shown as (Worker/Non-worker); M: Male F: Female

**Table 6**  
L.A. County major activity duration (minutes/day) comparison with JSD.

Demographic group (JSD)	Activity duration (minutes)							
	Home (1)	Work (2)	School (3)	Buy goods (5)	Buy meals (7)	Recreational (9)	Exercise (10)	Visit friends (11)
Young Male W. (0.0085)	843.7/837.5	209.5/224.5	117.7/134.7	19.6/15.0	52.5/38.5	100.7/53.3	23.8/18.6	61.5/53.4
Young Female W. (0.0159)	820.2/785.3	176.2/204.7	131.5/134.8	39.8/20.8	57.3/34.3	70.2/34.6	23.5/18.8	75.8/88.8
Mid-Aged Male W. (0.0189)	798.4/800.4	380.8/382.9	0.0/0.0	37.9/19.1	61.7/77.0	76.2/62.2	23.9/4.7	42.4/1.0
Mid-Aged Female W. (0.0185)	839.0/825.6	333.0/330.8	0.0/0.0	42.4/22.7	57.0/74.3	63.2/58.7	16.8/4.4	46.1/2.1
Elderly Male W. (0.0157)	875.0/886.0	295.6/284.5	0.0/0.0	34.8/30.0	69.2/77.7	83.2/61.2	20.5/4.4	35.6/0.8
Elderly Female W. (0.0136)	913.2/902.4	245.7/260.2	0.0/0.0	41.4/35.6	62.3/67.6	84.7/60.8	15.0/4.0	37.0/0.8
Young Male N. (0.0090)	903.8/960.4	0.0/0.0	170.7/203.1	21.3/14.5	54.0/68.7	97.7/84.6	57.9/41.1	52.1/44.7
Young Female N. (0.0103)	899.9/987.7	0.0/0.0	170.5/149.2	36.0/21.8	59.9/75.2	79.7/76.9	31.5/39.9	49.2/40.0
Mid-Aged Male N. (0.0064)	1013.8/1054.2	0.0/0.0	18.5/3.8	40.6/43.0	57.2/69.0	127.3/118.2	27.0/37.2	45.6/32.7
Mid-Aged Female N. (0.0078)	994.0/1083.5	0.0/0.0	11.9/3.9	44.2/22.5	58.3/64.7	90.6/67.7	16.9/11.3	54.4/57.2
Elderly Male N. (0.0092)	1022.0/1101.4	0.0/0.0	1.6/3.3	42.1/43.8	70.3/80.5	144.4/150.7	25.0/5.8	41.6/22.4
Elderly Female N. (0.0099)	1037.1/1115.6	0.0/0.0	1.6/3.3	37.4/41.4	63.5/88.9	120.8/128.9	15.1/4.7	46.9/31.5

Note: Values shown as TUS/Output. W: Worker, N: Non-worker. JSD values indicate distribution similarity (lower JSD means better result).

The generation demonstrates strong alignment with the Time Use Survey across all demographic groups. Quantitative evaluation using JSD shows excellent performance (detailed results in Table 6), with non-worker groups achieving particularly strong alignment (JSD 0.0064–0.0103) and worker groups showing comparable performance (JSD 0.0085–0.0189). Mid-aged male non-workers achieve the best performance (JSD = 0.0064) while maintaining realistic activity patterns, such as work duration differences across age groups and gender-specific activity preferences.

Beyond activity durations, we evaluate temporal patterns to ensure the demographic adapter preserves learned temporal dynamics while incorporating demographic variations. The temporal evaluation spans demographic groups (Fig. 11) and activity-specific distributions (Fig. 12). At the demographic level, mid-aged and elderly groups show consistent patterns across both worker and non-worker categories. The young group shows some discrepancies between TUS and HTS because the young worker group is insignificant in HTS (around 2% of the total population).

At the aggregate level, the model maintains strong agreement with both raw GPS data (JSD: 0.004) and NHTS patterns (JSD: 0.014), capturing key features like morning (8:00–10:00) and afternoon (16:00–18:00) peaks. At the activity level, major activities like work and school are well captured and the patterns learned from the base model were maintained throughout the complex transfer learning process. Notably, our generated patterns and GPS data capture early morning activities (before 6 AM) that are absent in HTS data, addressing known survey limitations such as underreporting and rounding effects, for example, people typically fail to report some emergency travel activities and tend to make the pattern more routine (Bricka et al., 2012). This demonstrates our framework’s ability to balance GPS and survey data characteristics while maintaining demographic representativeness.

At the activity level, as shown in Fig. 12, major activities like work and school are well captured, and the patterns learned from the base model were maintained throughout the complex transfer learning process. The model maintains high fidelity across both major mandatory and non-mandatory activities, with average JSD values consistently below 0.06. Notably, our generated patterns and GPS data capture early morning activities (before 6 AM) that are absent in HTS data, addressing known survey limitations such as underreporting and rounding effects, for example, people typically fail to report some emergency travel activities and tend to make the pattern more routine (Bricka et al., 2012). This demonstrates our framework’s ability to balance GPS and survey data characteristics while maintaining demographic representativeness.

### 6.3. Spatial validation

We utilized the generated output from the prior pipeline to assign locations in the context of LA County using Algorithm 3, with a spatial division of eight sub-regions (Fig. 13(a)). The model was first calibrated using 100,000 agents from the SCAG Activity-Based Model dataset, then scaled to generate mobility patterns for 1 million agents.

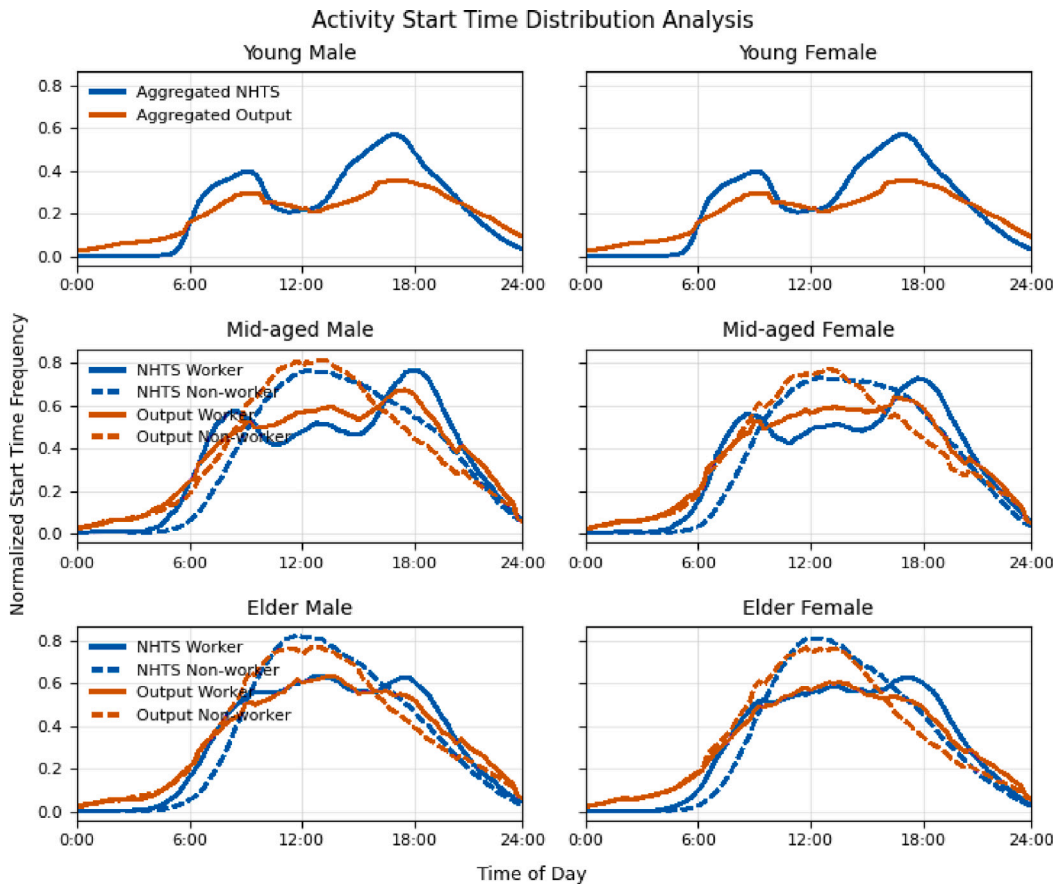


Fig. 11. Temporal performance evaluation per demographic group based on. Note: HTS lacks a Young Worker group; comparison for young individuals ignores employment status.

To ensure the reliability of our validation results, we implemented a rigorous three-stage calibration process following the methodology established in He et al. (2024). First, we calibrated the MATSim-based transportation network using an iterative proportional adjustment algorithm that minimizes the sum of squared differences between simulated and observed traffic volumes across 100+ Caltrans PeMS detector stations on the top 10 highest-traffic freeways in LA County. The calibration algorithm iteratively adjusts link capacity factors starting from initial values of 0.8, continuing until convergence criteria are met (typically 6 iterations achieving < 5% relative difference across time periods). Second, we validated spatial trip distributions against SCAG ABM benchmarks, achieving cosine similarity of 0.997 for origin–destination matrices. Third, we conducted temporal validation using real-world traffic patterns, confirming realistic peak-hour dynamics and congestion onset/dissipation patterns.

The spatial distribution of activities across sub-regions demonstrates the effectiveness of our location assignment method (Fig. 13(a)). The high cosine similarity (0.997) between the origin–destination matrices of our model and SCAG ABM validates the preservation of regional flow patterns.

Traffic simulation on the LA transportation network (Fig. 13(b)) reveals strong system-level performance. The hourly VMT and traffic speed patterns closely match the benchmark SCAG ABM results, with MAPEs of 4.97 and 1.16 respectively (Fig. 14(b)). These metrics indicate accurate reproduction of aggregate travel patterns.

We conducted detailed analysis on a major segment of Interstate 405, comparing our results with both SCAG ABM predictions and PeMS real-world observations. The analysis reveals that directional variations in congestion patterns are well-captured. For northbound traffic, we observed a MAPE of 5.85 for volume and 4.45 for speed. Similarly, for southbound traffic, we measured a MAPE of 9.32 for volume and 4.36 for speed. The corridor-level results (Fig. 15) demonstrate the model’s capability to reproduce fine-grained traffic dynamics, including peak-hour congestion and directional flow patterns.

## 7. Egypt case study

### 7.1. Regional adaptation in data-scarce contexts

To evaluate the regional adaptability of our proposed framework, we conducted a case study on Egypt, a data-scarce region. While the Los Angeles case demonstrated the framework’s ability to reflect demographic and temporal characteristics of human

Activity Start Time Distributions Comparison

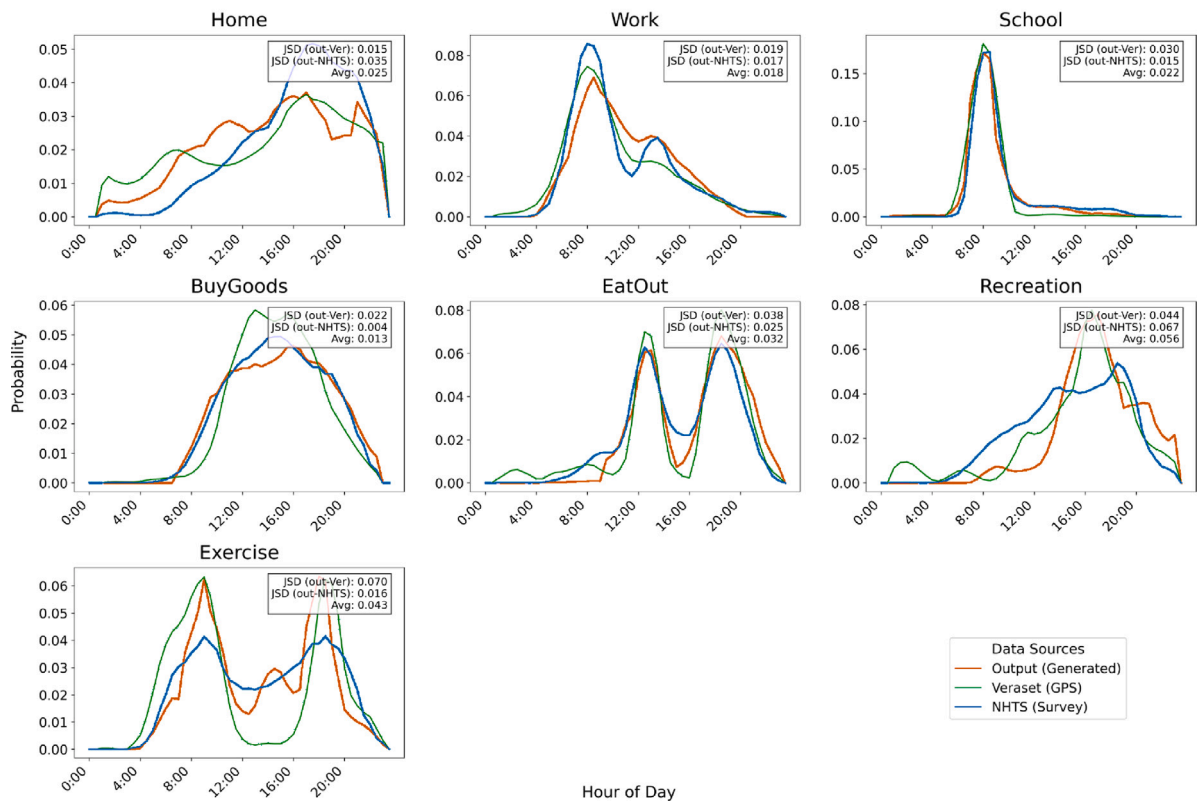


Fig. 12. Activity-specific temporal pattern comparison: start time distributions for major activities across different data sources.

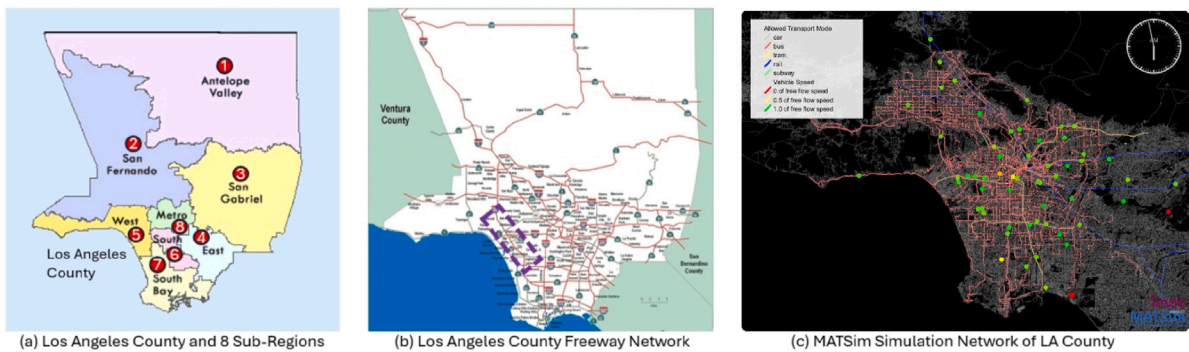


Fig. 13. LA county map and freeway network.

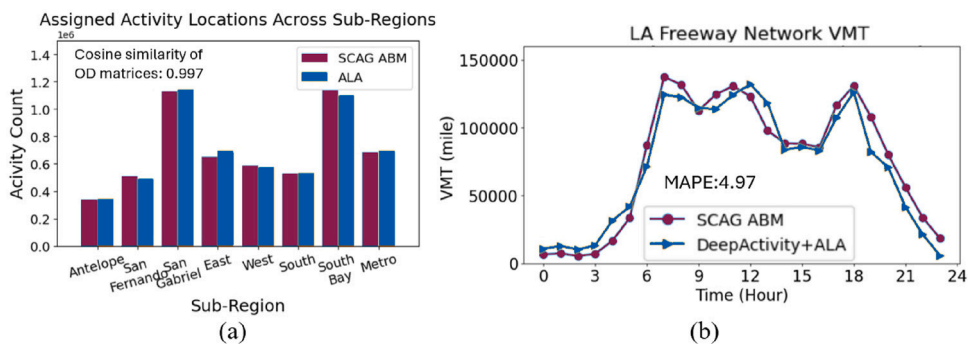


Fig. 14. Validation for location assignment and traffic loading at network level.

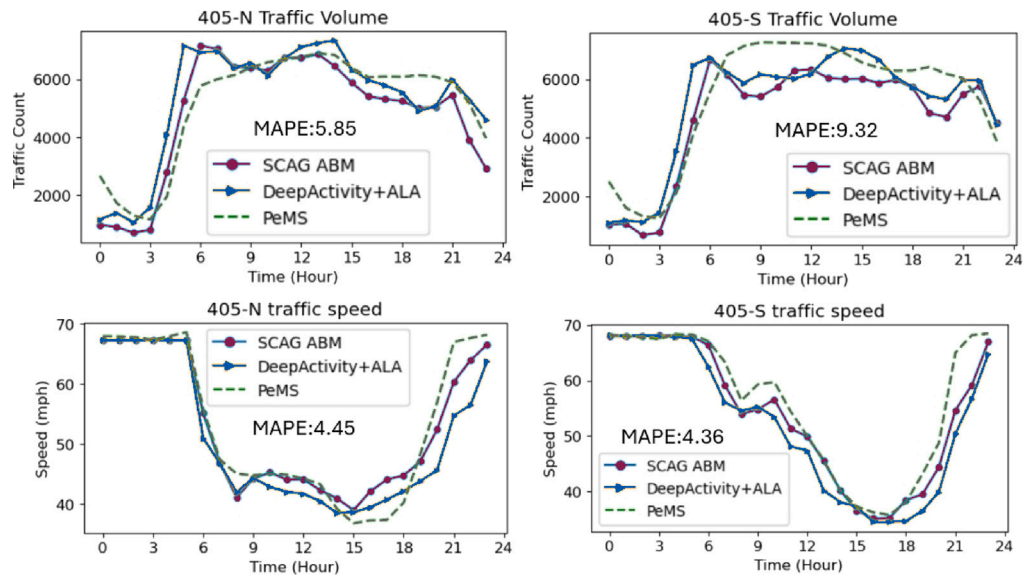


Fig. 15. Validation for location assignment and traffic loading at corridor level.

Table 7

Model evolution during iterative transfer learning.

JSD values	Activity type	Start time	End time
NHTS vs Egypt GPS (baseline)	0.0255	0.0637	0.0433
Iteration_1	0.0055 ± 0.0009	0.0052 ± 0.0004	0.0027 ± 0.0005
Iteration_2	0.0038 ± 0.0006	0.0035 ± 0.0003	0.0021 ± 0.0004
Iteration_3	0.0039 ± 0.0004	<b>0.0025 ± 0.0002</b>	<b>0.0021 ± 0.0003</b>
Iteration_4	<b>0.0027 ± 0.0002</b>	0.0034 ± 0.0003	0.0029 ± 0.0004
Iteration_5	0.0043 ± 0.0005	0.0042 ± 0.0004	0.0045 ± 0.0006
Iteration_6	0.0029 ± 0.0003	0.0049 ± 0.0006	0.0046 ± 0.0005

mobility, the Egypt case emphasizes its capability to adapt to new cultural and regional contexts. We utilized a dataset of 479,526 samples, rigorously curated to ensure quality by retaining only agents with at least 7 days of recorded activity and sufficient daily temporal coverage. This dataset was partitioned into training (80%) and validation (20%) sets.

Demographic information was assigned to the dataset based on Egypt's demographic distribution, incorporating two dimensions: *age* and *sex*. Each agent was categorized into one of six demographic groups: young, mid-aged, or elderly males and females. The GPS trajectory annotation follows the same methodology. Notably, the LLM demonstrated robust cross-linguistic capabilities, accurately matching Arabic-language POIs with activity types. This enabled the integration of diverse POI data to enhance the semantic context of stay points, a crucial step in regions like Egypt where linguistic diversity adds complexity to POI interpretation.

## 7.2. Transfer learning and model training

To adapt the base model trained on the NHTS dataset, we employed transfer learning to the Egypt case scenario, leveraging the semi-supervised iterative refinement process 1. The adaptation process retained universal temporal patterns while progressively aligning with the distinct cultural and regional mobility characteristics of Egypt. We then applied the progressive freezing method during the training process, as mentioned in the methodology. Initially, only the MLP and embedding layers were fine-tuned, allowing for rapid adaptation of activity classifications and temporal representations. In subsequent iterations, the transformer layers were gradually unfrozen to refine the model's understanding of intermediate and complex activity dependencies. Its evolution through the semi-supervised transfer learning using Egypt GPS data is visualized and quantified in Fig. 16 and Table 7.

The Egypt TUS data was aggregated to derive target activity duration distributions across the six demographic groups. Adapter training followed the same computational configuration as the case of LA County Case Study

The inference process uses the same setup as the prior case study, generating synthetic activity chains that closely aligned with the target distributions. As shown in Table 8, the JSD values for all demographic groups were low (ranging from **0.0054** for mid-aged males to **0.0302** for mid-aged females), indicating strong alignment between the generated and target distributions. This underscores the framework's ability to replicate realistic activity patterns even in a data-scarce context.

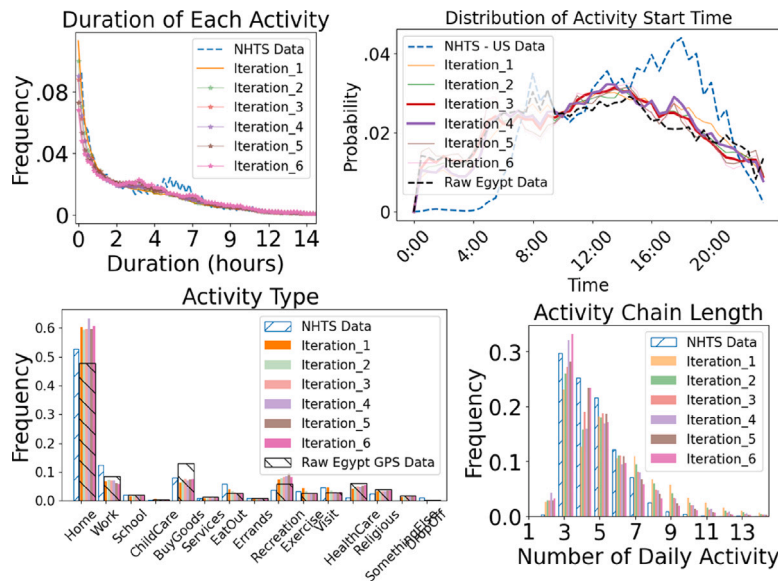


Fig. 16. Iterative adaptation from NHTS to Egypt via semi-supervised transfer learning with Egypt GPS data.

Table 8

Egypt major activity duration (minutes/day) comparison with JSD.

Demographic group (JSD)	Activity duration (minutes)							
	Home (1)	Work (2)	School (3)	Buy goods (5)	Buy meals (7)	Recreational (9)	Exercise (10)	Visit friends (11)
Young Male (0.0067)	842.2/790.6	95.0/90.7	347.6/308.2	4.2/15.8	15.0/1.4	33.5/38.9	43.0/31.7	32.4/56.2
Young Female (0.0172)	786.1/761.8	76.8/70.6	316.6/313.9	3.1/53.3	14.4/4.3	131.5/148.3	1.4/8.6	36.2/53.3
Mid-Aged Male (0.0054)	684.1/688.3	342.4/348.5	12.7/10.1	8.9/14.4	15.0/15.8	136.7/126.7	4.1/4.3	119.9/112.3
Mid-Aged Female (0.0302)	881.7/832.3	199.1/146.9	63.7/0.0	8.1/24.5	19.7/5.8	146.7/126.7	0.3/4.3	36.2/53.3
Elderly Male (0.0141)	752.7/764.6	80.2/56.2	2.2/2.9	4.6/18.7	17.4/47.5	186.9/197.3	1.3/5.8	135.4/108.0
Elderly Female (0.0198)	848.6/853.9	107.9/79.2	0.0/17.3	3.1/18.7	10.9/4.3	164.3/142.6	0.0/11.5	46.6/106.6

Note: Values shown as TUS/Output. Each demographic group’s JSD value indicates overall distribution similarity (lower is better).

### 7.3. Observations on generated patterns

The generated activity chains demonstrated high consistency with target distributions, particularly for mandatory activities. Mid-aged males exhibited significantly longer work durations compared to mid-aged females, reflecting the disparity in employment rates (65% for males vs. 18% for females, as reported by the UN). Work durations for females in Egypt were notably lower than those in Los Angeles, illustrating the framework’s sensitivity to cultural and socio-economic differences. Religious activities are prominent across all demographic groups, with peaks reflecting the significance of prayer times in Egyptian society. Recreational activities are observed to be less frequent, consistent with the conservative cultural context of Egypt compared to Western regions like Los Angeles.

The Egypt case study demonstrates the framework’s adaptability to data-scarce regions, effectively generating realistic activity chains aligned with local sociodemographic patterns and cultural contexts. By leveraging transfer learning, adapter training, and LLM-based semantic enrichment, the framework overcomes the challenges of sparse data and fragmented sources, showcasing its potential for mobility modeling in diverse global settings.

## 8. Discussion

### 8.1. Framework positioning and methodological considerations

Our framework represents the first cross-domain mobility foundation model that jointly addresses semantic trajectory enrichment, demographic-aware activity reconstruction, and cross-regional transfer learning. Unlike existing approaches that focus on either raw coordinate generation (Wang et al., 2024; Feng et al., 2020; Zhu et al., 2023) or trajectory recovery (Xia et al., 2021; Ren et al., 2021; Si et al., 2023), our three-stage pipeline unifies LLM-based semantic enrichment, transformer-driven reconstruction, and progressive transfer learning to generate comprehensive mobility datasets with activity semantics and demographic conditioning.

The absence of directly comparable baselines stems from our unique problem formulation that combines privacy-preserving data fusion with transferable demographic adaptation. While trajectory generation methods excel at spatiotemporal prediction, they lack the semantic labeling and cross-domain capabilities that form our core contribution. Our evaluation strategy appropriately emphasizes transfer learning effectiveness and real-world validation rather than traditional trajectory prediction metrics.

### 8.2. Limitations and bias considerations

**LLM-Based POI Classification Bias:** Our LLM approach (Section 4.1.1) demonstrates strong cross-linguistic performance (93.4% accuracy for Arabic POIs, 91.4% for English), but potential cultural biases remain from English-centric pre-training. Multi-functional POIs present semantic ambiguity challenges, which we address through probabilistic top-3 outputs rather than hard classifications. Future work should incorporate more diverse geographic training data and formal bias quantification methods.

**Multi-Source Data Disagreements:** Survey data (HTS/TUS) and GPS (Section 5) traces exhibit inherent contradictions due to underreporting of short trips, recall errors, and regional methodology differences (Stopher et al., 2010; Gruber and Stark, 2024). Our framework addresses these through demographic adapters that reconcile GPS precision with survey contextual richness, though generated patterns may lean toward GPS characteristics due to data volume dominance. The demographic scope (age, gender, employment) captures fundamental mobility variations while maintaining computational tractability for million-agent simulations, though income and household structure represent valuable future extensions.

**GPS Coverage Bias:** Commercial GPS data overrepresents younger, urban smartphone users, potentially excluding rural and low-income populations. Our mitigation strategies include census-guided demographic priors, large masking ratios (70%) during inference, and demographic reweighting, though systematic coverage gaps remain a limitation for comprehensive population representation.

### 8.3. Privacy and practical considerations

Our framework provides inherent privacy protection through complete synthetic data generation with no re-identifiable trajectory outputs and demographic aggregation, preventing individual identification. Following project completion, we plan to release de-identified synthetic datasets, complete model implementations, and evaluation protocols to facilitate community benchmarking and comparative studies.

The computational efficiency demonstrated enables practical deployment for transportation agencies, with scalability to 10-million-level megacities population through streamlined processing architectures.

### 8.4. Future research directions

Key areas for advancement include expanding demographic dimensions to incorporate income and household structure, developing lightweight adapter architectures for rapid regional customization, further enlarging the scale of trajectory generation, and implementing formal differential privacy guarantees. Cross-continental validation with diverse cultural contexts and integration of real-time mobility data streams represent additional opportunities for enhancing framework generalizability and practical impact.

## 9. Conclusion

This paper presents a framework that enables realistic mobility pattern generation through data fusion. Our approach addresses fundamental data challenges: semantically enriching and completing trajectories through LLM annotation and reconstruction, adapting to new regions through transfer learning, and incorporating demographic patterns without paired labels. We validate the generated patterns using real-world test cases through large-scale simulation and demonstrate strong performance. By fusing diverse data sources, our framework demonstrates the capability of overcoming individual data limitations. Future work could extend this framework to capture more complex social dynamics, including coordinated activities between individuals, large gathering events, and social network effects on mobility patterns, further enhancing the realism of generated trajectories.

### CRedit authorship contribution statement

**Haoxuan Ma:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Xishun Liao:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Yifan Liu:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Conceptualization. **Qinhua Jiang:** Validation, Methodology. **Chris Stanford:** Conceptualization. **Shangqing Cao:** Methodology. **Jiaqi Ma:** Writing – review & editing, Conceptualization.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported in part by the FHWA Center for Excellence on New Mobility and Automated Vehicles Program, and in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number 140D0423C0033. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

## Data availability

The datasets generated and/or analyzed during the current study are not publicly available due to privacy agreements with the data provider (Veraset) but are available from the corresponding author on reasonable request.

## References

- Afyouni, I., Al Aghbari, Z., Razack, R.A., 2022. Multi-feature, multi-modal, and multi-source social event detection: A comprehensive survey. *Inf. Fusion* 79, 279–308.
- Alexander, L., Jiang, S., Murga, M., González, M.C., 2015. Origin–destination trips by purpose and time of day inferred from mobile phone data. *Transp. Res. Part C: Emerg. Technol.* 58, 240–250.
- Anil, R., Dai, A.M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., et al., 2023. Palm 2 technical report. arXiv preprint arXiv:2305.10403.
- Bricka, S., Paleti, R., Bhat, C.R., 2012. An analysis of the factors influencing differences in survey-reported and GPS-recorded trips. *Transp. Res. Part C: Emerg. Technol.* 21, 67–78.
- Cai, T., Liu, Y., Zhou, Z., Ma, H., Zhao, S.Z., Wu, Z., Ma, J., 2024. Driving with regulation: Interpretable decision-making for autonomous vehicles with retrieval-augmented reasoning via LLM. arXiv preprint arXiv:2410.04759.
- Cao, J., Li, Q., Tu, W., Wang, F., 2019. Characterizing preferred motif choices and distance impacts. *Plos One* 14 (4), e0215242.
- Chen, G., Viana, A.C., Fiore, M., Sarraute, C., 2019. Complete trajectory reconstruction from sparse mobile phone data. *EPJ Data Sci.* 8 (1), 1–24.
- Federal Highway Administration, 2017. National household travel survey. <https://nhts.ornl.gov>.
- Feng, J., Yang, Z., Xu, F., Yu, H., Wang, M., Li, Y., 2020. Learning to simulate human mobility. In: *Proc. 26th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining*. pp. 3426–3434.
- González, M.C., Hidalgo, C.A., Barabási, A.L., 2008. Understanding individual human mobility patterns. *Nature* 453 (7196), 779–782.
- Gruber, C., Stark, J., 2024. Underreported trips, a non-negligible empirical effect of traditional survey methods. *Transp. Res. Procedia* 69, 312–320.
- Hasan, S., Schneider, C.M., Ukkusuri, S.V., González, M.C., 2013. Spatiotemporal patterns of urban human mobility. *J. Stat. Phys.* 151 (1), 304–318.
- He, B.Y., Jiang, Q., Ma, H., Ma, J., 2024. Multi-agent multimodal transportation simulation for mega-cities: Application of los angeles. In: *The 13th International Workshop on Agent-Based Mobility, Traffic and Transportation Models, Methodologies and Applications. ABMTRANS*.
- Howard, J., Ruder, S., 2018. Universal language model fine-tuning for text classification. arXiv preprint arXiv:1801.06146.
- Kapp, A., Mihaljevic, H., 2023. Reconsidering utility: unveiling the limitations of synthetic mobility data generation algorithms in real-life scenarios. In: *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*. pp. 1–12.
- Li, M., Gao, S., Lu, F., Zhang, H., 2019. Reconstruction of human movement trajectories from large-scale low-frequency mobile phone data. *Comput. Environ. Urban Syst.* 77, 101346.
- Li, Z., Gou, Z., Zhang, X., Liu, Z., Li, S., Hu, Y., Ling, C., Zhang, Z., Zhao, L., 2024. TEG-DB: A comprehensive dataset and benchmark of textual-edge graphs. arXiv preprint arXiv:2406.10310.
- Liao, X., He, B.Y., Jiang, Q., Kuai, C., Ma, J., 2024a. Deep activity model: A generative approach for human mobility pattern synthesis. arXiv preprint arXiv:2405.17468.
- Liao, X., Liu, Y., Kuai, C., Ma, H., He, Y., Cao, S., Stanford, C., Ma, J., 2024b. Reconstructing human mobility pattern: A semi-supervised approach for cross-dataset transfer learning. arXiv preprint arXiv:2410.03788.
- Liu, Y., Kuai, C., Liao, X., Ma, H., He, B.Y., Ma, J., 2024. Semantic trajectory data mining with LLM-informed POI classification. In: *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*. pp. 207–213.
- Liu, Y., Liao, X., Ma, H., He, B.Y., Stanford, C., Ma, J., 2024b. Human mobility modeling with limited information via large language models. arXiv preprint arXiv:2409.17495.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.
- Luca, M., Barlacchi, G., Lepri, B., Pappalardo, L., 2021. A survey on deep learning for human mobility. *ACM Comput. Surv.* 55 (1), 1–44.
- Ma, H., He, B.Y., Kaljevic, T., Ma, J., 2024. A two-sided model for EV market dynamics and policy implications. In: *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*. pp. 3439–3444.
- Ma, H., Liu, Y., Jiang, Q., He, B.Y., Liao, X., Ma, J., 2024b. Mobility AI agents and networks. *IEEE Trans. Intell. Veh.* 9 (7), 5124–5129.
- Merchant, A., Rahimtoroghi, E., Pavlick, E., Tenney, I., 2020. What happens to BERT embeddings during fine-tuning? arXiv preprint arXiv:2004.14448.
- Müller, M., 2007. Dynamic time warping. *Inf. Retr. Music. Motion* 69–84.
- OpenStreetMap contributors, 2017. Planet dump. retrieved from <https://planet.osm.org>, <https://www.openstreetmap.org>.
- Ounoughi, C., Yahia, S.B., 2023. Data fusion for ITS: A systematic literature review. *Inf. Fusion* 89, 267–291.
- Peters, M.E., Ruder, S., Smith, N.A., 2019. To tune or not to tune? adapting pretrained representations to diverse tasks. arXiv preprint arXiv:1903.05987.
- Ren, H., Ruan, S., Li, Y., Bao, J., Meng, C., Li, R., Zheng, Y., 2021. MTrajRec: Map-constrained trajectory recovery via Seq2Seq multi-task learning. In: *Proc. 27th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining*. pp. 1410–1419.
- Schneider, C.M., Belik, V., Couronné, T., Smoreda, Z., González, M.C., 2013. Unravelling daily human mobility motifs. *J. R. Soc. Interface* 10 (84), 20130246.
- Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., Gal, Y., 2024. AI models collapse when trained on recursively generated data. *Nature* 631 (8022), 755–759.
- Si, J., Yang, J., Xiang, Y., Wang, H., Li, L., Zhang, R., Tu, B., Chen, X., 2023. TrajBERT: BERT-based trajectory recovery with spatial-temporal refinement for implicit sparse trajectories. *IEEE Trans. Mob. Comput.* 1–12.

- Stanford, C., Adari, S., Liao, X., He, Y., Jiang, Q., Kuai, C., Ma, J., Tung, E., Qian, Y., Zhao, L., et al., 2024. NUMOSIM: A synthetic mobility dataset with anomaly detection benchmarks. In: Proceedings of the 1st ACM SIGSPATIAL International Workshop on Geospatial Anomaly Detection. pp. 68–78.
- Stopher, P., Xu, M., Fitzgerald, C., 2010. Assessing the accuracy of the sydney household travel survey with GPS. *Australas. Transp. Res. Forum* 33, 1–18.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al., 2023. Llama: Open and efficient foundation language models. arXiv preprint [arXiv:2302.13971](https://arxiv.org/abs/2302.13971).
- Uber Engineering, 2018. H3: Uber's hexagonal hierarchical spatial index. <https://www.uber.com/blog/h3/>. (Accessed 5 July 2025).
- U.S. BLS, 2024. American Time Use Survey (TUS) Data.
- US Bureau of Labor Statistics, 2023. County employment and wages in California — Fourth quarter 2023. URL [https://www.bls.gov/regions/west/news-release/countyemploymentandwages\\_california.htm](https://www.bls.gov/regions/west/news-release/countyemploymentandwages_california.htm). (Accessed 26 January 2025).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30.
- Veraset, 2024. Global mobility and location data provider. <https://www.veraset.com/>. (Accessed 17 June 2024).
- Wang, X., Fang, M., Zeng, Z., Cheng, T., 2024. Where would I go next? Large language models as human mobility predictors. arXiv preprint [arXiv:2308.15197](https://arxiv.org/abs/2308.15197).
- Wei, X., Shen, L., Du, X., Yang, Z., Guo, Z., Yin, Q., 2024. How can multi-source heterogeneous data contribute to assessing urban transportation carrying capacity? *Environ. Impact Assess. Rev.* 108, 107602.
- Xia, T., Qi, Y., Feng, J., Xu, F., Sun, F., Guo, D., Li, Y., 2021. AttnMove: History enhanced trajectory recovery via attentional network. In: Proc. AAAI Conf. on Artificial Intelligence, vol. 35, pp. 447–455.
- Yan, H., Liao, Y., Ma, Z., Ma, X., 2024. Improving multi-modal transportation recommendation systems through contrastive De-biased heterogeneous graph neural networks. *Transp. Res. Part C: Emerg. Technol.* 164, 104689.
- Yan, H., Ma, X., Liu, B., Tan, E., Li, Y., Ni, Z., Liu, T.-L., 2025. Enhancing public transit adoption through personalized incentives: a large-scale analysis leveraging adaptive stacking extreme gradient boosting in China. *Transp. Res. Part C: Emerg. Technol.* 171, 104992.
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H., 2014. How transferable are features in deep neural networks? *Adv. Neural Inf. Process. Syst.* 27.
- Zheng, Y., 2015. Methodologies for cross-domain data fusion: An overview. *IEEE Trans. Big Data* 1 (1), 16–34.
- Zheng, V.W., Zheng, Y., Xie, X., Yang, Q., 2010. Collaborative location and activity recommendations with gps history data. In: Proceedings of the 19th International Conference on World Wide Web. pp. 1029–1038.
- Zhu, Y., Ye, Y., Zhang, S., Zhao, X., Yu, J.J.Q., 2023. DiffTraj: Generating GPS trajectory with diffusion probabilistic model. arXiv preprint [arXiv:2304.11582](https://arxiv.org/abs/2304.11582).