

NUMOSIM: A Synthetic Mobility Dataset with Anomaly Detection Benchmarks

Chris Stanford
cstanford@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Suman Adari
sadari@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Xishun Liao
xishunliao@ucla.edu
University of California, Los Angeles
Los Angeles, California, USA

Yueshuai He
yueshuaihe@ucla.edu
University of California, Los Angeles
Los Angeles, California, USA

Qinghua Jiang
qhjiang93@ucla.edu
University of California, Los Angeles
Los Angeles, California, USA

Chenchen Kuai
kuai0407@ucla.edu
University of California, Los Angeles
Los Angeles, California, USA

Jiaqi Ma
jiaqima@ucla.edu
University of California, Los Angeles
Los Angeles, California, USA

Emmanuel Tung
etung@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Yinlong Qian
yqian@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Lingyi Zhao
lzhao@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Zihao Zhou
ziz244@ucsd.edu
University of California, San Diego
San Diego, California, USA

Zeeshan Rasheed
zrasheed@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

Khurram Shafique
kshafique@novateur.ai
Novateur Research Solutions
Ashburn, Virginia, USA

ABSTRACT

Collecting real-world mobility data is challenging. It is often fraught with privacy concerns, logistical difficulties, and inherent biases. Moreover, accurately annotating anomalies in large-scale data is nearly impossible, as it demands meticulous effort to distinguish subtle and complex patterns. These challenges significantly impede progress in geospatial anomaly detection research by restricting access to reliable data and complicating the rigorous evaluation, comparison, and benchmarking of methodologies. To address these limitations, we introduce a synthetic mobility dataset, NUMOSIM, that provides a controlled, ethical, and diverse environment for benchmarking anomaly detection techniques. NUMOSIM simulates a wide array of realistic mobility scenarios, encompassing both typical and anomalous behaviours, generated through advanced deep learning models trained on real mobility data. This approach allows NUMOSIM to accurately replicate the complexities of real-world movement patterns while strategically injecting anomalies to challenge and evaluate detection algorithms based on how effectively they capture the interplay between demographic, geospatial, and

temporal factors. Our goal is to advance geospatial mobility analysis by offering a realistic benchmark for improving anomaly detection and mobility modeling techniques. To support this, we provide open access to the NUMOSIM dataset, along with comprehensive documentation, evaluation metrics, and benchmark results.

CCS CONCEPTS

• **Computing methodologies** → **Modeling and simulation.**

KEYWORDS

mobility, geospatial data, anomaly detection

ACM Reference Format:

Chris Stanford, Suman Adari, Xishun Liao, Yueshuai He, Qinghua Jiang, Chenchen Kuai, Jiaqi Ma, Emmanuel Tung, Yinlong Qian, Lingyi Zhao, Zihao Zhou, Zeeshan Rasheed, and Khurram Shafique. 2024. NUMOSIM: A Synthetic Mobility Dataset with Anomaly Detection Benchmarks. In *1st ACM SIGSPATIAL International Workshop on Geospatial Anomaly Detection (GeoAnomalies'24)*, October 29–November 1 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3681765.3698455>

1 INTRODUCTION

Mobility data analysis plays a critical role in various domains, including urban planning, disaster management, and epidemiology. Detecting anomalies within mobility patterns — deviations from expected behaviors — can provide valuable insights into significant events or disruptions. For instance, identifying mobility anomalies is essential for understanding the spread of infectious diseases and

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

GeoAnomalies'24, October 29–November 1 2024, Atlanta, GA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1144-2/24/10...\$15.00

<https://doi.org/10.1145/3681765.3698455>

assessing the impact of public health interventions. However, research in geospatial anomaly detection is often constrained by the challenges associated with collecting and utilizing real-world mobility data. These challenges include stringent privacy regulations, high costs, logistical difficulties, and inherent biases in the data, all of which can impede the development and evaluation of robust anomaly detection algorithms.

Accurate annotation of anomalies in large-scale datasets presents another significant obstacle. Identifying subtle and complex patterns that indicate anomalies typically requires extensive domain expertise and meticulous effort. This process is not only time-consuming but also prone to inconsistencies, as human annotators may have different interpretations of what constitutes an anomaly. These limitations hinder access to reliable, well-annotated data, which is crucial for the rigorous evaluation and benchmarking of anomaly detection methodologies.

Given these constraints, synthetic data generation has emerged as a promising alternative for creating datasets that support the development and testing of anomaly detection algorithms. Synthetic datasets offer the flexibility to model a wide range of scenarios, including rare or extreme events that may be underrepresented in real-world data. However, generating synthetic mobility data that accurately captures the complexity of real-world movement patterns is challenging. A key difficulty lies in modeling the intricate interplay of demographic, geospatial, and temporal factors that shape human mobility behaviors. Traditional statistical approaches, typically employed in synthetic data generation, often fall short in capturing these complexities, resulting in datasets that lack the necessary variability and realism to effectively mimic real-world scenarios.

The remainder of this paper is organized as follows. Section 2 reviews prior work related to synthetic mobility data generation. Section 3 introduces the NUMOSIM data. Section 4 discusses the methodology used for generating NUMOSIM, including data modeling, model training, and the data generation processes. Section 5 details the anomaly injection processes, the types of anomalies introduced, and the rationale behind them. Section 6 describes the evaluation metrics and benchmark results provided with NUMOSIM. Section 7 presents our plans for expanding the dataset. Section 8 explains the data format and files released as part of the dataset. Finally, Section 9 concludes the paper with a discussion of contributions and future work.

2 PRIOR WORK

There are numerous datasets available for studying human mobility, encompassing both real and synthetic data. These datasets can be categorized into three main types: location-based service (LBS) data, synthetic data, and vehicle data.

2.1 Location-Based Services (LBS) Data

LBS data has been extensively used in recent studies of human mobility. For example, Yabe et al. [31] introduced YJMob100K, an open-source large-scale mobility dataset based on mobile phone data from Yahoo Japan Corporation. This dataset comprises 75 days of human mobility trajectories for 100,000 individuals, with location pings spatially and temporally discretized to protect user

privacy. Other notable real-world datasets include Gowalla, which captures location-based social network data [8], Foursquare, which includes check-ins from New York City and Tokyo from April 12, 2012, to February 13, 2013 [32], Meta’s travel pattern data, derived from the location history of Facebook users (Meta, n.d.), and OpenPaths from The New York Times Labs, which collects voluntary location data from iPhone and iPad users [10]. However, LBS data often have limitations, such as incomplete daily trajectories, as seen with Gowalla and Foursquare, and a lack of detailed point of interest information due to privacy controls, as with YJMob100K. Additionally, these datasets typically do not include underlying sociodemographic characteristics of the agents because of privacy concerns. Another critical limitation is that real-world LBS datasets generally do not contain annotated anomalies, making it challenging to use them directly for developing and benchmarking anomaly detection algorithms.

2.2 Synthetic Data

The second category includes synthetic datasets, which offer an alternative by simulating human mobility patterns while overcoming some of the limitations of real-world data. For instance, HumoNet [17] and the Patterns of Life Simulation [3, 4] are frameworks that generate synthetic human trajectory data by leveraging multiple layers of real-world data, including transportation networks, points of interest (POI), population data, and observed human trajectories. Another example is MetaPoL, a synthetic dataset created by simulating human behavior within an immersive virtual reality (VR) environment modeled after a real-world secure facility [16]. A significant advantage of synthetic datasets is their ability to provide fully annotated anomalies, allowing researchers to evaluate and benchmark anomaly detection algorithms effectively. This level of control and precision in data generation is difficult to achieve with real-world datasets. However, despite these benefits and innovative approaches, there is a notable lack of large-scale public synthetic datasets that can represent the complex dynamics of megacities with populations in the millions.

2.3 Vehicle Data

The final category is vehicle data, which primarily consists of datasets collected from GPS-equipped vehicles operating in urban environments. For instance, since 2009, the New York City Taxi & Limousine Commission (NYC TLC) has released data on pick-up and drop-off locations for taxis and for-hire vehicles in New York City [23]. Other notable datasets include taxi datasets from cities such as Porto [22], San Francisco [26], Rome [7], and Beijing [33]. The Porto dataset tracks 441 taxis over the course of one year [22], while the Rome dataset monitors 320 taxis over a 30-day period [6, 7]. In addition to taxi datasets, there are also bus datasets that provide further insights into urban mobility. The Rio bus dataset [5] and the Dublin bus dataset [13] capture the trajectories of buses across multiple stoppoints or bus stops, allowing for the detection of abnormal mobility patterns. The anomalies are clearly defined as deviations from the established bus routes. Although these datasets offer valuable insights into the mobility patterns of taxi and bus users within their respective cities, they

fall short of capturing complete daily trajectories and do not represent the broader urban population. This limitation restricts their applicability to more comprehensive mobility studies.

3 PRESENT WORK

To address the limitations of existing datasets, we introduce NUMOSIM (Novateur/UCLA Mobility Simulation), a novel synthetic mobility dataset designed to provide a controlled, ethical, and diverse environment for benchmarking anomaly detection techniques. NUMOSIM is generated using advanced generative deep learning models trained on real travel survey data (and optionally sparse mobility data, when available), enabling it to simulate a wide array of realistic mobility scenarios that encompass both typical and anomalous behaviors. Generative deep learning models are particularly well suited for this task, as they have demonstrated the ability to capture the complex interplay of demographic, geospatial, social, and temporal factors that drive human behavior, resulting in highly realistic synthetic data [20]. Our approach leverages these capabilities to simulate the trajectories of a large population of synthetic agents in urban areas. Additionally, NUMOSIM strategically injects a small quantity of generated anomalies into the dataset, facilitating rigorous benchmarking of anomaly detection algorithms by testing how effectively they capture the intricate dynamics of mobility patterns.

To further support the development and refinement of these algorithms, we provide a comprehensive set of evaluation metrics and benchmark results alongside the NUMOSIM dataset. These resources are intended to help researchers assess the performance of their models and compare them with established baselines.

The initial release of NUMOSIM focuses on a large-scale simulation of Los Angeles, resulting in NUMOSIM-LA — a dataset that reflects the unique mobility dynamics of this urban area. However, our vision for NUMOSIM extends beyond a single city. We plan to expand the dataset by continuing to add more cities in the future, beginning with NUMOSIM-CAIRO, to increase its utility and applicability across different geographic contexts. We aim to advance geospatial mobility analysis by offering a realistic and versatile benchmark for testing and improving anomaly detection and mobility modeling techniques.

4 METHODOLOGY

The Deep Activity Model (DeepAM) [20] introduces a novel approach to the synthesis of human mobility patterns by harnessing the capabilities of deep learning techniques and the widespread availability of household travel survey (HTS) data. This model employs a transformer-based architecture, featuring an encoder-decoder structure, to generate realistic activity chains that capture the complex interdependencies among household members and their activities.

The input of the model consists of sociodemographic data of individuals and their household members. These diverse data types are concatenated and embedded to form a comprehensive feature vector. The encoder processes this input through self-attention layers, effectively capturing the intricate relationships within the data. Subsequently, the decoder generates activity predictions in

an autoregressive manner, predicting the activity type, start time, and end time for each activity in an activity chain.

DeepAM is trained on 180,000 samples from the 2017 National Household Travel Survey (NHTS) [1]. The activities used in the NHTS informed our activity classification into the 16 categories shown in Table 1. The model’s loss function incorporates multiple components: cross-entropy loss for activity-type prediction, soft-label loss for temporal predictions, and specialized losses for temporal order and sequential timing to ensure logically consistent chronological sequences.

Activity Types	
Transportation	Errands
Home	Recreation
Work	Exercise
School	Visit
ChildCare	HealthCare
BuyGoods	Religious
Services	SomethingElse
EatOut	DropOff

Table 1: List of the 16 modeled activity types.

Given the inherent uncertainty of human behavior, the performance and realism of DeepAM are assessed at a systemic level by comparing the distributions of generated and real-world activity patterns. Specifically, Jensen-Shannon divergence (JSD) is used to evaluate five key aspects of the generated activity chains: activity frequencies, distributions of activity start and end times, the number of daily activities, and activity durations. In addition, the logical consistency of the model is evaluated through the activity-transition probability similarity. The results of these evaluations can be found in [20].

For this release, activity type chains were generated for a synthetic subpopulation of Los Angeles consisting of 200,000 agents. An example of an activity chain for a single agent is shown in Table 2.

The list of available Points of Interest (POIs) for the simulation was curated using data from Planetsense[19] and USA Structures[14] for the Los Angeles region, as defined by the boundary in Figure 1. POIs were matched to nearby structures, with unmatched structures assumed to be residential locations. Each location was then assigned a set of valid activity types based on a manually defined mapping. For example, locations labeled "education" by Planetsense were assigned activity types such as Work, School, and DropOff. Residential locations were assigned activity types like Home and Visit. The total number of POIs available for each activity type is presented in Table 3.

Following the generation of activity type chains by DeepAM, a POI assignment procedure was employed to assign each activity to a valid POI for that activity. This procedure considered available data and literature on LA, including average commute times [18],

Day	Activity	Start time	End time
1	Home	00:00	08:30
1	Work	09:00	17:00
1	Home	17:30	24:00
2	Home	00:00	09:00
2	Work	09:30	12:00
2	EatOut	12:15	12:45
2	Work	13:00	17:30
2	Home	18:00	24:00
...
56	Home	17:45	24:00

Table 2: Example of a single agent’s 8-week activity chain output from the Deep Activity Model.

Activity Type	Number of POIs
Transportation	449
Home	2509756
Work	409920
School	10904
ChildCare	33821
BuyGoods	108496
Services	17028
EatOut	165442
Errands	4414
Recreation	17685
Exercise	30520
Visit	2509756
HealthCare	3100
Religious	7255
SomethingElse	2054
DropOff	2838192

Table 3: Number of POIs in the area of interest that correspond to each activity type. A given POI can correspond to multiple activity types.

total distance traveled per day [2, 24], radius of gyration [15, 25], number of locations visited per day [27], and Zipf’s law on human movement [15, 34]. Once agents were assigned POIs, they navigated the Los Angeles road network, obtained via Open Street Map




Figure 1: Area of interest covered in the simulation. All agents stay within the boundary during the 8 week period. Map data from OpenStreetMap [11].

(OSM) [11] to perform their activities. Their arrival times are results of their navigation through the road network using the speed limits supplied by OSM. Although we did not include second-by-second trajectories in the data release, this step is necessary to generate realistic arrival and departure times.

The final output is a schedule of POI visits for each agent, spanning two consecutive 4-week periods: `train` and `test`. An example of the format for the train POI visits is shown in Table 4.

5 ANOMALY INJECTION

The NUMOSIM dataset includes two primary types of anomalies: non-recurring anomalies and recurring anomalies. These anomalies are strategically injected into the dataset to simulate different deviations from expected agent behavior patterns, providing a robust benchmark for evaluating the effectiveness of anomaly detection algorithms in geospatial mobility data.

Non-recurring anomalies are singular, isolated events that disrupt an agent’s typical sequence of activities during the test period. These anomalies are designed to represent unexpected, one-time deviations from an agent’s routine. For example, an agent might unexpectedly leave a stay point that they typically visit at a specific time, or they might visit a new location that fundamentally differs from their usual pattern of stay points. These anomalies challenge detection models to identify significant deviations without being misled by the inherent variability of human behavior.

Recurring anomalies, on the other hand, represent patterns of behavior that deviate from the norm on a regular basis. These anomalies are injected to simulate shifts in an agent’s established behavioral patterns, such as changes in the frequency, timing, or location of visits to Points of Interest (POIs). For example, an agent who previously visited a particular location once a month might now visit it every weekend. While each visit might seem normal in isolation, the recurring nature of these visits marks a significant

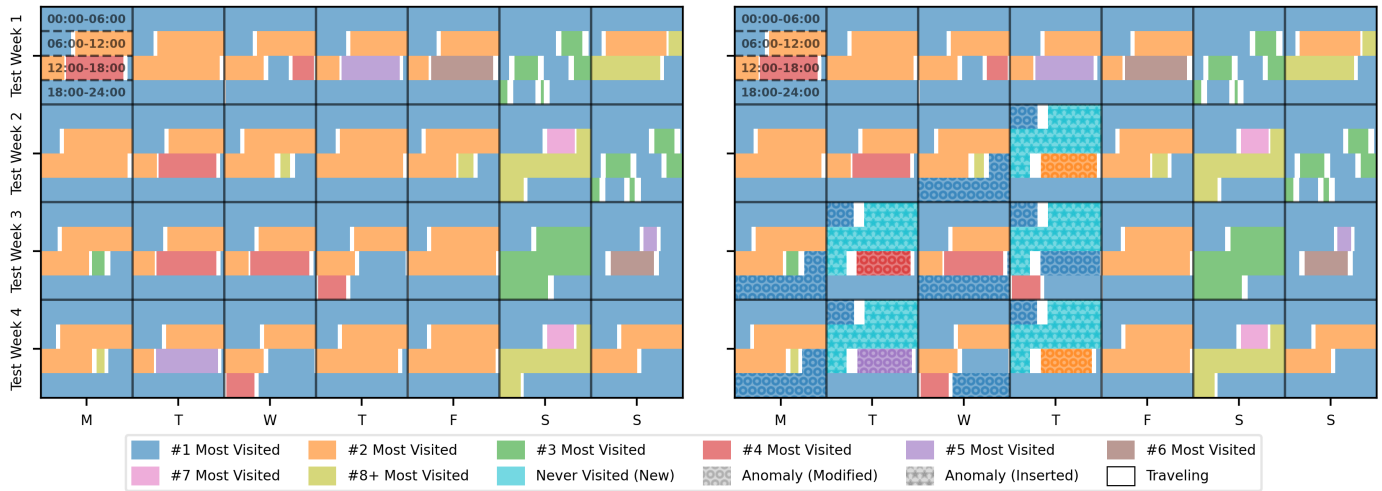
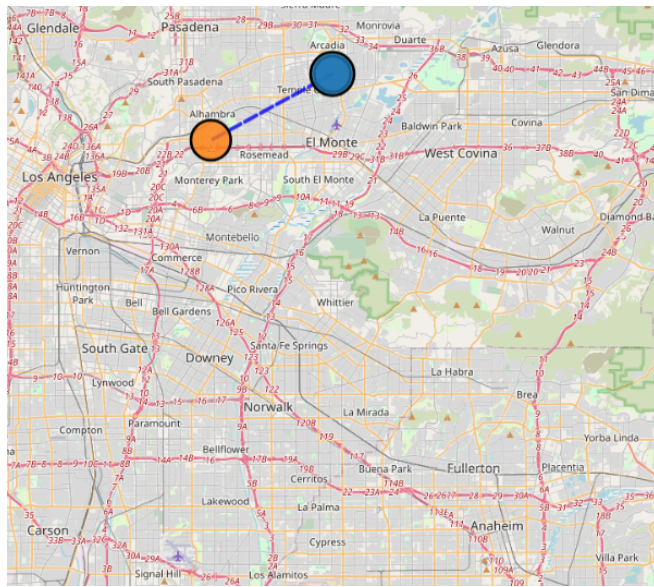
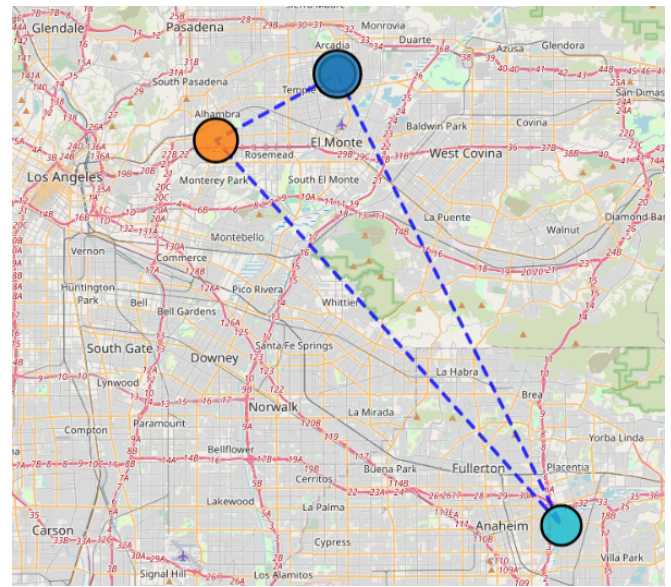


Figure 2: Example of a recurring anomaly. Left: the unaltered, normal test period for an agent, represented as a calendar. Each box represents one day, and each row within a box represents a six-hour period (as demonstrated in the top-left box). Right: the altered, anomalous test period for the same agent. The injected anomalies (marked with the star hatch pattern) are a visit to a new location (cyan) for that agent, recurring at the same time of day each time. The surrounding visits are also altered temporally to accommodate and are therefore also considered anomalous.



(a) Normal daily activity sequence. The agent starts at home (blue), then goes to a frequently visited location (orange), and then returns home.



(b) Anomalous daily activity sequence. A visit at a new location (cyan) is inserted, making that visit anomalous, as well as the surrounding points which are altered temporally to accommodate.

Figure 3: Example of normal (Left) and anomalous (Right) daily stay point sequence for the same agent as Figure 2 on the day of Test Week 2, Day 4 (Thursday).

agent	poi_id	start_datetime	end_datetime
1	215521	2024-01-01 T00:00:00-08:00	2024-01-01 T08:29:09-08:00
1	89362	2024-01-01 T09:01:01-08:00	2024-01-01 T17:05:11-08:00
1	215521	2024-01-02 T08:35:10-08:00	2024-01-0 2T09:00:55-08:00
1	89362	2024-01-02 T09:31:05-08:00	2024-01-02 T12:00:05-08:00
1	99721	2024-01-02 T12:12:44-08:00	2024-01-02 T12:41:31-08:00
1	89362	2024-01-02 T13:02:00-08:00	2024-01-02 T17:36:52-08:00
1	215521	2024-01-02 T18:03:20-08:00	2024-01-03 T08:16:30-08:00
...
200000	788111	2024-01-28 T17:34:50-08:00	2024-01-29 T08:16:21-08:00

Table 4: Example of the `stay_points_train.parquet` data format, containing the stay points for 200,000 agents for 4 weeks.

departure from the agent’s typical behavior, thereby classifying it as an anomaly. Another instance of a recurring anomaly could involve an agent routinely visiting a new location instead of their usual spot.

Overall, in NUMSIM-LA there are 381/200,000 anomalous agents and 3468/17,067,273 anomalous stay points. Approximately 10% of the inserted stay points are to random locations in the anomalous agent’s own history. The remainder are to random locations in other agents’ histories.

Figure 2, Figure 3, and Table 5 illustrate an example of a recurring anomaly. Figure 2 shows a calendar-based visualization of a normal 4-week period compared to an anomalous 4-week period. Figure 3 shows a map-based visualization of the spatial activity of a particular day from Figure 2, which is further specified by Table 5.

The objective of introducing these anomalies is to assess how well models differentiate between normal variations and true anomalies. An effective model should recognize both subtle and overt deviations, minimizing false positives by distinguishing irregular behaviors from real anomalies. This capability is crucial for identifying real disruptions caused by emergencies or unexpected events. Non-recurring anomalies test a model’s sensitivity to singular disruptions, while recurring anomalies evaluate its ability to recognize significant, repetitive shifts. Together, these anomalies provide a comprehensive assessment of model robustness, making NUMOSIM a powerful tool for benchmarking anomaly detection in geospatial mobility data.

6 BENCHMARKS

To demonstrate the utility of NUMOSIM as a comprehensive testbed for geospatial anomaly detection, we implement a set of baseline methods for benchmarking. These benchmarks are designed not only to validate the detectability of the injected anomalies but also to establish a performance baseline for future work in this domain.

NUMOSIM, with its labeled anomalous and non-anomalous data, provides a versatile platform for developing and evaluating anomaly detection models. The dataset’s flexibility allows for various configurations, including both unsupervised and semi-supervised approaches. The anomaly prevalence rate during evaluation can be adjusted by selecting specific subsets of the `test` data. Users can introduce various types of noise, such as Gaussian spatial/temporal noise, missing staypoints, or agent ID switches to tailor the data to specific application needs.

Beyond data partitioning and augmentation, NUMOSIM allows for training and evaluation at different levels of granularity: agent-level, staypoint-level, trip-level, or point-level. For instance, an agent may be classified as anomalous if it has a certain number or proportion of anomalous staypoints; this classification can then be used to score models at the agent-level. Similarly, models that output scores for staypoints or trips (i.e., staypoint-pair transitions) can be evaluated at those levels, or aggregated to the agent-level using strategies such as max-pooling staypoint scores. For applications sensitive to duration, evaluations can be conducted on a second-by-second or point-by-point basis, allowing for fine-grained analysis of model performance.

Our benchmarking methodology proceeds as follows. First, we train the selected methods on the `train` dataset, which is free of anomalies, using an unsupervised learning approach. Next, we perform inference on the `test` period data to generate anomaly scores. Finally, we evaluate each model using metrics such as the Average Precision (AP) and Area Under the Receiver Operating Characteristic curve (AUCROC). Although other metrics like Maximum F1-Score or Average Precision Recall are viable, they not are included in this analysis. It is important to note that no forms of noise or data augmentation were applied during this benchmarking process.

We selected three baseline anomaly detection methods for this evaluation: RioBusData [5], Spatial-Temporal Outlier Detector (STOD) [12], and Gaussian Mixture Variational Sequence AutoEncoder (GM-VSAE) [21]. RioBusData [5] is an interactive analysis tool with a CNN-based anomaly detection model, initially applied to bus mobility data. STOD [12] is a GRU-based anomaly detection model, also originally designed for bus mobility data. GM-VSAE [21] is a VAE-based anomaly detection model, applied to taxi mobility data. To adapt these baseline methods to NUMOSIM, we made a few straightforward modifications, which are detailed in Appendix A. In addition to the baseline methods, we created a simple statistical test, described in Appendix B, based on the fact that the rate at which agents visit each POI should not change substantially between train and test.

The results of our baseline benchmarking experiments are presented in Table 6. In this table, we observe a striking trend: most

Original			Anomalous		
poi_id	start_datetime	end_datetime	poi_id	start_datetime	end_datetime
2364463 (residence)	2024-02-07 16:38:04	2024-02-08 07:17:01	2364463 (residence)	2024-02-07 16:38:04	2024-02-08 01:48:01
-	-	-	295533 Silver Screen Sight and Sound	2024-02-08 02:28:56	2024-02-08 13:19:47
399191 Ramona Elementary School	2024-02-08 07:33:18	2024-02-08 17:41:43	399191 Ramona Elementary School	2024-02-08 14:01:34	2024-02-08 17:41:43
2364463 (residence)	2024-02-08 17:58:47	2024-02-09 06:55:00	2364463 (residence)	2024-02-08 17:58:47	2024-02-09 06:55:00

Table 5: Tabular example of an injected anomaly, corresponding to Figure 3 and Figure 2 (Test Week 2, Thursday). Anomalies (red) are created as a result of the injected visit, as well as the temporal adjustments to surrounding visits.

existing, sophisticated anomaly detection models are underperforming, while a basic model that merely accounts for the visit rate of agents achieves significantly higher scores.

This outcome, while surprising, is not unprecedented in the field of anomaly detection. In Wu and Keogh [30], they demonstrate that with just a single line of code, they can replicate the performance of complex, many-parameter deep learning models on popular benchmark datasets.

This observation underscores the broader principle that the success of anomaly detection, particularly in the context of complex mobility patterns, hinges not just on the sophistication of the models employed but on the careful selection of features that truly capture the underlying dynamics of the data. In mobility analysis, where patterns can be influenced by a myriad of factors such as time, location, sociodemographic characteristics, and behavioral trends, identifying the most relevant features is crucial. A well-chosen feature, like the visit rate, can often encapsulate significant behavioral signals, providing clarity where more complex models might struggle to disentangle noise from meaningful patterns.

7 ONGOING RELEASES

In addition to the initial Los Angeles dataset, we have plans for several future releases that will be made available as they are completed. Our next planned release will focus on Cairo, Egypt, with additional regions to follow over time. Each new dataset will be tailored to reflect the unique mobility patterns and cultural contexts of its respective region.

To achieve this level of specificity, we will employ transfer learning techniques, enabling our deep activity model to adapt to the local customs and behaviors of each area. By incorporating local human mobility data into the training process, our model will be fine-tuned to capture the distinct activity patterns that characterize different regions, ensuring that each dataset offers a realistic and region-specific representation of human movement.

These ongoing releases are part of our broader commitment to providing comprehensive, high-quality datasets that can be used to benchmark and refine geospatial mobility models and anomaly detection techniques. By expanding the geographic scope of our

Level	Method	Anomaly Prevalence	AP	AUCROC
Staypoint	RioBusData		0.000189	0.505
	STOD	0.000203	0.00024	0.594
	GM-VSAE		—	—
	Visit Rate		0.0165	0.906
Agent	RioBusData		0.00164	0.501
	STOD	0.00191	0.00182	0.518
	GM-VSAE		0.00192	0.507
	Visit Rate		0.0164	0.646

Table 6: Staypoint-level and agent-level results from baseline methods on NUMOSIM-LA. GM-VSAE concatenates an agent’s entire series of stay points, so there are no staypoint-level scores for this method.

datasets, we not only aim to support a wider range of research applications—from urban planning and transportation management to public health and security—but also to enhance the generalizability of anomaly detection models across different geospatial regions [28]. This approach will allow researchers to test and validate their models in diverse environments, ensuring that the developed techniques are robust and applicable across various contexts. By offering datasets that span multiple regions, we hope to contribute to the creation of more versatile and adaptable models capable of handling the complexities of human mobility on a global scale.

8 FILE DESCRIPTIONS

NUMOSIM data can be accessed on the Open Science Framework at (<https://osf.io/sjyfr/>).

This release contains 6 files, split into two categories: supplemental and stay point.

8.1 Supplemental Files

- (1) **readme.txt**: A text file describing the files in the data release.
- (2) **demographics.parquet**: Table containing the demographic information for each agent. Descriptions of the columns and their values can be found in the readme file.
- (3) **poi.parquet**: Table containing information for points of interest. Columns are `poi_id`, `name`, `latitude`, and `longitude`, and `act_types`.

8.2 Stay point files

The stay point files contain the stay point information for the agents. The 4-week `train` stay points and 4-week `test` stay points are split into separate files, but together comprise an 8-week sequence of stay points for 200,000 simulated agents, spanning a period from 2024-01-01T00:00:00-08:00 to 2024-02-25T23:59:59-08:00. Any activity that overlapped the beginning of `train` or the end of `test` was truncated. Any activity that overlapped the transition between `train` and `test` was included in the `train` file in full rather than being split across the two files.

- (1) **stay_points_train.parquet**: Table containing stay point information for all agents for a 4-week long training period. Columns are `agent_id`, `poi_id`, `start_datetime`, and `end_datetime`.
- (2) **stay_points_test_truth.parquet**: Table containing ground truth stay point information for all agents for a 4-week long period that immediately follows the training period. Columns are `agent_id`, `poi_id`, `start_datetime`, and `end_datetime`.
- (3) **stay_points_test_anomalous.parquet**: Table containing stay point information for all agents for a 4-week long period that immediately follows the training period, with injected anomalies. Columns are `agent_id`, `poi_id`, `start_datetime`, `end_datetime`, `anomaly`, and `anomaly_type`. The `anomaly` column is a boolean indicating whether the stay point has been injected or altered. The `anomaly_type` column is a integer where 1 indicates that the anomaly is a modified staypoint and 2 indicates that the anomaly is an injected staypoint.

9 CONCLUSION

In this paper, we introduced NUMOSIM, a synthetic mobility dataset with embedded anomalies, designed as a benchmark for evaluating geospatial mobility models and anomaly detection techniques. By leveraging deep activity models trained on real-world survey data, we aim to bridge the gap between synthetic data generation and the realistic representation of human mobility patterns.

This dataset represents a significant step forward in the field of geospatial mobility analysis, providing researchers with a robust tool to develop and validate more sophisticated models that can generalize beyond the limitations of traditional datasets. Our future work will focus on expanding the dataset to encompass different geographical regions, incorporating more complex and rare behavioral patterns and social interactions, and refining the anomaly generation process to further enhance the realism and utility of the dataset.

By making this dataset publically available, we aim to drive progress in mobility modeling and anomaly detection, ultimately

contributing to more accurate and effective applications in areas such as urban planning and transportation management.

ACKNOWLEDGMENTS

Supported 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.

REFERENCES

- [1] Federal Highway Administration. 2017. National Household Travel Survey. <https://nhts.ornl.gov>.
- [2] Laura Alessandretti, Piotr Sapiezynski, Sune Lehmann, and Andrea Baronchelli. 2017. Multi-scale spatio-temporal analysis of human mobility. *PLOS ONE* 12, 2 (02 2017), 1–17. <https://doi.org/10.1371/journal.pone.0171686>
- [3] Hossein Amiri, Will Kohn, Shiyang Ruan, Joon-Seok Kim, Hamdi Kavak, Andrew Crooks, Dieter Pfoser, Carola Wenk, and Andreas Zufle. 2024. The Patterns of Life Human Mobility Simulation. arXiv:2410.00185 [cs.MA] <https://arxiv.org/abs/2410.00185>
- [4] Hossein Amiri, Ruochen Kong, and Andreas Zufle. 2024. Urban Anomalies: A Simulated Human Mobility Dataset with Injected Anomalies. arXiv:2410.01844 [cs.SI] <https://arxiv.org/abs/2410.01844>
- [5] Aline Bessa, Fernando de Mesentier Silva, Rodrigo Frassetto Nogueira, Enrico Bertini, and Juliana Freire. 2016. RioBusData: Outlier Detection in Bus Routes of Rio de Janeiro. arXiv:1601.06128 (Jan. 2016). <http://arxiv.org/abs/1601.06128> arXiv:1601.06128 [cs].
- [6] Marco Bonola, Lorenzo Bracciale, Pierpaolo Loreti, Raul Amici, Antonello Rabuffi, and Giuseppe Bianchi. 2016. Opportunistic communication in smart city: Experimental insight with small-scale taxi fleets as data carriers. *Ad Hoc Networks* 43 (2016), 43–55.
- [7] Lorenzo Bracciale, Marco Bonola, Pierpaolo Loreti, Giuseppe Bianchi, Raul Amici, and Antonello Rabuffi. 2014. CRAWDAD dataset roma/taxi (v. 2014-07-17).
- [8] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego, California, USA) (KDD '11). Association for Computing Machinery, New York, NY, USA, 1082–1090. <https://doi.org/10.1145/2020408.2020579>
- [9] Kenneth Ward Church. 2017. Word2Vec. *Natural Language Engineering* 23, 1 (Jan. 2017), 155–162. <https://doi.org/10.1017/S1351324916000334>
- [10] The New York Times Company. 2011. OpenPaths. <https://nytlabs.com/projects/openpaths.html> Accessed: 2024-08-23.
- [11] OpenStreetMap contributors. 2024. OpenStreetMap. <https://www.openstreetmap.org> Accessed: 2024-08-10.
- [12] Michael Cruz and Luciano Barbosa. 2020. Learning GPS Point Representations to Detect Anomalous Bus Trajectories. *IEEE Access* 8 (2020), 229006–229017. <https://doi.org/10.1109/ACCESS.2020.3046912>
- [13] Michael Oliveira Cruz and Luciano Andrade Barbosa. 2023. Applying Transformers for Anomaly Detection in Bus Trajectories. 4633755 (Nov. 2023). <https://doi.org/10.2139/ssrn.4633755>
- [14] FEMA and Esri. 2024. USA Structures. <https://gis-fema.hub.arcgis.com/pages/usa-structures>. Accessed: 2024-08-10.
- [15] Marta C. González, César A. Hidalgo, and Albert-László Barabási. 2008. Understanding individual human mobility patterns. *Nature* 453, 7196 (2008), 779–782. <https://doi.org/10.1038/nature06958>
- [16] Chathika Gunaratne, Mason Stott, Debraj De, and Gautam Malviya Thakur. 2024. MetaPoL: Immersive VR based Indoor Patterns of Life (PoL) and Anomalies Data Generation for Insider Threat Modeling in Nuclear Security. In *2024 25th IEEE International Conference on Mobile Data Management (MDM)*. 255–258. <https://doi.org/10.1109/MDM61037.2024.00053>
- [17] Joon-Seok Kim, Gautam Malviya Thakur, Licia Amichi, Annetta Burger, Chathika Gunaratne, Joseph Tuccillo, Taylor Hauser, Joseph Bentley, Kevin Sparks, Debraj De, Chance Brown, Elizabeth McBride, Jesse McGaha, James Gaboardi, Xiuling Nie, and Steven Carter Christopher. 2024. HumoNet: A Framework for Realistic Modeling and Simulation of Human Mobility Network. In *2024 25th IEEE International Conference on Mobile Data Management (MDM)*. 185–194. <https://doi.org/10.1109/MDM61037.2024.00042>
- [18] Yardi Kube. 2024. Commuting in the US: Best Metros for a Short and Eco-Friendly Commute. <https://www.yardikube.com/blog/commuting-in-the-us-best-metros>

for-a-short-and-eco-friendly-commute/. Accessed: 2024-08-10.

[19] Oak Ridge National Laboratory. 2024. PlanetSense. <https://planetsense.ornl.gov/webps/about>. Accessed: 2024-08-10.

[20] Xishun Liao, Brian Yueshuai He, Qinhuo Jiang, Chenchen Kuai, and Jiaqi Ma. 2024. Deep Activity Model: A Generative Approach for Human Mobility Pattern Synthesis. *arXiv preprint arXiv:2405.17468* (2024).

[21] Yiding Liu, Kaiqi Zhao, Gao Cong, and Zhifeng Bao. 2020. Online Anomalous Trajectory Detection with Deep Generative Sequence Modeling. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. 949–960. <https://doi.org/10.1109/ICDE48307.2020.00087>

[22] Luis Moreira-Matias, Joao Gama, Michel Ferreira, Joao Mendes-Moreira, and Luis Damas. 2013. Predicting taxi–passenger demand using streaming data. *IEEE Transactions on Intelligent Transportation Systems* 14, 3 (2013), 1393–1402.

[23] New York City Taxi and Limousine Commission. 20XX. Taxi and Limousine Commission Trip Data. <https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page> Accessed: 2024-08-23.

[24] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. 2012. A Tale of Many Cities: Universal Patterns in Human Urban Mobility. *PLOS ONE* 7, 5 (05 2012), 1–10. <https://doi.org/10.1371/journal.pone.0037027>

[25] Luca Pappalardo, Filippo Simini, Salvatore Rinzivillo, Dino Pedreschi, Fosca Giannotti, and Albert-László Barabási. 2015. Returners and explorers dichotomy in human mobility. *Nature Communications* 6, 1 (2015), 8166. <https://doi.org/10.1038/ncomms9166>

[26] Michal Piorkowski, Natasa Sarafijanovic-Djukic, and Matthias Grossglauser. 2009. CRAWDAD dataset epfl/mobility (v. 2009-02-24).

[27] Christian Schneider, Vitaly Belik, Thomas Couronné, Zbigniew Smoreda, and Marta C. Gonzalez. 2013. Unravelling daily human mobility motifs. *Journal of the Royal Society, Interface / the Royal Society* 10 (07 2013), 20130246. <https://doi.org/10.1098/rsif.2013.0246>

[28] Mark Tenzer, Zeeshan Rasheed, and Khurram Shafique. 2023. The Geospatial Generalization Problem: When Mobility Isn't Mobile. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems (Hamburg, Germany) (SIGSPATIAL '23)*. Association for Computing Machinery, New York, NY, USA, Article 47, 4 pages. <https://doi.org/10.1145/3589132.3625590>

[29] Inc. Uber Technologies. 2018. H3: A Hexagonal Hierarchical Spatial Index. <https://h3geo.org/>. [https://h3geo.org/Version 3.7.2](https://h3geo.org/Version%203.7.2).

[30] Renjie Wu and Eamonn J. Keogh. 2020. Current Time Series Anomaly Detection Benchmarks are Flawed and are Creating the Illusion of Progress. *CoRR* abs/2009.13807 (2020). [arXiv:2009.13807](https://arxiv.org/abs/2009.13807) <https://arxiv.org/abs/2009.13807>

[31] Takaaki Yabe, Kiyohito Tsubouchi, Takuro Shimizu, Yoshihide Sekimoto, Kaoru Sezaki, Esteban Moro, and Alex Pentland. 2024. YJMob100K: City-scale and longitudinal dataset of anonymized human mobility trajectories. *Scientific Data* 11, 1 (2024), 1–8. <https://doi.org/10.1038/s41597-024-03237-9>

[32] Dingqi Yang. 2014. Foursquare Dataset. <https://sites.google.com/site/yangdingqi/home/foursquare-dataset> Accessed: 2024-08-23.

[33] Jing Yuan, Yu Zheng, Chengyang Zhang, et al. 2010. T-Drive: Driving Directions Based on Taxi Trajectories. In *Proceedings of 18th ACM SIGSPATIAL*.

[34] George Kingsley Zipf. 1946. The P1 P2/D Hypothesis: On the Intercity Movement of Persons. *American Sociological Review* 11, 6 (1946), 677–686. <http://www.jstor.org/stable/2087063>

Appendices

Appendix A BASELINE METHOD ADAPTATIONS

In this section, we describe how we modify the Rio-Bus Model [5], the STOD [12] and the GM-VSAE [21] to use with our multi-agent stajpoint trajectories dataset.

A.1 RioBusData

The original RioBusData model [5] uses a Convolutional Neural Network (CNN) to analyze bus trajectories in Rio de Janeiro. The CNN takes as input a sequence of GPS coordinates and predicts the bus ID, learning to distinguish between normal and outlier trajectories by assuming outlier bus trajectories’ IDs are harder to predict. The architecture of the adapted model is presented in Figure 4.

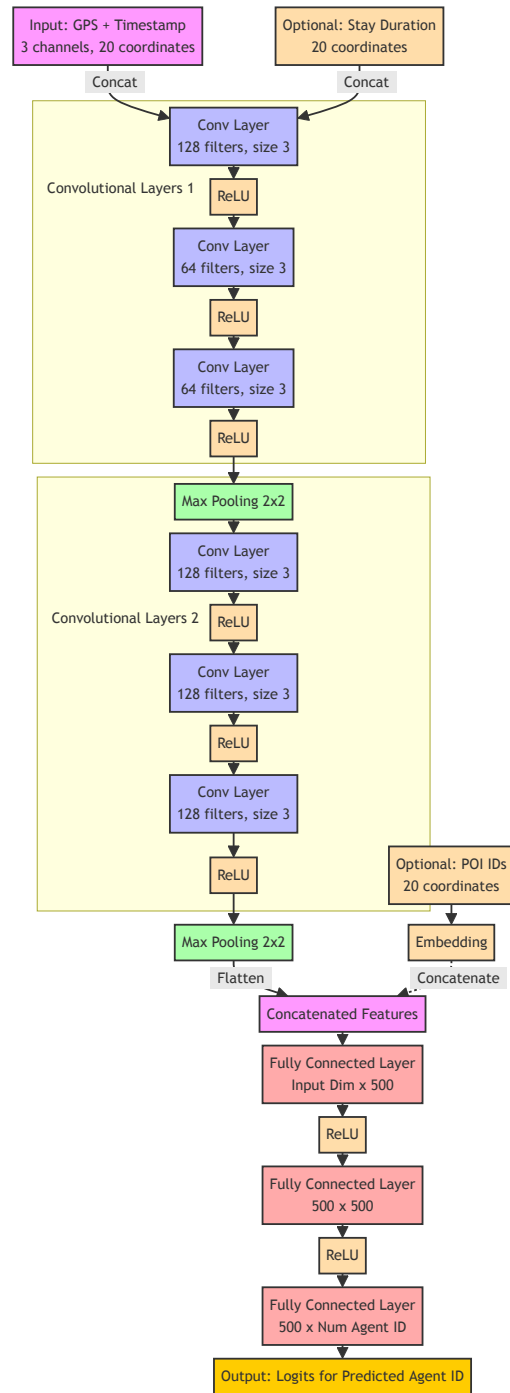


Figure 4: Illustration of the RioBUS CNN architecture adapted to the multi-agent GPS stajpoint trajectory. The input to the network is 20 GPS coordinates of a single agent that are coded in three channels, which correspond to the timestamp, latitude, and longitude of each GPS coordinate. The input can be optionally extended to four channels by the additional stay duration feature. This input is fed to two convolutional layer blocks. The CNN output is then concatenated with an optional POI embedding. The last layer predicts the agent ID to which the input 20 GPS coordinates belong.

Agent ID Replacement. The bus ID in the original model is replaced by an agent ID. This allows the model to build agent behaviors' classification instead of bus routes' classification.

Additional Features. Points of Interest (POI) embeddings are concatenated with the CNN output before the final prediction layer. This provides context about the surrounding environment of each GPS coordinate. Additionally, a new channel is added to the input, representing the stay duration at each GPS coordinate. This helps capture the temporal aspect of staypoints.

Input Structure. The adapted model still takes as input 20 GPS coordinates for a single agent. These are encoded in three or four channels: timestamp, latitude, longitude, and optionally, stay duration.

Network Architecture. The core CNN structure remains similar to the original RioBusData model. It consists of three initial convolutional layers, followed by a max pooling layer, then three more convolutional layers, another max pooling layer, and finally three fully connected layers.

Output Score. We use the negative probability of the agent ID being correctly predicted as the anomaly score. If the agent's sub-trajectory is less likely to be correctly identified, it is more likely to be anomalous. The staypoint-level anomaly score is computed as the anomaly score of the 20-coordinates sub-trajectory ended with the staypoint. The agent-level anomaly score is computed from all 20-coordinate sub-trajectories of a specific agent. It is equal to the maximum score among these sub-trajectories.

A.2 STOD

The original STOD model [5] applies a three-stage approach. The model first employs a bidirectional GRU (Gated Recurrent Unit) to process contextual information. This includes historical trajectory, future trajectory, and other features of the coordinate of interest. The GRU then predicts the bus state at the coordinate of interest. The bus state is one of "in route, bus stop, traffic signal, and other stops". The classifier is called PAC (Point Activity Classifier). The high dimensional input to the last layer of the PAC is considered a contextual representation of a point in the trajectory.

Meanwhile, it maps all GPS sequences to H3 [29] hexagonal grids. All GPS coordinates are converted to discrete words representing the grid index. A Word2Vec [9] model is trained on the H3 sequences to gather the association between H3 grids. The Word2Vec embeddings of a GPS coordinate are considered a geographical representation of a point.

Then, for a given GPS trajectory, it uses the concatenated PAC embeddings and Word2Vec embeddings of all points in the trajectory as input to a GRU to finally predict the bus ID. Similar to Rio-Bus, it assumes outlier bus trajectories' IDs are harder to predict. The architecture of the adapted PAC classifier is presented in Figure 5.

Word2Vec modification. The Word2Vec representation is compatible with agent trajectories, so we leave it unchanged. For training the Word2Vec model, we use the full trajectory of each agent as a sentence.

Agent ID Replacement. To output the score per staypoint, we remove the GRU at the final stage and directly use a 2-layer MLP with ReLU activation to predict the agent ID, instead of the bus ID.

PAC modification. We mainly modify the Point Activity Classifier, see Figure 5. We remove the route ID, velocity and acceleration from the GRU input since they are not available in our dataset. Instead, we concatenate the GRU input with Points of Interest (POI) embeddings. This provides context about the surrounding environment of each GPS coordinate. Staypoint durations are also directly concatenated with the GRU input. This helps capture the temporal aspect of staypoints. Instead of predicting the bus state, we predict the POI type of the GPS coordinate of interest and make sure the POI type of the point is removed from the input.

Output Score. We use the negative probability of the agent ID being correctly predicted as the anomaly score. If the agent sub-trajectory is less likely correctly identified, it is more likely to be anomalous. The agent-level anomaly score is, again, computed as the maximum of all staypoint-level anomaly score.

A.3 GM-VSAE

The Gaussian Mixture Variational Sequence AutoEncoder (GM-VSAE) [21] captures complex sequential information in trajectories, discovers different types of normal routes via Gaussian mixture, and represents them in a continuous latent space. This method requires gridified sequences. Therefore, we only retain the longitude and latitude from the input GPS sequences. We then convert these coordinates into grid series for use with the model. We treat each agent as a separate trajectory.

This model does not support staypoint-level scoring. We calculate the negative likelihood of each test agent trajectory and use this value as the agent-level anomaly scores.

Appendix B VISIT RATE MODEL

The visit rate model assumes that the number of times an agent i visits a POI p in a 4-week time window is a Poisson process with a rate parameter $\lambda_{i,p}$. We further assume that this parameter does not change between `train` and `test` for non-anomalous agents and stay points.

We then compute staypoint-level anomaly scores $a_{i,p}$ for each (agent_id, poi_id) tuple simply by comparing the observations $\hat{\lambda}_{i,p}$ between between `train` and `test`, normalized by the standard deviation in `train`.

$$a_{i,p} = \frac{\text{abs}(\hat{\lambda}_{i,p,\text{train}} - \hat{\lambda}_{i,p,\text{test}})}{\sqrt{\hat{\lambda}_{i,p,\text{train}}}} \quad (1)$$

If no visits occurred in `train`, a baseline value of 0.5 was assumed. This results in a single visit in `test` receiving a slightly anomalous score rather than causing division by zero.

Finally, since this method only picks up inserted anomalies, which are typically preceded and followed by altered anomalies, we re-assigned each staypoint's anomaly score to be the max of itself, the preceding, and the following score.

Agent-level scores were then computed by taking $\max_p(a_{i,p})$ for each agent.

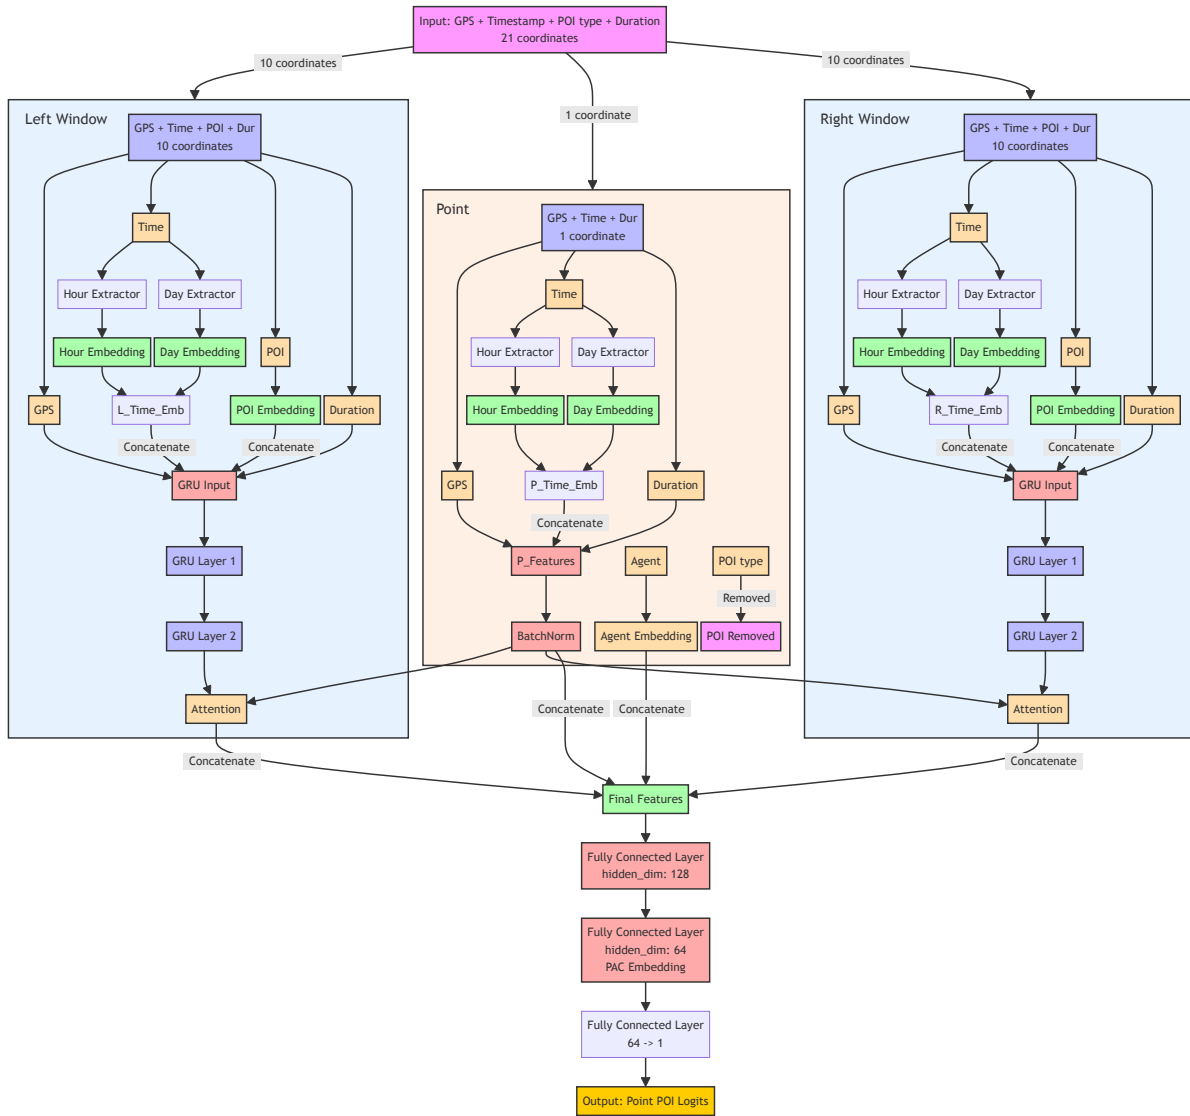


Figure 5: Illustration of the Point Activity Classifier (of STOD) adapted to the multi-agent GPS staypoint trajectory. The input to the network is 21 GPS coordinates with some extra features, which correspond to the timestamp, latitude, longitude, POI types, and stay duration of each GPS coordinate. The input is split into the left window, the coordinate of interest, and the right window. All other inputs are used to predict the POI type of the POI type of the coordinate of interest.