Large Language Models for Spatial Trajectory Patterns Mining

Zheng Zhang*
Emory University, Atlanta, GA, USA
zheng.zhang@emory.edu

Hossein Amiri*
Emory University, Atlanta, GA, USA
hossein.amiri@emory.edu

Zhenke Liu
Emory University, Atlanta, GA, USA
zhenke.liu@emory.edu

Liang Zhao
Emory University, Atlanta, GA, USA
liang.zhao@emory.edu

Andreas Züfle
Emory University, Atlanta, GA, USA
azufle@emory.edu

Abstract

Identifying anomalous human spatial trajectory patterns can indicate dynamic changes in mobility behavior with applications in domains like infectious disease monitoring and elderly care. Recent advancements in large language models (LLMs) have demonstrated their ability to reason in a manner akin to humans. This presents significant potential for analyzing temporal patterns in human mobility. In this paper, we conduct empirical studies to assess the capabilities of leading LLMs like GPT-4 and Claude-2 in detecting anomalous behaviors from mobility data, by comparing to specialized methods. Our key findings demonstrate that LLMs can attain reasonable anomaly detection performance even without any specific cues. In addition, providing contextual clues about potential irregularities could further enhances their prediction efficacy. Moreover, LLMs can provide reasonable explanations for their judgments, thereby improving transparency. Our work provides insights on the strengths and limitations of LLMs for human spatial trajectory analysis.

CCS Concepts

• **Information systems → Geographic information systems**; **Location based services**.

Keywords

Geolife, Patern of Life, Simulation, Trajectory, Dataset, LLM

ACM Reference Format:
Zheng Zhang, Hossein Amiri, Zhenke Liu, Liang Zhao, and Andreas Züfle. 2024. Large Language Models for Spatial Trajectory Patterns Mining. 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, 4 pages. https://doi.org/10.1145/3681765.3698467

1 Introduction

The widespread adoption of location-enabled mobile devices has led to a massive collection of human mobility data [24, 29], comprising diverse trajectory types from individual app usages to public transportation systems. These mobility traces can be modeled as dynamic graphs, representing sequences of location visits with associated semantics [25]. Analyzing these dynamic graphs enables valuable insights for applications like transportation mode classification and detecting spatiotemporal patterns [7, 10, 14, 19, 22, 31, 36, 39, 42]. A particularly difficult task is identifying anomalous mobility patterns within an individual's semantic trajectories, where the trajectory significantly deviates from their historical patterns. Finding such anomalous patterns of individuals may indicate a change in behavior which has many important applications. For instance, infectious disease monitoring [16, 27, 40] or tracking elderly behaviors [30].

In recent times, there has been a surge in progress with large language models (LLMs) [21, 26] like Transformers [34], BERT [12], GPT [11], among others. These LLMs act as foundational models, which can be easily adapted for various downstream applications with minimal adjustments [11, 17, 21, 41]. Notably, breakthroughs in design and training techniques have enabled emerging abilities in LLMs, distinguishing cutting-edge models like GPT-3.5 [11], GPT-4 [3], Claude-2 [8], BARD [1], LlaMA [32], and LlaMA-2 [33] from earlier versions. For example, features such as in-context training [23] and zero-shot learning [17, 35] allow these models to adapt to tasks they were not explicitly trained for.

Despite the remarkable progress LLMs have made in diverse NLP tasks like question answering (QA) and machine translation, their potential in analyzing human mobility patterns remains largely unexplored. Human mobility data, unlike typical language sequences, presents with intricate spatial-temporal dynamics and rich topological connections between entities. Detecting anomalous behaviors are especially difficult due to the intrinsic property of unknown nature of anomalies. Existing methods typically rely on creating hand-crafted features such as the total traveled distance and use heuristic rule to determine outliers, which limits their capability to generalize effectively to detect unseen outlier patterns. In contrast, LLMs has natural advantage since they can directly perceive natural language input. As LLMs have shown powerful reasoning ability and generalization capabilities directly from the input prompt, it becomes intriguing to assess to what extent LLMs can detect diverse anomaly behaviors under the human mobility patterns.

To systematically study the capabilities of LLMs on detecting outliers (anomalies) in human mobility trajectories, we conduct a series of empirical experiments with leading LLMs on diverse datasets. By comparing their performance to specialized human mobility anomaly detection methods, we aim to assess the potential strengths and limitations of LLMs in this domain. Critically, by altering the input prompt formats, we aim to evaluate how effectively LLMs can extract and leverage the underlying structural information from the dynamic mobility patterns to enhance their performance in subsequent tasks. Moreover, we delve

*These authors contributed equally to the paper and are named in random order.

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

ACM ISBN 979-8-4007-1144-2/24/10
https://doi.org/10.1145/3681765.3698467

into both the effectiveness and interoperability of LLMs' predictions. The source code of the model and the datasets used in this study are available at https://github.com/onspatial/LLM-outlier-detection, https://github.com/onspatial/geolife-outlier-dataset and https://github.com/onspatial/pol-outlier-dataset, respectively.

The rest of the paper is organized as follows. Section 2.2 describes the simulated and real-world datasets used in this study. Section 2 presents the methodology, including the input prompt formats and the comparison methods. Section 3 details the experimental setup and results. Finally, Section 4 concludes the paper.

2 Methodology

In this section we present the methodology used to evaluate the performance of large language models (LLMs) in detecting anomalous human mobility patterns. We first describe the datasets used in this study, followed by the comparison methods and the input prompt formats used to evaluate the LLMs.

2.1 Experimental Design and Implementation

Encoding Trajectories to Sequence. The original dataset comprises check-in records, each represented as a row containing a timestamp, user ID, and location coordinates, along with the type of location. For example, a typical row might appear as follows: latitude 39.90523006334842, longitude 116.39160929411752, start time 2009-10-11,07:02:49, end time 2009-10-11,07:27:34, location type: Pub, user ID: 46. We then transform this data for each user into a structured format encoding the time, location type, and distance from the previous check-in, e.g., "Sun 07:27, Pub, 2.1 km." This encoded data is subsequently utilized to generate the input prompts for large language models (LLMs).

Input Prompt Formats. To systematically study the research questions above, we design two different dimensions to create input prompts for the LLMs. Specifically, (1) **With/Without Hint**: Given that an anomaly begins at a specific time point in the data, it is crucial to evaluate the performance of the LLMs whether this information is provided or not. Notably, for all comparative methods, this hint is used to divide the data into training and testing sets.; (2) **Separate vs Combine**: It would also be interesting to assess whether there is a significant difference in performance when presenting all the trajectories in a single prompt versus in separate prompts. This is because placing them all in one prompt might allow the model to consider interactions between different trajectories. On GitHub we present examples to illustrate the details of prompt.

Choices of LLMs. We opted to utilize OpenAI's state-of-the-art models, GPT-3.5 and GPT-4, via their API system, and Claude-2 by Anthropic. Specifically, we use gpt-3.5-turbo-16k-0613 and gpt-4-0613 for 'Separate' prompt, and Claude-2 for 'Combine' prompt due to its capability to hold long input prompt up to 100K input context window size.

Evaluation Metrics. We also compared the performance of LLMs with two state-of-the-art LLMs, **GPT-3.5** and **GPT-4**, and a state-of-the-art LLM by Anthropic, **Claude-2**. We used the **Top-10 Hits**, **Top-25 Hits**, **AP score**, and **AUC score** as evaluation metrics. The detailed experimental settings are described in the following subsections.

Outlier Type	#Agents	Source	Period	#Outliers
hunger	1000	POL	450+14 days	90
work	1000	POL	450+14 days	30
social	1000	POL	450+14 days	30
combined	3000	POL	450+14 days	150
imposter	69	GeoLife	4 years	20

Table 1: Specification of the datasets utilized in this paper.

2.2 Simulation and Dataset Generation

This subsection describes both the simulated using the Patterns-of-Life Simulation [4, 6, 15, 16] and real-world dataset based on the GeoLife dataset [38, 43]. Specifications of the datasets, including details and key attributes, can be found in Table 1. The source code of the simulation and data processing of the GeoLife dataset is accessible through the GitHub repository: https://github.com/onspatial/geolife-outlier-dataset. In addition, the dataset is available for download at https://osf.io/rxnz7/.

2.2.1 Simulation of Patterns of Life . The patterns of life simulation was designed to emulate human needs and behavior in an urban environment [45]. Within the simulated environment, virtual entities referred to as agents perform actions that mirror human activities. These include attending work, forming friendships, engaging in social gatherings, and more. The agents' existence is crafted to resemble human life in a real-world environment (roads, buildings) obtained from OpenStreetMap [2]. Throughout their simulated lives, agents navigate to diverse locations, including restaurants, workplaces, residential apartments, and recreational venues. A salient feature of the simulation is the generation of comprehensive log files. These logs contain extensive data regarding the agents, including their location and current state information, thus allowing for in-depth analysis and research.

In our study, we generated data by running simulations over four distinct maps, namely Fairfax County, Virginia, USA (FVA); the French Quarter of New Orleans, Louisiana, USA (NOLA); Atlanta, Georgia, USA (ATL); and Beijing, China (BJNG). The simulations were conducted over a period of 450 days to replicate normal life, followed by an additional 14 days to incorporate abnormal behavior into the regular patterns. We introduced three specific types of abnormal behavior that define outliers trajectories:

- **Hunger outlier:** An agent under this category becomes hungry more quickly. Such agents have to go to restaurants or their homes much more often.
- **Social outlier:** This type of agent randomly selects recreational sites to visit when needed, rather than being guided by their attributes and social network.
- **Work outlier:** Agents in this category abstain from going to work on workdays.

We further divided these abnormalities into three intensity levels: red, orange, and yellow. Red outliers exhibit extremely abnormal behavior, orange outliers act moderately abnormal, and yellow outliers display abnormal behavior less frequently. For example, a work outlier will decide not to go to work 100%, 50%, or 20% of the time when classified as red, orange, or yellow, respectively. We divide the simulation into 450 simulation of days of normal behavior, followed by 14 days of a small number of agents exhibiting outlier behavior. Details are in Table 1, with an extended dataset in [5].

2.2.2 Real World Dataset (GeoLife). The real-world dataset for this study was created using the Microsoft Research Asia's GPS

Table 2: Outlier detection performance for all datasets. *We report Top-25 Hits instead of Top-100 for Geolife dataset due to their size constraints on datasets. (-) denotes the absence of experiments due to the API cost issue.

	Geolife				Patterns-of-Life			
Model	Top-10 Hits	Top-25 Hits*	AP score	AUC score	Top-10 Hits	Top-100 Hits	AP score	AUC score
OMPAD	1	4	0.1665	0.1697	0	0	0.0079	0.4512
MoNav-TT	0	7	0.2849	0.3989	0	0	0.0094	0.4798
TRAOD	4	7	0.1060	0.5498	0	1	0.0030	0.4390
DSVDD	7	15	0.6246	0.7714	1	2	0.0120	0.5398
DAE	5	12	0.4627	0.6234	0	1	0.0089	0.4649
GPT-3.5	5	8	0.4014	0.4979	0	6	0.0365	0.7572
GPT-3.5-with-hint	4	12	0.3741	0.5917	0	2	0.0176	0.6220
GPT-4	3	9	0.2732	0.4417	-	-	-	-
GPT-4-with-hint	5	8	0.3181	0.4818	-	-	-	-
Claude-2	4	13	0.4756	0.7474	-	-	-	-
Claude-2-with-hint	7	16	0.6879	0.8875	-	-	-	-

Trajectory dataset [43]. Since the original data did not conform to a check-in format, we employed the method outlined in [20] to extract stay points, thereby transforming the data to fit the check-in pattern used in life simulation studies. Next, we utilized OpenStreetMap to categorize locations into four groups: apartments, workplaces, pubs, and restaurants. Given that OpenStreetMap encompasses a broad array of categories and types, we manually classified them into these four distinct groups. Upon preprocessing the data, we eliminated agents with fewer than 50 records, resulting in a final count of 69 agents with a total of 14,080 training trajectories and 3,552 test trajectories. Within the context of the GeoLife dataset, we introduced a specific outlier type called the "imposter outlier". An agent acting as an imposter outlier by switching the trajectories with another agent after a specific time point. The dataset was then divided into two segments: 80% of the stay points for training and introduced outliers into the remaining 20% for test.

3 Experimental Results

3.1 Experimental Settings

We conducted the experiments on two human mobility benchmark datasets: GeoLife [43] and Patterns-of-Life [6, 15]. We compared the performance of LLMs with several unsupervised trajectory outlier detection methods, including three non-deep learning methods **OMPAD** [9], **MoNav-TT** [37] and **TRAOD** [18], and two state-of-the-art deep learning methods **DSVDD** [28] and **DAE** [13, 44].

3.2 LLM Detection Results

Broadly, this paper focuses on studying the central question of investigating the capabilities of LLMs on identifying anomalous behaviors within human mobility patterns from three perspectives:

- **Can LLMs effectively detect anomalous behaviors within human mobility patterns without any indicative information?** It is intriguing to assess whether LLMs can attain substantial predictive performance on anomaly detection tasks, even in the absence of any clue about the anomalies, e.g. such as temporal occurrence or the nature of the anomaly.
- **Can providing indicative clues about the anomaly enhance the detection efficacy of LLMs?** Incorporating specific clues or hints about potential anomalies might bolster the LLM's ability to identify irregularities more accurately. By offering contextual information, it could guide the LLM to focus on certain aspects of the data and make more informed predictions.
- **Can LLMs provide reasonable explanation to their judgements?** Beyond mere classification, it is imperative to observe whether LLMs can elucidate the fundamental reasoning behind their determinations. Specifically, can these models articulate the underlying rationale when predicting human mobility patterns as anomalous or normal, thereby enhancing the transparency and trustworthiness of their judgments?
- **LLMs can effectively detect anomaly behaviors without any indicative information.** We observed that the LLM demonstrates commendable detection results on both datasets. For the Geolife dataset, Claude-2 surpasses all non-deep learning methods, achieving performance on par with the deep learning method. Both GPT-3.5 and GPT-4 also produce results that are comparable to those of other methods. This might suggest that presenting all mobility trajectories in a single prompt may lead to better performance than using separate prompts. As for the PoL dataset, the GPT-3.5 model significantly outperforms all the methods it was compared against.
- **Providing additional indicative information can further enhance the detection efficacy of LLMs.** We observed that by incorporating a 'hint' into the LLMs, detection performance consistently improved across all models when tested on the Geolife dataset. Notably, Claude-2-with-hint demonstrated a significantly superior detection rate, surpassing all other comparison methods. On the other hand, there was a slight dip in performance on the PoL dataset when adding the hint. This could arguably be attributed to the LLM's ability to manage longer input temporal trajectories, as evidenced by the average length of trajectories being 52.2 for Geolife and 182.0 for PoL.
- **LLMs is capable to provide reasonable explanation to their judgements.** Examples of generated explanations alongside predictions can be found on GitHub for the Geolife dataset. Notably, we observed that the LLMs are capable of providing cogent explanations for their prediction results. Such clarity is pivotal for ensuring transparency in anomaly detection methods.

4 Conclusion

In this work, we conduct empirical studies to provide insights on the strengths and limitations of large language models (LLMs) for detecting anomalous behaviors from mobility data, by comparing LLMs to specialized anomaly detection methods. Our key findings show that LLMs can achieve promising anomaly detection performance even without any specific cues about potential anomalies. Furthermore, providing contextual information about possible irregularities can enhance the prediction accuracy of LLMs. In addition, LLMs can provide explanations for their anomaly judgments, thereby improving model transparency. Our results suggest that LLMs can be a valuable tool for detecting anomalies in human mobility data, offering a new perspective on the application of LLMs in the field of anomaly detection.

For future work, we plan to study the effectiveness of open source LLMs such as Llama-2 models to improve model transparency. We also aim to address the issue that LLMs have difficulty processing long mobility trajectories due to the limited context window size. Moreover, we intend to evaluate our approach on additional mobility datasets. This work represents an initial exploration of applying LLMs for the important and promising task of mobility anomaly detection. We hope it will inspire more research in this direction.

5 Acknowledgements

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0025. 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] 2023. *Google Gemini*. https://gemini.google.com

[2] 2023. *OpenStreetMap*. https://www.openstreetmap.org

[3] Josh Achiam, Steven Adler, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).

[4] Hossein Amiri, Will Kohn, et al. 2024. The Patterns of Life Human Mobility Simulation. (2024). arXiv:2410.00185

[5] Hossein Amiri, Ruochen Kong, and Andreas Zufle. 2024. Urban Anomalies: A Simulated Human Mobility Dataset with Injected Anomalies. (2024). arXiv:2410.01844

[6] Hossein Amiri, Shiyang Ruan, et al. 2023. Massive Trajectory Data Based on Patterns of Life. In *SIGSPATIAL'23*. ACM, 1–4.

[7] Guangji Bai, Ziyang Yu, et al. 2023. Staleness-Alleviated Distributed GNN Training via Online Dynamic-Embedding Prediction. *arXiv preprint arXiv:2308.13466* (2023).

[8] Yuntao Bai, Andy Jones, Kamal Ndousse, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862* (2022).

[9] Arslan Basharat, Alexei Gritai, and Mubarak Shah. 2008. Learning object motion patterns for anomaly detection and improved object detection. In *2008 IEEE conference on computer vision and pattern recognition*. IEEE, 1–8.

[10] Filip Biljecki, Hugo Ledoux, and Peter Van Oosterom. 2013. Transportation mode-based segmentation and classification of movement trajectories. *International Journal of Geographical Information Science* 27, 2 (2013), 385–407.

[11] Tom Brown, Benjamin Mann, Nick Ryder, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.

[12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[13] Dario Dotti, Mirela Popa, and Stylianos Asteriadis. 2020. A hierarchical autoencoder learning model for path prediction and abnormality detection. *Pattern Recognition Letters* 130 (2020), 216–224.

[14] Xiaojie Guo, Shiyu Wang, and Liang Zhao. 2022. Graph neural networks: Graph transformation. *Graph Neural Networks: Foundations, Frontiers, and Applications* (2022), 251–275.

[15] Joon-Seok Kim, Hyunjee Jin, et al. 2020. Location-based social network data generation based on patterns of life. In *2020 21st IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 158–167.

[16] Will Kohn, Hossein Amiri, and Andreas Züfle. 2023. EPIPOL: An Epidemiological Patterns of Life Simulation (Demonstration Paper). In *SIGSPATIAL SpatialEpi'23 Workshop*. ACM, 13–16.

[17] Takeshi Kojima, Shixiang Shane Gu, et al. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems* 35 (2022), 22199–22213.

[18] Jae-Gil Lee, Jiawei Han, and Xiaolei Li. 2008. Trajectory outlier detection: A partition-and-detect framework. In *2008 IEEE 24th International Conference on Data Engineering*. IEEE, 140–149.

[19] Jae-Gil Lee, Jiawei Han, and Kyu-Young Whang. 2007. Trajectory clustering: a partition-and-group framework. In *SIGMOD'07*. 593–604.

[20] Quannan Li, Yu Zheng, Xing Xie, Yukun Chen, Wenyu Liu, and Wei-Ying Ma. 2008. Mining user similarity based on location history. In *Proceedings of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems*. 1–10.

[21] Chen Ling, Xujiang Zhao, Jiaying Lu, et al. 2023. Beyond one-model-fits-all: A survey of domain specialization for large language models. *arXiv preprint arXiv:2305.18703* (2023).

[22] Yueyang Liu, Lance Kennedy, Hossein Amiri, and Andreas Züfle. 2024. Neural Collaborative Filtering to Detect Anomalies in Human Semantic Trajectories. arXiv:2409.18427

[23] Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2021. Metaicl: Learning to learn in context. *arXiv preprint arXiv:2110.15943* (2021).

[24] Mohamed Mokbel, Mahmoud Sakr, Li Xiong, et al. 2024. Mobility Data Science: Perspectives and Challenges. *ACM Transactions on Spatial Algorithms and Systems* 10 (2024). Issue 2.

[25] Christine Parent, Stefano Spaccapietra, et al. 2013. Semantic trajectories modeling and analysis. *ACM Computing Surveys (CSUR)* 45, 4 (2013), 1–32.

[26] Xipeng Qiu, Tianxiang Sun, Yige Xu, et al. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences* 63, 10 (2020), 1872–1897.

[27] Sirisha Rambhatla, Sepanta Zeighami, et al. 2022. Toward accurate spatiotemporal COVID-19 risk scores using high-resolution real-world mobility data. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 8, 2 (2022), 1–30.

[28] Lukas Ruff, Robert Vandermeulen, et al. 2018. Deep one-class classification. In *International conference on machine learning*. PMLR, 4393–4402.

[29] Chris Stanford, Suman Adari, et al. 2024. NUMOSIM: A Synthetic Mobility Dataset with Anomaly Detection Benchmarks. *arXiv preprint arXiv:2409.03024* (2024).

[30] Thanos G Stavropoulos, Asterios Papastergiou, et al. 2020. IoT wearable sensors and devices in elderly care: A literature review. *Sensors* 20, 10 (2020), 2826.

[31] Kevin Toohey and Matt Duckham. 2015. Trajectory similarity measures. *Sigspatial Special* 7, 1 (2015), 43–50.

[32] Hugo Touvron, Thibaut Lavril, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).

[33] Hugo Touvron, Louis Martin, Kevin Stone, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).

[34] Ashish Vaswani, Noam Shazeer, et al. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[35] Jason Wei, Maarten Bosma, et al. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652* (2021).

[36] Di Yao, Chao Zhang, Zhihua Zhu, Jianhui Huang, and Jingping Bi. 2017. Trajectory clustering via deep representation learning. In *2017 international joint conference on neural networks (IJCNN)*. IEEE, 3880–3887.

[37] Jianting Zhang. 2012. Smarter outlier detection and deeper understanding of large-scale taxi trip records: a case study of NYC. In *Proceedings of the ACM SIGKDD International Workshop on Urban Computing*. 157–162.

[38] Zheng Zhang, Hossein Amiri, et al. 2024. Transferable Unsupervised Outlier Detection Framework for Human Semantic Trajectories. (2024). arXiv:2410.00054

[39] Zheng Zhang and Liang Zhao. 2021. Representation learning on spatial networks. *Advances in Neural Information Processing Systems* 34 (2021), 2303–2318.

[40] Zheng Zhang and Liang Zhao. 2022. Unsupervised deep subgraph anomaly detection. In *ICDM*. IEEE, 753–762.

[41] Wayne Xin Zhao, Kun Zhou, Junyi Li, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).

[42] Yu Zheng, Yukun Chen, Quannan Li, Xing Xie, and Wei-Ying Ma. 2010. Understanding transportation modes based on GPS data for web applications. *ACM Transactions on the Web (TWEB)* 4, 1 (2010), 1–36.

[43] Yu Zheng, Xing Xie, Wei-Ying Ma, et al. 2010. GeoLife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.* 33, 2 (2010), 32–39.

[44] Chong Zhou and Randy C Paffenroth. 2017. Anomaly detection with robust deep autoencoders. In *KDD '17*. 665–674.

[45] Andreas Züfle, Carola Wenk, et al. 2023. Urban life: a model of people and places. *Computational and Mathematical Organization Theory* 29, 1 (2023), 20–51.