

Kinematic Detection of Anomalies in Human Trajectory Data

Lance Kennedy
Emory University, USA
lance.kennedy@emory.edu

Andreas Züfle
Emory University, USA
azufle@emory.edu

Abstract

Historically, much of the research in understanding, modeling, and mining human trajectory data has focused on where an individual stays. Thus, the focus of existing research has been on *where* a user goes. On the other hand, the study of *how* a user moves between locations has great potential for new research opportunities. Kinematic features describe how an individual moves between locations and can be used for tasks such as identification of individuals or anomaly detection. Unfortunately, data availability and quality challenges make kinematic trajectory mining difficult. In this paper, we leverage the Geolife dataset of human trajectories to investigate the viability of using kinematic features to identify individuals and detect anomalies. We show that humans have an individual “kinematic profile” which can be used as a strong signal to identify individual humans. We experimentally show that, for the two use-cases of individual identification and anomaly detection, simple kinematic features fed to standard classification and anomaly detection algorithms significantly improve results.

CCS Concepts

• Information systems → Geographic information systems.

Keywords

Human Trajectory Data, Trajectory Outlier Detection

ACM Reference Format:

Lance Kennedy and Andreas Züfle. 2024. Kinematic Detection of Anomalies in Human Trajectory Data. 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/10.1145/3681765.3698464>

1 Introduction

Most often, when working with human trajectory data, research has focused on mining and analyzing the locations where a user arrives and stays (or what we may call ‘staypoints’). Using these staypoints, one can learn about a user’s behavior and daily/weekly/monthly patterns. In fact, having just 4 of an individual’s staypoints is sufficient to uniquely identify them [2]. This makes an individual’s staypoints Personally Identifiable Information (PII), requiring certain privacy measures to be taken when working with them. However, staypoints comprise only a part of human trajectory data. We call the movement of users on roads or paths between these staypoints ‘trips.’ Very little research has been done on identifying

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

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

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1144-2/24/10

<https://doi.org/10.1145/3681765.3698464>

individuals based on their trips, since studies are generally more focused on *where* an individual goes, not how they get there.

That said, if a user’s staypoints can be used to uniquely identify them, then it stands to reason that their trips could as well. Consider two individuals who drive in their car from one staypoint to another at the exact same time. Based solely on their staypoints, they would appear the exact same, but in fact they likely *moved* between the staypoints with completely different kinematic behavior. Some individuals drive faster, while some speed up or slow down more suddenly. Each individual should have a *kinematic profile* which is unique to them and how they move from place to place.

These kinematic profiles may also be PII. Additionally, they could be used to identify kinematic anomalies. Consider, for example, if an individual’s phone was stolen. The phone, which tracks the GPS location of the user, would know the kinematic profile of how the owner moves. Sensing a kinematic anomaly in the kinematic profile of the thief, it could take security measures to protect the owner’s data. Alternatively, a ship may veer off-course to partake in illegal fishing activities, but falsify its trajectory during that time to appear that it stayed on course. The falsified trajectory could be flagged as kinematically different from the ship’s normal trajectory. While this is all idealistic, it is a challenging problem for several reasons. First, it is hard to collect high-quality human trajectory data. For reasons already discussed, these data are protected for privacy [15], so it needs to be collected from anonymous volunteers. One of the most well-known examples is the Geolife GPS trajectory dataset [16, 17], which collected GPS data from 182 users in Beijing China, over a period of several years. Unfortunately, due to the volunteers selectively using the GPS tracking, inconsistent sampling rates, and GPS location errors, it is still difficult to mine accurate kinematic data from this dataset. However, our results show that we can still do so sufficiently to both classify users by their kinematic data, and identify kinematic anomalies.

This paper is organized as follows: In Section 2, we highlight related works. In Section 3, we describe the dataset, then explain how we mine the kinematic features in Section 4. In Section 5, we present our experimental results. We provide case studies for users on whom the models worked well and conclude in Section 6.

2 Related Works

Human trajectory data is spatio-temporal data which records the locations of humans at different timestamps. It has been shown to be PII, as it can be used accurately for user classification [2, 12].

Meng et al. [11] surveys and categorizes various methods of anomaly detection trajectory data. Lee et al. [7], for instance, partitions trajectories into sub-trajectories and focuses on grouping these, rather than looking at them in whole. This serves as the basis for many trajectory anomaly detection algorithms.

Due to privacy concerns, much recent research uses simulated data such as Stanford et al. [13]. Some such methods for trajectory outlier detection employ deep neural networks [3], large language

models [14], or more explainable approaches such as collaborative filtering [9]. However, simulations usually do not model realistic kinematic features, so we focus on real-life data for this work.

Some trajectory anomaly detection algorithms which consider kinematic features concern maritime vessels [4, 8], but only consider speed and direction. Lu & Xia [10] also consider speed, acceleration, and direction for trajectory classification with great success.

3 Description of Data

The Geolife GPS trajectory dataset was collected from 182 individuals in Beijing, China between 2007 and 2012. It records, for each user, a sequence of timestamps, latitudes, longitudes, and other spatial data such as altitude over some period of time.

While many of the users carried their GPS trackers over long periods of times, about 24% only did so for less than a week, and only about 42% of them did so for more than a month. This reduced the amount of users with sufficient, usable data. Furthermore, many of the users exhibit inconsistent sampling rates, where the amount of time between timestamps can vary wildly. This not only makes it more difficult to mine kinematic features from the users, but along with GPS errors, makes it difficult to determine when a user is at a staypoint, trip, or simply not using the tracker. Fortunately, the dataset also provides a list of trips, annotated by their *modality*, or method of transportation (e.g. train, car, walk), for 73 of the users. These trips include a start time, end time, and modality for the user’s movements. Using these, we no longer need to analyze the raw trajectories to extract trips, although it again reduces the amount of available data. Therefore, for the following study and experiments, we select GeoLife users having (1) annotated trip modality labels, (2) trips which do not appear to have extreme GPS errors, and (3) at least 30 such trips.

4 Methodology

Using the trajectory data, and the start and end times for the trips, we can extract kinematic features per trip. We approximate the speed of the user by calculating the distance between the coordinate pairs at consecutive timestamps, and further approximate the acceleration of the user by taking the differences of consecutive speeds. Of course, neither of these can be perfectly accurate: calculating speed as distance over time assumes that the user was moving in a straight line, and calculating acceleration in this way assumes that changes in speed happen instantaneously at each timestamp. Additionally, as the sampling rate decreases, the more rough these approximations become. Using the speeds and accelerations, we extract 10 kinematic features per trip, listed in Table 1. After extracting these features, we notice that many trips include impossible values, such as trips spanning a period of over 100 days, or with speeds of up to 3000000 m/s. Therefore, we reduce the set of trips by removing those which constitute an anomaly for any of our 10 features. These anomalies are determined using the interquartile range. Finally, we take only those users which have a sufficient amount of trips after this reduction, a bound which we select as 30. This leaves only 6145 trips over 26 users which are usable for our experiments. Furthermore, the data is considerably imbalanced. The number of trips per user range from 31 to 748. On average, a user has approximately 236 trips, but with great variance; 12 of the users have less than 100 trips.

Kinematic Feature	Mean	Std Dev
Trip Duration in s	20144.587	23781.262
Maximum Speed in $\frac{m}{s}$	29.006	16.737
Minimum Speed in $\frac{m}{s}$	0.0	0.0
Max. Positive Acceleration in $\frac{m}{s^2}$	21.845	16.439
Min. Negative Acceleration in $\frac{m}{s^2}$	-20.038	15.606
Mean Speed in $\frac{m}{s}$	4.095	2.361
Mean Absolute Accelerations in $\frac{m}{s^2}$	0.946	0.340
Std Dev of Speed in $\frac{m}{s}$	3.889	1.963
Std Dev of Acceleration in $\frac{m}{s^2}$	1.978	0.942
Std Dev of Absolute Accelerations in $\frac{m}{s^2}$	1.714	0.930

Table 1: Summary of Kinematic Features Used in this Work.

Model	Accuracy	ROC-AUC	F1 score
Decision Tree	.303 ± .008	.589 ± .004	.208 ± .006
Weighted Guess	.079 ± .005	.500	.041 ± .005
Random Guess	.035 ± .003	.500	.026 ± .002

Table 2: Results compared to random classifiers

5 Experiments and Results

We provide two experiments: in Section 5.1, we use the kinematic features to perform user-wise classification of the trips. This is done with a simple decision tree over the kinematic features. In Section 5.2, we perform anomaly detection by taking the trips of a single user and adding trips randomly selected from other users as anomalies. This detection is done using the Local Outlier Factor (LOF) algorithm [1]. In both experiments, we provide a case study for one user who provides particularly promising results. The code to generate the data and run these experiments is available at <https://github.com/lancek23/trajectory-kinematics>.

5.1 Classification

Trips are classified user-wise over 5 stratified folds, helping the decision tree to train and be tested on a sufficient amount of data per user. We evaluate the results of classification by the following metrics: accuracy, area under the receiver operating characteristic curve (ROC-AUC), and F1 score. The ROC-AUC helps us ensure that the model works better than a random classifier. The F1 score uses “macro” aggregation, taking the unweighted mean of scores across the different class labels, and serves as a conservative metric for class-imbalanced problems based on precision and recall.

We provide two random classifiers against which we measure our results. The first classifier uses a “weighted random guess” in which it predicts a label with a probability equal to the proportion of that label in the dataset. The second classifier uses a “true random guess,” and predicts each label with equal probability. The results can be seen in Table 2, aggregated over the 5 folds. By all metrics, our model is superior to any random classifier, correctly classifying around 30% of all trips in the dataset.

We now provide an example confusion matrix in Figure 1. The user labels are organized in this matrix from most to least trips, for convenience. We can observe that the decision tree lacks the bias that the weighted random guess will have towards those users with more trips, and is still able to correctly predict some of the trips for those users with relatively few trips. We observe that the kinematics based user identification is substantially better than random guessing. We also observe that for some users, the classification is very accurate. For example, for User 068, we observe a Recall of 36%




Figure 1: Confusion matrix for the decision tree model. Users are sorted by their number of trips.




Figure 2: A plot of kinematic features of 5 users. Note the dense purple cluster formed by one of the users. and a Precision of 38%. For User 128, we observe a Recall of 69% and a Precision of 71%. These are substantially higher than the Precision/Recall values of $1/26 \sim 4\%$ that we would expect using random guessing. We also observe that the model confuses some users. For example, User 126 is frequently misclassified as User 167, having a Recall of only 6% and a Precision of only 5%. What makes some users easier to detect kinematically than others?

5.1.1 Case Study. To investigate this question, we provide a case study for a user we call "User 128." User 128 has noticeable kinematic patterns, particularly in their maximum speed and in their standard deviation of absolute accelerations per trip. The trips of User 128 make up the dense, purple cluster seen in Figure 2. We can again observe the kinematic profile of User 128 in Figure 3, which plots the maximum speed vs the average speed of the trips for a few users. We see that despite changes in average speed, the maximum speed achieved by User 128 appears consistent. Alternatively, some of the other users plots look nearly uniform, with less apparent trends. In reality, it is likely that the other users achieve more consistent maximum speeds during different trips, but data collection errors and low sampling frequencies lead to issues in capturing this.

User 128 was consistently easy for the decision tree to classify. In Figure 1, we see that more of the true positives we find are from User 128 than any other, despite having only the third-most trips in




Figure 3: A plot of kinematic features of 5 users. Note the consistency of max speed for the user plotted in purple.

	LOF	Random Classifier
Mean	0.087	0.046
Std	0.164	0.057
Min	0.014	0.011
Median	0.042	0.030
Max	1.000	0.518

Table 3: PR-AUC summary statistics

the dataset. While this is partially explainable from decision trees having a natural bias towards labels with more observations in a dataset, User 128 is still much easier to identify than the other users with a large number of trips. This is because users 128’s trips are more well-clustered than the trips of other users. User 128 provides an example of what is possible in using kinematics to discriminate between human trajectories. With more accurate datasets in the future, we hope to show that the results for User 128 can be replicated on a greater number of human trajectories.

5.2 Anomaly Detection

To perform anomaly detection, we first inject anomalous mobility by swapping user IDs of trips, thus assigning a trip of one user to another user. For each user, we randomly select trips from the other users up to roughly 3% of the number of trips the original user has in the dataset. These trips together form our new dataset.

Using the LOF algorithm, we assign an anomaly score to each trip based on their local density [1]. We expect that the trips of the normal user will be kinematically similar and will form a dense cluster, while the anomaly trips will be far from this cluster with few nearby neighbors. Furthermore, the different modalities (walk, car, etc) of the normal user should each form their own cluster, with the inserted trips apart from each of them. Sorting the trips by their anomaly score, we find the area under the precision-recall curve (PR-AUC), which is a useful indicator for binary data.

We allow each of our 26 users to serve as the “normal” user for 10 trials, randomly adding anomalies for each trial as described, for 260 total trials. We also run the same experiments but assign anomaly scores using a random classifier rather than the LOF algorithm as a baseline. The results are summarized in Table 3.

These results are, unfortunately, not as good as we had hoped. We’re still able to outperform a random baseline, which uniformly at random selects 3% of trips as anomalies. But the results of the LOF having a mean AUC-PR of 8.7% is not practically useful. This is surprising, as we found that the decision tree classifier was able




Figure 4: A plot of kinematic features of 4 users. Note the consistency of average speed for the user plotted in red.

to accurately identify the user that generated a trip. Yet, the LOF struggles to detect an agent changing their identity. A problem might be the large number of kinematic features that we consider, which may cause the LOF, which computes distances between multi-dimensional points, to run into the Curse of Dimensionality [5]. It may be possible to automatically find the right subspaces to detect outliers in as proposed in [6], but we leave this study for future work. Additionally, we believe that much of this is due to the previously mentioned issues with data collection. If there are errors in mining the kinematic features of a user, then it may be impossible to find the true patterns for that user and build an accurate kinematic profile. However, there are some users which we are able to perform better on than others. We believe that these users are not necessarily *easier* to discriminate from others, but rather that these users have more accurate data, allowing us to build accurate kinematic profiles of them. We provide a case study of one such user in section 5.2.1.

5.2.1 Case Study. The user we call “User 147” consistently performed well in anomaly detection. Figure 4 shows this user, along with some others, plotted by two kinematic features. We notice that User 147, plotted in red, is very consistent in its average speed compared to the other users. This explains why User 147 is so easy to find anomalies with: any trip with an average speed outside of its typical range can be marked as an anomaly. In the classification problem, User 147 wasn’t as easy to identify due to so many other trips having similar average speeds, but when it represents the majority “normal” class, this is no longer an issue. In fact, in the 10 trials where User 147 was “normal,” the LOF algorithm had an average PR-AUC of 0.29, even achieving 0.792 in its best trial. User 147 demonstrates an important principle: when the user has a consistent kinematic profile, anomaly detection becomes easy.

6 Conclusion and Future Work

This paper provides an early look into the potential of kinematic data mined from human trajectory data. Using relatively simple models, we are able to classify trajectories by user, and identify anomalous trips from a user of interest.

This has important implications in privacy, as kinematic data may be PII, and therefore important to keep under privacy restrictions. It also has practical uses in identifying unusual or falsified

trajectories. These may signify a phone being stolen as the kinematic profile of the user changes, or a ship falsifying their trajectory to cover up illegal behavior. With more accurate datasets in the future, the ability to perform classification and anomaly detection tasks using kinematic data could be even greater. In addition, more complex models may be able to detect new patterns in the kinematic profiles, yielding improved performance.

7 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] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. 2000. LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*. 93–104.
- [2] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleyesen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3, 1 (2013), 1–5.
- [3] Minxuan Duan, Yinlong Qian, et al. 2024. Back to Bayesics: Uncovering Human Mobility Distributions and Anomalies with an Integrated Statistical and Neural Framework. (2024). arXiv:2410.01011 <https://arxiv.org/abs/2410.01011>
- [4] Peigu Fu, Haozhou Wang, Kuien Liu, Xiaohui Hu, and Hui Zhang. 2017. Finding abnormal vessel trajectories using feature learning. *IEEE Access* 5 (2017), 7898–7909.
- [5] Piotr Indyk and Rajeev Motwani. 1998. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. 604–613.
- [6] Hans-Peter Kriegel, Peer Kröger, Erich Schubert, and Arthur Zimek. 2009. Outlier detection in axis-parallel subspaces of high dimensional data. In *Advances in Knowledge Discovery and Data Mining: 13th Pacific-Asia Conference, PAKDD 2009 Bangkok, Thailand, April 27–30, 2009 Proceedings 13*. Springer, 831–838.
- [7] Jae-Gil Lee, Jiawei Han, and Kyu-Young Whang. 2007. Trajectory clustering: a partition-and-group framework. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. 593–604.
- [8] Bao Lei and Du Mingchao. 2018. A distance-based trajectory outlier detection method on maritime traffic data. In *2018 4th International Conference on Control, Automation and Robotics (ICCAR)*. IEEE, 340–343.
- [9] Yueyang Liu, Lance Kennedy, Hossein Amiri, and Andreas Züfle. 2024. Neural Collaborative Filtering to Detect Anomalies in Human Semantic Trajectories. arXiv:2409.18427 [cs.LG] <https://arxiv.org/abs/2409.18427>
- [10] Sichong Lu and Ying Xia. 2020. Dual Supervised Autoencoder Based Trajectory Classification Using Enhanced Spatio-Temporal Information. *IEEE Access* 8 (2020), 173918–173932. <https://api.semanticscholar.org/CorpusID:222096600>
- [11] Fanrong Meng, Guan Yuan, Shaoqian Lv, Zhixiao Wang, and Shixiong Xia. 2019. An overview on trajectory outlier detection. *Artificial Intelligence Review* 52 (2019), 2437–2456.
- [12] Erik Seglem, Andreas Züfle, Jan Stutzki, Felix Borutta, Evgheniy Faerman, and Matthias Schubert. 2017. On privacy in spatio-temporal data: User identification using microblog data. In *SSTD 2017*. Springer, 43–61.
- [13] Chris Stanford, Suman Adari, et al. 2024. NUMOSIM: A Synthetic Mobility Dataset with Anomaly Detection Benchmarks. *arXiv preprint arXiv:2409.03024* (2024).
- [14] Zheng Zhang, Hossein Amiri, Zhenke Liu, Andreas Züfle, and Liang Zhao. 2023. Large language models for spatial trajectory patterns mining. *arXiv preprint arXiv:2310.04942* (2023).
- [15] Yu Zheng. 2015. Trajectory data mining: an overview. *ACM Transactions on Intelligent Systems and Technology (TIST)* 6, 3 (2015), 1–41.
- [16] Yu Zheng, Quannan Li, Yukun Chen, Xing Xie, and Wei-Ying Ma. 2008. Understanding mobility based on GPS data. In *UbiComp ’08* (Seoul, Korea). Association for Computing Machinery, New York, NY, USA, 312–321. <https://doi.org/10.1145/1409635.1409677>
- [17] 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.