Exploring Spatio-temporal Movements for Intelligent Mobility Services

Tobias Grüner[1][a], Sören Frey[1][b], Jens Nahm[1] and Dirk Reichardt[2]

[1]*Independent Researcher, 70563 Stuttgart, Germany*

[2]*Baden-Wuerttemberg Cooperative State University Stuttgart (DHBW Stuttgart), Jägerstraße 56, 70174 Stuttgart, Germany*

Keywords: Mobility Services, Machine Learning, Prediction, Classification, POI Extraction, Clustering, Location Data.

Abstract: Mobility services can substantially benefit from incorporating movement behavior information. Models of daily travel routines can facilitate intelligent recommendations of suitable car sharing, ride pooling, or Mobility as a Service (MaaS) offerings, for instance. However, existing approaches that infer regular travel activities from historical location data exhibit several limitations. For example, they often have an insufficient resolution in the spatial and temporal dimension or are restricted to predicting only the next location visit. This paper presents an activity-based approach to model daily travel routines and predict regularities with the help of machine learning (ML). We first extract points of interest (POIs) and corresponding visits from historical location data. Then, regularities for these visits are identified with the help of classification. We validate our work in progress approach using data from voluntary, consenting test subjects (CTS) who agreed to track their movements. They labeled their own data for each activity with corresponding regularity information. We show that POI visits can already be predicted reliably for the first classes of movements.

1 INTRODUCTION

Over the last years, the automotive industry has established an ample service portfolio that augments and enhances classical transportation schemes. Corresponding mobility services address many different areas and needs. They range from trip planning, congierge, and sharing services to further offerings subsumed under the Mobility as a Service (MaaS) (Jittrapirom et al., 2017) umbrella that also encourage multi-modal transportation modes. To keep up and advance valuable assistance and support, the services have to become increasingly intelligent and personalized. Utilizing location information and understanding consenting individuals' movement behavior opens up many ways to increase convenience by optimizing their daily routines.

For example, the presence of nearby car sharing vehicles, the pre-booking of a ride pooling tour, or the most cost-efficient MaaS-based mix of transport carriers can be proactively suggested based on movement information inferred from historical location data. However, existing approaches for inferring movement behavior are less suited in the context of mobility services. For example, they often address different target domains with lower resolutions in the spatial and temporal dimension or are restricted to predicting only the next location visit.

This paper presents our activity-based work in progress approach (Grüner, 2019) to describe daily travel routines and predict regularities with the help of machine learning (ML). We introduce a movement behavior model that covers regularity and irregularity in both the temporal and spatial dimension. In our pre-processing pipeline we first extract points of interest (POIs) and corresponding visits from historical location data with the help of clustering. The identification of the movement classes for these visits is then approached as a classification problem. A group of voluntary, consenting test subjects (CTS) agreed to track their movements. They labeled the data for each of their activities with corresponding movement information. The resulting dataset is split and used for training and validating the classification model. We show that our work in progress approach can already predict POI visits reliably for the first types of movement classes. In summary, the main contributions of our work are:

- An activity-based approach to model spatio-temporal movement behavior.

- A pre-processing pipeline for the extraction of POI visits from historical location data.

[a] https://orcid.org/0000-0002-0913-0472
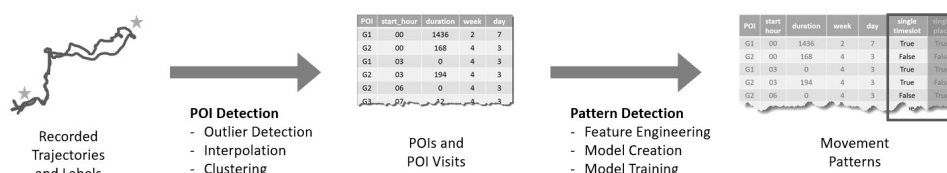[b] https://orcid.org/0000-0002-4087-6117

123

Figure 1: High-Level approach for learning and predicting spatio-temporal movement patterns.

- A validated classification model for the prediction of first types of movement information.

The remainder of the paper is structured as follows. The related work is described in the next Section 2. Our approach for modeling movement behavior, the extraction of POI visits, and the classification model used for prediction is described in Section 3. The evaluation of this model is then presented in Section 4. Section 5 discusses the experiment results before Section 6 draws the conclusions.

2 RELATED WORK

The investigation and usage of movement patterns is relevant in many areas. For example, synthetic Daily Activity-Travel Patterns (DAPs) were generated based on a household travel survey from 1991 to simulate and forecast travel demands (Kitamura et al., 1997). While DAPs rely on a sample of a single day in order to forecast travel demand for a large group of people, our work focuses on the detection of actual movement patterns of a single person.

Global Positioning System (GPS) data is used in (Ashbrook et al., 2002) to learn POIs and predict user movements with the help of a Markov model. From a current location, the model can be asked for a user's next most likely significant location (POI). Vintan et al. also propose an approach that tries to determine a person's next movement (Vintan et al., 2004). They use multi-layer perceptron neural predictors with and without pre-training. In comparison, our work utilizes a gradient boosting machine learning approach (XG-Boost). It is not restricted to predicting the immediately next POI but rather identifies inherent movement patterns independent from the current state.

Vukovic et al. discuss the prediction of movements using a hybrid solution based on user movement statistics and neural networks to identify movement regularities (Vukovic et al., 2007). The position data is gained from mobile network cell information as opposed to the GPS positioning with a higher resolution used in our approach.[1] Instead of just distinguishing regular from irregular movements, our movement patterns can represent regularity and irregularity in the spatial and temporal dimensions.

A sequential patterns data mining approach to extract frequent movement patterns of vehicles in vehicular ad-hoc networks (VANETs) is proposed in (Merah et al., 2013). The movement patterns are used to generate movement rules with associated probabilities. In contrast, we utilize machine learning (XGBoost) and are not restricted to VANETs.

3 PREDICTION OF SPATIO-TEMPORAL MOVEMENT PATTERNS

3.1 Overview

Processing Phases. This work is structured into two key phases (see Figure 1). In the *POI Detection* phase, the recorded and labeled historic location data of individuals is interpolated and POIs and visits at such are detected (see Section 3.3). The phase results in a list containing all recorded visits at all POIs of an individual. This list of visits is extended in the following *Pattern Detection* phase in order to find activities and their related movement patterns based on the regularity and similarity of visits at POIs (see Section 3.4).

To achieve this, an appropriate set of features is designed and an ML model is trained to classify the activities with their corresponding movement patterns. Basic concepts which constitute the basis for the two phases are introduced in Section 3.2

Data Basis. We recruit a small group ($n = 13$) of voluntary CTS (referred to as P1-P13) to create a data basis for developing and evaluating our proposed approach. Following EU's General Data Protection Regulation (GDPR) is of primary concern. We make sure to comply with all corresponding rules. We are interested in a comprehensive set of actual movements independent of the used means of transportation. This allows for an extensive analysis of all movements in order to recommend, for example, multi-modal

---

[1]Our work uses GPS but the approach can also be utilized with other geo-spatial positioning systems

routes. Hence, the CTSs rather record their location data using a smartphone app[2] instead of relying solely on vehicle telematics systems, for instance.

The density of the recorded tracking positions varies per CTS (see Figure 2) as they all use different devices and visit places with different GPS signal strengths. This emulates a realistic setting with heterogeneous types of user devices and behaviors. The historical location data is recorded over a period of approx. three months. For each day, the CTSs label their data for each of their activities by defining the corresponding (1) movement pattern type, (2) timeslot(s), and (3) POI(s). These concepts are described in Section 3.2. The emerging dataset is split up into a modeling dataset D_M, which is required to find a proper pre-processing pipeline configuration, and an evaluation dataset D_E, on which the approach is evaluated (see Section 4). Due to the numerous parameters which have to be determined in order to configure the pre-processing steps, D_M consists of ten (n=10) and D_E consists of three (n=3) CTSs.

3.2 Basic Concepts

This work uses a movement behavior model that utilizes an activity-based approach[3] to describe an individual's travel behavior and to link each activity to a movement pattern. A movement pattern is defined by a spatial and temporal regularity. The spatial regularity is a set of locations (POIs) that are visited at the given timeslot(s). A timeslot T defines the temporal regularity at which the activity is performed at the given location(s). There are many examples for activities like working, buying weekly groceries, and sport activities that take place at specific weekdays on weekly periodicities. Hence, a timeslot T consists of a set of weekdays ($D_v(T)$) for which a location is visited on a $p_w(T)$-weekly periodicity (see Equation 1).

$$T = (D_v(T), \ p_w(T)),$$
$$D_v(T) \subseteq \{Monday, Tuesday, \ldots, Sunday\} \wedge \quad (1)$$
$$p_w(T) \in \mathbb{N}^+, \ n \leq 53$$

Movement patterns are categorized into four types based on their spatial and temporal regularity. For each regularity dimension the types are differentiated into a regular and an irregular case. The types are RT-RL, IT-RL, RT-IL, and IT-IL as displayed in Table 1.

For example, if an individual performs an activity always on Mondays and Tuesdays every week at

---

[2]https://gpslogger.app

[3]Activity-based approaches consider travel as a consequence of individuals pursuing different activities at distributed locations (Hall, 2012).
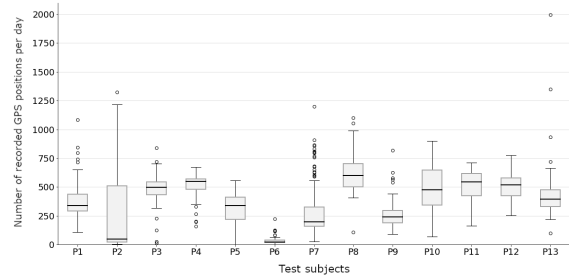
Figure 2: Density of recorded GPS positions.

Table 1: Movement pattern types.

		Temporal	
		Regular	Irregular
Spatial	Regular	Regular Time, Regular Location (*RT-RL*)	Irregular Time, Regular Location (*IT-RL*)
	Irregular	Regular Time, Irregular Location (*RT-IL*)	Irregular Time, Irregular Location (*IT-IL*)

a specific location, the activity has a movement pattern of type RT-RL and is linked to one timeslot $T = (\{Monday, Tuesday\}, 1)$. If the exemplary individual performs the same activity alternating each week (e.g. for even weeks on Monday and for uneven weeks on Tuesday), then the movement pattern of the activity is still of type RT-RL but it is now linked to two timeslots $T_1 = (\{Monday\}, 2)$ and $T_2 = (\{Tuesday\}, 2)$. If the activity takes place at different weekdays but no temporal regularity determines which weekday is visited, the movement pattern of the activity is of type IT-RL and is therefore linked to no timeslot.

3.3 POI Detection

In order to detect POIs and POI visits, the location data is filtered for outliers, interpolated, and clustered.

Outlier Removal. There is a variety of outliers in the location traces of the dataset due to inaccurate measures of the GPS devices. By interpolating the traces the impact of these errors on the data quality increases. Therefore, the location traces of the CTS group have to be filtered before they are interpolated. We use the Isolation Forest algorithm (Liu et al., 2008) to remove outliers in the data basis.

Interpolation. As a GPS signal is not always available, e.g. inside buildings (Kjærgaard, M. B. et al., 2010), the recorded GPS positions of the CTS group can be sparse and therefore include gaps (see Figure 3a). The GPS traces are linearly interpolated in order to create a continuous history (see Figure 3b).
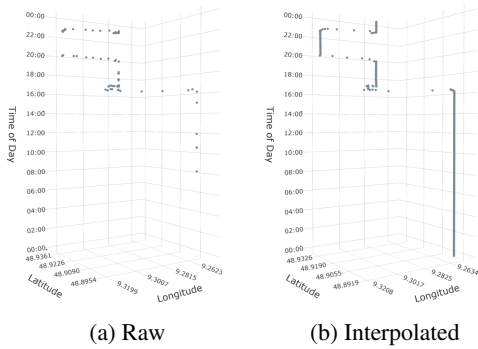
(a) Raw   (b) Interpolated

Figure 3: The sparse location data of an exemplary day from a single CTS (a) is interpolated every 120 seconds (b). Four visits at three POIs are recorded for this day.

The interpolation interval depends on the use case, since a longer interpolation interval eliminates short stops at locations. We use a 120 second interval to interpolate the location data of the CTSs in order to also detect visits of short duration.

Clustering. The interpolated location traces are clustered using the Density-based Spatial Clustering of Applications with Noise (DBSCAN) algorithm (Ester et al., 1996). The maximum point distance ε and the minimum number of cluster points *min_pts* for DBSCAN are empirically determined with the help of D_M. Here, the parameters $\varepsilon = 0.2$ and *min_pts* = 40 identify the most POIs according to the F1 score (see Figure 4). To measure the correct identification of POIs via DBSCAN, the identified cluster centers are compared to the known POIs for a set of test days. An uncertainty radius r_u has to be determined within which locations cannot be distinguished due to the accuracy of the GPS signal. We estimate this radius by calculating the median of all standard distances (Bachi, 2005) of all GPS points for those days on which the CTS stays at one POI for the whole day. On our data set D_M we approximate the radius with $r_u \approx 300m$. The cluster centroids are considered as POIs. A POI can be visited multiple times a day. The visits are extracted from the POI clusters as follows. The corresponding points of each cluster are ordered temporally. Every time the timespan between two points of a cluster exceeds a certain threshold[4] the following points are considered as another visit.

3.4 Temporal Pattern Detection

To detect temporal patterns two feature types are created and an ML model is trained. For the detection

---

[4]As the GPS traces are interpolated there are no time gaps between consecutive tracking points of a cluster.
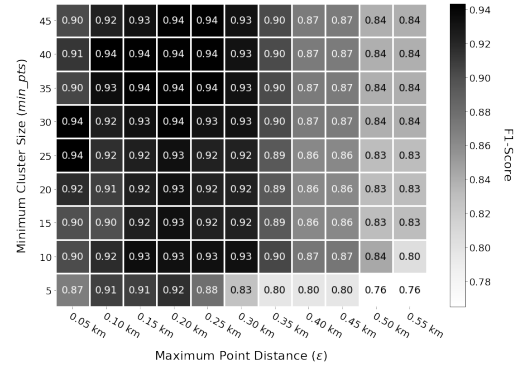
Figure 4: POI detection performance of DBSCAN depending on ε and *min_pts* parameters with $r_u = 300m$.

of spatial patterns the similarity of visits at different locations has to be measured. As only an insufficient number of activities are labeled with the RT-IL pattern type by our CTSs, the feature types for the detection of this pattern will be addressed separately in our future work.

Feature Engineering. To indicate whether visits at a certain location are occurrences of a temporally regular or irregular activity, the feature types $cv_{vpd}(T)$ (visits per day coefficient of variation) and $c_{dpw}(T)$ (days per k-th week coefficient) are introduced. Each feature type represents a set of features. Each feature corresponds to a specific timeslot. In order to determine the likeliness of a specific timeslot to be present, each place and its visits are considered separately. If a timeslot T is present, the corresponding place should be equally often visited for all of its weekdays ($D_v(T)$) over the entire observed time. To measure the uniformity of the distribution of weekdays at a given place, the visits per day coefficient of variation ($cv_{vpd}(T)$) is defined as the quotient of the standard deviation between the visits per weekday ($s_{vpd}(T)$) and the mean visits per weekday ($\overline{n_{vpd}}(T)$, see Equation 2) of the given timeslot T.

$$cv_{vpd}(T) = min(\frac{s_{vpd}(T)}{\overline{n_{vpd}}(T)}, 1) \qquad (2)$$

To determine the dividend ($s_{vpd}(T)$) and the divisor ($\overline{n_{vpd}}(T)$), only the weekdays $D_v(T)$ of the timeslot T are considered. The coefficient of variation is restricted to the value range $cv_{vpd}(T) \in [0, 1]$. The closer $cv_{vpd}(T)$ is to 0 the more likely a movement pattern with a temporal regularity for the given timeslot is present. An example for calculating $cv_{vpd}(T)$ is provided later in this section.

The feature $c_{dpw}(T)$ (see Equation 3) indicates how close the average number of visited days per

$p_w(T)$-th week ($\overline{n_{dpw}}(T)$) is to the overall number of different days on which a place is visited ($|D_v(T)|$). The place should be visited on each of the timeslot's weekdays in every $p_w(T)$-th week for the timeslot T to be present. The closer $c_{dpw}(T)$ is to the maximum 1.0, the more likely a movement pattern with a temporal regularity for the given timeslot T is present.

$$c_{dpw}(T) = \frac{\overline{n_{dpw}(T)}}{|D_v(T)|} \qquad (3)$$

Example: In Table 2, exemplary visits for an activity with a movement pattern of type RT-RL is present and a single (correct) timeslot $T_c = (\{Tuesday, Thursday\}, 1)$ at a specific place is displayed. Two outliers from the timeslot are present in the example: On Monday in the third week an additional visit is recorded on a non-regular day and on Thursday in the fifth week a visit is missing. To compare different instances of the feature types (features of a specific timeslot), both feature types are calculated for a second (incorrect) timeslot $T_i = (\{Monday, Tuesday, Thursday\}, 1)$.

The feature type $cv_{vpd}(T)$ for both timeslots is calculated in Equation 4 and Equation 5. The low $cv_{vpd}(T_c) \approx 0.1$ (see Equation 4) indicates that the exemplary place is visited similarly often for $D_v(T_c)$, wherefore the timeslot T_c is likely to be present.

$$\overline{n_{vpd}}(T_c) = \frac{6+5}{2} = 5.5$$
$$s_{vpd}(T_c) = \sqrt{\frac{(6-5.5)^2 + (5-5.5)^2)}{2}} = 0.5 \quad (4)$$
$$cv_{vpd}(T_c) = min(\frac{s_{vpd}(T_c)}{\overline{n_{vpd}}(T_c)}, 1) = \frac{0.5}{5.5} \approx 0.1$$

For the incorrect timeslot T_i, $cv_{vpd}(T)$ is no longer close to 0 ($cv_{vpd}(T_i) \approx 0.54$, see Equation 5).

$$\overline{n_{vpd}}(T_i) = \frac{1+6+5}{3} = 4 \quad s_{vpd}(T_i) \approx 2.16$$
$$cv_{vpd}(T_i) = min(\frac{s_{vpd}(T_i)}{\overline{n_{vpd}}(T_i)}, 1) = \frac{2.16}{4} \approx 0.54 \quad (5)$$

By also considering feature type $c_{dpw}(T)$ (in the given example $c_{dpw}(T_c) \approx 0.92$, see Equation 6), the presence of the investigated timeslot can be confirmed as it is close to the maximum 1. The feature type $c_{dpw}(T)$ decreases if an outlier is introduced to the timeslot. For the incorrect timeslot T_i, the feature value decreases to $c_{dpw}(T_i) \approx 0.67$.

$$\overline{n_{dpw}}(T_c) = \frac{2+2+2+2+1+2}{6} \approx 1.83$$
$$c_{dpw}(T_c) = \frac{\overline{n_{dpw}}(T_c)}{|D_v(T_c)|} \approx \frac{1.83}{2} \approx 0.92 \quad (6)$$

Table 2: Calendar with exemplary visits at a place with a single timeslot. Visited days are marked gray. For both timeslots (T_c, T_i), the number of visits at the place for each day of the week ($n_{vpd}(T)$) is counted below and the number of visits for the place in every single week is counted on the right ($n_{dpw}(T)$).

Week	Mo.	Tu.	Wed.	Thu.	Fri.	Sa.	Su.	$n_{dpw}(T_c)$	$n_{dpw}(T_i)$
1								2	2
2								2	2
3								2	3
4								2	2
5								1	1
6								2	2

$n_{vpd}(T_c)$	0	6	0	5	0	0	0
$n_{vpd}(T_i)$	1	6	0	5	0	0	0

Model Creation and Training. The described features are used to train an XGBoost (Chen et al., 2016) model. The model is trained to determine if the visits at a POI relate to a movement pattern which is regular or irregular in the temporal dimension by using $cv_{vpd}(T)$ and $c_{dpw}(T)$. The parameter configuration $n\_estimators = 1000$, $max\_depth = 100$, $\lambda_{reg} = 5$ performed best with an F1 score of 0.81 for D_M.

4 EVALUATION

The experiments evaluate (1) the capabilities of our pre-processing pipeline for detecting visited POIs and (2) the performance of our classification model regarding the differentiation between temporal regular and irregular movement patterns.

POI Detection. The evaluation dataset D_E (see Section 3.1) is used for evaluating the POI detection quality. We compare the labeled POIs with the POIs identified by the pre-processing pipeline by checking whether an identified POI is within the uncertainty radius $r_u \approx 300m$ (see Section 3.3) of a labeled POI.

Table 3 shows the results with the metric F1 as the main outcome of the experiment execution. The mean F1 score from all three CTSs is approx. 0.82 with a standard deviation of approx. 0.05.

The number of POIs that have been labeled but are not identified by the pre-processing pipeline (FN) is higher than the number of points that have been incorrectly identified as POIs (FP) for all three CTSs (see Table 3). Therefore the recall with a mean of approx. 0.78 is lower than the precision with a mean of approx. 0.86.

Model Performance. The model is trained on the dataset D_M and tested using D_E. The evaluation of the model performance yields an F1 score of 0.81 and 0.86 for D_M and D_E, respectively.

Table 3: POI detection quality for three CTSs (TP, FP, FN correspond to the number of true positives, false positives, and false negatives, respectively).

No.	TP	FP	FN	Precision	Recall	F1
1	165	43	52	0.79	0.76	0.78
2	110	17	35	0.87	0.76	0.80
3	106	8	24	0.93	0.82	0.87

5 DISCUSSION

The distributions of the movement pattern types of D_M and D_E are skewed. Hence, F1 is our main metric of interest as it combines precision and recall. We also use F1 for assessing the POI detection, which performs reasonably well with a mean F1 score of approx. 0.82 (D_E). Our classification model performs even better with an F1 score of approx. 0.86 (D_E).

However, there exist several threats to validity. The number of available CTSs and, as a consequence, the dataset size are rather low. This also leads to a limited size of the evaluation dataset D_E, which might explain that F1 for D_M is lower than for D_E. For the POI detection, the clustering technique DBSCAN provided the best results for the given dataset. Other clustering techniques might be better suited for larger datasets. Furthermore, in order to minimize the labeling effort, the CTS group only had to label the activities they considered as regular. Not labeling presumably non-regular activities can lead to more errors as each CTS might not be aware of all her actual regularities.

Moreover, the location data is recorded by a rather homogeneous group of CTSs that are very similar in terms of worksite affiliation, working hours, and age. In contrast, the travel behavior of distinct user groups differs, e.g. between home-based persons (like homemakers) and persons who travel to their workplace (Kutter, 1973). The proposed features may therefore be not as effective for different group compositions. Furthermore, our work in progress approach was only validated for timeslots with 1-week periodicities. We will further investigate the robustness of the features and our approach, especially for additional timeslot types, in our future work.

6 CONCLUSIONS

Digitization in the automotive industry causes the change from car manufacturers to mobility service providers. For example, to propose meaningful MaaS offerings to interested and consenting individuals,

movement regularities have to be identified. Our proposed approach can model daily travel routines and predict regularities using the machine learning algorithm XGBoost. We demonstrate that already small datasets enable acceptable performance for POI detection and future movement prediction ($F1 > 0.8$).

In our future work, we will investigate the remaining movement pattern types, further temporal periodicities, and sub-weekday time unit granularity. Moreover, we will explore how the evolution of movement behavior over time can be incorporated.

REFERENCES

Ashbrook, D. et al. (2002). Learning significant locations and predicting user movement with GPS. In *Proc. 6th Int'l Symp. on Wearable Computers*, pages 101–108.

Bachi, R. (2005). Standard Distance Measures and Related Models for Spatial Analysis. *Papers in Regional Science*, 10:83–132.

Chen, T. et al. (2016). XGBoost: A Scalable Tree Boosting System. *Proc. of the 22nd ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining - KDD '16*.

Ester, M. et al. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proc. of the Second Int'l Conf. on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press.

Grüner, T. (2019). *Detecting Complex Spatio-temporal Movement Patterns Using Historical Location Data*. DHBW Stuttgart, Germany. Bachelor's Thesis.

Hall, R. (2012). *Handbook of Transportation Science*. Int'l Series in Operations Research & Management Science. Springer US.

Jittrapirom, P. et al. (2017). Mobility as a Service: A Critical Review of Definitions, Assessments of Schemes, and Key Challenges. *Urban Planning*, 2(2):13–25.

Kitamura, R. et al. (1997). Generation of Synthetic Daily Activity-Travel Patterns. *Transportation Research Record*, 1607(1):154–162.

Kjærgaard, M. B. et al. (2010). Indoor Positioning Using GPS Revisited. In *Pervasive Computing*, pages 38–56.

Kutter, E. (1973). A Model for Individual Travel Behaviour. *Urban Studies*, 10(2):235–258.

Liu, F. T. et al. (2008). Isolation Forest. In *2008 Eighth IEEE Int'l Conf. on Data Mining*, pages 413–422.

Merah, A. F. et al. (2013). A Sequential Patterns Data Mining Approach Towards Vehicular Route Prediction in VANETs. *Mobile Networks and Applications*, 18:788–802.

Vintan, L. et al. (2004). Person Movement Prediction Using Neural Networks. In *First Workshop on Modeling and Retrieval of Context*, pages 1–12.

Vukovic, M. et al. (2007). Predicting user movement for advanced location-aware services. In *2007 15th Int'l Conf. on Software, Telecommunications and Computer Networks*, pages 1–5.