Poster: Exploring the Usefulness of Bluetooth and WiFi Proximity for Transportation Mode Recognition

Vlad C. Coroamă, Can Türk, Friedemann Mattern
vcoroama@ethz.ch
ETH Zurich, Department of Computer Science
Zurich, Switzerland

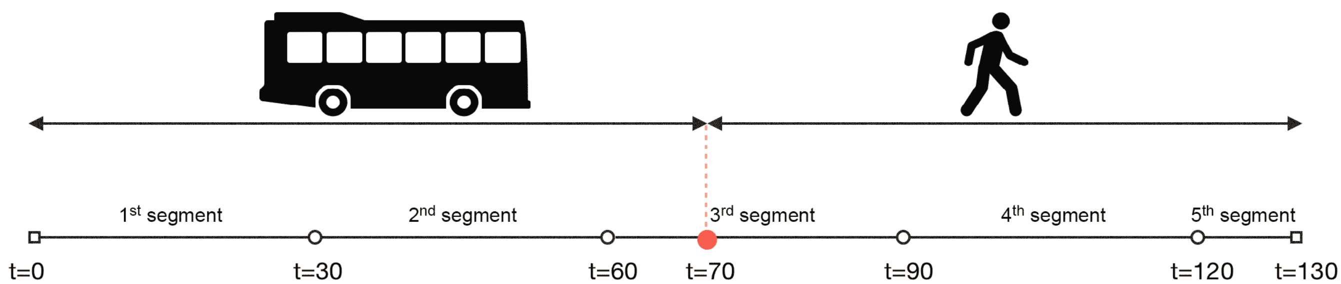

Figure 1: A simple trip example consisting of two trip legs, and a classification window size of 30 seconds.

ABSTRACT

Understanding the mobility patterns of large groups of people is essential in transport planning. Today's assessments rely on questionnaires or self-reported data, which are cumbersome, expensive, and prone to errors. With recent developments in mobile and ubiquitous computing, it has become feasible to automate this process and classify transportation modes using data collected by users' smartphones. Previous work has mainly considered GPS and accelerometers; however, the achieved accuracies were often insufficient. We propose a novel method which also considers the *proximity patterns of WiFi and Bluetooth (BT)* devices in the environment, which are expected to be quite specific to the different transportation modes. In this poster, we present the promising results of a preliminary study in Zurich.

CCS CONCEPTS

• **Applied computing → Transportation**; • **Computing methodologies** → *Supervised learning*; *Classification and regression trees.*

UbiComp/ISWC '19 Adjunct, September 9–13, 2019, London, United Kingdom

ACM ISBN 978-1-4503-6869-8/19/09.
https://doi.org/10.1145/3341162.3343847

KEYWORDS

transportation mode recognition; classification; random forest; transport planning

ACM Reference Format:
Vlad C. Coroamă, Can Türk, Friedemann Mattern. 2019. Poster: Exploring the Usefulness of Bluetooth and WiFi Proximity for Transportation Mode Recognition. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2019 International Symposium on Wearable Computers (UbiComp/ISWC '19 Adjunct), September 9–13, 2019, London, United Kingdom.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3341162.3343847

1 INTRODUCTION

Nowadays, the assessment of transportation practices relies on either questionnaires or self-reported data [6]. Both methods are cumbersome, expensive, prone to errors, and their data quickly becomes outdated. Given this situation, a growing body of literature examines how smartphone sensors can be used to automatically infer the transportation mode. Such an approach could potentially yield large-scale and yet fine-grained data across regions, groups of users, and means of transportation. Previous work has considered several sensors such as GPS or inertial measurement units (e.g., accelerometers) for this task. Although promising, this approach could so far only be explored in academic or industrial research. The main challenge to date was the insufficient accuracy in distinguishing among modes with similar speeds, acceleration, or routes. While high accuracies could be achieved in classifying distinct modes such as walking, cycling, or driving, previous work has not been able to satisfyingly distinguish cars from buses in urban environments or among different types of public transportation such as buses and trams.

Our project explores a new method that exploits proximity patterns of WiFi and BT devices in the environment. By regularly scanning the environment for WiFi and BT devices, it was expected to uncover distinctive proximity patterns, which could be quite specific to the different transportation modes. In urban public transportation, for example, every couple of minutes (i.e., at each stop) some BT devices could disappear, and others appear. In cars, by contrast, one or two BT devices might be constantly reachable over longer periods of time, e.g. the BT car stereo or the phone of another passenger. We present the preliminary results of a study with data collected around the city of Zurich. The study considered 7 different transportation modes among which four vehicle modes *car*, *bus*, *tram*, and *train*.

2 BACKGROUND

While not yet deployed in real applications, a growing body of academic literature investigates the use of data from mobile phone sensors in transportation planning. Early studies using sensor data from mobile devices for transport mode classification used mainly GPS data from the GeoLife dataset. Geolife is a GPS trajectory dataset collected by Microsoft Research Asia over 4.5 years between 2007 and 2011 [5]. Its more than 17,000 trajectories collected by 178 users comprise a total distance of over 1.2 million kilometers and span more than 48,000 hours. The first study to use the GeoLife dataset for transportation mode classification was [8]. Via rigid boundaries for the average speed, it first classifies the data as *walk* or *non-walk*, and postulates that walking segments must always exist between the other transportation modes. The study then uses features such as *heading change rate* or *stop rate* to distinguish between three further modes: *drive*, *bus*, or *bike*, reaching an overall accuracy among all four modes of 76%. Building on the same Geolife dataset, [9] first partitions trips into uni-modal segments. Assuming a correct partition, and deploying a deep neural network (DNN) with a stacked auto-encoder, the study then reaches a precision of 93.5% among five modes: *walking*, *cycling*, *bus*, *subway*, and *driving*; however, the needed segmentation prerequisite is notoriously error-prone, and cannot be realistically assumed.

Further frequently deployed sensors, either alone or in conjunction with GPS, are inertial measurement units such as accelerometers or gyroscopes. A dataset comprising data from 224 volunteers who jointly collected over 8000 hours of readings was released by HTC research [7]. The dataset comprises data from three sensors: accelerometer, gyroscope, and magnetometer; however, no GPS data. The manually recorded ground truth includes many more types of transportation than previous datasets: next to *still*, *walk*, *run*, *bike*, there are six vehicle modes: *motorcycle*, *car*, *bus*, *metro*, *train*, and *high speed rail*. Investigating the appropriateness of a variety of machine learning techniques using this data set, [2] reaches accuracies from around 60% and up to 86%. The analysis, however, does either distinguish between *still*, *walk*, *run*, *bike*, and a generic *vehicle* mode, or only between vehicle modes for data already known (from the ground truth) to be some type of vehicle. It does not address the far more difficult topic of classifying vehicle and non-vehicle transportation modes in one step without a-priori knowledge. Consequently, in a later work the authors no longer aimed at classifying individual vehicle modes, but only classified between *still*, *walk*, *run*, *bike*, and the generic *vehicle* mode [1].

Finally, some of the literature employed either WiFi or BT for transportation mode determination. An early study [4] used the rate of change of both the GSM cell a phone is connected to, and the WiFi base stations in reach, to infer the approximate speed of the user, and subsequently distinguish between *dwelling*, *walking*, and *driving*. By contrast, we do not want to infer the user's speed, but are interested in proximity patterns (e.g., the share of WiFis that stay in range for longer periods of time, their rate of change, etc). A single study took an approach similar to ours, and informed many of our design choices. To classify between the *walk*, *bike*, *vehicle*, and *rail* modes, [3] uses data from several sensors: GPS, accelerometer, WiFi, and BT. For the latter two, it defines the *recurrent address count* feature, i.e. the share of devices in range for two consecutive scans, hypothesizing that this share is greater for specific transportation modes (e.g., for public transportation) than for others. We expand this concept with several further features reflecting proximity patterns.

3 METHOD

An Android application was developed to collect both sensor and ground truth data. The app periodically reads a variety of sensors, i.e., gyroscope, accelerometer, gravity, GPS, Bluetooth and WiFi. Table 1 gives an overview of the sampling period for each sensor. Depending on the system resources available, the actual sampling frequency may be lower than the values shown in the table. A minimum of 30 seconds between WiFi scans is imposed by the Android system.

Training: As shown in Fig. 2, the user interface (UI) is fairly basic, allowing users to choose one of the seven different transportation modes to collect the ground truth, and cancel or finish a trip. The application starts recording sensor data as soon as the first transportation mode is chosen and records until a trip is either finished or cancelled. A mode change is reflected by selecting a new transportation mode, which automatically deselects the previous one. The wait

Table 1: Sampling period for each sensor

Sensor	Sampling Period
Gyroscope Accelerometer Gravity	50 Hz
GPS	1 Hz
Bluetooth	For 12 seconds every 30 seconds
Wi-Fi	For 15 seconds every 30 seconds

mode is used to indicate periods of inactivity between transportation modes during a trip. Users were asked to select this mode when e.g. waiting for a vehicle to arrive. To guarantee the accuracy of the ground truth, we asked the volunteers to cancel a trip if any inaccuracies appeared while recording it, e.g., if they forgot to indicate a mode change.

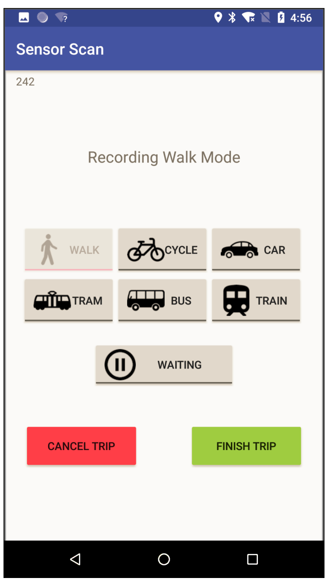

Figure 2: The UI of the data gathering app.

Classification: Based on the analysis of results from previous studies, Random Forest was chosen for the classification of transportation modes. Random Forest has several advantages: i) it makes few statistical assumptions about the dataset, ii) it has few hyperparameters and does not need exhaustive parameter tuning, iii) it reveals how well the model generalizes without the need for cross-validation or a separate test set thanks to the built-in out-of-bag (OOB) error calculation, iv) it runs efficiently on large datasets, v) it yields probabilities of belonging to different classes, vi) it supports changing the weight of prediction errors individually for each class, vii) it measures the relative importance of each feature easily, and viii) it supports randomization, such as searching for the best features among a random subset of features while splitting a node, thus allowing diversity.

4 DATA

The app was installed on Nexus 5X devices that were handed out to volunteers who recorded their everyday commutes for a couple of days each. A total of 20 volunteers gathered jointly over 150 hours of data in and around Zurich. There were 385 trips comprising 1111 *trip legs*, i.e., continuous periods of the same transportation mode within a trip. Table 2 shows the distribution of the individual transportation modes in the data. As seen in the table, transportation modes were not evenly distributed among the recordings, *train*, *bicycle*, and *tram* dominating the data. We thus deployed two oversampling techniques: Synthetic Minority Oversampling Technique (*SMOTE*) and Adaptive Synthetic Sampling Approach (*ADASYN*). From the recorded trips, 70% (269 trips) were used for training and the remaining 30% (116 trips) for testing.

Table 2: Duration and mode distribution of the collected data in the Zurich metropolitan area.

	Number of trip legs	Duration
Tram	150	31h:15m
Bus	49	10h:07m
Walk	298	28h:17m
Train	78	31h:55m
Car	43	13h:56m
Bicycle	126	24h:55m
Wait	367	12h:11m
TOTAL	1111	152h:39m

5 EVALUATION

For classification, the window length is one of the key parameters, and entails the following trade-off: too short a window is unlikely to provide enough useful information. Too long a window, as also shown in Figure 1, is likely to include two or more different modes, thus confuse the classification.

In our analysis, for each window we extracted 87 features from the available sensors (28 for GPS, 16 for accelerometer, 28 for BT proximity, and 15 for WiFi proximity). As soon as it started to be meaningful, we started to extract these features also for longer window sizes: *prevMinute*, *prev2Minutes*, *prev4Minutes*, and *wholeTrip*. In a longer trip with 10-second windows, for example, we would extract the features for each 10-second window, but after one minute into the trip, also for the last minute, after two minutes into the trip for the last 2 minutes, and so on. The 'short' windows are not overlapping, while the 'long' ones are sliding, always with the increment of the short ones. Thus, a maximum of 5 ∗ 87 features were used for classification.

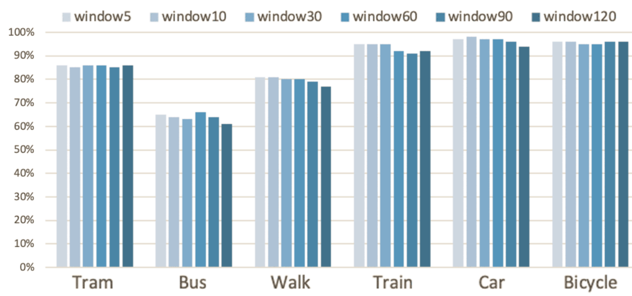

Figure 3: Classification precision of individual transportation modes, for different window sizes, and including the consideration of past windows.

Through this combination of short and long windows, we aimed at solving the aforementioned trade-off, capturing the mode changes better reflected by short windows as well as the continuity and information richness provided by longer ones. Figure 3 shows the precision scores for short window sizes varying between 5 and 120 seconds. Each analysis also deployed the appropriate longer sliding windows, i.e., those longer than the respective short window size (e.g., 1, 2, 4 minutes and entire trip for a 5 or 30 seconds short window, but only 4 minutes and entire trip for a 120 seconds short window). The variation between different window sizes is relatively small. Perhaps counter-intuitively, smaller windows yield slightly better results.

The confusion matrix (not shown due to space constraints) shows precision values of 88.1% *tram*, 64.9% *bus*, 83.1% *walk*, 95.8% *train*, 98.1% *car*, and 96.6% *bicycle*. Recall values are similar, 87.1% *tram*, 67.5% *bus*, 88.7% *walk*, 89.8% *train*, 97.3% *car*, and 92.7% *bicycle*, and the OOB score is 99.2%. Most confusions are between *tram* and *bus*. When merging them into a *public transport* mode, precision values are all between 82.6 and 98.7% and recall values between 88.4 and 95.2%.

6 DISCUSSION

Achieved accuracies using features based on the GPS and inertial sensors were on par with previous work using similar study conditions (e.g., not confining users to hold the smartphone still while en route). Additionally deploying features based on BT and WiFI proximity increased the accuracy for all transportation modes studied. Figure 4 shows the classification precision for BT and WiFi only, GPS and accelerometer only, and all combined. For trains with a relatively stable collection of BT devices in proximity and not so stable GPS signal, a classification based just on BT and WiFi proximity even outperforms a more traditional one based on GPS and accelerometer. By improving the average precision from 80% to 87%, BT and WiFi proximity appear a promising path for further investigation in order to improve the automatic classification of transportation modes based on smartphone sensor data.

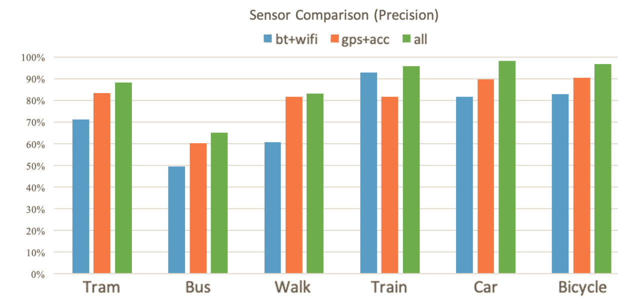

Figure 4: Contribution of BT and WiFi proximity to the classification.

REFERENCES

[1] Shih Hau Fang, Yu Xaing Fei, Zhezhuang Xu, and Yu Tsao. 2017. Learning Transportation Modes from Smartphone Sensors Based on Deep Neural Network. *IEEE Sensors Journal* 17, 18 (2017), 6111–6118.

[2] Shih Hau Fang, Hao Hsiang Liao, Yu Xiang Fei, Kai Hsiang Chen, Jen Wei Huang, Yu Ding Lu, and Yu Tsao. 2016. Transportation Modes Classification Using Sensors on Smartphones. *Sensors* 16, 8 (Aug. 2016).

[3] David Montoya, Serge Abiteboul, and Pierre Senellart. 2015. Hup-me. *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '15* (2015), 1–4.

[4] Min Y. Mun, Deborah Estrin, Jeff Burke, and Mark H. Hansen. 2008. Parsimonious mobility classification using GSM and WiFi traces. In *Proc. of the 5th Workshop on Embedded Networked Sensors (EmNets)*. 1–5.

[5] Microsoft Research. 2011. Geolife GPS trajectory dataset – User Guide. https://www.microsoft.com/en-us/research/publication/geolife-gps-trajectory-dataset-user-guide/.

[6] Leon Stenneth, Ouri Wolfson, Philip S Yu, and Bo Xu. 2011. Transportation mode detection using mobile phones and GIS information. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 54–63.

[7] Meng-Chieh Yu, Tong Yu, Shao-Chen Wang, Chih-Jen Lin, and Edward Y. Chang. 2014. Big Data Small Footprint: The Design of a Low-power Classifier for Detecting Transportation Modes. *Proc. VLDB Endow.* 7, 13 (Aug. 2014), 1429–1440. https://doi.org/10.14778/2733004.2733015

[8] 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* 4, 1 (2010), 1–36. https://doi.org/10.1145/1658373.1658374 arXiv:arXiv:1011.1669v3

[9] Xiaolu Zhu, Jinglin Li, Zhihan Liu, Shangguang Wang, and Fangchun Yang. 2016. Learning transportation annotated mobility profiles from GPS data for context-aware mobile services. *Proceedings - 2016 IEEE International Conference on Services Computing, SCC 2016* (2016), 475–482. https://doi.org/10.1109/SCC.2016.68