Characterizing Driving Environments Through Bluetooth Discovery

Walter Bronzi, Sébastien Faye, Raphael Frank, Thomas Engel
University of Luxembourg, Interdisciplinary Centre for Security, Reliability and Trust (SnT)
6 avenue de la Fonte, L-4362 Esch-sur-Alzette, Luxembourg
E-mail: {firstname.lastname}@uni.lu

***Abstract*—Within the world of wireless technologies, Bluetooth has recently been at the forefront of innovation. It is becoming increasingly relevant for vehicles to become aware of their surroundings. Therefore, having knowledge of nearby Bluetooth devices, both inside and outside other vehicles, can provide the listening vehicles with enough data to learn about their environment. In this paper, we collect and analyze a dataset of Bluetooth Classic (BC) and Low Energy (BLE) discoveries. We evaluate their respective characteristics and ability to provide context-aware information from a vehicular perspective. By taking a look at data about the encountered devices, such as GPS location, quantity, quality of signal and device class information, we infer distinctive behaviors between BC and BLE relative to context and application. For this purpose, we propose a set a features to train a classifier for the recognition of different driving environments (i.e. road classes) from Bluetooth discovery data alone. Comparing the performance of our classifier with different sampling parameters, the presented results indicate that, with our feature selection, we are able to predict with reasonable confidence up to three classes (Highway, City, Extra-Urban) by using only discovery data and no geographical information. This outcome gives promising results targeted at low energy and privacy-friendly applications and can open up a wide range of research directions.**

***Index Terms*—Bluetooth, Bluetooth Low Energy, Sensing Systems, Context-aware Services, Machine Learning**

I. Introduction

We live in a world where everything is increasingly connected, at a time when transportation companies are trying to offer personalized services to their users. Although the availability of opportunistic systems is extremely promising, the required data sources are currently not used to their full potential. The downside of standard sensors that collect the required data, such as GPS, is that they often consume a lot of energy and compromise user privacy. In this article, we take on the challenge of exploiting the natural growth in the connectivity of our everyday devices using Bluetooth traces, one of the least expensive data sources, to analyze how this technology could complement current systems.

Two types of Bluetooth currently exist: 1) Bluetooth Classic, which relies on an association mechanism between two clients (i.e. key exchange); and 2) Bluetooth Low Energy or Bluetooth Smart (Bluetooth version 4.0 and higher). The latter dispenses with the association mechanism and directly allows devices to broadcast information that can be quickly retrieved by other devices. The generated packets of Bluetooth Low Energy (BLE) are therefore lighter, and the overall transmission consumes less energy than Bluetooth Classic (BC) whilst reaching ranges of nearly 100 meters, compared to a few meters for BC. In general, it is widely accepted that both BC and BLE consume significantly less power than other communication protocols (e.g. Wi-Fi, GPS, etc.) [1]–[4]. Moreover, the inexpensive nature of BC and BLE makes these technologies easy and advantageous for manufacturers to implement.

BC can be found mainly in smartphones, cars, computers, etc. whilst BLE in wearables, bicycles, key-rings, smart TVs and many more [1]. Furthermore, both these technologies are found in objects that can be identified as either static (e.g. computers, televisions) or mobile (e.g. key-rings, wearables).

In previous studies, we benchmarked BLE for Vehicle-to-Vehicle (V2V) communications [5]. Results have shown how this technology is best used in dense traffic situations and for simple non-safety related data exchanges. Its performance in terms of delivery ratio and round-trip time make this technology a perfect candidate for general low-to-medium speed V2X communications.

In this paper, we go one step further by gathering knowledge on the quantity and type of devices we can discover whilst driving in diverse scenarios.

A vehicular scenario allowed us to quickly cover a geographical area in Luxembourg by means of a data collection campaign. Participants were asked to use a Android mobile application while driving in order to passively discover Bluetooth (Classic and Low Energy) devices. Moreover, we propose a classifier that is capable of outputting a certain driving context giving as input the raw Bluetooth discovery data. The purpose of this classification is to benchmark the precision with which we can approximately estimate *in what type of environment* we are, by looking at *which devices* are around us. By the means of machine learning algorithms we score different configurations using diverse sampling methods.

The potential of using Bluetooth as a support technology to offload other communication means is threefold. First, it is energy-efficient if compared to technologies such as GPS, Wi-Fi and ZigBee. Secondly, it can be appropriately duty-cycled depending on the applications and circumstances. Finally and above all, by not requiring a complex client-server

[1]https://bluetooth.com/what-is-bluetooth-technology/bluetooth-devices

architecture, this approach is privacy-friendly as every analysis and contextual estimation is done locally on the device.

II. Related Work

Bluetooth has been standardized as a short to medium range communication technology operating in the 2.4 GHz band. Its Low Energy (or Smart) standard was subsequently introduced with the Bluetooth 4.0 specification, adding a low power mode of operation to the original protocol. Despite being similar in implementation, BC and BLE use different controllers (i.e. physical and link layer) and are thus not compatible. Most devices that are on the market today, unless they have a specific use, support both protocol stacks in dual-mode.

Although the use of Bluetooth is wide-ranging, this article is focusing on its use in sensing systems. Such systems generally allow the collection of vast amounts of data, where the most conventional parameters are a sampling and recording frequency that vary according to the metrics collected and the sensors used. This type of system is primarily passive (or opportunistic) and allows the collection, analysis, processing and storage of a set of data.

Recent studies have shown that the use of network traces makes it possible to understand different contextual and environmental characteristics. For example, Wi-Fi traces can be used to identify places [6] or to understand different mobility patterns [7]. Bluetooth traces, at the heart of this article, are also promising as they allow not only to measure interactions between devices, but also to retrieve details (varying depending on the version) of the nature and properties of devices.

Using Bluetooth for context-awareness was experimented by Chen et al. [8], before the introduction of its Low Energy version, for detecting daily activities. Aditya et al. also used Bluetooth to build context specific privacy profiles for mobile communication [9].

In previous research we focused on the usage of BLE, for active communication in a vehicular context. Our findings showed how this technology behaves in different scenarios such as Highway or City. The protocol proved to be resilient to many situations but found its best case scenario, in its current implementation, to be at urban travel speeds [5].

While the aforementioned work focused on actively using Bluetooth for vehicular communications, we decided in this paper to experiment with passive scanning from a vehicle perspective. With BLE, once discoverability is enabled, nearby devices can be tracked even if they are not being actively used.

The aim of this article is to outline ways in which a technology as simple as Bluetooth can be used to identify complex environmental characteristics in a vehicular context. Our claim is that the quantity and the quality of traces obtained by Bluetooth are strongly dependent on the mobility (or speed) of the user. If this assumption is verified, it is certainly possible to identify these locations or certain mobility characteristics based solely on Bluetooth traces, which do not consume energy and which have the great advantage of preserving privacy.

Currently, there are different studies that target vehicular communication with VRUs [10] with a particular focus on pedestrians and bicycles [11] but, to the best of our knowledge, none using BLE.

III. Field Experiment

The first step to achieve our goal is to set up a sensing system deployable on a small scale. In our case, in order to discover Bluetooth devices in a vehicular context, we set up an experiment involving data collection from smartphones inside cars. We developed a mobile application able to run passively and scan the environment while the user is driving. In order to minimize polluting the sessions and to keep the user experience fast and seamless, we developed a logic able to recognize the driving activity from our participants and automatically start and stop the scanning procedure.

We deployed the application on 20 Android phones from different brands belonging to the participants of the study. For a period of two months, we collected data from daily drives. The mobile app, named *BlueScanner*, is available for free under a MIT Licence on GitHub [2].

A. Sensing System

Our implementation, although not comparable with more expensive sniffers, provides enough precision to prove our claim and to enable a concrete and easily deployable application which can be installed on a multitude of devices.

Since we wanted to keep the user interaction to a minimum, so as to avoid unnecessary distraction at the wheel, our application uses the Android activity recognition API, with some added logic, to automatically start discovering for devices once a *driving activity* is detected with a certain precision. Furthermore, the upload of traces to our server was encrypted and automated. The mentality behind these design choices is to reduce user interaction to a minimum, customize user settings for different usages and standardize a way of starting, stopping and logging sessions in the same way for everyone. Furthermore, we instructed the participants to keep the phone on a holder (that we provided if none was available) on the windshield of their vehicle in order to provide optimal discovery.

Once a *driving activity* is detected, a session is started. By listening for both BC and BLE devices we can thus discover nearby objects. BC devices are discovered during a 12 sec discovery cycle (Android platform specific [3]) which is then looped for the length of the session. For BLE, because of the protocol differences, we keep the scanning device in discovery mode which in return will log every received advertisement at any time. Throughout the session, we log the GPS location of the device which we will use further on for labeling the dataset for our classification, i.e. to determine the road class (also in relation of the speed) of the discovering device. Furthermore, Table I shows the recorded information for every discovery.

[2]https://github.com/wbronzi/bluescanner

[3]https://dev.android.com/reference/android/bluetooth/BluetoothAdapter.html

Table I: Discovery Data

Type	Example	Details
MAC	AA:BB:CC:DD:EE:FF	Used for cross identification (unique only for BC)
RSSI	-50 dBm	Received Signal Strength
Device Name	Audi_MMI_XXX	Assigned device name if defined
Class Identifier (BC only)	200408	Type of device (e.g. hands-free kit)
Manufacturer Information (BLE only)	4C 00	Manufacturer (and sometimes model) of the discovered device (e.g. 4C corresponds to the Apple).

Table II: Dataset Summary

Number of participants	20
Number of recorded sessions	794
Location entries	284,975
Total coverage	13,486 km
BC entries	64,099 (16,201 Unique)
BLE entries	184,497 (6,241 Unique)
Average session duration	24.93 min

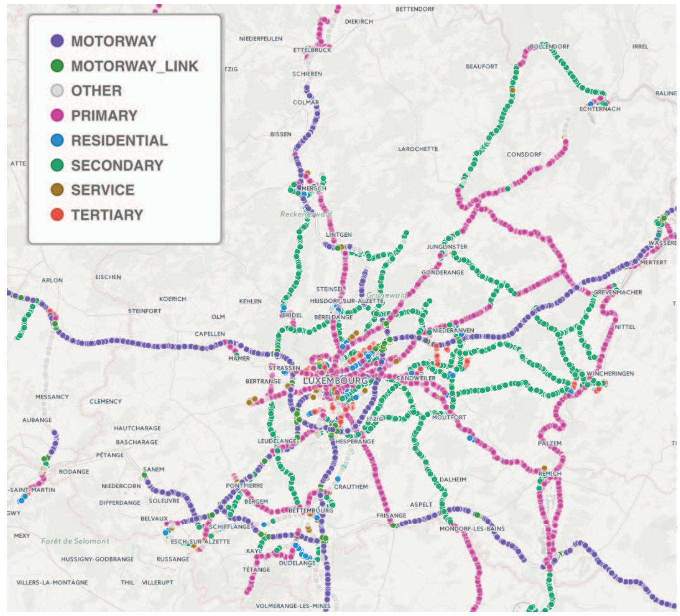

Figure 1: Experiment coverage

B. Data Collection Campaign

Because of the nature of this study, sensitive information about the participants had to be logged to conduct our analysis. Information such as locations and details about the users' device was securely stored on our back-end to be later utilized for building the classifier. All our users were prompted to read a terms and condition form and accept the privacy policy of the study.

Furthermore, no specific conditions were imposed on the participants other than having the application running with the device on a phone holder while driving. Typical driving routes include to and from work plus extra activities that require taking the car.

C. Dataset Statistics

Table II provides an overview of the final statistics after our two-month campaign. As the discovery behavior is different for the two Bluetooth technologies, it is reflected here in the difference of total number of entries between them. For more information on how the two protocol stacks handle discovery, interested readers can refer to the work of Gomez at al. [12].

For an initial validation of the dataset, we plotted the recorded latitudes and longitudes with *osm2po*[4] to a more readable label representing the road class [13]. As we can see from Figure 1, the spatial distribution of trips provides a large coverage of diverse road types within the country of Luxembourg.

The mapping between GPS locations and OSM road types is required for defining labels when building the classifier.

An in-depth statistical analysis can be found in our previous publication [14]. We found BC to be mainly present in primary and motorway roads, because of the large quantity of hands-free kits, whilst BLE is generally found in primary, residential and secondary roads. BLE is mostly implemented in objects that pedestrians and drivers are likely to be carrying and that therefore are more likely to be encountered on roads with slower speed limits such as residential and secondary roads.

The definition of primary roads is sometimes ambiguous. At times this road type, at least within Luxembourg, is represented by both extra-urban roads that connect villages within the country and roads usually very crowded with pedestrians inside the city. Generalizing a scenario for different cities can therefore be challenging. We considered such conclusion when building the scenarios for our classification study.

From these findings we conclude that the dataset is consistent and large enough to be exploitable. The presented metrics and the knowledge of discovered devices can lead to classification techniques for estimating the environment (road class) and in the future the travel speed range through the BC and BLE discovery behavior.

IV. Classification

Combining our conclusions from the statistical analysis with the additional information we gather during discovery, such as the device class (BC only) and manufacturer information (BLE only), we can further expand our knowledge of the nature of devices discovered nearby (e.g. smartphones, hands-free from vehicles, smartwatches, TV, bicycles).

Using this information as part of a feature set in a classifier can prove useful for different applications that require devices to contextualize themselves in a privacy-friendly and low energy way (no access to GPS). In the scope of this work, we aim at classifying a generic environment by providing as input Bluetooth data alone.

[4]http://osm2po.de/

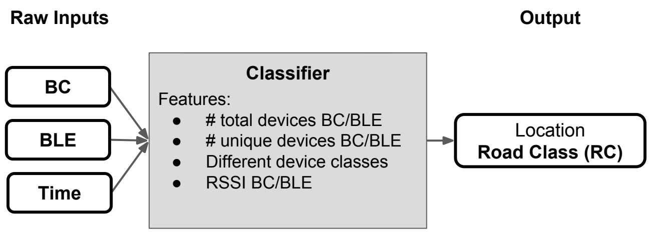

Figure 2: Classifier description

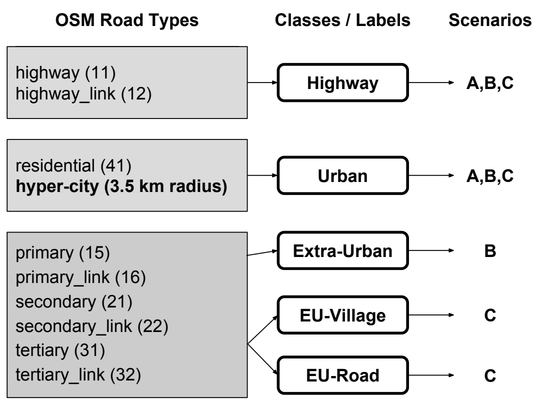

Figure 3: Open Street Map to Scenario Mapping

A. *Classifier Description*

The raw discovery entries taken as inputs from our dataset (Figure 2) are preprocessed for our classification in order to provide features for different training timeslots. More precisely, the features are: total number of nearby discovered devices; unique number of nearby discovered devices; device Class or Manufacturer Information (see Table I); and RSSI.

For each feature, we consider both the BC and BLE metrics. All different configurations of technologies and feature sets were evaluated but ultimately the combination of both protocols proved to perform better.

To correctly homogenize the data we selected the most relevant timeframes that reflect our findings. We ordered all entries temporally and only considered rush hours (from 8:00 to 10:59 and from 16:00 to 19:59). This selection correctly matches the typical traffic demand during a work day in Luxembourg, as previously observed by Codeca et al. [15]. Moreover, we divided our training and test sets respectively 80% and 20%.

B. *Scoring metrics*

As is the case in many classification approaches, we consider the accuracy as a starting point for scoring our classifier. This metric is produced by dividing the number of correct predictions made by the total number of predictions. As we are dealing with an imbalance of classes in our dataset, other performance measures must be considered (On average highway and primary roads cover 60% of the total amount locations, cf. [14]).

The sensitivity, also called recall or true positive rate, can be used to measure the integrity of a classifier. This score is obtained by dividing the number of true positives with the sum of true positives and false negatives. A low sensitivity implies a high number of false negatives. The specificity or true negative rate, is represented by dividing true negatives with the sum of true negatives and false positives. Again, a low specificity implies a high number of false positives.

By combining these two metrics we obtain the area under the curve (AUC) which is a performance measure between 0 and 1: $\text{AUC} = \frac{\text{Sensitivity + Specificity}}{2}$. A classifier with an AUC score > 0.5 performs better than random guessing [16]. By adopting this metric, we can give unbiased scores to our classifier by correctly representing minority classes.

C. *Classifying Environments*

We wanted the output of our classifier to provide a low precision context at first. Doing so provided us with a proof-of-concept of what we could achieve with our particular dataset and feature selection. For this purpose we propose three different scenarios built upon the data we collected.

These scenarios, as illustrated in Figure 4 have different labels and can describe a driving environment with different precision.

As Figure 3 shows, our simplest configuration (Scenario A) only has two classes indicating that only highways and urban environments are taken into consideration. For every scenario, we utilize an *osm2po* class-to-label mapping to determine the road type label.

In scenario A, we can find the following OpenStreetMap road classes: *highway* and *highway_link* represented by the label Highway whilst *residential* goes under Urban. Additionally, for the Urban label, we also consider a hyper-city class which includes locations within a certain radius (3.5 km in our case). This class is fine tuned for our scenario but can easily be adopted for other cities with similar road topology.

For scenario B, we have a similar configuration with an added class representing everything that is not highway or residential roads (mainly primary, secondary and tertiary roads).

Finally, for our last configuration (Scenario C), we further divide the Extra-Urban (EU) group into two sub-groups: Village and Road. An EU location is labeled as EU-Village when at least two *residential* roads are found within 350 m. The rest is labeled as EU-Road. We found this tuning to be the best at representing the road network structure.

For our classification we benchmarked two of the most popular and well-known set of models and multi-class learning algorithms; Support Vector Machines (SVM) and Random Forest (RF) [17]. In our case case RF provides, overall, better accuracies than SVM other than better handling of imbalance between classes. RF is a method used for classification and regression tasks that makes use of decision trees at training

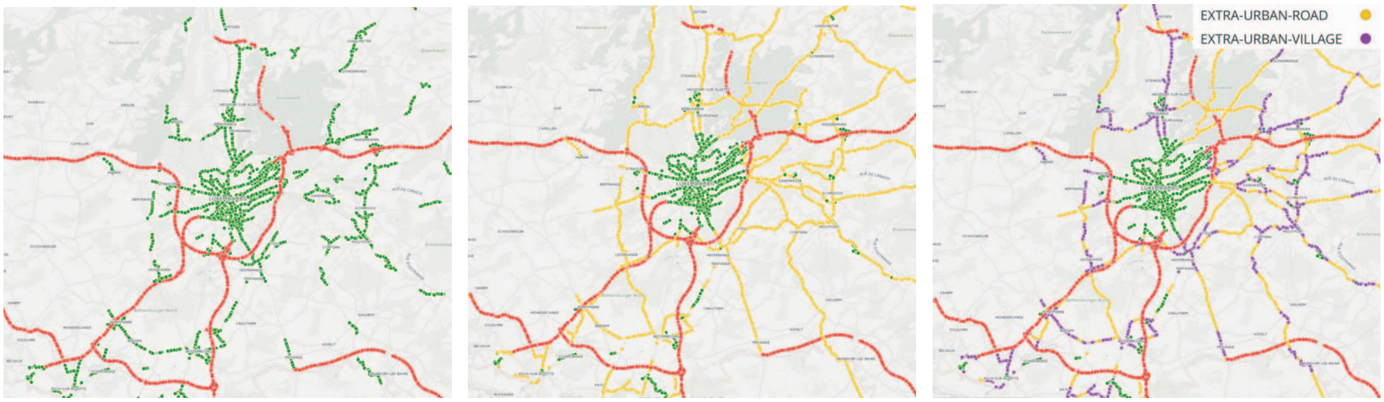

(a) Scenario A - 2 Classes (b) Scenario B - 3 Classes (c) Scenario C - 4 Classes

Figure 4: Scenarios overview

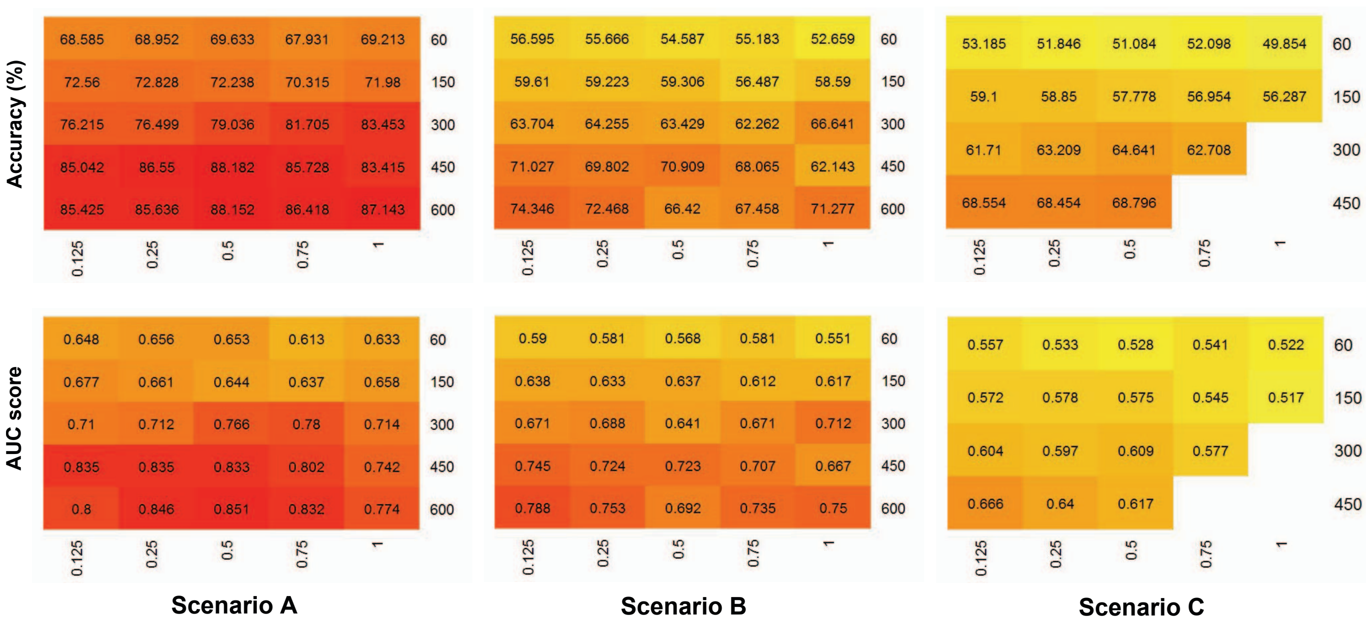

Figure 5: Classification results
(X axes represent sliding windows and Y axes represent timeframes)

time [18]. We chose this particular algorithm because, other than being fast and scalable, it can handle multiple classes (up to 4 in our case) and behaves well with non-linear features. Using this algorithm allowed us to avoid hyperparameter tuning and instead to focus on resampling our dataset with different timeframes and sliding windows to observe the behavior of the classifier. Moreover, for all our tests, we used 100 trees for each forest.

Having exported and ordered our dataset temporally as described in Section IV-A, we then resampled it for every scenario using five different timeframes (60s, 150s, 300s, 450s, 600s) and 5 different sliding windows (12.5%, 25%, 50%, 75%, 100%), giving us a total of 75 possible configurations (3 scenarios x 5 timeframes x 5 sliding windows). As the average of a session is $\approx$ 25 minutes, it did not make sense to try and sample windows bigger than 10 minutes.

In Figure 5 we show the entirety of the results for all timeframe/sliding window configuration and scenario. For every single combination, we provide the average results from 50 runs of the Random Forest algorithm.

We notice, as expected, that our Scenario A, with two road classes, is our best performing scenario. With a maximum accuracy of up to 88% and up to 0.85 AUC score, we can say that our feature set can be successfully used to predict the road category. The more labels, hence the more road classes, we try to predict the more we see the accuracy of classifier decline. With three labels (Scenario B) we reach up to 74% and up to 0.78 AUC score.

With the last scenario, trying to predict 4 road classes, we observe that we cannot fully complete our result matrix. This means that in the empty cases one or more of our classes was so rare that, even though we could achieve an overall high accuracy for the other classes, it could not be classified at all. This is the result of a very limited number of devices discovered in rural areas which makes it impossible to perform classification in these cases. The higher the time window is, the more a class can get *diluted* when averaging the majority class for that particular timeslot.

Although we could use, in specific cases, our classifier to predict up to 4 road categories, we find that with our particular dataset, 3 classes achieve the best trad-off between number of potential predicted classes and overall accuracy of prediction.

V. Conclusion and Perspectives

In this work, using a sensing system made freely available to the community, we characterized Bluetooth discovery in relation to a proof-of-concept vehicular scenario, with the aim of identifying different environmental characteristics by showing that the quantity and the quality of traces obtained are strongly dependent on the mobility of the user and on the places he or she visits. In particular, we proposed a mechanism that allows to identify different types of environments, such as Highway, Urban and Extra-Urban roads, by proposing different sampling methods capable of obtaining promising accuracies.

This mechanism can open the way to a wide variety of applications and can easily be extended to more than vehicular systems (e.g. pedestrians). More and more applications are using passive, opportunistic sensing systems to automatically characterize different aspects of their users. Most of the time, these sensing systems aim to provide personalized services for the user, for instance through a recommendation system. Apart from sensing systems, the mechanism described in this paper can be used as an intermediate step in opening the way to user mobility and environment profiling applications.

An interesting example of this is the estimation of passenger flow in public transportation systems. If a user is moving surrounded by other users, on a bus for example, it is possible to transcribe this behavior by observing Bluetooth devices that remain in the vicinity of the user for a certain period of time, and comparing this value with the speed estimated by the other devices (i.e. only crossed a few times). Another application example is driver behavior profiling by observing travel patterns. By detecting changes of environment (e.g. Highway to Urban and back to Highway again), it is possible to identify recurring patterns like from home to work and vice versa. Combining such knowledge with traffic data could allow the system to provide recommendation on route alternatives when needed. For instance, advising for a new route, at the discretion of the user, when a certain driving pattern is recognized within a known congested area.

The application developed for this article can be implemented as part of existing services that detect user contextual information and can easily be extended to cover more than just driving.

In future work, we plan to go one step further by proposing new methods of identifying other user characteristics, including current traffic state and type of vehicle as well as finding new classifiers for different mobility scenarios such as bikes and pedestrians.

Finally, exploring the possibilities of Bluetooth 5.0, a new standard that will provide devices with a far wider communication range, has the potential to significantly increase the classification process.

Acknowledgment

The authors would like to thank Telindus Luxembourg for their support and their industrial insight on the topic.

References

[1] Matti Siekkinen, Markus Hiienkari, Jukka K Nurminen, and Johanna Nieminen. How low energy is bluetooth low energy? comparative measurements with zigbee/802.15. 4. In *Wireless Communications and Networking Conference Workshops (WCNCW), 2012 IEEE*, pages 232–237. IEEE, 2012.

[2] Jin-Shyan Lee, Yu-Wei Su, and Chung-Chou Shen. A comparative study of wireless protocols: Bluetooth, uwb, zigbee, and wi-fi. In *Industrial Electronics Society, 2007. IECON 2007. 33rd Annual Conference of the IEEE*, pages 46–51. IEEE, 2007.

[3] Roy Friedman, Alex Kogan, and Yevgeny Krivolapov. On power and throughput tradeoffs of wifi and bluetooth in smartphones. *IEEE Transactions on Mobile Computing*, 12(7):1363–1376, 2013.

[4] Aaron Carroll, Gernot Heiser, et al. An analysis of power consumption in a smartphone.

[5] Walter Bronzi, Raphael Frank, German Castignani, and Thomas Engel. Bluetooth low energy performance and robustness analysis for inter-vehicular communications. *Ad Hoc Networks*, 37:76–86, 2016.

[6] Hongbo Liu, Jie Yang, Simon Sidhom, Yan Wang, Yingying Chen, and Fan Ye. Accurate wifi based localization for smartphones using peer assistance. *IEEE Transactions on Mobile Computing*, 13(10):2199–2214, 2014.

[7] Timothy Ellersiek, Gennady Andrienko, Natalia Andrienko, Dirk Hecker, Hendrik Stange, and Marc Mueller. Using bluetooth to track mobility patterns: depicting its potential based on various case studies. In *Proceedings of the Fifth ACM SIGSPATIAL International Workshop on Indoor Spatial Awareness*, pages 1–7. ACM, 2013.

[8] Yiqiang Chen, Zhenyu Chen, Junfa Liu, Derek Hao Hu, and Qiang Yang. Surrounding context and episode awareness using dynamic bluetooth data. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 629–630. ACM, 2012.

[9] Paarijaat Aditya, Viktor Erdélyi, Matthew Lentz, Elaine Shi, Bobby Bhattacharjee, and Peter Druschel. Encore: private, context-based communication for mobile social apps. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*, pages 135–148. ACM, 2014.

[10] José Javier Anaya, Pierre Merdrignac, Oyunchimeg Shagdar, Fawzi Nashashibi, and José E Naranjo. Vehicle to pedestrian communications for protection of vulnerable road users. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 1037–1042. IEEE, 2014.

[11] General Motors. GM Developing Wireless Pedestrian Detection Technology. URL: http://media.gm.com/media/us/en/gm/news.detail.html/content/Pages/news/us/en/2012/Jul/0726_pedestrian.html.

[12] Carles Gomez, Joaquim Oller, and Josep Paradells. Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology. *Sensors*, 12(9):11734–11753, 2012.

[13] Mordechai Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12–18, 2008.

[14] W. Bronzi, T. Derrmann, G. Castignani, and T. Engel. Towards characterizing bluetooth discovery in a vehicular context. In *2016 IEEE Vehicular Networking Conference (VNC)*, pages 1–4, Dec 2016.

[15] Lara Codeca, Raphael Frank, Sébastien Faye, and Thomas Engel. Luxembourg sumo traffic (lust) scenario: Traffic demand evaluation. *IEEE Intelligent Transportation Systems Magazine*, 2016.

[16] Patrick O Glauner, Andre Boechat, Lautaro Dolberg, Radu State, Franck Bettinger, Yves Rangoni, and Diogo Duarte. Large-scale detection of non-technical losses in imbalanced data sets. *arXiv preprint arXiv:1602.08350*, 2016.

[17] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.

[18] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.