

Where am I? Using Mobile Sensor Data to Predict a User’s Semantic Place with a Random Forest Algorithm

Elisabeth Lex, Oliver Pimas, Jörg Simon, and Viktoria Pammer-Schindler

Know-Center, Graz, Austria

Abstract. We use mobile sensor data to predict a mobile phone user’s semantic place, e.g. at home, at work, in a restaurant etc. Such information can be used to feed context-aware systems, that adapt for instance mobile phone settings like energy saving, connection to Internet, volume of ringtones etc. We consider the task of semantic place prediction as classification problem. In this paper we exploit five feature groups: (i) daily patterns, (ii) weekly patterns, (iii) WLAN information, (iv) battery charging state and (v) accelerometer data. We compare the performance of a Random Forest algorithm and two Support Vector Machines, one with an RBF kernel and one with a Pearson VII function based kernel, on a labelled dataset, and analyse the separate performances of the feature groups as well as promising combinations of feature groups. The winning combination of feature groups achieves an accuracy of 0.871 using a Random Forest algorithm on daily patterns and accelerometer data.

A detailed analysis reveals that daily patterns are the most discriminative feature group for the given semantic place labels. Combining daily patterns with WLAN information, battery charging state or accelerometer data further improves the performance. The classifiers using these selected combinations perform better than the classifiers using all feature groups. This is especially encouraging for mobile computing, as fewer features mean that less computational power is required for classification.

1 Introduction

Smartphones currently hold a handheld market share of over 30% - and this market share is rising¹. Because of their built-in sensors, smartphones are a particularly suitable tool for capturing people’s activities in a physical environment as opposed to people’s interactions with electronic devices or interactions within virtual environments. Such mobile sensor data can be used to analyse behavioural patterns, or within user- and context-adaptive systems. Given the wide spread of smartphones, such systems have the potential to reach an incredible amount of users. In this paper, we describe how to use mobile sensor data to predict a mobile phone user’s semantic place, i.e. *home*, *work*, *restaurant* etc.

¹ <http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats/a#smartphone-shipments>

Semantic place information exceeds geographic location information in that it gives a meaning to a user's location. Location-aware systems that exploit the geographic location or just the uniqueness of places (e.g., based on WLAN IDs) are state-of-the-art. Recommender systems like Yelp for restaurants or Friends for finding friends in the vicinity use geographic location information. The Llama App executes location-specific rules w.r.t. device system settings, and uses cell tower information to identify locations. In such Apps, semantic categories are assigned to places by users, but not exploited by the system. Systems that exploit place semantics are now cutting edge. Only recently, a recommender system for advertisements has been described that depends on geographic locations, identified via WLAN ID, but distinguishes places also via their semantics, e.g., fashion shop, restaurant, cinema etc. [8]. The mapping between geographic location and place semantics is not automated, but used by the system.

2 Dataset

The work described in this paper was carried out in the context of the Nokia Mobile Data Challenge 2012. The challenge provided a data set, the MDC dataset, collected by the NRC/Lausanne Data Collection Campaign 2009-2010 [9]. Smartphone data has been collected by almost 200 participants in the course of at least one year [10]. For each user, data about telephone usage, media usage, motion (accelerometer data), telephone status (bluetooth, battery charging) etc. has been collected [9].

In the MDC dataset, each data record contains the data of a sensor (e.g., battery charging status) and a timestamp. Each record has been collected by a single user, and is assigned to a place ID p_{ID} that defines a geographic location. However, it cannot be related back to geographic coordinates, and it corresponds to a circle of 100m radius. Each place ID p_{ID} is associated to a single user. Since geographic coordinates of place IDs p_{ID} are unknown, it is unknown whether place IDs p_{IDs} from different users correspond to the same geographic location. A subset of records in the MDC dataset has been labelled with one of the predefined semantic place labels as ground truth. The full list of predefined semantic place labels is given below in Table 1.

Discussion. The MDC dataset is very unbalanced, in that much more labelled records exist for instance for the semantic place label *Home* than for *Holiday resort or vacation spot*. We do not know whether the distribution of labels is representative, and as we will discuss below, labelling behaviour may have influenced the classification results.

The classification problem that we tackle based on the MDC dataset is based on unique, but in terms of geographic location, unknown place IDs which define circles of 100m radius. Such an accuracy is plausible, if for instance cell tower triangulation is used, whilst with GPS or assisted GPS the location information should be more accurate². However, in a scenario of real application, any semantic place prediction algorithm would probably have access to geographic location

² <http://technowizz.wordpress.com/2010/01/03/lbs-technologies-part-1/>

Table 1. Semantic place labels

| | |
|----|---|
| 1 | Home |
| 2 | Home of a friend, relative or colleague |
| 3 | My workplace or school |
| 4 | Location related to transportation (e.g., bus stop, metro stop, train station, parking lot, airport) |
| 5 | Workplace or school of a friend, relative or colleague |
| 6 | Place for outdoor sports (e.g., walking, hiking, skiing) |
| 7 | Place for indoor sports (e.g., gym) |
| 8 | Restaurant or bar |
| 9 | Shop or shopping center |
| 10 | Holiday resort or vacation spot |

information. This would be an important piece of complementary information that could be used for instance in map lookups.

3 Problem Statement

In this work, we tackle the following problem:

Given an unlabeled place p_{ID} that has a number of features (computed from the data records associated with p_{ID}), predict the semantics of p_{ID} out of a list of predefined semantic place labels.

We consider this task as a supervised classification problem in combination with resampling to address the unbalanced nature of the dataset, and feature selection. As classifier, we used the Weka [5] implementation of a Random Forest algorithm [1] and of a Support Vector Machine (SVM) with the Pearson VII function (PuK) kernel [15] and as well as of an SVM with an Radial Basis Function (RBF) kernel.

4 Features

We used five feature groups: (i) daily patterns, (ii) weekly patterns, (iii) WLAN information, (iv) battery charging state, and (v) accelerometer information. Each feature group consists of multiple features described in more detail below.

4.1 Daily and Weekly Patterns

A daily pattern is a behavioural pattern that changes with the time of the day. We use the term “weekly pattern” in analogy to denote patterns of behaviour that change with the day of the week.

In the MDC dataset, we found strong evidence for daily and weekly patterns of users. The strongest evidence exists for the semantic place labels “Home”, “Home

of a friend, relative or colleague”, “My workplace or school” and “Workplace or school of a friend, relative or colleague” for daily patterns and “Place for outdoor sports” and “Place for indoor sports” for weekly patterns. Based on these insights, we used daily and weekly patterns as features.

For instance, between 1am and 4am, people are most often at home or at the home of a friend, relative or colleague (cf. Fig. 1 for “Home”). Figure 1 depicts the probability that the visited place is “Home” or “My workplace or school” on the time of the day. The probability is calculated as follows:

$$p_s = \left(\frac{\#records(s, t_1)}{\#records(t_1)} \dots, \frac{\#records(s, t_{24})}{\#records(t_{24})} \right) \quad (1)$$

where s is a semantic place label, t_i with $i = 0 \dots 23$ is the timespan of an hour starting at the time denoted by t_i , $\#records(s, t_1)$ is the number of records in the timespan t_i that are labelled with the semantic place label s in the MDC dataset and $\#records(t_1)$ is the number of records in the timespan t_1 for which a semantic place label exists in the MDC dataset.




Fig. 1. Daily pattern for the semantic place label “Home” (left) and “My Workplace or school” (right)

Given a place ID p_{ID} and all data records associated with p_{ID} , the daily pattern feature group consists of 24 features, each for a timespan t_i , $i = 0 \dots 23$ such that t_0 corresponds to the interval between midnight and 1am, t_1 corresponds to the interval between 1am and 2am etc. Each of the 24 features is computed as follows:

$$\frac{\#records(p_{ID}, t_i)}{\#records(t_i)} \quad (2)$$

where $\#records(p_{ID}, t_i)$ is the number of records at place p_{ID} in the timespan t_i , and $\#records(t_i)$ is the number of records taken in the timespan t_i by the user associated with p_{ID} .

The weekly pattern feature group consists of 3 features: The number of records taken at place ID p_{ID} on weekdays (wd), the number of records taken at place ID p_{ID} on weekend days (we), and the ratio $\frac{we}{we+wd}$.

4.2 Accelerometer Data

We hypothesized that information about users' movements would be discriminative for the given semantic place labels, e.g., to identify sports and transportation related places. In the MDC dataset, such information could be derived from accelerometer data.

The MDC dataset was collected from Nokia N95 phones which have an accelerometer with a sensitivity of $\pm 2G$ and a bandwidth of 35Hz [17]. The main challenge for using accelerometer data from the MDC dataset was how to compute velocity information without orientation and position information. In smartphones newer than the N95, a gyroscope is used to deliver orientation information. With accelerometer data from N95 smartphones, a normalization of the coordinate system is necessary to calculate average velocity given accelerometer data only. However, in [16], it has been shown that such a normalization is subject to errors which lead to an inaccurately computed direction of velocity. Therefore we do not normalize accelerometer data w.r.t. gravity at all, but simply integrate acceleration information to get average, approximate, velocity. We hypothesized that for the very short time intervals over which we integrate this approximation is sufficient to distinguish between semantic place labels.

In the MDC dataset, an accelerometer record is an array that consists of single accelerometer measurements within a timeframe. Each accelerometer measurement consists of the x, y, z acceleration in mG ($10^{-3}G$) and the time difference to the start of the timeframe. Each accelerometer record is also associated with a unique place ID p_{ID} .

We computed the average velocity within a time frame (i.e. for one *record*) by first integrating the x, y, z values separately. This gives us the velocity in the x, y, z direction of the accelerometer within the time frame of the *record*:

$$v_x(record) = \int_x record * dt \quad (3)$$

$$v_y(record) = \int_y record * dt \quad (4)$$

$$v_z(record) = \int_z record * dt \quad (5)$$

The Euclidean distance gives the velocity within the timeframe of the record:

$$v(record) = \sqrt{v_x(record)^2 + v_y(record)^2 + v_z(record)^2} \quad (6)$$

All velocities, computed over timeframes, are aggregated to give the average velocity and its standard deviation at a unique place p_{ID} . Both values are normalised to values between 0 and 1 with a min-max normalization.

$$av(p_{ID}) = \frac{\sum_i v(record_i(p_{ID}))}{\#records(p_{ID})} \quad (7)$$

is the average velocity at place p_{ID} , and $i = 1 \dots \#records(p_{ID})$.

$$std_v(p_{ID}) = \sqrt{\frac{1}{\#records(p_{ID})} \sum_i (v(record_i(p_{ID})) - av(p_{ID}))^2} \quad (8)$$

is the standard deviation of the average velocity with $i = 1 \dots \#records(p_{ID})$.

The accelerometer data feature group for p_{ID} consists of two features, namely the min-max normalised average velocity (min-max normalised Eq. 7) and the min-max normalised standard deviation of velocity (min-max normalised Eq. 8).

4.3 WLAN Information

We assumed that users connect to WLAN more frequently at some semantic places than at others. We therefore defined a feature that indicates the frequency of WLAN usage at a unique place p_{ID} :

$$fv_{p_k}^{\rightarrow} = \frac{\#connections(place_k)}{\sum_{i=0}^n \#connections(place_i)} \quad (9)$$

Given the fact that the N55 mobile phones that have been used to record the dataset, deactivate the WLAN connection when not currently used [12], the number of connections resembles the intensity of the WLAN usage.

The WLAN feature group consists of this single feature.

4.4 Battery Charging State

We assume that users charge their phones at selected semantic places. The MDC dataset provides four different charging states: (i) charger not connected (s_0), (ii) device is charging (s_1), (iii) charging completed (s_2), and (iv) charging continued after brief interruption (s_3). The feature vector for a unique place ID p_{ID} has four dimensions, i.e. one dimension for each charging state. Each dimension has a value between 0 and 1, denoting the number of records with the corresponding charging state at p_{ID} divided by the number of all charging state records of the user associated with p_{ID} :

$$\frac{|\{records(p_{ID}) | record(p_{ID}) = s_i\}|}{|\{records(p_{ID})\}|} \quad (10)$$

for $i = 0 \dots 3$ where $record(p_{ID})$ is a charging state record at place p_{ID} .

The battery charging state feature group thus consists of four features.

5 Experiments and Results

The Random Forest (RF) algorithm is an ensemble classifier consisting of multiple randomized decision trees that are combined using bagging [14]. We used it since it is known to be a highly accurate and fast classification algorithm [2]. Besides, the algorithm is not very sensitive to outliers, is able to deal with missing values, and, as stated in [1], avoids overfitting. Due to their wide use for classification problems we also evaluated the performance of Support Vector Machines on the problem at hand. We used two Support Vector Machines with two different kernels. Both SVMs implement John C. Platt’s sequential minimal optimization algorithm for training a support vector classifier using polynomial or RBF kernels³. The first SVM has a Radial Basis Function kernel (RBF-SVM). The second SVM has an SVM with the Pearson VII [15] function as a universal kernel function (PuK-SVM). The Pearson VII function is an alternative to the standard SVM kernels for which studies in the field of remote sensing suggest that it outperforms standard kernels [13].

To extract the best performing features from each feature group, we applied *feature selection* with a CfsSubsetEval⁴ filter from Weka. This filter evaluates features with respect to their individual predictive ability along with the degree of redundancy between the features [6]. Since the training data is unbalanced, we applied a *resampling filter* from Weka to introduce a bias towards a uniform class distribution. If features, for instance WLAN information, are not available, we treat them as missing values. All evaluation results have been computed with a *10-fold cross validation* on the MDC dataset.

5.1 All Feature Groups for All Semantic Place Labels

We evaluated the performance of the Random Forest algorithm and both types of Support Vector Machines, the SVM with an RBF kernel and the SVM with the PuK kernel with all feature groups for all semantic place labels, which gives a 10 class multi-class problem with 5 feature groups and 32 single features. This experiment resulted in an average f-measure of 0.854 for the Random Forest, an average f-measure of 0.764 for the SVM with PuK kernel, and an average f-measure of 0.366 for the SVM with RBF kernel. Detailed results for each semantic place label are given in the next section.

5.2 All Feature Groups for Each Semantic Place Label

We evaluated both the Random Forest algorithm and the SVM with all features for each semantic place label separately. Their performances are given in terms of the F-Measure in Table 2 next to each other.

³ <http://weka.sourceforge.net/doc/weka/classifiers/SMO.html>

⁴ <http://wiki.pentaho.com/display/DATAMINING/CfsSubsetEval>

Table 2. F-measure of the Random Forest (RF) algorithm and the SVM using the Pearson VII function kernel (PuK-SVM) and the SVM using the RBF kernel (RBF-SVM). All three algorithms used all feature groups and classified each semantic place label separately.

| Semantic Place Label | RF | PuK-SVM | RBF-SVM |
|---|------|---------|---------|
| 1: Home | .766 | .714 | .769 |
| 2: Home of friend, relative or colleague | .566 | .638 | .364 |
| 3: My workplace or school | .8 | .659 | .737 |
| 4: Place related to transportation | .825 | .538 | .105 |
| 5: Workplace or school of a friend, relative or colleague | .875 | .793 | .305 |
| 6: Place related to outdoor sports | .871 | .712 | .061 |
| 7: Place related to indoor sports | .866 | .788 | .103 |
| 8: Restaurant or bar | .962 | .938 | .354 |
| 9: Shop or shopping center | .918 | .857 | .449 |
| 10: Holiday resort or vacation spot | .966 | .894 | .425 |

Discussion. The SVMs underperformed the Random Forest algorithm for all semantic place labels except *Home of a friend, relative or colleague*. Therefore we carry out further analyses only with the Random Forest algorithm. The performance of the last three classes (8-10) is deceptively high. This is an artefact of the MDC dataset, which contains only few examples for these classes. Therefore we cannot assume a good generalisation ability of the classifiers for these classes (semantic place labels). This also holds true for all experiments below.

5.3 Single Feature Groups for Each Semantic Place Label

Next, we created Random Forest classifiers such that each classifier detects only one semantic place label and uses only one feature group (results shown in Table 3 on the next page). Such experiments lead to insights on the relevance of a particular feature group for different semantic places.

Discussion From Table 3 we can see that daily patterns perform very well on nearly all semantic place labels, except *Home of a friend, relative or colleague* and *Place related to transportation*. Especially the latter is surprising. One might assume that people travel very regularly for instance to and from work. In [3,4], strong daily patterns have been found in transportation networks. One possible solution for this discrepancy of results lies in the unknown labelling behaviour of the study participants who created the MDC labels within the MDC dataset. Did they tend to label places related to transportation maybe rather for unusual transportation paths and not for their daily routes to and from work, to and from supermarkets etc? The confusion matrix supports this interpretation as it shows that based on daily patterns, the class *Place related to transportation* is often confused with the class *Holiday resort or vacation spot*. We also expected

Table 3. F-measure of Random Forest classifiers that detect a single semantic place label using a single feature group. In the table below, Daily P. abbreviates Daily Patterns, Weekly P. abbreviates Weekly Patterns, Charging abbreviates Battery Charging State, and Accel. abbreviates Accelerometer Data.

| Semantic Place Label | Daily P. | Weekly P. | Charging | WLAN | Accel. |
|---|----------|-----------|----------|-------|--------|
| 1: Home | 0.776 | 0.826 | 0.72 | 0.627 | 0.4 |
| 2: Home of a friend, relative or colleague | 0.536 | 0.68 | 0.582 | 0.593 | 0.431 |
| 3: My workplace or school | 0.836 | 0.737 | 0.698 | 0.485 | 0.426 |
| 4: Place related to transportation | 0.646 | 0.468 | 0.794 | 0.806 | 0.889 |
| 5: Workplace or school of a friend, relative or colleague | 0.966 | 0.655 | 0.889 | 0.918 | 0.903 |
| 6: Place related to outdoor sports | 0.862 | 0.566 | 0.867 | 0.746 | 0.813 |
| 7: Place related to indoor sports | 0.875 | 0.787 | 0.866 | 0.828 | 0.907 |
| 8: Restaurant or bar | 0.883 | 0.763 | 0.949 | 0.95 | 0.884 |
| 9: Shop or shopping center | 0.881 | 0.75 | 0.897 | 0.833 | 0.907 |
| 10: Holiday resort or vacation spot | 0.785 | 0.672 | 0.977 | 0.966 | 0.966 |

the daily pattern feature group to perform better on the class *Home*. The main class of confusion is *Home of a friend, relative or colleague*.

Weekly patterns perform worse overall compared to daily patterns. However, they outperform all feature groups with respect to the *Home* and the *Home of a friend, relative or colleague* class.

The typical mobile sensor data, i.e. battery charging state, WLAN information and accelerometer data serve very well to predict the semantic place labels 4-10. The battery charging state feature group for instance performs exceptionally well on the classes *Restaurant or bar* and *Holiday resort or vacation spot*.

5.4 Feature Group Combinations for Each Semantic Place Label

We evaluated combinations of the winning feature group of daily patterns with the battery charging state, WLAN information and accelerometer data feature groups. The results are shown in Table 4, next to the results of the daily pattern feature group.

Discussion. Combining the daily pattern feature group with other feature groups improves the classification performance, except for the classes *My workplace or school* and *Workplace or school of a friend, relative or colleague*. However, even

Table 4. F-measure of the Random Forest classifiers that detect a single semantic place label using a combination of feature groups. In the table below, DP abbreviates Daily Patterns, Charging abbreviates Battery Charging State, and Accel. abbreviates Accelerometer Data. The last column repeats the results when using only the daily pattern feature group.

| Semantic Place Label | DP + Charging | DP + WLAN | DP + Accel. | DP |
|---|---------------|-----------|-------------|------|
| 1: Home | .826 | .783 | .8 | .776 |
| 2: Home of a friend, relative or colleague | .667 | .604 | .667 | .536 |
| 3: My workplace or school | .75 | .825 | .743 | .836 |
| 4: Place related to transportation | .836 | .794 | .844 | .646 |
| 5: Workplace or school of a friend, relative or colleague | .935 | .889 | .935 | .966 |
| 6: Place related to outdoor sports | .867 | .844 | .871 | .862 |
| 7: Place related to indoor sports | .862 | .892 | .879 | .875 |
| 8: Restaurant or bar | .916 | .962 | .962 | .883 |
| 9: Shop or shopping center | .93 | .876 | .941 | .881 |
| 10: Holiday resort or vacation spot | .988 | .966 | .955 | .785 |

here, the differences in performance are very small. The different combinations of feature groups perform approximately equally well, with the combination “Daily Patterns and Accelerometer” slightly in the lead. The results shown in Table 4 for combinations of the daily pattern feature group with the battery charging state, WLAN or accelerometer feature group are even better than when using all feature groups (results shown in Table 2).

6 Comparison with Other Work on the MDC Dataset

Other authors, such as [7,11,18] have also worked on the MDC dataset but used different machine-learning approaches.

In [11], the authors develop one binary classifier for each semantic place label, and binary classifiers are either k-Nearest Neighbour or Support Vector Machines, using very similar features than the ones used within this paper. Results given within the paper are derived using a 2-fold cross-validation. The relevance of this paper lies in a newly developed multi-coded class based multiclass evaluation rule that combines classification results of the binary classifiers. However,

the overall accuracy of the developed multi-class classifier is, with 73.26% significantly lower than what we show can be achieved within this paper.

In [7], the authors use a multi-level classification approach to address the fact that the dataset is unbalanced, and that the semantic place labels form sub-groups of semantic places that are distinguishable by different features. The authors use multiple classification algorithms like SVM, J48, etc. and combine diverging results of different algorithms by a fusion model. The authors have evaluated a very broad range of features (54) and identified movement behaviour, phone usage behaviour, communication behaviour as well as temporal behaviour of users and WLAN- and bluetooth information as good features. Using a 10-fold cross-validation, the authors reach an accuracy of 65.77%.

In [18], the authors compare the performance of Logistic Regression, SVMs, Gradient Boosted Trees, and Random Forests; the latter corresponds to our approach. The authors selected features automatically from the space of all combinations of all possible features. While many features correspond to raw sensor data, some are also preprocessed sensor data such as variance of accelerometer data etc. The best achieved result, with 10-fold cross validation, lies at 65.3%, achieved with the gradient boosted tree algorithm.

In terms of accuracy, our approach thus compares extremely favourable with algorithmically more complex approaches.

7 Conclusion

We have shown that for the problem of predicting semantic places based on mobile sensor data, a Random Forest algorithm outperforms both an SVM with an RBF kernel and an SVM with a Pearson VII function kernel, as well as other algorithmically more complex approaches. The performance of different feature groups, daily patterns, weekly patterns, WLAN information, battery charging state and accelerometer data, was analysed for the Random Forest algorithm. The single best performing feature group however was the daily pattern feature group. Its performance could be further improved by combining it with WLAN, battery charging state or accelerometer data. These combined feature groups result in a better performing classifier than even the classifier that uses all feature groups. This is highly encouraging for mobile computing, as fewer features mean that less computational power is needed to perform the classification.

Acknowledgements. The Know-Center is funded within the Austrian COMET Program under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

References

1. Breiman, L.: Random forests. *Machine Learning* 45, 5–32 (2001)
2. Caruana, R., Karampatziakis, N., Yessenalina, A.: An empirical evaluation of supervised learning in high dimensions. In: *Proceedings of the 25th International Conference on Machine Learning, ICML 2008*, pp. 96–103. ACM (2008)
3. Chen, C., Chen, J., Barry, J.: Diurnal pattern of transit ridership: a case study of the new york city subway system. *Journal of Transport Geography* 17(3), 176–186 (2009)
4. Grieco, M., Urry, J.: *Mobilities: new perspectives on transport and society*. Ashgate (January 2012)
5. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA data mining software: An update. *SIGKDD Explor. Newsl.* 11(1), 10–18 (2009)
6. Hall, M.A.: *Correlation-based Feature Subset Selection for Machine Learning*. PhD thesis, University of Waikato, Hamilton, New Zealand (1998)
7. Huang, C.-M., Ying, J.J.-C., Tseng, V.S.: Mining Users Behaviors and Environments for Semantic Place Prediction. In: *Proceedings of the Mobile Data Challenge by Nokia Workshop, Co-Located with Pervasive 2012* (2012)
8. Kim, B., Ha, J.-Y., Lee, S., Kang, S., Lee, Y., Rhee, Y., Nachman, L., Song, J.: Adnext: a visit-pattern-aware mobile advertising system for urban commercial complexes. In: *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications, HotMobile 2011*, pp. 7–12. ACM, New York (2011)
9. Kiukkonen, N., Blom, J., Dousse, O., Gatica-Perez, D., Laurila, J.: Towards rich mobile phone datasets: Lausanne data collection campaign. In: *Proceedings of the ACM International Conference on Pervasive Services, ICPS* (2010)
10. Laurila, J.K., Gatica-Perez, D., Aad, I., Bornet, B.J.O., Do, T.-M.-T., Dousse, O., Eberle, J., Miettinen, M.: The mobile data challenge: Big data for mobile computing research. In: *Pervasive Computing* (2012)
11. Montoliu, R., Martinez-Usó, A., Martínez Sotoca, J.: Semantic place prediction by combining smart binary classifiers. In: *Proceedings of the Mobile Data Challenge by Nokia Workshop, Co-Located with Pervasive 2012* (2012)
12. Oliver, E.: A survey of platforms for mobile networks research. *SIGMOBILE Mob. Comput. Commun. Rev.* 12(4), 56–63 (2009)
13. Pal, M.: Kernel methods in remote sensing: a review. *ISH. J. Hydraulic Eng. (Special Issue)*, 194–215 (2009)
14. Saffari, A., Leistner, C., Santner, J., Godec, M., Bischof, H.: On-line random forests. In: *2009 IEEE 12th International Conference on Computer Vision Workshops (ICCV Workshops)*, September 27–October 4, pp. 1393–1400 (2009)
15. Üstün, B., Melssen, W.J., Buydens, L.M.C.: Facilitating the application of Support Vector Regression by using a universal Pearson VII function based kernel. *Chemometrics and Intelligent Laboratory Systems* 81(1), 29–40 (2006)
16. Wang, S., Chen, C., Ma, J.: Accelerometer Based Transportation Mode Recognition on Mobile Phones. In: *Proceedings of the 2010 Asia-Pacific Conference on Wearable Computing Systems, APWCS 2010*, pp. 44–46. IEEE Computer Society, Washington, DC (2010)
17. Yang, J.: Toward physical activity diary: Motion recognition using simple acceleration features with mobile phones. *Data Processing*, 1–9 (2009)
18. Zhu, Y., Zhong, E., Lu, Z., Yang, Q.: Feature Engineering for Place Category Classification. In: *Proceedings of the Mobile Data Challenge by Nokia Workshop, Co-Located with Pervasive 2012* (2012)