



DAILY ROUTINE CLASSIFICATION FROM MOBILE PHONE DATA

Katayoun Farrahi ¹ Daniel Gatica-Perez ¹

IDIAP-RR 07-62

JULY 2008

TO APPEAR IN
MLMI'08

¹ IDIAP Research Institute, CP 592, 1920 Martigny, Switzerland and Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, {kfarrahi,gatica}@idiap.ch

DAILY ROUTINE CLASSIFICATION FROM MOBILE PHONE DATA

Katayoun Farrahi

Daniel Gatica-Perez

JULY 2008

TO APPEAR IN
MLMI'08

Abstract. The automatic analysis of real-life, long-term behavior and dynamics of individuals and groups from mobile sensor data constitutes an emerging and challenging domain. We present a framework to classify people's daily routines (defined by day type, and by group affiliation type) from real-life data collected with mobile phones, which include physical location information (derived from cell tower connectivity), and social context (given by person proximity information derived from Bluetooth). We propose and compare single- and multi-modal routine representations at multiple time scales, each capable of highlighting different features from the data, to determine which best characterized the underlying structure of the daily routines. Using a massive data set of 87000+ hours spanning four months of the life of 30 university students, we show that the integration of location and social context and the use of multiple time-scales used in our method is effective, producing accuracies of over 80% for the two daily routine classification tasks investigated, with significant performance differences with respect to the single-modal cues.

1 Introduction

Human activity modeling from large-scale sensor data is an emerging domain relevant to many applications, such as determining the behaviour and habits of individuals and the structure and dynamics of organizations [1, 3, 12]. This could be useful for social science research and self-awareness tools. Given the massive amount of data captured by ubiquitous sensors over long periods of time and involving many people, fundamental questions to address through automatic analysis include: Do people follow similar routines? Do certain people not follow other’s routines? Are routines useful in group discovery?

Recent research has attempted to analyze complex, real-life activities from indoor sensors such as cameras, microphones, proximity, or motion sensors [6, 12, 10, 8]. The limitations with indoor spaces are that the sensors are often fixed and only those activities that occur in the (local) physical space covered by the sensors can be recognized. Other recent approaches use wearable devices carried by people, which collect various types of evidence of their activities, including motion in dynamic environments [7] and audio in face-to-face conversations [1, 13]. However, these wearable devices are not always practical for multiple users over long periods of time. In this paper, we study human routines from sensors that have become an integral part of our daily lives, mobile phones. The functionality of this ubiquitous infrastructure of mobile devices is dramatically increasing [2, 3], not requiring users to modify their daily behavior for data collection.

We define routines to be temporal regularities in people’s lives. A routine often involves patterns of locations (e.g. being at work or at home, or going from work to home) and human interactions (e.g. as reflected by proximity information) over time, possibly over different time scales. Automatic routine classification and discovery are in general challenging tasks as people’s locations and interactions often vary from day to day and from individual to individual, and data from sensors can frequently be incomplete or noisy.

The problems addressed in this work are as follows: given a day in someone’s life, measured solely in terms of the noisy location and proximity information obtainable from a mobile phone, would this day more closely resemble a weekend or a weekday? If the person analyzed was a student, would a day in her life reveal potential group affiliations? More concretely, looking at the visualizations of location and proximity days in Figure 1, does a given day (a row in each of the visualizations) more closely resemble a weekend or a weekday? And do the day’s routines appear more like an engineering or a business student’s typical rituals? Answering such questions is difficult as users often work on weekends and the similarity in routines over days is often high. We would like to know how well we can automate these tasks. Using real-life data from the Reality Mining dataset [3], involving a large group of people over thousands of hours of activity, our work provides answers to these questions. This domain of research has been reviewed as a very promising technology [11].

The first contribution of this work is the novel investigation of a set of discriminant representations of location (measured from cell tower connection information) and proximity (measured from Bluetooth information) within a supervised learning framework. We investigated various representations characterizing proximity and location features in a day, such as multiple time-scales, proximity identity, quantity of proximate people, and representations with and without time considerations, to determine which best represented the underlying structure of the daily routines. The second contribution is the investigation of location-driven and proximity-driven day-type classification from a single day in the life of a user. The third contribution is the investigation of location-driven and proximity-driven group-type classification from single users’ days. The fourth contribution is the comparison of single-modal versus multi-modal (i.e. multiple information sources) representations for location and proximity data for the two activity recognition tasks at hand.

Overall, we found that integrating information at multiple time-scales is useful, that fusing proximity and location information is beneficial compared with individual cues, and that the targeted daily routines (day-type and student-type) can be recognized with good accuracy (80.3% and 89.6%, respectively) even though the sensor data is partly incomplete.

There are many difficulties inherent to the activity recognition tasks at hand, complicating the

already challenging dataset we used. Issues with mobile phone sensor data include poor indoor reception, incorrect data entries (due to the phone being left behind) and Bluetooth errors, to name a few. Further, proximity data is not always available, leaving many days without any information. Besides difficulties with the dataset itself, other challenges include the facts that students do not follow strict schedules, for example, they work on weekends regularly complicating the day type classification and that students might work or take classes in different buildings or offices, share offices or other spaces infrequently, etc and none of this is known a priori. Further, the dataset contains various types of students (undergraduate, graduate), which may follow different routines.

The paper is organized as follows. Section 2 presents the data set and highlights its inherent challenges. Section 3 describes our approach. Experiments and results are discussed in Section 4. Conclusions are given in Section 5.

2 Sensing Activity with Mobile Phones

The most widely deployed and used mobile computing device today is the mobile phone [9]. Current mobile phones can capture data related to the daily routines of large numbers of people over a large period of time. More specifically, their locations, such as being at work or home, can be captured from cell tower connections. Interactions can be captured by Bluetooth, which detects other Bluetooth devices within a small radius. Phone call and SMS activities can further be recorded. Phone application usage can be saved including the camera, calendar, games, and web browser usage [3]. Finally, content, including photos and video, can also be collected [2]. From the potential options, in this paper we examine both location and Bluetooth data.

Recent work has been done using coarse-grained Global System for Mobile communications (GSM) data from mobile phones to recognize high-level properties of user mobility (walking versus driving), as well as daily step count for a very small number (3) of users over the course of one month [9]. Both coarse and fine-grained location systems have been used to perform location-driven activity inference [9, 5]. In work by Eagle and Pentland [4], which is the closest to ours, student type affiliations are determined by clustering location information aggregated over a period of nine months. All of the works described used location-driven activity inference. In this work, we investigate the student-type task considering proximity-driven inference, in addition to location-driven inference. Further, we investigate an additional task of day type classification. In addition, we evaluate several representations for the dataset, inferring class types from *single days* of data, as opposed to aggregated intervals of data.

There are many challenges and sources of noise inherent in mobile phone data. They can be forgotten, turned off, or out of battery. There are also issues with cell tower connections such as poor indoor reception and fluctuating connections. Bluetooth errors include detection between certain types of walls, recording people who are not physically proximate. There is also a small probability Bluetooth will not discover other proximate devices [3]. Further, ground truth collection is a difficult task especially over long periods of time. Users labeling is often incomplete, unclear and often unavailable, sometimes due to privacy concerns. All these issues lead to noisy, partly incomplete and partly inaccurate data with very little ground truth to rely on.

We use the Reality Mining dataset [3], collected by N. Eagle at MIT. The activities of 100 subjects were recorded by Nokia 6600 smart phones over the 2004-2005 academic year. This comprises over 800 000 hours of data on human activity; if we take into account the location, proximity, and phone call information, this corresponds to over 2 million hours of collected data. This dataset has been built respecting the privacy concerns of individuals in the study. The subjects in the study are students and staff of MIT that live in a large geographical area covered by over 32000 cell towers. They work in offices with computers that have Bluetooth devices which can sense in a 5-10m radius [3]. The public location information available for all subjects in the study includes the cell tower ID, as well as the date and duration of connection. All of the subjects labeled the cell tower ID's which correspond to their homes. We obtained a list of MIT work cell towers which correspond to the Media Lab and

the Sloan Business school. The Bluetooth proximity data collected contains the IDs of two proximate devices as well as the date and duration of interaction. The list of work cell tower IDs obtained from MIT was incomplete as many students never connected to any of the cell towers in this list and thus were never considered to be at work. To resolve this issue, additional work labels were inferred from being in proximity to each person’s computer; we did not consider being in proximity to one’s laptop as being at work due to the mobile nature of the device. We assign a location label of HOME(H), WORK(W), or OTHER(O) to the 32000 cell towers. Towers which are not labeled as H or W are categorized as O. We have a fourth location label, NO RECEPTION(N), when there is no tower connection recorded for a person for a given time (eg. no battery, phone off or no reception).

3 Classifying Daily Routines

We address two classification tasks for daily routines: weekday vs. weekend routines, and engineering student-like vs. business student-like routines. In both cases, the input data is one day of location and/or proximity information.

3.1 Data Representation

The goal is to represent a day using location and proximity information that is discriminant to daily pattern classification. A day can be represented at multiple time scales, and people’s routines usually follow block-type schedules. In this paper, we quantify location and proximity information at two levels (one fine-grain at 30 minutes and one coarse-grain at 3-4 hours). These two time scales provide a simple model of time management that is appropriate, in our opinion, to characterize many people’s lives. For location data, keeping in mind the H, W, O, N labels, in addition to time considerations, useful information may be contained in the quantity of these locations present in a day, or the dynamics in which they occur (for example, work often follows home). Further, for proximity data, sources of useful data include the identity of the person with whom a user was proximate, the number of proximate people (quantity of proximity disregarding the user’s identities), as well as time considerations. These features motivated the various location-driven and proximity-driven representations presented next.

3.1.1 Location Representation

L_a *Fine-Grain Location.* For the fine-grain location representation, visualized in Figure 1 and 2 a), a day is divided into 30 minute non-overlapping time intervals, resulting in 48 blocks per day. We assign a location label of W, H, O, or N to each 30 minute block. For classification purposes, this 48 element vector was transformed to binary format. Note that over a 30 minute interval, typically several cell tower connections are made, often with continuous fluctuations between a few. To address this source of noise we select the cell tower with the maximum connection time over each 30 minute interval.

L_b *Bag of Location Transitions.* This representation was built from the fine-grain location representation considering 8 coarse-grain timeslots in a day. A location word contains 3 consecutive location labels presented for the fine-grain representation corresponding to 1.5 hour intervals followed by one of the 8 timeslots in which it occurred. Thus a location word has 4 components, 3 location labels followed by a timeslot. We take overlapping 1.5 hour sets of labels to make a location word, so that if we had a pattern HHHOW in timeslot 1, we would have the following location words: HHH1, HHO1, and HOW1. The bag of location transitions is the histogram of the present location words in the day.

L_c *Coarse-Grain Location.* For this representation, visualized in Figure 2 b), a day is also divided into 8 coarse timeslots. For each timeslot, there is a binary element representing the four location labels (H, W, O, N). If one of these labels was present within the given timeslot, it is counted as one, if this location was not present, it is counted as zero. This is a simplification of the bag of

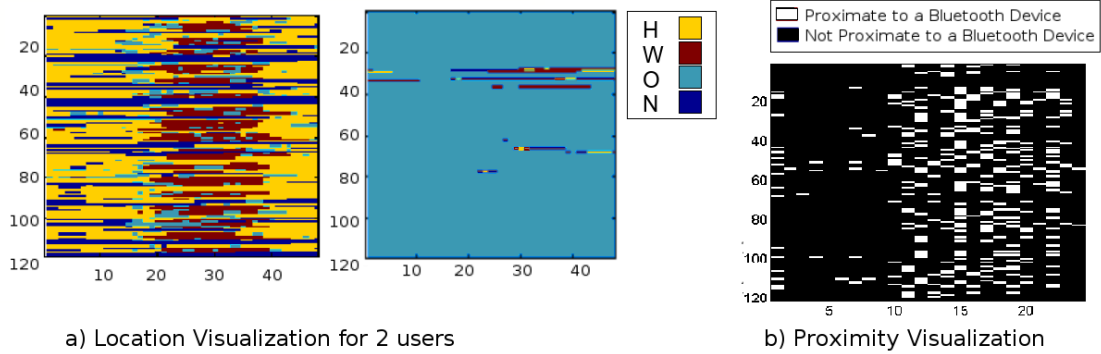


Figure 1: a) Visualization of location patterns using the fine-grain location representation, L_a , for 2 users over 121 days. Each row in the graph represents a day in the life of the individual. The labels H, W, O, N represent home, work, other, and no reception respectively. The day is divided into 48 fine-grain (30 minute) timeslots, each with a location label. The user on the left has a rich set of routines visible in the location patterns, whereas the user on the right is mostly incomplete due to lack of celltower labels. b) The proximity representation, P_b , is visualized for a user. Only proximity with users in the group are considered. Each row in the graph represents a day in the life of the user. The day is divided into 8 timeslots, each with 3 elements indicating the quantity of proximate users for that timeslot. For this user, most proximity activity occurs later in the day for most days.

location transitions, in which the dimensionality was reduced to be comparable to some of the proximity representations described in the next subsection.

L_d *Two-Feature Location*. This representation is the simplest, in which the number of 30-minute H and W labels are counted without taking into account when exactly they occur in a day.

3.1.2 Proximity Representation

P_a *UserID Proximity*. The userID proximity representation is also illustrated in Figure 2 c). There are 31 binary components for a given day, reflecting the 30 people considered in this study (see Section 4.1), and the last component for the case when no one is in proximity. If the person was in proximity with one of the 30 individuals, the value for that component will be one; for days when the person is not in proximity with anyone, the last component will be one. Thus, we only consider proximity within the set of 30 people. We do not consider a person to be in proximity with oneself.

P_b *Coarse-Grain Proximity*. The coarse-grain proximity representation, visualized in Figure 1 b), contains again the same 8 timeslots in a day. In this description of proximity, the identities of people are disregarded and only the quantity of proximate people for a timeslot is considered. In the first timeslot, the first element is one if 1 to 2 people are in proximity, the second if 3 to 4 people are, and the last if 5 or more people are in proximity. The resulting representation contains 8 timeslots, each with 3 elements. This idea of binary quantization is repeated over the 8 timeslots giving a quantification of interaction within the total set of people over different times in the day.

P_c *One-Feature Proximity*. This is the simplest representation for proximity. We count the number of proximate people for a person within a day, and use this value.

3.1.3 Combined Representation

For the combined representation, we concatenate one of the location representations with one of the proximity representations. In this paper, we only consider cases with comparable location and proximity dimensionality. Feature extraction techniques (e.g. PCA) could have been applied on the joint representations but were not explored here.

3.2 Classification

The classification was performed using a support vector machine (SVM) with a Gaussian kernel. For both daily routine classification tasks (days as weekends or weekdays, or days as belonging to business students or engineering students), the training strategy was leave-one-user-out, specifically testing on all the days for one unseen person while training on the data for all other people (note: proximity features are by definition relational involving pairs of people); we tested on each of the people and averaged the results. We optimized the kernel parameter on one data split for a randomly chosen person.

4 Experiments and Results

4.1 Data set

From the Reality Mining data set, we experimented with 30 people and 121 consecutive days, resulting in approximately 3600 data points. Our choice was guided by the goal of analyzing people and days for which data was reasonably available. The exact dates in the experiment were August 26, 2004 to December 24, 2004. The people selected had the most number of days with at least one W or H label. We removed days which were entirely N (no reception) labels since these had no useful information, which resulted in approximately 2800 data points. To select the interval of 121 days, we found the time interval with the most number of useful days (i.e., days with W, H, or O labels) over all 30 people. The resulting subset is still massive, amounting to over 87 000 hours, or about 10 years of data, and remains quite challenging in terms of noise, incompleteness, and complexity. This is illustrated in Figure 2 where it might be very difficult for a human to differentiate days as weekends/weekdays, or whether the day corresponds to a business student or engineering student.

For the student-type daily routine classification task, a subset of 23 of these 30 people were considered based on their student type labels. There were 6 business school students, and 17 engineering students. The engineering students covered a broader scope, including both undergraduate and graduate levels.

4.2 Weekday/Weekend Routine Classification

The weekend/weekday classification results are presented in Table 1 and reveal the difficulty of the task solely based on location or proximity information. In each table, the classification accuracy averaged over all people is presented first, and the average accuracy for each class is presented later. Generally, weekdays are more easily identified with location as input, and weekends are characterised better by proximity data. We can understand this by identifying weekdays with WORK cell towers, and weekends by not being in proximity with colleagues. However, in this dataset, students appearing to be in W locations on weekends complicate the classification task, resulting in at best 44.1% weekend classification accuracy by the bag of location transitions (L_b), which performs overall better than the others, also having the highest dimensionality. The coarse-grain approach L_c (fused bag of location words) performs slightly worse for weekends with a significantly smaller dimension. The fine-grain location representation, L_a , performs the worst for WE, the best for WD, and slightly better than the two-feature location case. All methods perform better than a ‘naive’ guess that assumes all days are weekdays, which results in $5/7 = 71.4\%$ accuracy.

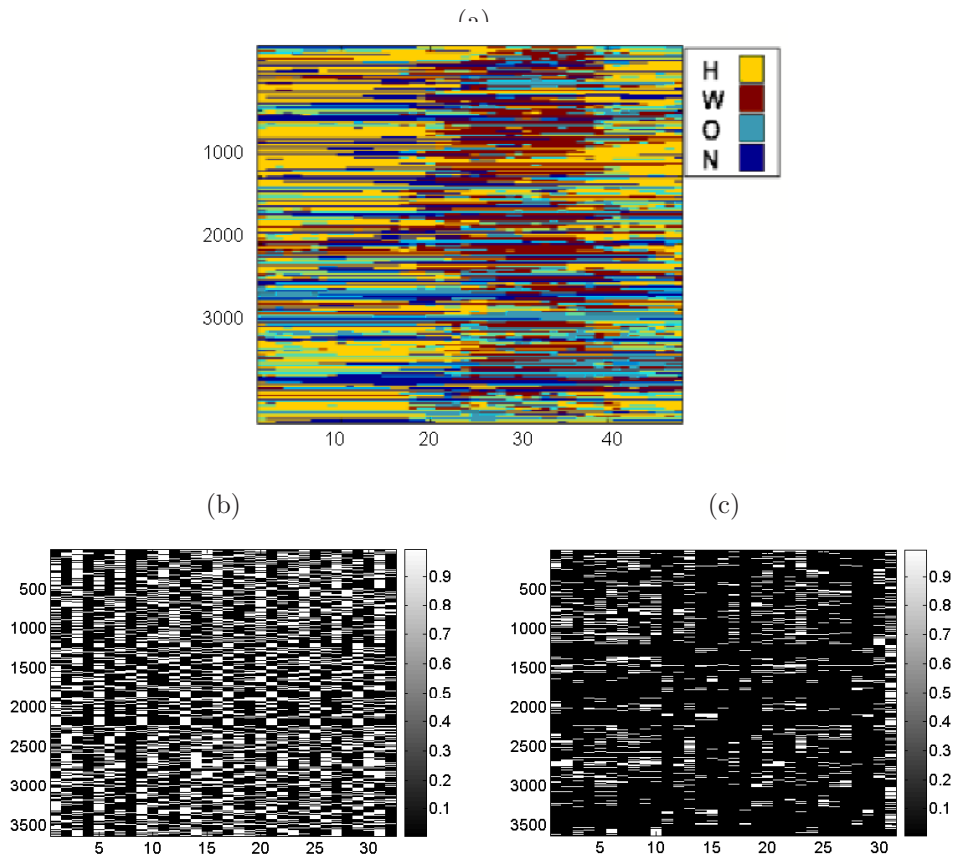


Figure 2: a) Fine-grain location representation, L_a , visualized over the entire set of days and users in the study. The x axis corresponds to the 48 half hour time intervals in a day. The y axis corresponds to a given day of a user in the dataset. b) Coarse-grain location representation, L_c , visualized over all days and users. c) UserID proximity, P_a , displayed over all users and days.

Table 1: Weekend (WE) and Weekday (WD) daily routine classification accuracy. The top table shows the difficulty in determining weekends based on location information alone. Proximity data is more deterministic of weekend routines. Classification obtained by combining location and proximity results in the best performance. Significance values are shown for the most significant results.

Location Accuracy (%)				Proximity Accuracy (%)			
Method	Overall	WE	WD	Method	Overall	WE	WD
L_a	74.2	19.3	95.3	P_a	74.3	70.7	75.8
L_b	76.8	44.1	89.1	P_b	72	54.2	78.7
L_c	76	36.6	90.8	P_c	74.6	67.9	77.1
L_d	75.7	30	93.1				
Combined Accuracy (%)							
Method	Overall	Eng	Bus				
(L_d, P_c)	76.9	47.35	88.1				
(L_c, P_a)	80.3	65.8	85.8				
(L_c, P_b)	79	53.4	89.3				
(L_a, P_a)	76.5	60.2	82.8				

Proximity information alone is useful in characterizing weekends, but does not perform as well as location data for identifying weekdays. There are many weekdays with little group interaction, resulting in higher confusion with weekends. The userID proximity and one-feature cases (P_a and P_c) reveal about 2% difference between their weekend and weekday performances, overall resulting in the highest performance of approximately 74%.

The lower panel in Table 1 shows the improvement in classification with the combination of proximity and location data. Note that in all cases the overall performance of the joint representations improved over that of the singleton case. We achieved over 80% accuracy with the combined representation (L_c, P_a) trading-off 2-3% weekday accuracy for improved weekend classification. In Figure 3, we visualize the days for which the proximity-alone data (columns 33-56) was misclassified, however when we added the location data (columns 1-32), the resulting 56-component vectors were correctly classified. In both figures, the first 32 columns visualize the location representation L_c and the last 24 columns illustrate the proximity representation P_b , so each row displays a day of the combination (L_c, P_b) . Figure 3a) are weekends which performed incorrectly for proximity-alone data due to the abundance of proximity interactions, which are not typical of weekends. In contrast, Figure 3b) shows weekdays which were mistaken for weekends due to the sparsity in interactions, not typical of weekdays. The addition of the location information in both cases resulted in correct classification, thus illustrating cases for which the combination of information improved classification performance.

The performance difference between the best location only method (L_b) and the best combined method (L_c, P_a) is statistically significant at the 1% level. The same is true for the performance difference between the best proximity only method (P_c) and the best combined method (L_c, P_a) .

4.3 Business/Engineering Student Routine Classification

Effectively classifying daily routines as belonging to business students or engineering students based on proximity-only observations was representation-dependent. Proximity representation P_c , the one-feature case, was inadequate in differentiating between student types, suggesting that the overall quantity of proximity within each group is on average the same. If the business students had much more proximity within the total set of people, or vice versa, we could expect the one-feature case to have higher accuracy. The coarse-grain proximity representation P_b improved the accuracy of business student classification, however, the userID proximity representation proved to be the best, with almost 99% accuracy in engineering student classification and 61% for business students. The knowledge of identity from proximity is the key for discriminating student disciplines.

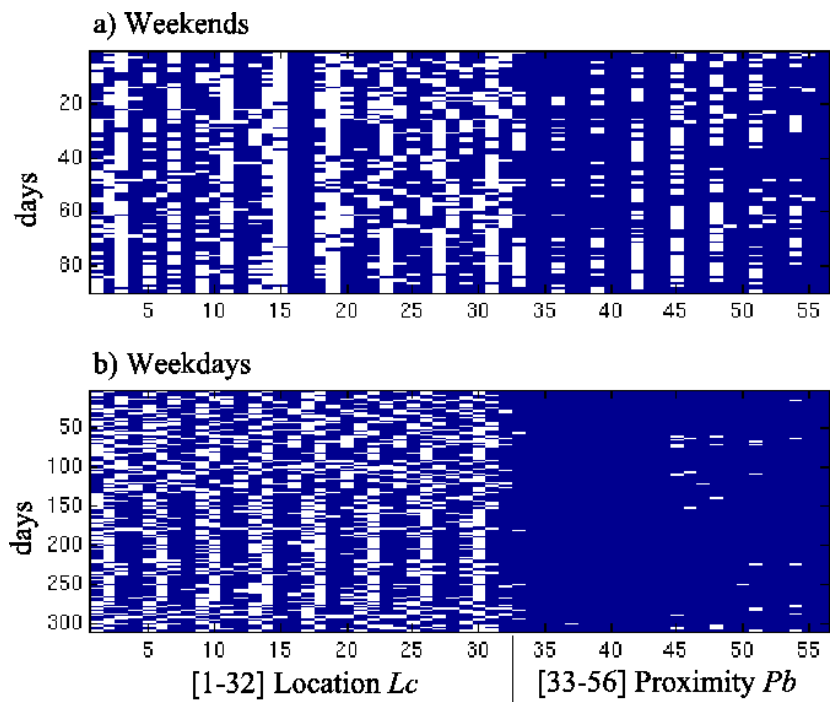


Figure 3: Advantages of the joint location-proximity representation (L_c, P_b) . Visualization of a) weekends and b) weekdays for which the proximity-only data was misclassified, but for which the location-only data and the combined proximity-location data were correctly classified. The sparsity of the weekday proximity-only data (columns 33-56 in b)), resulted in incorrect classification since sparsity in interaction is typical weekend behavior. However, when we added the location information, the resulting combined representation was correctly classified. The opposite phenomena can be observed in plot a), for which weekends have abundant proximity data, typical of weekday behavior.

Table 2: Engineering (Eng) vs. Business (Bus) student daily routine classification results. Proximity within the specific group is most representative of student type, especially when student identity is retained. The joint location and proximity data improves classification performance for the (L_c, P_a) combination. However, the other combinations generally perform as well as the singleton cases. Significance values are shown for the most relevant results.

Location Accuracy (%)				Proximity Accuracy (%)			
Method	Overall	Eng	Bus	Method	Overall	Eng	Bus
L_a	66.8	90.4	0	P_a	89.1	98.9	61.2
L_b	74.54	94.3	19	P_b	78.1	96	28.1
L_c	74.5	94.8	17.1	P_c	50.2	95.3	0
L_d	74.8	99.6	4.5				
Combined Accuracy (%)							
Method	Overall	Eng	Bus				
(L_d, P_c)	73.3	97.6	4.5				
(L_c, P_a)	89.6	99	62.9				
(L_c, P_b)	78.76	93.4	37.4				
(L_a, P_a)	84.5	95	54.7				

Location knowledge was inadequate in student type determination for the most part. This is likely due to the simplified representation used where the 32 000 cell tower IDs have been reduced to four location classes. It is expected that a representation more precisely identifying the location of a student would perform better. However, the representation used here is useful in understanding whether student types differ in the amount of time spent at school, home, or out and about. The two-feature location case, L_d , having low accuracy, indicates that the amount of time spent at school and home is not indicative of student type. The most effective characteristics in differentiating, which can be observed by the highest performance with the bag of location transitions representation, might be patterns of “going to work” in a timeslot, or “coming home” in a timeslot, or other similar routines which are captured by this representation.

The performance difference between the best location only method (L_d) and the best combined method (L_c, P_a) is statistically significant at the 1% level. The performance difference between the best proximity only method (P_a) and the best combined method (L_c, P_a) is not statistically significant.

5 Conclusion

We presented a method to classify daily life routines from massive, complex data collected with mobile phones. Using over 87 000 hours of data, we achieved over 80% accuracy in identifying whether a given day more closely resembles a weekend or weekday. This is not an easy task as students spend many weekends in work locations and have many weekdays with few group interactions. We showed that the integration of location and proximity data performed significantly better than the single observation sources, and that using representations that consider multiple time scales was beneficial. We further succeeded in identifying whether a user is an engineering or business student with over 89% accuracy based on a single day pattern of activity. The identity of individuals, measured by proximity, was key in this case, which confirms that social context is very helpful to identify people’s routines. We plan to further exploit this concept for other daily routines relevant for the analysis of mobile social networks.

Acknowledgments This research has been supported by the Swiss National Science Foundation through the MULTI project. We thank Nathan Eagle (MIT) for sharing the data and helping with various aspects of the collection structure.

References

- [1] Choudhury, T., Pentland, A.: Sensing and Modeling Human Networks using the Sociometer. In: Proceedings of ISWC, (2003).
- [2] Davis, M., House, N.V., Towle, J., King, S., Ahern, S., Burgener, C., Perkel, D., Finn, M., Viswanathan, V., Rothenberg, M.: MMM2: Mobile media metadata for media sharing. In: Proceedings of ACM CHI, Portland (2005).
- [3] Eagle, N., Pentland, A.: Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing* 10(4), 255-268 (2006).
- [4] Eagle, N., Pentland, A.: Eigenbehaviors: Identifying structure in routine. *Behavioral Ecology and Sociobiology* (in submission), (2007).
- [5] Hariharan, R., Krumm, J., Horvitz, E.: Web-Enhanced GPS. In Proc. of the International Workshop on Location and Context Awareness (LoCA), (2005).
- [6] McCowan, I., Gatica-Perez, D., Bengio, S., Lathoud, G.: Automatic Analysis of Multimodal Group Actions in Meetings. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)* 27(3), 305-317 (2005).
- [7] Munguia Tapia, E., Intille, S.S., Haskell, W., Larson, K., Wright, J., King, A., Friedman, R.: Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart monitor. In: Proc. Int. Symp. on Wearable Comp., Boston (2007).
- [8] Oliver, N., Horvitz, E., Garg, A.: Layered Representations for Learning and Inferring Office Activity from Multiple Sensory Channels. In: Proceedings of Int. Conf. on Multimodal Interfaces (ICMI), Pittsburgh, PA. (2002).
- [9] Sohn, T., Varshavsky, A., LaMarca, A., Chen, M., Choudhury, T., Smith, I., Consolvo, S., Hightower, J., Griswold, W.G., Lara, E.: Mobility Detection Using Everyday GSM Traces. In: Proceedings of UbiComp 4206, 212-224 (2006).
- [10] Stiefelhagen, R., Bernardin, K., Ekenel, H.K., McDonough, J., Nickel, K., Voit, M., Woelfel, M.: Audio-Visual Perception of a Lecturer in a Smart Seminar Room. *Signal Processing - Special Issue on Multimodal Interfaces Vol. 86* (12), Elsevier, (2006).
- [11] Technology Review,
<http://www.technologyreview.com/specialreports/specialreport.aspx?id=25>
- [12] Wren, C., Ivanov, Y., Kaur, I., Leigh, D., Westhues, J.: SocialMotion: Measuring the Hidden Social Life of a Building. *Third International Symposium on Location- and Context-Awareness* 4718, 85-102, 2007.
- [13] Wyatt, D., Choudhury, T., Kautz, H.: Capturing Spontaneous Conversation and Social Dynamics: A Privacy-Sensitive Data Collection Effort. In: Proceedings of International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Honolulu (2007).