

Supporting Location-Aware Services for Mobile Users with the Whereabouts Diary

Nicola Biccocchi, Gabriella Castelli, Marco Mamei, Alberto Rosi, Franco Zambonelli
Dipartimento di Scienze e Metodi dell'Ingegneria – Università di Modena e Reggio Emilia
Via Amendola 2 – 42100 Reggio Emilia – Italy
name.surname@unimore.it

ABSTRACT

Modern handheld devices provided with localization capabilities could be used to automatically create a diary of user's whereabouts, and use it as a complement of the user profile in many applications. In this paper we present the Whereabouts diary, an application/service to log the places visited by the user and to label them, in an automatic way, with descriptive semantic information. In particular, Web-retrieved data and the temporal patterns in which places are visited can be used to define such meaningful semantic labels. In this paper, we describe the general idea at the basis of our service and discuss our implementation and the associated experimental results. In addition, we illustrate an application that can fruitfully exploit the whereabouts diary as a supporting service, and discuss areas for future work.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Systems]: Distributed Systems; I.2.3 [Artificial Intelligence] Probabilistic Reasoning.

General Terms

Algorithms, Location-based Services

Keywords

Location services, Localization algorithms, Bayesian networks.

1. INTRODUCTION

The recent diffusion of handheld devices and smart phones equipped with localization capabilities is opening new scenarios in the development of context-aware services. In this paper we present the ideas and a first prototype implementation of the *whereabouts diary*: an application running on a GPS-equipped handheld device that records the list of relevant places visited by the user. The diary runs autonomously without requiring user's interaction and is able to classify *semantically* the places being visited in an unsupervised way. Relevant places can be extracted by considering clusters and dropouts in the GPS signal (typically indicating the user staying in a place or entering in a building). Semantic information can be added by exploiting the structure of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Mobilware'08, February 12-15, 2008, Innsbruck, Austria.
Copyright © 2008 ACM 978-1-59593-984-5/08/02... \$5.00

people daily routine. For example, the place where the user usually spends night-time can be tagged semantically as "home", while the place where the user usually goes from 8am to 6pm can be tagged as "work". Specifically, we realized a set of Bayesian networks to diagnose the kind of place given the temporal pattern of user visits. Further information can be extracted by geocoding the place and mining the Web in search for relevant information.

The result is a diary describing the user daily life that could be used either as a stand-alone application allowing a user to browse through his past locations, or as a tool supporting other applications.

From the latter perspective the Whereabouts diary can be regarded as a sort of middleware-level service to support location-aware mobile services. While most of the proposed middleware infrastructures focus on "low-level" issues like communication and interaction among nodes, the Whereabouts diary focuses on high-level issues such as context-representation and its understanding (e.g., it encodes what places matter to the users, and what they mean). In our opinion, such kind of infrastructures will become more and more important in the future. As context information will become more detailed and relevant, there will be the need of managing them at the middleware level, providing application with a pre-digested and understandable view of such data [3].

2. THE WHEREABOUTS DIARY

The construction of the whereabouts diary is an incremental process. In this section we describe (i) how to extract the places visited by the users in terms of their geographical locations, (ii) how to derive addresses and business' activities from such places, (iii) and finally how to add personal semantic labels (e.g., home, work, etc.) to the places.

2.1 Diary based on Coordinates

At the most basic level, the diary is an application that continuously collects and stores the user location. In our implementation the diary collects a log of GPS reading sampled at 0.1 Hz. This creates a log like the one presented in Fig. 1.

Longitude	Latitude	Time
-73.974	40.763	July, 4, 2006, 4:35:00 pm
-73.973	40.766	July, 4, 2006, 4:35:10 pm
-73.974	40.765	July, 4, 2006, 4:35:20 pm

Fig. 1 Log of GPS coordinates.

Starting from such a log of GPS readings, it is possible to run segmentation and clustering algorithms to infer the places where the user spends most of his time [6]. Following an approach

similar to the one proposed in [1], we tagged as relevant those places for which one of the following conditions apply:

1. the GPS signal is lost for at least T seconds and it is re-acquired later on at a distance of less than L meters from where it was lost. This reflects the situation in which a user enters a building and leaves it after some time. Some empirical evaluations let us to set T = 20 minutes, L = 20 meters. The constraint on time is important to wash out GPS signal glitches, the constraint on space is useful to avoid those situation in which the GPS has been shut down and the user moves away.
2. The GPS readings over a time window of W seconds are clustered within a radius of R meters from each other. This reflects the situation in which the user stays for a long time in a place like a park or a square. Some empirical evaluations let us to set W = 20 minutes, R = 100 meters.

The list of relevant places is built online and incrementally. When a set of coordinates meets one of the above criteria, the system looks in the list of the already discovered places for one closer than L = 10 meters to the coordinates. If such a place does not exist, a new place is created and the time of visit is recorded. If the place exists, the place coordinates are averaged with the new coordinates, and if enough time has passed since the previous visit (30 min), the time of the new visit is recorded.

The result of this operation is a list of places described in terms of longitude and latitude, and a list of time intervals associated to each of the coordinates indicating when the user has been there. This process creates a diary like the one presented in Fig. 2.

Longitude	Latitude	Time
-73.974	40.763	July, 4, 2006, 4:35pm-5:41pm
-73.774	40.563	July, 4, 2006, 6:35pm-7:00pm

Fig. 2 Diary based on GPS coordinates.

2.2 Diary based on Addresses and Businesses

A simple list of coordinates is only partially informative and the need of translating from positions to places (i.e., adding semantic meaningful tags to the discovered coordinates) has been widely recognized [5]. A diary containing information like “the user was at home” rather than “the user was at coordinates (10.873, 44.630)” would be naturally much more informative and easy to use in context-aware applications.

A step in the process of adding semantic information would be to translate from coordinates to addresses. This can be done via standard tracking and geocoding services (as common GPS navigators do). However, because of errors in GPS localization and errors in the process of segmenting and clustering the GPS readings to identify relevant places, in most of the situations, it will not be possible to identify the unique address where the user is located, and only a partial estimate can be given. This step converts the diary in the one depicted in Fig. 3.

In our implementation, the coordinates associated to places have been translated into addresses using a custom geocoding service. Most of the geocoding services available online (e.g., that provided by the Google Maps API) translate addresses into coordinates. Instead, the diary needs the reverse operation: from coordinates to addresses. We developed a “reverse” geocoding service for our region, on the basis of maps available from a

commercial navigator software. To take into account GPS and geocoding inaccuracies, and the errors introduced by the place retrieving process, the diary application tries to reverse geocode all the addresses within a radius of 10m from the place being segmented. Thus, the diary actually creates a list of candidate addresses where the user has been.

Place	Time
123, 5 th Ave, NY, USA	July, 4, 2006, 4:35pm-5:41pm
4,5,...,21, 26 th St., NY, USA	July 6, 2006, 7:00am – 8:00am

Fig. 3 Diary based on addresses. Because of GPS errors multiple addresses can be associated to a single place.

In another step, the diary can try to mine the Web to identify what is in a particular address. However, due to the aforementioned localization errors, this process will return in some of the cases a list of all the businesses performed in the geocoded addresses. Still, in some situations a single exact match could be retrieved like in the case of the user being in a big stadium or entering a big shopping mall. Even more semantic information could derive by searching relevant events that happened in that place at that time. For example, it could be possible to extract from the Web the fact that “the 4th of July parade” took place near the geocoded location at the same time the user was there. This process could create a diary like the one depicted in Fig. 4.

Place	Time
4 th July Parade	July, 4, 2006, 4:35pm-5:41pm
126, 13th St., NY, USA	July 5, 2006, 11:00pm – 7:00am
Starbucks Coffee Uno’s Pizza	July 6, 2006, 7:00am – 8:00am

Fig. 4 Diary based on places. Because of errors in the previous phases, multiple businesses can be associated to a single place.

To perform this operation, in our implementation, we screen-scraped information coming from a widely used online white-pages service (www.paginebianche.it) in our region allowing to query for who is at a given address. In particular, each geocoded address belonging to a given place (as provided by the previous step) is looked up in the white-pages and the corresponding business is retrieved. The result of this process is a set of entries labeled with the possible businesses found in that place. This translation process is not completely accurate, since several addresses are not listed in the white-pages (mainly due to privacy constraints). Still, the fact that most public businesses (like shops, etc.) are listed, while several private houses are not, allows to prune out a lot of unlikely addresses being discovered by the previous step. Private spaces like “home” – that are not listed in the white-pages – can be derived from the following analysis.

It is finally worth pointing out that the operations described in this subsection could be automated and improved given the availability of a complete spatial database like the one integrated into commercial navigator systems.

2.3 Diary based on Personal Places

Given that the user activities are profiled in some way (e.g., the diary may know a priori that the user tends to stay at home at night), the diary application can give labels to places by looking

at the temporal patterns in which places are visited. For example, the place most visited at night during weekdays can be meaningfully labeled as “Home”. Such kind of analysis can be also used in combination with commonsense information [8] to disambiguate between alternative retrieved places. For example, in the table in Fig. 4 the ambiguity among “Starbucks Coffee” and “Uno’s Pizza” can be resolved (at least from a probabilistic point of view) in favor of the former, in consideration of the fact the place has been visited from 7:00 am to 8:00 am.

To implement this step, for each place being identified in the first phase, the diary creates a Bayesian network to analyze the temporal pattern in which the place has been visited by the user (see Fig. 5). The Bayesian network is composed of 4 nodes.

1. The *weekend* node represents a boolean variable used to represent whether a given observation takes place in the weekend or not. This node is always observed on the basis of the information stored in the GPS signal. This information represents the variability in people behavior between weekdays and weekends.
2. The *hour* node is a 24-values discrete node storing the time of day. This node is always observed on the basis of the information stored in the GPS signal.
3. The *kind of place* node is a discrete node modeling what a given place is. In our implementation, we try to classify among 5 different kind of places: home, work, restaurant (to indicate any kind of dining place), pub (to indicate any kind of evening entertainment), and disco (to indicate any kind of late-night entertainment). This classification is rather arbitrary, and each user of the diary should provide the kinds of place that best match his habits. This node is never observed, and is inferred by probability computations.
4. The *happens* node is a boolean variable expressing whether the user visits that place at that time. This node is always observed on the basis of the outcomes of the diary localization phase.

The role of the Bayesian network is to encode the routine of the user daily life. This is done by compiling the probability distribution associated to the fact that the user, in a given moment, is in a certain kind of place. For example, the probability of the user being at home during weekdays is in the table in Fig. 6.

Similar tables can be created for other kind of places. In our current implementation, these tables are compiled by hand by the users that are asked to self-report the likelihood of being in a given kind of place at a given time. Such kind of data could be derived automatically also by a labeled trace of user’s past whereabouts, using standard learning algorithms [10]. Once the tables are filled in, basic inference operations in Bayesian networks will be used to derive the most likely kind of place given the visit pattern.

Specifically, when the diary previous phases identify that the user is visiting a place, the corresponding Bayesian network is retrieved, and the *weekend*, *hour*, *happens* nodes are set to their actual values (the *happens* node is trivially set to *true* to indicate that there is a visit). Then, the diary computes the probability distribution of the *kind of place* node. The newly computed distribution will be used as a prior for subsequent visits. This naturally allows evidences to add up, actually enabling the Bayesian network to classify the places on the basis of the visit temporal pattern.




Fig. 5. Bayesian network to classify places. White nodes are those that will be provided as evidence.

Weekend = false, Kind of Place = home						
time	11pm-6am	7am-8am	9am-1pm	2pm-5pm	6pm-7pm	8pm-10pm
P(happens)	0.8	0.6	0.2	0.2	0.4	0.5

Fig. 6. Conditional probability table describing the probability of the user being at home during weekdays.

In its final form the diary represents a powerful source of context information allowing to extrapolate user’s habits, preferences and routine behavior.

Of course, should other kind of sensing devices be available (e.g., RFID and NFC – Near Field Communication – readers), classification could use such information to better identify the places. For example, a powerful source of data could come from credit card transaction records that would identify not only in which shop the user has been, but also what he has bought. In the end, the combination of all the above steps leads to a diary close to the one in Fig. 7.

Place	Time
4 th July Parade	July, 4, 2006, 4:35pm-5:41pm
Home	July 5, 2006, 11pm – 7am
Starbucks Coffee	July 5, 2006, 7am – 8am

Fig. 7 Diary based on personalized places.

It is fair to point out that the idea of the whereabouts diary is not completely new. One of the first systems trying to log the places visited by the user is has been proposed in 1991 [9]. Still, the recent availability of cheap localization devices and the recognition of location as a primary source of context information has stimulated a wealth of related works ranging from localization algorithms [6] to semantic analysis of the discovered places [10].

3.EXPERIMENTS

To test the effectiveness of the whereabouts diary, we collected GPS traces for three weeks from three members of our research team (among the authors) as they went about their normal lives. Each member carried either an i-mate PDA 2K smart phone, or a HP IPAQ RX3700 pda, connected with a Bluetooth GPS reader. GPS signal has been acquired at 0.1Hz and processed on the fly by the handheld device. Overall, we acquired about 90000 GPS poses amounting at 360 MB of data. Overall, this resulted in 25 places being identified as relevant. During the data collection weeks, data collectors recorded ground-truth information about the places they have been. Such information has been collected with a simple notepad application running on the PDAs and

allowing to write a textual description of where the user has been at a given time. In the following, we present some results obtained by comparing the whereabouts diary entries, after some weeks of usage, with recorded ground-truth information.

In a first set of experiments, we tried to verify the accuracy of the algorithm to identify relevant places on the basis of the GPS trace log. Following an approach similar to [6], we classify the incorrect results into: (i) *wrong*: the user is in a place, but the diary reports he is in a different place, (ii) *false negative*: the user is in a place, but the diary reports he is moving, (iii) *false positive*: the user is moving, but the diary reports he is in a place. The results of this experiment are reported in Fig. 8, and they actually show the average of the results obtained by the data collectors. The results we obtained show that the algorithm is correct in 80.6% of the cases. This figure is coherent with the results presented in [6] with regard to the A-S algorithm [1] that is the one closer to our implementation. The high-percentage of false negatives (compared to the other cases) is mainly due to the fact sometimes the GPS takes a long time before acquiring the signal. Thus, it can happen that a user leaves a building, and the trace of the GPS is acquired only when he is already far away. In such a situation the place is not detected given the constraint described in Sect. 2.1.

Correct	Incorrect		
	Wrong	False neg.	False pos.
80.6%	0%	17.3%	2.1%

Fig. 8. Errors in the algorithm to identify relevant places on the basis of the GPS trace log.

In a second set of experiments, we tried to verify the results of the (reverse) geocoding service. Basically, the idea is to verify the impact of localization errors in the process of geocoding. It is worth noticing that the maps we used to perform this operation record only the first and the last number of a street segment and span, uniformly, all the other numbers among the segment. This of course introduces further errors in that it does not take in to account the differences in the sizes of the buildings.

Since the place discovery algorithm clusters together points that are closer than 10 meters, we counted the number of addresses retrieved within a circle of 10 meters radius centered at the relevant place. The results of this operations are displayed in Fig. 9, and highlight two aspects. On the one hand, the address of almost half of the places can be retrieved uniquely (this is the case of large buildings – like the departments of our university). On the other hand, some places produce more than 10 associated addresses. This is the case of small buildings in the center of the city. It is fair to report that these distributions are rather preliminary since they are based on a dataset of only 25 places (those identified by the diaries of the 3 data collectors). We are currently conducting a more extensive data collection process that would allow us to identify more stable distributions.

In a third group of experiments, we tried to evaluate the performance of retrieving the businesses performed on a given place via the white-pages service. Comparing the results with the ground-truth annotations, we first tried to determine whether the correct place is retrieved. With disappointment, we verified that the actual place can be retrieved in only 40% of the cases. This is either due to localization or white-pages errors. Moreover, due to

the multiplicity of addresses being discovered several businesses can be assigned to a given place. In Fig. 9, we report the distribution of the number of businesses found for a given place. It is easy to see that some addresses are not listed in the white-pages, since there are no places with more than 10 retrieved businesses. In addition, It is worth reporting that the number of businesses being retrieved is almost independent of whether the correct place has been found or not. Some places, in fact, return a long list of candidate entries not containing the correct one. The main source of errors of this phase is related to the white-pages interface and how it handles street numbers.

In our future work we will try to improve this result in many directions. First, we will try to integrate our software with commercial spatial databases (e.g., <http://www.tomtom.com/pro>). This would notably improve the coverage of the addresses and businesses being mapped. Second, we will try to implement more advanced localization algorithms in order to reduce the uncertainty about user location [6]. Finally, we will try to embed commonsense reasoning (e.g., the user does not go to a disco in the morning) to cut off unlikely possibilities. The last resort would be to ask the user. In uncertain situation the diary could pop up and let the user solve possible ambiguities.

Finally, the last set of experiments verified the results of the Bayesian classification. Overall, our approach classifies the places correctly in 64% of the cases. In order to better analyze the results we tried to assess the confidence of the diary in its own classification – most probable estimate (MPE). To this end, we compute the information entropy of the resulting distributions. The lower the entropy, the more the system is confident about the MPE (i.e., the distribution peaks on the MPE value). More in detail, we separated the entropies related to the distributions that produced a correct MPE from those that produced a wrong MPE. For each of these two categories, we averaged together the entropies of the distributions producing the same MPE (see Fig. 10). Looking at the graph it is possible to notice that entropies related to wrong MPEs are higher than entropies related to correct MPEs. This is good and reflects the fact that the distributions associated to wrong estimates are less peaked, and thus the diary is less confident about its own classification. In fact, examining the wrong distributions in a lot of circumstances the distribution is bimodal: one peak is the correct one and another slightly more probable is incorrect. In such circumstances, it is likely that more observations on the temporal pattern of visits will correct the classification outcome.




Fig. 9. Percent of relevant places corresponding to a given number of addresses/businesses.

Another interesting remark about the result in Fig. 10, relates to the average entropy associated to the different kind of places. Not surprisingly, *home* and *work* places have lower entropies since their associated temporal pattern of visits is defined more precisely. On the contrary, places like *pubs*, *restaurants* and *discos* have a more flexible pattern of visits and thus they are classified less precisely and tend to produce higher entropies.




Fig. 10. Information entropy of the resulting distributions for the various kind of places. The information entropy for a 5-value discrete distribution ranges from 0 to 2.32 (flat distribution).

4.Applications

Once the diary is available, it can be used as a complement of the user profile in several applications. In this section we present a tourist diary application that we already started developing, other examples can be found in [4].

A tourist diary application could match the places visited by the user across a database of tourist attraction and suggest possible next visits. In addition, the exchange of diaries among tourists could allow to exchange experiences, hints and recommendations [2]. Also in this case, the extraction of semantic information (e.g., retrieving shops and other businesses in a given area) could provide a notable added value to the information being communicated.

To test the whereabouts diary, we started implementing a guide application of our city offering advices and recommendations to tourist on the basis of their past travel pattern. For example, if the diary indicates that the user likes visiting museums, the guide suggests visiting other museums of a city.

The guide has been implemented as a Web application, based on Google Maps and accessible via PDA and smart-phones. The application works on the basis of these points:

- All the attractions in the city are stored in a database with their type (e.g., name="Raffaello" type="cinema").
- The diary runs the above clustering algorithms and uploads to the server the coordinates of relevant places where the user has been.
- For each place identified by the diary, the list of all the tourist attractions within 80 meters is retrieved. Each attraction is associated with the probability that it was the one visited by the user among the list of possibilities. We simply calculate it with:

$$\frac{1}{\text{number of retrieved attractions for that place}}$$

- This information is stored in a database associating to each place all the retrieved attractions together with their type and with their probability (see Fig. 11).

date	duration	name	Type	p
28-9-2007:19:00	24000	Ariosto	Cinema	0.25
28-9-2007:19:00	24000	Il Pozzo	Bar	0.25
28-9-2007:19:00	24000	Dante	Cinema	0.25
28-9-2007:19:00	24000	Grande	Square	0.25

Fig. 11. database associating places with attractions together with their type and their probability.

Once a user requires a recommendation to the guide, the guide computes the most likely kind of attraction the tourist may be interested in. In particular, for each type of attractions the diary computes a score by adding all the time the user spent in a given kind of place multiplied by the associated probability:

$$score_{kind_i} = \sum_{places\ of\ kind_i} time_{of\ visit} \cdot probability$$

The kind of attraction with the highest score wins and is proposed to the user. More specifically, the guide proposes the list of attractions of the winning kind, ordered by the distance from the user, and excluding those that the user has already visited (i.e., having a probability > 0.7 in the table in Fig. 11).

We are currently evaluating the performance of our guide application in terms of precision and recall. First results indicate that the whereabouts diary provides a rather effective support to the identification of places the user may be interested in.

5.OPEN ISSUES

We identified several research directions to improve our work.

First, by exploiting data mining techniques and commonsense knowledge, one could think at more effective ways to discriminate among several candidate places [8]. For example, if a person went to a restaurant at noon, it is very unlikely that he will go to another restaurant at 2p.m. As another example, a deeper knowledge of human activities could also suggest that a middle-aged woman is more likely to go to a market after work than a young girl. Due to these considerations, we envision an additional module able to learn users' habits and to automatically assign, with an increasing precision level, a probability value to all the candidate places.

A well-known and annoying drawback of this kind of systems is the training phase. To reduce the effort required by users to tune the system to their habits, we propose an additional distributed layer, which can be considered a sort of social knowledge repository. We envision, in fact, to collect and aggregate thousands of users profiles over a central location. Indeed, using data mining techniques, a synthetic profile of each kind of user could be generated. This "socially derived profile" can be used as a starting point for each user of our service. For example a middle-aged housewife will not start with an empty profile, but will be provided with a profile precompiled with those common patterns of her social group. An interesting aspect of this approach is its own intrinsic incremental mode of operation. In fact, in every moment, every user can upload his own updated profile to

the central location. In this way each user's experience will contribute to increase the accuracy of the "socially derived profile" of his own social group. Moreover, this approach is not really prone to privacy problems, because of a user is not asked to exchange the full log of its diary, but only a synthetic description of its general behavior [7].

The system described in this paper is mainly based on the verified assumption that most of the users usually spend the most of time in a small number of key places (home, work, favorite pub) [10]. Because of this, a future improvement of our diary is related to the ability of inferring, with a good level of accuracy, not only well known places but also less predictable locations. With this regard, it is also important to understand how the overall accuracy of our system will increase by recognizing places in which the users usually spend a little amount of their time. Is the overall complexity increase justified by an higher accuracy in recognizing only those rare events? It mostly depends on the metric that we chose. To properly evaluate this kind of systems, a cost-based metric should be used. Accordingly, the overall quality of the place-recognition engine is not only related to the ratio between right and wrong classifications, but also to the cost that a wrong decision could have in different situations. For example, if a predictive system is built upon our diary log, a slight variation in the places visited by the user in the afternoon can cause relevant errors, and then costs, in the prediction for the evening. In this context, we are also concerned about the actual feasibility of this task: are little fluctuations manageable? With which kinds of mechanisms? Can the aforementioned "social" approach be useful in this task? Answering these questions is worth of further researches.

Another aspect to be investigated regards the use of other kinds of sensing devices. For example, Wi-Fi and GSM could naturally complement GPS localization. Moreover, we would like to make our system able to use not only location oriented services, but also others services, belonging to different domains. For example, powerful sources of data can be credit card transaction records, Bluetooth meetings, phone working mode and the ratio of missed calls. Our goal is to increase the accuracy of the diary with this data.

The last, but maybe most complex and fascinating future perspective of our system is to develop an additional module able, not only to detect the current location, but also to predict future ones. This work, beside its scientific value, can lead to several relevant applications. For example, the model of each user, including his favorites places and habits, could be used as a portion of user's profile in a bunch of other applications. For example, a pervasive advertisement service showing personalized commercials to the users in front of a wide screen could show restaurant advertisements if the diary predicts that the users are going to have lunch in a while [4]. Another interesting application can be some sort of personal agenda which is able to inform the user that he is risking to fail some planned activities due to his current, and then predicted, behavior (e.g., if you stay awake longer, you risk of missing tomorrow morning meeting as it happened last month). Moreover we envision to build a system that is able to publish the predicted user position autonomously. Using such a system, friends and family, can for example be informed about the user behavior in an automatic way.

6. CONCLUSIONS

The Whereabouts diary can represent a useful middleware service to support location-aware activities for mobile users. While most of the proposed middleware infrastructures focus on "low-level" issues like communication and interaction among nodes, the Whereabouts diary makes available high-level knowledge about the context and the user activities, and could be fruitfully exploited in a variety of application scenarios.

Although the presented results are very promising, and the current version of the Whereabouts Diary is usable, we have also identified several directions along which to improve and enrich it with further features.

7. ACKNOWLEDGEMENTS

Work supported by the project CASCADAS (IST-027807) funded by the FET Program of the European Commission".

8. REFERENCES

- [1] D. Ashbrook, T. Starner, "Using GPS to Learn Significant Locations and Predict Movement Across Multiple Users", *Personal and Ubiquitous Computing*, 7(1):275-286, 2003.
- [2] J. Axup, S. Viller, I. MacColl, R. Cooper, "Lo-Fi Matchmaking: A Study of Social Pairing for Backpackers", *International Conference on Ubiquitous Computing*, Orange County (CA), USA, 2006.
- [3] G. Castelli, A. Rosi, M. Mamei, F. Zambonelli, A Simple Model and Infrastructure for Context-aware Browsing of the World, *International Conference on Pervasive Computing and Communication*, White Plains (NY), USA, 2007.
- [4] G. Castelli, M. Mamei, A. Rosi, "The Whereabouts Diary", *International Symposium on Location- and Context-Awareness*, LNCS 4718, Munich, Germany, 2007.
- [5] J. Hightower, "From Position to Place", *Workshop on Location-Aware Computing*, Seattle (WA), USA, 2003.
- [6] J. Hightower, S. Consolvo, A. LaMarca, I. Smith, and J. Hughes, "Learning and Recognizing the Places We Go", *International Conference on Ubiquitous Computing*, Tokyo, Japan, 2005.
- [7] V. Kostakos, E. O'Neill, A. Shahi, "Building Common Ground for Face to Face Interactions by Sharing Mobile Device Context", *International Workshop on Location and Context-Awareness*, Dublin, Ireland, 2006.
- [8] H. Liu, P. Singh, "ConceptNet: a practical commonsense reasoning toolkit", *BT Technology Journal* 22(4):211-226, Kluwer Academic Publishers, 2004.
- [9] W. Newman, M. Eldridge, M. Lamming, "PEPYS: Generating Autobiographies by Automatic Tracking", *European Conference on Computer Supported Cooperative Work*, Amsterdam, Netherlands, 1991.
- [10] D. Patterson, L. Liao, K. Gajos, M. Collier, N. Livic, K. Olson, S. Wang, D. Fox, H. Kautz, "Opportunity Knocks: a System to Provide Cognitive Assistance with Transportation Services". *International Conference on Ubiquitous Computing*, Nottingham, UK, 2004.