Locator: A Self-adaptive Framework for the Recognition of Relevant Places

Paul Baumann
Wireless Sensor Networks Lab
TU Darmstadt, Germany
paul.baumann@wsn.tu-darmstadt.de

Johannes Klaus
Bensheim, Germany
johannes.klaus@gmx.net

Silvia Santini
Wireless Sensor Networks Lab
TU Darmstadt, Germany
santinis@wsn.tu-darmstadt.de

UbiComp '14 Adjunct, September 13-17 2014, Seattle, WA, USA
ACM 978-1-4503-3047-3/14/09.
http://dx.doi.org/10.1145/2638728.2638751

Abstract
A high number of algorithms for the recognition of users' relevant places exist. However, none of them provide an optimal solution across all users and scenarios. We present a preliminary design of LOCATOR – a self-adaptive framework for recognizing users' relevant places. LOCATOR learns for different contextual situations, combinations of algorithms and location sensor data that achieve the best performance in recognizing relevant places. We conducted a 5-weeks study and collected sensor and ground-truth data from 6 users. Our preliminary results indicate the shortcomings of relying on one algorithm and sensor for recognizing places and thus motivates the rational behind our approach.

Author Keywords
Localization, relevant places, self-adaptiveness

ACM Classification Keywords
H.m [Information systems]: Miscellaneous.

Introduction
The wide spread of mobile personal devices opens new opportunities for a plethora of new applications that support users' daily-life. The capabilities of modern mobile devices like smartphones go far beyond the possibility to perform voice calls and text messages. For the most time of the day users' smartphone is in the close proximity of its owner

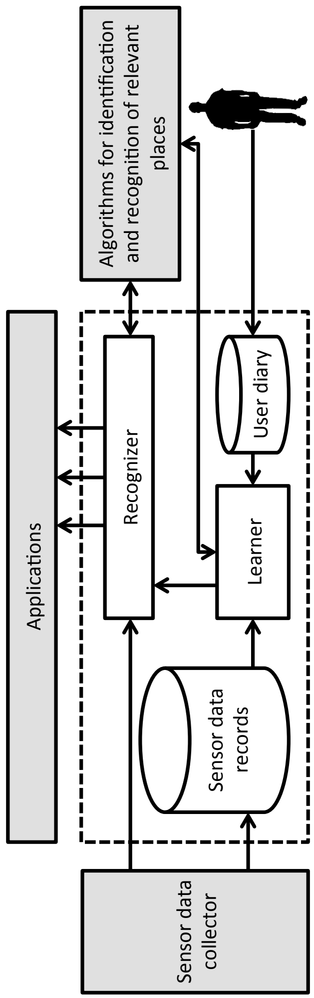

Figure 1: A preliminary architectural overview of LOCATOR.

[7] and thus allows inferring her context. An example for applications that leverage information about users' relevant places are location-based services. We refer to a relevant place as a place: "*where the user spends a substantial amount of time and/or visits frequently*" [1]. For instance, this information is then used for the prediction of users' next relevant place [5], or home automation [3].

In order to identify relevant places and to recognize when users' visit them again, many algorithms have been proposed [6, 1, 2]. The majority of these algorithms recognize places by leveraging one of the smartphones' location sensor, e.g., GPS, Wi-Fi, GSM, or microphone. The right choice of the algorithm and sensor has implications on the achievable performance. First, all the mentioned location sensors differ in their spatial coverage. For instance, in order to recognize a relevant place with a large geographical area like campus, location data with a high spatial coverage is required. One example of such a location data is the GSM cell identifier. Alternatively, an algorithm that relies, e.g., on the Wi-Fi data, will recognize the campus as a set of different places. This set needs to be identified in order to treat all these recognized places as one users' relevant place. At the same time, location data with a high spatial coverage does not allow recognizing places that are close to each other. In this case, data from a location sensor with a low spatial coverage, e.g., Wi-Fi or Bluetooth, is required. The second implication is the data availability, e.g., GPS does not work indoor. The third implication are the differences in the energy consumption across the sensors. Despite the fact that none of the existing solutions provide an optimal solution across all users and scenarios, there are only few attempts for overcoming the existing shortcomings by combining algorithms and sensors [8].

To overcome these shortcomings, we present a preliminary design of LOCATOR – a self-adaptive framework for recognizing users' relevant places. Our framework leverages location sensors, state-of-the-art algorithms for the recognition of relevant places, and users' diary records about the visited relevant places. LOCATOR learns which algorithm in the combination with which location sensors achieves the best performance for the given contextual situation, e.g., time of the day. Building upon the gained knowledge, LOCATOR achieves its self-adaptiveness by automatically switching to the combination that will achieve the best performance. We conduct a 5-weeks study with 6 participants to motivate the necessity of our approach by highlighting the drawbacks that LOCATOR aims to overcome.

We first describe LOCATOR's preliminary design, then present our first results to motivate the ongoing work, and finally outline our next steps.

A Self-adaptive Framework for the Recognition of Relevant Places

LOCATOR aims to recognize users' relevant places in a self-adaptive manner. The self-adaptiveness is achieved by learning the best *configuration* for each user individually and then switching to the best performing configuration for the given context. A *configuration* consists of a specific combination of algorithms, their parameters, and sensors.

To this end, LOCATOR leverages users' diary with places they have visited in the association with the sensor data records. This step helps to learn which configuration performs best in the particular situation. Our framework involves two phases: a *learning* and a *recognizing* phase. LOCATOR's architecture is depicted in Figure 1 and involves the LEARNER and RECOGNIZER component and SENSOR DATA RECORDS and USER DIARY databases.

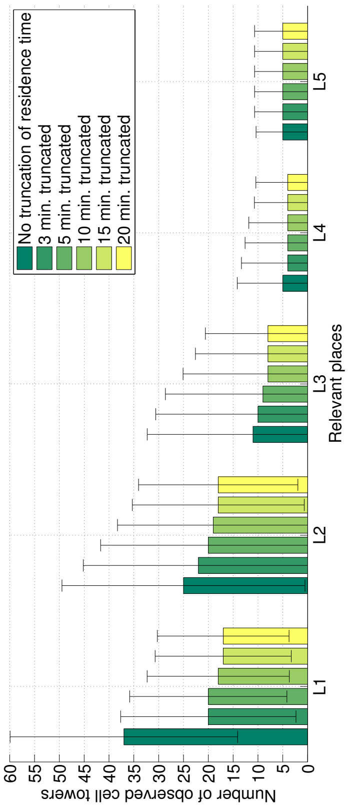

Figure 2: The median and the standard deviation of the number of observed GSM cell identifiers at the five most visited relevant places. In order to cope with users' inaccuracies in keeping record or their visits, x minutes after arrival and before departure events have been truncated.

The LEARNER component is responsible for learning best configurations for the particular context. To this end, LEARNER accesses the sensor data in the SENSOR DATA RECORDS database and the corresponding diary records in the USER DIARY database. The diary records are provided by users during the *learning* phase. Building upon this data, LEARNER configures the RECOGNIZER component.

After the *learning* phase is completed, the *recognizing* phase is triggered. It aims to recognize places based on the learned configurations. Building upon these configurations, RECOGNIZER recognizes users' places by leveraging the configuration that has been identified as the best performing solution for the given context. The output of the RECOGNIZER component is the user's current place.

The components in grey color outside of the dotted box in Figure 1 are not part of LOCATOR. However, these components provide a support function for LOCATOR.

Recognition of Relevant Places with GSM

After briefly sketching the frameworks' preliminary design, we now describe its first realization and the preliminary results that support the idea of our approach.

We have implemented our framework as an Android-based application. The SENSOR DATA COLLECTOR component builds upon the open sensing framework *funf* [4]. To capture users' diary, we provide a built-in feature that allows users marking their relevant places on a map and keeping record of their place visits. For the further analysis steps, we conducted a 5-weeks study with 6 participants that were provided with Android phones and the implementation of LOCATOR. The outcome is a rich data set that contains sensor records for GSM, Wi-Fi, Bluetooth, and GPS. However, the collected ground-truth data and its completeness is the most valuable part of our data set.

Number of GSM Cell Identifiers observed at Relevant Places

We first analyze the number of observable unique GSM cell identifiers at each relevant place. The goal is to investigate the feasibility of using a single GSM cell identifier to recognize a relevant place. We indicate with L_j, $j = 1 : N_L$, the j-th relevant place and define the sorted set $\mathcal{L} = \{L_1, L_2, \ldots, L_{N_L}\}$ as the set of N_L places relevant to a user i. The places are sorted according to the amount of time a user i spends there, i.e., L_1 is the place where a user i spends most of her time.

Figure 2 shows median and standard deviation of the number of observed unique GSM cell identifiers during the conducted 5-weeks study. We cope with users' inaccuracy in capturing ground-truth data, e.g., a user left her home at 14:03 but indicated 14:00 as the time of the departure event. To this end, we truncate the first and last x minutes of the data for each visit. We ran our analysis for $x = \{3, 5, 10, 15, 20\}$ minutes.

We observe a high number of unique observed GSM cell identifiers at the five most visited places across all users. The results indicate the inability to recognize places by only considering one single GSM cell identifier.

Defining Sets of GSM Cell Identifiers

Building upon our initial results, we analyze whether GSM cell identifiers can directly be mapped to the corresponding relevant place. Therefore, for each relevant place we compute the percentage of time each GSM cell identifier was observed over the total amount of time a user was visiting a particular place. We then define for each place three sets of cell identifiers. The first set contains GSM cell identifiers that have been observed at the particular place only. We refer to these cells as *distinct cells*. The second set contains GSM cell identifiers that have been observed at multiple places but with the highest relative time at this

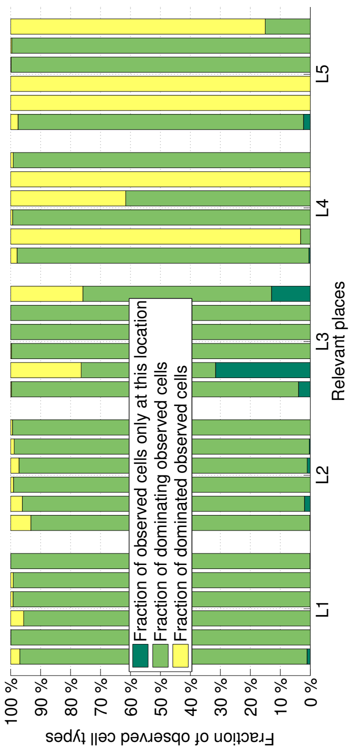

Figure 3: The fraction of time each set of GSM cell identifiers has been observed by each participant at the five most relevant places.

place. We refer to these GSM cell identifiers as *dominating cells*. Finally, the last set contains GSM cell identifiers that have been observed at multiple places, however, at this place with a lower relative time as at another place. We refer to these cell identifiers as *dominated cells*.

We now investigate the amount of time each set of GSM cell identifiers covers the five most relevant places. We observe in Figure 3 that the majority of the time is covered by *dominating cells*. It further indicates that most of the observed GSM cell identifiers are usually observable at at least two places, e.g., at home and at the bus station in front of the house door. With the decrease of places' relevance the fraction of *dominated cells* for these places increases. This observation highlights the aspect of the high spatial coverage leading to the inability to directly link the observed GSM cell identifier to the particular relevant place.

Conclusions and Future Work

We have presented our preliminary design of a self-adaptive framework for the recognition of users' relevant places. The self-adaptiveness is achieved by learning the combinations of algorithms and sensors that achieve best performance in recognizing places. The learned combinations are used as configurations to automatically switch between the algorithms and sensors. We conducted a 5-weeks study and collected data from 6 users. Our preliminary results indicate the shortcomings of relying only on one algorithm and sensor and thus motivate the rational behind our approach. The results support the necessity of switching between algorithms and sensors in order to recognize relevant places.

Building upon these results, we plan the following steps. First, we will extend our framework with further algorithms that leverage location sensors, e.g., Wi-Fi, Bluetooth, or GPS. We then will extend the LEARNER component with the ability to learn strategies for switching between different algorithms and location sensors with the goal to optimize the performance. Finally, we will evaluate our framework by leveraging the collected data as well as a number of publicly available data sets.

Acknowledgements

This work has been partially supported by the Collaborative Research Center 1053 funded by the German Research Foundation and by the Priority Program Cocoon funded by the LOEWE research initiative of the state of Hesse, Germany.

References

[1] D. Kim *et al.* Discovering Semantically Meaningful Places from Pervasive RF-Beacons. In *UbiComp* (2009).

[2] J. Hightower *et al.* Learning and Recognizing the Places We Go. In *UbiComp* (2005).

[3] J. Scott *et al.* PreHeat: Controlling Home Heating Using Occupancy Prediction. In *UbiComp* (2011).

[4] N. Aharony *et al.* Social fMRI: Investigating and Shaping Social Mechanisms in the Real World. *Pervasive and Mobile Computing* (2011).

[5] P. Baumann *et al.* The Influence of Temporal and Spatial Features on the Performance of Next-place Prediction Algorithms. In *UbiComp* (2013).

[6] R. Montoliu *et al.* Discovering Places of Interest in Everyday Life from Smartphone Data. *Multimedia Tools and Applications 62* (2013), 179–207.

[7] S. Patel *et al.* Farther Than You May Think: An Empirical Investigation of the Proximity of Users to their Mobile Phones. In *UbiComp* (2006).

[8] Y. Chon *et al.* Mobility prediction-based smartphone energy optimization for everyday location monitoring. In *SenSys* (2011).