Employing User Feedback for Semantic Location Services

Donnie H. Kim, Kyungsik Han, Deborah Estrin
UCLA CSD CENS
{dhjkim, kshan, destrin}@cs.ucla.edu

ABSTRACT
Just as coordinate-oriented location-based applications have exploded recently with mapping services, new semantic location services will be critical for the next wave of killer applications. People are going to want everyday applications to have location-awareness that goes beyond simple numerical latitude and longitude. Loci is a new semantic location service layer that employs user feedback to bridge the gap between machine-learned and human-defined places. Advances in place learning techniques have provided us the tools to detect nearly 95% of the visits we make to places and the distances we travel. The difficulty of recovering the remaining 5% comes from designing parameters that work for every user in every place. Based on a user study with 29 participants over three weeks, we show that the level of user feedback required by the service is acceptable and most of the users are willing to provide help to improve their experiences with the service. Our results suggest that user feedback has the potential to significantly improve semantic location services, but requires well-timed prompting mechanisms to improve the quality of the feedback.

Author Keywords
Semantic Location, User Feedback

ACM Classification Keywords
H.5.m Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Design, Human Factors, Management

INTRODUCTION
Mobile phones are increasingly expected to do more than make phone calls. A growing number of mobile devices can now determine their own physical location, and various location-based mobile applications are on the rise. In addition to applications that require instantaneous locations (*e.g.*, navigation, emergency response, and local search), many applications can benefit from continuously understanding the user's location context. Ranging from personal applications such as geo-reminders and location-diaries [27], to social applications such as whereabout-sharing and ride-sharing [3, 29], numerous emerging applications want colloquial places and paths rather than just coordinates. As people normally think and speak locations in terms of places like "my home" or "Ed's office", location-aware applications will encounter the problem of translating coordinates provided by today's positioning systems to semantic places. Track-based applications for map building, traffic estimation, and trace sharing, can also benefit from automatically parsed paths that are sequences of coordinates between places.

Unfortunately, current location-aware applications must handle the coordinates directly. Absence of a universal place database forces developers to ask users to define places one-by-one by drawing circles (or polygons) on a map for each application. Paths are also parsed by post-processing algorithms from day-long location traces or by depending on users to manually start and stop location tracking. However, such schemes fail in describing many interesting indoor places and poorly scale. Creating indoor localization systems in every building is currently far-fetched, and relying on users to delineate places from scratch can omit many interesting places. Moreover, continuously estimating absolute positions significantly impacts battery-life on the mobile device, preventing these applications from penetrating into a wider population. Current applications either cope with a fixed sampling interval with a lower resolution or depend on users to manually check-in to places and track paths.

Place learning is one way to move from coordinates to semantic locations. Several algorithms have been developed for extracting significant locations by post-processing complete location traces offline [1, 18]. More recently, algorithms that learn and recognize semantic places from pervasive Wi-Fi access points enable discovering many indoor places in real-time [12, 14]. Furthermore, researchers have shown that semantically interpreting location context also helps designing an energy-efficient location tracking mechanism by leveraging people's tendency of staying at indoor places for the majority of the time [15]. The state-of-the-art technique can accurately detect over 95% of the place visits and the travel distance while using significantly less energy than collecting coordinates at a fixed rate. These techniques

This work was supported in part by Intel PhD Fellowship.

UbiComp'11, September 17–21, 2011, Beijing, China.

have been shown to be effective in capturing people's concept of a place most of the time. However, machine-learned places are not always same as human-defined places.

This paper presents Loci, a service layer that abstracts a user's location context as places and paths. The service learns new places, suggests potentially meaningful places, recognizes registered places, and tracks paths connecting places using previously proposed techniques [15]. To overcome the problem of when machine-learned places are different from user-defined places, Loci employs user feedback. User feedback comes in two forms: 1) managing suggested places, and 2) telling when the service detects places incorrectly. We sought to understand how often user feedback is needed, what helps a user manage suggested places, and how we can encourage a user to provide feedback to the service.

To answer these questions, we conducted a three-week user study with 29 participants. We found that the amount of user feedback required by Loci is acceptable and many of our participants were willing to provide feedback to the service. Comments made by study participants suggest that the places found by Loci were easily understood and were close to how they perceive places most of the time. Our results also suggest that well-timed notifications and system feedback motivate users to provide more input. Overall, our study points towards the feasibility of engaging users for improving semantic location services.

RELATED WORK

Place Learning

Place learning algorithms extract semantic places from raw sensor signals and can be divided into two classes by input data: geometry and fingerprint.

Geometry-based algorithms find significant locations by clustering coordinates from location traces. For example, Ashbrook *et al.* and Toyama *et al.* used a variant of the k-means clustering algorithm and found clusters by GPS signal losses [1, 28]. Liao *et al.* introduced a supervised learning approach using a hierarchical CRF model to recognize high level activities and places from GPS traces. Parameters were fit by training on a labeled data set [18]. Zhou et al. proposed a density-based clustering algorithm that can discover clusters of arbitrary shape [30]. Harihanran *et al.* and Kang *et al.* used a temporal and spatial stay threshold to find places from location traces [11, 13]. Unlike other approaches that required an entire trace, their heuristics based on time and distance are computationally inexpensive and can be used in real-time. These approaches extract building-level or outdoor places, as GPS is usually not available inside buildings. Marmasse *et al.* defined a place as an area where GPS is unavailable [22].

Fingerprint-based algorithms use ambient signals to identify a specific place, and work better than coordinate-based approaches in discovering indoor places. RF-based fingerprints (*e.g.*, cell tower, Wi-Fi, Bluetooth) are the most common while other ambient signals such as sound, color, light texture information are also used [21, 2]. RF-based algorithms infer a single visit to a place by detecting stable radio environments. Laasonen *et al.* found cliques by clustering cell towers while Froehich *et al.* identified places by prompting users when the currently connect cell tower changes [17, 9]. Hightower *et al.* used newly detected Wi-Fi access points to determine a place visit [12]. Kim *et al.* further improved the place detection accuracy by using selected strong access points to detect a place visit [14]. These algorithms can be applied to any RF beacon signals but relying on Wi-Fi provided better resolution than cell towers.

Place Labeling

Several research efforts have been made to understand how people assign place names. Zhou *et al.* analyzed how people describe places and suggested five categories of descriptions [31]. Barkhuus *et al.* proposed four different types of location labels namely geographic references, personally meaningful labels, activity-related labels, and hybrid [3]. Lin *et al.* refined the categories by imposing a hierarchy with more fine-grained subcategories [19]. Patterson *et al.* studied the use of place labels as meaningful status messages for instant messaging (IM) applications [26]. Monibi *et al.* used collaborative filtering techniques for recommending place labels for IM status based on history data from multiple users [23].

Continuous Localization

Tradeoff between energy consumption and location accuracy has been extensively studied for continuous localization. Several approaches assumed that a set of positioning systems with varying accuracy and energy consumption are available while others used low power sensors to find GPS sleep opportunities. Constandache *et al.* intermixed positioning systems by solving an optimization problem using dynamic programming [8]. Lin *et al.* assumed that the accuracy requirement changes as a user moves around and used an HMM model to select positioning systems [20]. Kjaergaard *et al.* assumed that applications define their distance error limits, used accelerometer to turn off GPS, and predicted next GPS sampling time based on speed estimation [16]. Paek *et al.* and Zhuang *et al.* also similarly used accelerometer data to adaptively sample GPS coordinates [24, 32]. These approaches tried to continuously estimate absolute locations within a distance error-bound while reducing the energy cost. Kim *et al.* took a different approach by interpreting location context in a more semantically meaningful way which helped to reduce the energy cost for continuously sensing locations [15].

User Feedback to create Indoor Location Systems

Over the years, several localization techniques relying on existing Wi-Fi infrastructure have been proposed. Due to the complex nature of indoor signal propagation, signal mapping has been a more practical and accurate solution than signal modeling. However, constructing a signal map (often called as *training phase*) requires a large amount of human labor for measuring signal strengths in every corners of a building. To bypass this, a couple of systems depending on user input to create the database have been proposed.

ActiveCampus uses a hill-climbing algorithm to approximate

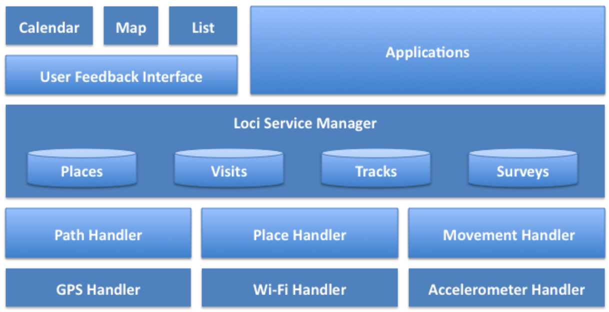

Figure 1. Loci semantic location service architecture

a user's location and employs a correction mechanism to improve its accuracy [10, 5]. When a user's location is incorrectly estimated, the user corrects the location by clicking on the right spot on a map. RedPin also avoids a costly offline training phase and generates location fingerprints from user input [6]. In contrast to ActiveCampus's geometric approach, RedPin generates symbolic location identifiers with room-level accuracy. The authors later proposed detecting intervals of immobility (using an accelerometer) to collect fingerprints while stationary and allow users to label data in a more appropriate time referring to the starting time and duration of the interval [7]. This interval method reduced the influence of second-by-second signal fluctuations that may be severe when fingerprints are instantaneously collected when a user dedicates to label a fingerprint on-site. Barry *et al.* conducted a year-long study of a similar crowd-sourcing localization system and showed its utility [4]. OIL additionally conveys graphical feedback about spacial uncertainties derived from discrete Voronoi diagrams and weeds out erroneous user contributions [25]. It prompts the user for feedback when the spatial uncertainty is too high (large Vornonoi region), and the user can postpone or turn off prompting.

However, these systems instantaneously localize a user in a particular building as opposed to continuously understanding a user's everyday location context.

THE LOCI SEMANTIC LOCATION SERVICE

Loci runs as a background service on a mobile device and abstracts a user's location context as places and paths rather than just coordinates. Visits to a place are recognized by detecting enter and exit events at a particular place. Paths are tracked by collecting a series of position fixes (a bundle of time, latitude, longitude, and optionally altitude, speed, accuracy, heading) that a user creates while moving from one place to another. Figure 1 shows the architecture of our service, which is comprised of a user feedback interface, a service manager, and three components handling place visits, path travels, and user movements. We first give an overview of how the service is used by a user and then explain the details of each component.

Places are managed by the user through a single interface directly provided by Loci and shared with various applications. When the user begins using Loci, no places are registered with the service. However, as the user carries around the mobile device running Loci, new places are gathered as potential everyday places waiting to be reviewed and confirmed. Places are suggested when the user visits them for the first time and spends a substantial amount of time (in our case, defined as more than 5 minutes). Loci automatically infers entrance to and exit from a place using Wi-Fi fingerprints collected during the stay. These Wi-Fi fingerprints are stored and used to identify places in the future. Recent visit times (as enter and exit times), approximate position of the place, and a list of neighboring Wi-Fi access point (AP) names are provided as hints to assist the user in deciding if a place is valid and should be saved or is spurious and should be discard.

At a convenient time, the user can decide which suggested places to register, add to existing place, delete, or block. As Loci uses Wi-Fi fingerprints to recognize a place, registering a place involves assigning a set of Wi-Fi fingerprints to a place. A Wi-Fi fingerprint is defined as a list of APs visible by the device and their received signal strength (RSS). Places are suggested when the Wi-Fi environment stabilizes for a period of time indicating a stay but does not match with any previously registered places. Adding a suggested place to an existing place allows user to register several Wi-Fi fingerprints to a single place. Some places may have more than one fingerprint. A user can define a place as a large space (e.g., markets, warehouses, etc.) where a single Wi-Fi fingerprint cannot cover the entire place, or the Wi-Fi environment of a place can change over time due to indoor signal interference and network reconfigurations. Loci also estimates the approximate position of the place using position fixes acquired around the visit time while tracking paths[1]. If a place is blocked, it will never be suggested as a place again, while a deleted place can be. Once a place is registered, Loci recognizes it using Wi-Fi scans. The user can also manually register a place when a visit to a new place was not successfully suggested (*e.g.*, brief visits or incorrect recognitions).

In the following sections, we describe Loci's user interfaces, embedded survey tools we designed for our user study, and its operational architecture.

User Interface

Loci offers a set of user interfaces for managing places, viewing location histories, and providing feedback. The user interface is organized primarily around a tab bar at the top of the screen that provides instance access to three main views: a *calendar view*, a *map view*, and a *list view*.

Calendar View

Calendar view allows a user to view place visit histories in a calendar which consists of a *monthly view* (Figure 2(a)) and a *daily view* (Figure 3(g)). The user can click on a date presented on the *monthly view* (which is shown when the *calendar view* is selected) to open a *daily view*. This takes the

[1]Path tracking is initiated when a user leaves a place and is suspended as the next destination is reached.

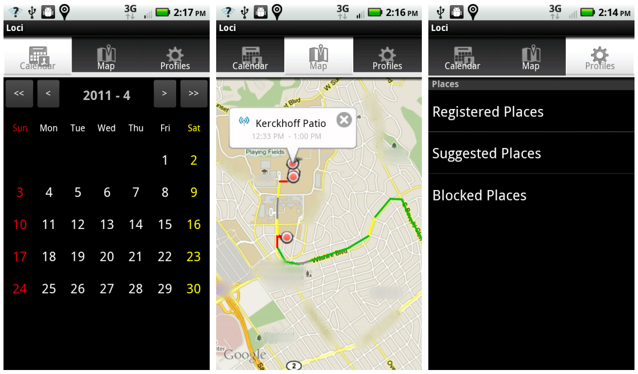

(a) Calendar (b) Map (c) List

Figure 2. Loci's user interface consists of three main tabs. *Calendar view* shows visit histories in time-order. *Map view* shows daily place visits and travel paths. *List view* categorizes places as registered, suggested, and blocked.

user to a screen which presents all visit activities of that particular day on a list. A list entry displays information about a single visit including enter time, exit time, and place name. When a new place is visited and suggested by Loci, the place name is presented as "New Place #X". The user can click on a list entry (representing a visit instance) to view more information about that particular visit (Figure 3(b)) such as stay duration and the Wi-Fi fingerprint collected during the visit (Figure 3(c)). This screen also provides an interface to register a new place (if recognized as a visit to a new place), edit a place's profile (if recognized as a visit to a registered place), or correct a misclassification (if recognized as a visit to a wrong place) through a single button labeled as "Edit". Correcting misclassification allows users to either register a new place using the Wi-Fi fingerprint collected during the visit, or add it to an existing place. Short visits to an unrecognized place is hidden by default in a daily view as they are mostly transient locations. However, the user can reveal these hidden visits and select a visit instance to manually register a new place using the saved fingerprint (Figure 3(g)).

Map View

Map view lets a user view location trajectories of a selected day on a map (both in satellite and map mode) (Figure 2(b)). The user can either select to display every single GPS point or a subset of filtered points (to reduce loading overhead) of a particular day. When the filtering option is selected, GPS points are filtered based on a selected time interval t, and displays one position per every t. Color of a line connecting each pair of points is defined by speed provided by the GPS module. Green, yellow, red, and gray are used when the estimated speed is above 10 m/s, between 5-10 m/s, between 1.5-5 m/s, and below 1.5 m/s respectively. A place is displayed as a dot on the map and uses the coordinate associated with the place when the place was registered. The user can zoom and pan the map and navigate the locations. Clicking on a place dot populates a popup bubble showing the place name and every visit instance during the day (showing enter and exit times). Clicking on the bubble takes the user to a *place info* screen (Figure 3(a)). If the selected place dot is a new place (suggested and waiting to be registered), the user can handle the suggested place instead.

List View

List view allows a user to view places by its status: *registered*, *suggested*, and *blocked* (Figure 2(c)). The three place lists are shown by choosing the *Profiles tab* (figure 2(c)). The *suggested place list* displays places suggested by Loci but not reviewed by the user yet (Figure 3(e)). Suggested places are given a temporary name such as "New Place #X". Selecting a new place will take the user to a screen where visit times, a position, and neighboring Wi-Fi APs are presented as hints (*calendar* and *map view* also arrive to the same place view). Tapping the map allows the user to enlarge the map for a better view (Figure 3(d)). The user can register, add to existing place, delete, or block the suggested place. Registering a place requires assigning a unique name and optionally correcting the estimated position of the place to the exact location. A place is shown as a circle, and a map tool is provided so the user can move the center or adjust the radius. The circle is currently not used to detect places, but is presented as an additional information about the place[2]. When the *registered place list* is viewed, every registered place are listed in an alphabetical order and the time elapsed since the last visit is shown (Figure 3(f)). Selecting a place from the list takes the user to a screen where the user can view and edit the place (Figure 3(a)). The *blocked place list* displays all blocked places. The user can unblock places by selecting a blocked place on this list.

Embedded Survey Tools

We embedded two survey tools in Loci to understand its performance in the field and the amount of feedback about place failures a user may provide to the service in everyday use. We directly asked users to tell us when a place found by Loci differs from their perception of a place using these tools.

Place Surveys

Place surveys are used to characterize places based on user feedback. Users can take a survey about a particular place by choosing a place from the *registered place list*. On a *place survey*, a user can select a name category of the chosen place name [19], add keywords, report place-detection failures that happened during visits, and provide free-text comments. *Place survey* is initiated by the user. We have primarily used this tool to understand place characteristics that lead to a specific place-detection failure and discuss potential solutions. Users can edit the place survey anytime they want. The embedded nature of the survey tool allows the users to provide *in-situ* feedback. We categorized place failures as 1) *missed* visits, 2) one visit *divided* into multiple visits, 3) two or more visits to different places *merged* into a single visit to a single place, and 4) *wrong* place recognition. When a machine-logged visit duration is significantly shorter than human-remembered duration, we asked the user to categorize them as a *divided* visit. Users can report on these four different failure types by answering how often each failure

[2]Circles can be used for detecting large outdoor places such as parks, beaches, or large superstores in low-rise buildings.

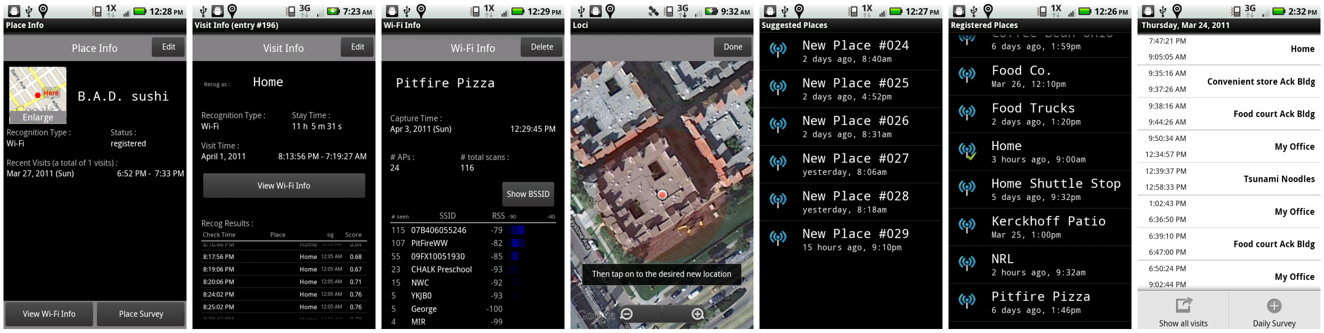

(a) (b) (c) (d) (e) (f) (g)

Figure 3. User interfaces for viewing and managing places. Screen shots of place info screen, visit info screen, Wi-Fi info screen, map edit screen, suggested place list, registered place list, and daily view are shown from left to right.

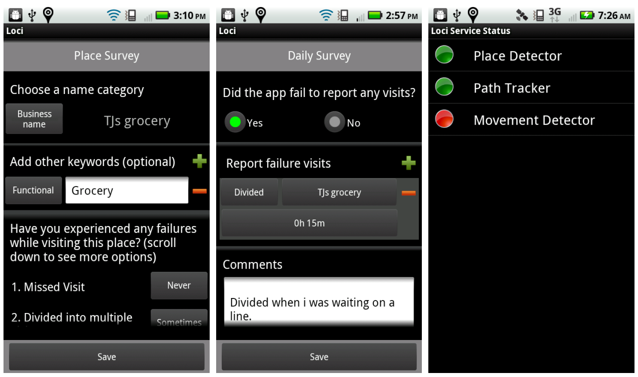

(a) (b) (c)

Figure 4. Embedded survey tools and status screen. Daily surveys are used to report place detection failures by day. Place surveys are used to report failures by place. Status screen illustrates activated components.

occurred. Frequency options include *never*, *sometimes* (less than a half of the visits), *often* (equal or more than a half of the visits), and *always*. Users can also additionally provide contextual information about the failure in the comment box.

Daily Surveys

To understand how often a user may need to report about failures, *daily surveys* are used to get feedback from the user. If any discrepancies between what Loci detects and what the user remembers are noticed during a day, they are reported in the *daily surveys*. On a *daily survey*, the user can report a failure by selecting a place from a list of registered places, a failure type, and a visit duration. Any number of failures can be added to the survey. If a place suggestion fails, the user can select "new place" from the place list and provide information about the detection failure during the first visit. A comment box is used to collect any additional information the user may wish to share about the failures.

Service Architecture

The service consists of three main building blocks to keep track of locations in a semantically meaningful way while reducing its energy expenditure: place handler, path handler, and movement handler (Figure 1). Loci service manager on top of these blocks controls the components, manages location history data, and provides a gateway to the user interface modules and applications. Our service architecture is built upon previously proposed ideas [15]. We briefly describe how each component works and interacts.

Place Handler

Place handler periodically scans neighboring Wi-Fi APs to detect place visits (for every 10 seconds in our case). Once the neighboring radio environment stabilizes (receiving c=3 consecutive similar scans), an entrance to a place is declared, and the service manager consults with the local database to recognize the place. If the Wi-Fi fingerprint similarity of the visited place against every registered places is lower than a threshold, the place is saved as a potentially new personal place in *place database*. Similarity of two Wi-Fi fingerprints is measured by computing a *Tanimoto coefficient* of the two signal strength vectors. *Place handler* continuously checks if the latest scan and the accumulated scans during a visit are similar. To suppress noisy signals misleading to prematurely declaring a departure, a subset of steady beacons (based on the response rate[3]) are used to measure similarity. If the scans are similar for longer than a predefined amount of time (5 minutes), place handler triggers the movement handler to find opportunities to sleep. If no movement is detected, place detector suspends until a movement is detected again. A departure from a place is detected when the similarity score is below a threshold for more than c scans. The visit information (a place name, enter and exit times, and a Wi-Fi fingerprint) is saved and *path handler* is initiated as the user leaves the place.

Path Handler

Path handler starts tracking positions using the GPS module on the mobile device as the user leaves a place. The sampling interval for acquiring position fixes is a tunable parameter defined by the application needs (we set as 15 seconds). Every position fixes acquired while the path handler is on are coupled with a timestamp and stored in *track database*.

[3]Response rate is the ratio of the detection count and the total number of scans.

Path handler signals *place handler* to suspend itself through the service manager when it detects that the user is traveling in high speed and is unlikely to arrive at a place any time soon. This avoids unnecessary Wi-Fi scans. When the average speed over a non-overlapping time window (30 seconds) is higher than 2 m/s[4], *place handler* is notified to sleep until the average speed drops below the same threshold. The speed information of each position fix is supplied by the GPS module.

Movement Handler

Movement handler is turned on to find opportunities to save energy while the device is staying at a place and immobile. Hinted by *place handler* that the user is staying at a place and the surrounding radio environment is stable, it powers the accelerometer on board and starts to monitor the motion of the device. The accelerometer is duty-cycled 50% (using a duration and a period set to 3 and 6 seconds) and is set at a low sampling rate (3-7Hz) to further save energy used for detecting movements. The variance of acceleration magnitude (to tolerate random orientation) is computed during the 3 seconds on-time window to determine a motion state. *Place handler* is put into sleep when the variance is consistently below a threshold for more than 30 seconds, and awaken when its above the threshold for more than 12 seconds. A shorter wait time is used for resuming *place handler* as the energy used for a couple of more Wi-Fi scans is negligible while a more conservative waiting time can lead to missing an exit event.

Data Storage

Loci manages four databases to maintain places, visit histories, tracks, and survey answers. *Place database* stores suggested places, registered places, and blocked places. For each place, a place name, a geographical area (as a circle on a map), a list of Wi-Fi fingerprints, and additional information associated with the place is saved. A new place is added to the database when *place handler* learns a new place (marked as suggested until the user confirms) or when a user manually registers a place using a particular Wi-Fi fingerprint. A place can be manually created by selecting a Wi-Fi fingerprint collected during a particular visit[5] and saved in *visit database*. When a user reviews a potential place to register (both from suggestions and manual entries), positions collected within the visit time (saved in *track database*) are used to estimate the geographical location of the place. Loci uses the most accurate position fix among the coordinates acquired during a stay. If no position with a timestamp that is within the stay time is available, it uses a position with the closest timestamp. *Visit database* maintains visit histories including enter and exit times, place name, and a Wi-Fi fingerprint collected during the stay. *Track database* stores positions collected while *path handler* was active. A position is a tuple of timestamp, latitude, longitude, altitude, speed, accuracy, and bearing (which are provided by the GPS module). Finally, survey answers provided by the user are saved in *survey database* until they are uploaded.

[4] Average human walking speed is 1.3-1.5 m/s.

[5] A short visit disqualified to be suggested as a place or a visit incorrectly labeled to another place can be used.

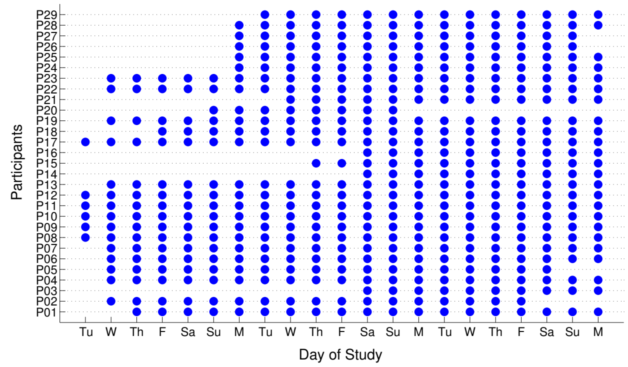

Figure 5. User study participation dates

Upload Handler

Data upload is triggered when the phone is plugged-in to a power outlet to conserve energy. Currently, upload handler is specifically designed for remotely collecting user study data, but can be extended to back-up the above described databases.

USER STUDY RESULTS

We conducted a user study with 29 participants living in four major cities and three continents to understand how people use Loci and adopt an interactive semantic location service.

Three research questions guided our study design:

1. How often does a user need to provide feedback?
2. What helps a user manage suggested places?
3. How can we motivate a user to provide feedback?

To address our research questions, we asked participants to use Loci for at least two weeks which preferably ran on their primary mobile phone, but we provided a secondary smart phone when necessary. The entire user study spanned for three weeks while each participant had different start dates and periods. However, a couple of participants experienced device failures or made unintentional errors during the study (*e.g.,* uninstalling before an upgrade), thus, couldn't provide a full two weeks of data (Figure 5). We present these partial experiment results as well as the results from the other participants who collected more than two weeks worth of data.

Procedure

Study sessions began by gathering a study consent form and providing a short tutorial about Loci. Users were required to enable GPS and Wi-Fi on their phone, then to download and install Loci (provided as an Android app) by visiting a URL. Once the participants successfully installed and used Loci for a couple of days, we requested that they complete an entrance survey posted online. The entrance survey probed the participants' commute distance, average number of visited places during weekdays and weekends, and their interests in location-based applications. Participants were expected to carry the phone most of the time, clear all suggested places everyday, make sure to recharge the phone every night, and use the embedded survey tools to report all place detection failures. We explained the three different user interfaces (calendar, map, and list) that can be used to access suggested places, and encouraged them to use the most convenient method. We also informed participants how to manually register places that were not successfully suggested (including brief visits or incorrectly recognized visits). The participants were advised to continue the study for at least two weeks. When the study concluded, we asked the participants to complete an exit survey. The exit survey asked about their experience and desired new features.

Participants

We collected data from 29 participants (21 males and 8 females) aged between 25 and 55 (median 30) who volunteered to participate. Among the participants, 11 were grad students, 6 were undergraduates, and the other 12 had different professions. Most of the participants used their own mobile phone and five participants were provided a secondary smart phone (Samsung 18, HTC 6, Motorola 5). Based on the entrance survey, commute distance for 79% of our participants were within 10 miles. For both weekdays and weekends, most of the participants selected between 3 and 6 as their number of unique places they visit per day (86% for weekdays and 71% for weekends). We also asked about their interests in location-based services. Location-based reminders were the most popular application (71%), followed by location diary apps (67%), place suggestion apps (57%), contributing data for urban planning (38%), and social apps (33%). Multiple choices were allowed.

Results

To measure how often a user needs to provide feedback to the service, we first investigate the suggested places a user receives everyday. We also examine how often users may report about place detection failures and help the service meet their needs. We then take a look at what factors help users to recognize a suggested place and provide correct user input. Finally, we discuss about possible methods for encouraging users to provide feedback to the service, and briefly analyze the resource usage of Loci.

User Feedback Demands

Loci requests user input for 1) handling suggested places (deciding to either register, add to existing place, delete, or block), 2) assigning a name when registering a place, and 3) manually registering a place when necessary. Since the service automatically recognizes registered places, user input is required primarily when a new place is visited. Another occasion is when Loci has trouble detecting a place as a user expects. The answers, gathered through the embedded survey tools, help to highlight when a gap exists between machine-learned and human-defined places. As we analyze the survey answers, many failures can be recovered by tuning the parameters of a place detection algorithm at challenging places. The difficulty comes from designing the parameters so that detections work everywhere. Having user feedback allows the parameters to be sub-optimally tuned so that the detection works reasonably well without overly burdening the user. We examine these two types of user inputs to measure how much feedback Loci demands from its user.

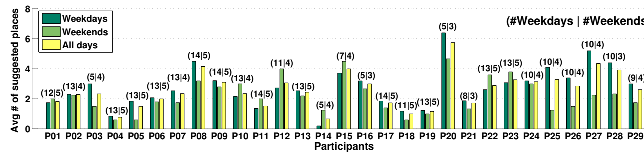

Figure 6. Average number of suggested places a user received by weekdays, weekends, and all days. About one third of the participants got more new place suggestions during weekends than weekdays.

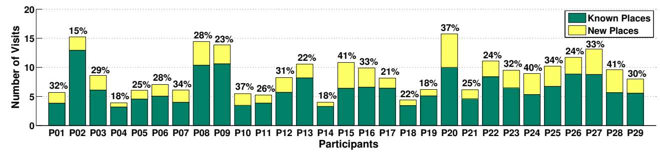

Figure 7. Average number of visits per day by user. The percentage on top of each bar indicates the percentage visits that require user input (visits to new places).

A) Registering new places.

Figure 6 illustrates the average number of suggested places a participant received for a day by weekdays, weekends, and all days. The number on top of each bar depicts the number of days used to calculate the average (weekdays and weekends). Participation dates varied from 8 days to 21 days. Overall, participants received 2.51 new places each day on average (weekends 2.18 and weekdays 2.64), and the number varied by users between 0.26 and 6.40. The maximum number of suggested places a participant received within a day was 11 and the minimum was 0. Among 29 participants, 9 users got more suggestions during weekends compared to weekdays.

To examine what influenced the number of suggestions a participant received, we first compare the average number of place visits per day by user (Figure 7). A visit is a single instance of a user staying at one place, and a place can be visited several times during a day. The percentage provided on top of each bar indicates the ratio of visits that required user input as the visited place was new to the service. P15 and P20 show a high ratio as the participant visited many new places over a weekend, but only provided week-long data sets. P23 exhibits a relatively high ratio due to a travel to another city during the first week of the study. In general, people with less place visits had to review less newly suggested places, but having more place visits did not always lead to a higher ratio (*e.g.*, P02). Some people occasionally traveled between a smaller set of places, increasing the number of visits but not new places, while others visited slightly more diverse places and less repeated places.

Next, we selected three participants and plotted how the number of suggested places changed over time. P08 registered the largest number of places during the study (over 85 unique places for three weeks), and visited 7-12 unique places each day. In the first few days, the percentage of new places were

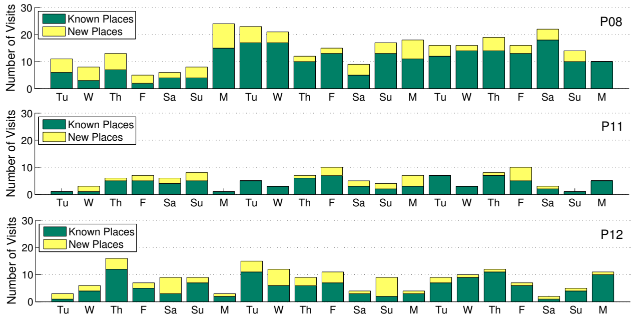

Figure 8. The number of place suggestions by participation date. P08 registered the largest number of places. P11 had place visits fewer than average. P12 visited average number of places. All participants generally experienced decreasing number of suggested places over time.

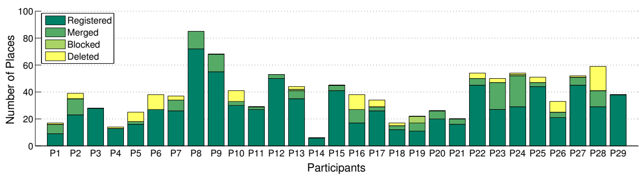

Figure 9. Quality of the suggested places. Overall, 75% of the suggested places were registered, 16% were merged, and 8% were deleted.

comparatively higher than the following days. However, regardless of the participant's high level of activity, many visits were recognized and new places were gradually learned over time. P11 made less place visits than the average and didn't have to deal with new places for 6 days out of 21 days. New places were mostly found over the weekends. P12 had average place visits and the number of suggested places went down as the participant entered the third week. Most participants experienced decreasing number of suggested places in general.

Finally, we analyzed how the participants decided to process each suggestion to understand the quality of our suggestions. Figure 9 breaks down the suggested places by its latest status: *registered*, *merged*, *deleted*, and *blocked*. Overall, 75% of the suggested places were registered, 16% were merged, and 8% were deleted. A *registered* suggestion is confirmed as a valid new place by the participant while a *merged* one is added to an existing place. P23 and P24 had noticeably more merged places as P23 merged many suggested places into "Disneyland" and two airport names while P24 defined places as buildings and merged many rooms. When a place was *deleted*, we asked why and saved the answers. Participants deleted suggestions mostly when Loci picked up a congested road as a place or when they couldn't recognize the suggested place. Some people blocked an open space where they often walked.

B) Reporting place failures.

When a place visit detected by Loci is incorrect, we asked participants to report such failures on both the *daily survey* and the *place survey*. We used *daily surveys* to gauge how often a user experienced place detection failures. We collected

Freq.	Missed		Divided		Merged		Wrong	
Never	845	(97.3%)	793	(91.4%)	844	(97.2%)	850	(97.9%)
Sometimes	20	(2.3%)	54	(6.2%)	20	(2.3%)	16	(1.8%)
Often	3	(0.4%)	17	(1.9%)	3	(0.4%)	1	(0.1%)
Always	0	(0%)	4	(0.5%)	1	(0.1%)	1	(0.1%)
Total	868	(100%)	868	(100%)	868	(100%)	868	(100%)

Table 1. Distribution of the place detection failures reported from the *place surveys*. Overall, *divided* failures were most common, followed by *merged*, *missed*, and *wrong* failures.

a total of 112 *daily surveys* from 29 participants during the study. However, we found that some failures were partially reported to one of the survey tools and not both. We manually fixed the discrepancies between the data sets from the two survey tools. As a result, the total increased to 176 *daily surveys*. First, we calculated the average number of *daily surveys* per day by individual user. About 51% of the participants submitted the *daily survey* every other day or less often, 27% reported once a day more, and 21% submitted none. However, this included repeating failures occurring while revisiting the same place. Thus, assuming that a failure can be fixed once it is reported, we excluded repeating reports and calculated the average again. In this case, every participant would report less than one failure per day. We also projected that the number will even get lower as the participant uses the service for a longer period. However, we also admit that the number we provide here may underestimate the failures as the participants may not have reported some failures.

With the *place survey*, the participants reported on how often a particular failure occurred at a place and, optionally, provided the context of the failure using a text box. We further analyzed place detection failures using these reports. For the 868 places registered by 29 participants during the study, *divided* failures happened most often, followed by *merged*, *missed*, and *wrong* failures (Table 1). Place visits were *divided* at relatively large spaces where a single Wi-Fi fingerprint couldn't cover the entire space (*e.g.*, department store, music hall). Next common places were in homes and offices where the participants stayed for extended time. Instantaneous failures occurred when strong APs briefly disappeared or radio signals fluctuated enough to trigger premature exits. *Divided* visits can be reduced by using a less sensitive similarity threshold or using several Wi-Fi fingerprints as a single place. Places were *merged* especially when a user directly moved from one place to another separated by a single floor or a hallway, These were also often incorrectly recognized places, resulting *wrong* failures. A solution for discriminating proximate places is manually separating the fingerprints for challenging places with a more aggressive similarity threshold. Finally, places were often *missed* when Wi-Fi wasn't available. These places are usually outdoors or in rural areas, where GPS may work better to detect places.

Managing Suggested Places

To understand how users dealt with suggested places, we logged which user interface was used to review suggested places and received additional feedback from exit surveys.

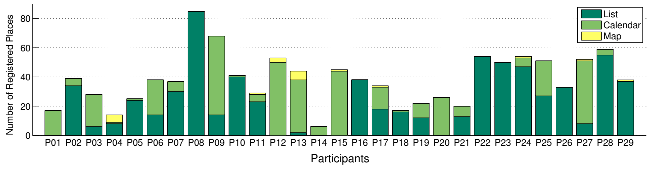

Figure 10. User interfaces selected for viewing and registering suggested places. Calendar view and list view were commonly used while map view was rarely used.

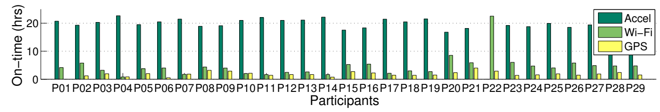

Figure 11. Daily average sensor on-time by participant

As illustrated in Figure 10, participants usually selected either *calendar view* or *list view* to review new places and stuck with one interface. *Map view* was used rarely. People who chose *calendar view* thought viewing place visits in a chronological order was helpful in remembering places, while others found it easier to view every suggested places in one place using the *list view*. About 38% of our participants answered that they preferred *calendar view* and 48% favored *list view*. We asked when the participants usually reviewed suggested places. 45% said they reviewed suggested places within a couple of hours after they left a new place and 31% reviewed new places at the end of the day. Others cleared the items right after they left the place (3%) or a couple of days later (3%). When we asked in what order they consulted the hints we provided (visit times, a map position, and visible Wi-Fi APs), 42% answered time-map-wifi and 31% answered map-time-wifi. While viewing the surrounding APs helped in some cases, participants found time (54% as the first consultant) and map (43%) more helpful to remember places. However, 71% of the participants thought that managing places became difficult when they couldn't recognize the suggested place, while assigning a name to each place (36%) or taking the time to handle them (18%) were less of an issue. It was also mentioned that map failed to provide enough information in cases where the GPS provided largely inaccurate indoor position estimations.

Encouraging User Feedback

Finally, we asked participants about their experience during the user study through exit surveys and sought for suggestions. We wanted to highlight two key implications to the semantic location services that resulted from the field study results: 1) people are generally open to providing feedback to improve the service, and 2) well-timed notifications and informative feedback motivates users to provide more input.

Based on the post-study exit survey, participants were mostly positive about receiving suggested places and managing personal places, M=7.32, (based on a Likert scale from 1 to 10, from annoying to fun). We also asked the maximum number of suggested places a participant will handle per day before giving up. 32% of the participants said that the number of suggested places doesn't really matter as long as they are reminded to review suggestions before the place is forgotten. 50% chose four or above and 28% chose three or below. People suggested that providing other contextual information about the visit (*e.g.*, phone calls made) or adding a social component (*e.g.*, sharing places with friends) may encourage them to engage more. However, the most desired feature was getting reminders about new suggestions. 75% of the participants thought they would register more places if Loci prompted when new places appeared. 46% wanted to be prompted while they were at the new place and 36% preferred to be prompted when arriving at home in the evening. Others favored being prompted when staying at some place for a long time (11%) or while in transit (7%).

Participants were also willing to provide feedback when Loci fails to correctly detect place visits as they expected. About 96% of the participants answered that they can tell Loci when it's not correctly detecting places so that it can do a better job for the next visit. Based on their experience using Loci, 25% thought they will have to provide place failure feedback once every other day, 39% as twice a week or less, and 35% as once a day or more. Lastly, 69% told us that visualizing and providing feedback about how Loci detects enter and exits as well as showing how it recognizes known places will have positive influence on getting more user feedback.

Resource Usage

Lastly, we analyze Loci's daily energy consumption and memory usage, which are greatly affected by the user's mobility pattern. We logged the time each sensor was activated to estimate power usage as different phones have different power profiles. As illustrated in Figure 11, on average, accelerometer, Wi-Fi, and GPS were activated for 19.9, 3.9, and 1.7 hours, respectively. Samsung Galaxy S lasted for about 15 hours while Motorola Droid X lasted for about 36 hours. As the accelerometer had the longest active time, in general, we suspect that the two models have significantly different power profile for their accelerometers. Loci required 0.25 MB per day on average (min: 0.05 MB, max: 0.53 MB) to save location histories.

CONCLUSIONS AND FUTURE WORK

Loci is a personal semantic location manager that helps users to identify everyday places and paths. Location-based applications can take advantage of the personal places registered by the service, much as the contact list of the user is currently used by social applications. Loci employs a place learning algorithm to help users identify semantic places. Places are gradually learned and suggested as users visit new places, instead of asking them to define places one-by-one ahead of time. Paths between places can also be traced for track-based applications. We investigated the amount of user feedback required by Loci when used everyday. Our three-week study with 29 participants indicates that the number of suggested places could vary from 0 to 10 or more per day depending on the user's activity level but gradually decreased overall as the service was used for a longer period of time. About 75% of the suggested places were registered as a new place while 16% were added to an existing place. Place detection failures were reported once every other day on average, assuming that a failure at a particular place can be fixed once it is reported. Most of the user study participants felt positive about providing feedback to the service for better performance. Finally, on a Likert scale from 1-10, where 1 stands for "Never" and 10 for "Absolutely", response to "Will you use a service like this?" had a mean of 7.28 and standard deviation of 2.36.

Our results provided us with a foundation for two new research agendas. First, we plan to investigate how user feedback can be transformed into system parameters and improve place detection performance. Adaptively intermixing place learning techniques based on Wi-Fi and GPS may also allow us to cover more places. Designing user interfaces that intuitively convey the cause of place detection failures and requests minimum user effort in providing feedback is also critical. This will help users to define and tune places as they want. Next, we plan to focus on identifying a rich class of location-based applications and implementing the service at a larger scale.

REFERENCES

1. D. Ashbrook and T. Starner. Using gps to learn significant locations and predict movement across multiple users. *Personal Ubiquitous Comput.*, 7(5):275–286, 2003.
2. M. Azizyan, I. Constandache, and R. Roy Choudhury. Surroundsense: mobile phone localization via ambience fingerprinting. In *MobiCom '09: Proceedings of the 15th annual international conference on Mobile computing and networking*, pages 261–272, New York, NY, USA, 2009. ACM.
3. L. Barkhuus, B. Brown, M. Bell, S. Sherwood, M. Hall, and M. Chalmers. From awareness to repartee: sharing location within social groups. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 497–506, New York, NY, USA, 2008. ACM.
4. A. Barry, B. Fisher, and M. L. Chang. A long-duration study of user-trained 802.11 localization. In *Proceedings of the 2nd international conference on Mobile entity localization and tracking in GPS-less environments*, MELT'09, pages 197–212, Berlin, Heidelberg, 2009. Springer-Verlag.
5. E. S. Bhasker, S. W. Brown, and W. G. Griswold. Employing user feedback for fast, accurate, low-maintenance geolocationing. In *Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications (PerCom'04)*, PERCOM '04, pages 111–, Washington, DC, USA, 2004. IEEE Computer Society.
6. P. Bolliger. Redpin - adaptive, zero-configuration indoor localization through user collaboration. In *Proceedings of the first ACM international workshop on Mobile entity localization and tracking in GPS-less environments*, MELT '08, pages 55–60, New York, NY, USA, 2008. ACM.
7. P. Bolliger, K. Partridge, M. Chu, and M. Langheinrich. Improving location fingerprinting through motion detection and asynchronous interval labeling. In *Proceedings of the 4th International Symposium on Location and Context Awareness*, LoCA '09, pages 37–51, Berlin, Heidelberg, 2009. Springer-Verlag.
8. I. Constandache, S. Gaonkar, M. Sayler, R. R. Choudhury, and L. P. Cox. Enloc: Energy-efficient localization for mobile phones. In *INFOCOM*, pages 2716–2720. IEEE, 2009.
9. J. Froehlich, M. Y. Chen, I. E. Smith, and F. Potter. Voting with your feet: An investigative study of the relationship between place visit behavior and preference. In *Ubicomp '06*, pages 333–350, 2006.
10. W. Griswold, P. Shanahan, S. Brown, R. Boyer, M. Ratto, R. Shapiro, and T. Truong. Activecampus: experiments in community-oriented ubiquitous computing. *Computer*, 37(10):73–81, 2004.
11. R. Hariharan and K. Toyama. Project lachesis: Parsing and modeling location histories. In M. J. Egenhofer, C. Freksa, and H. J. Miller, editors, *GIScience*, volume 3234, pages 106–124. Springer, 2004.
12. J. Hightower, S. Consolvo, A. LaMarca, I. E. Smith, and J. Hughes. Learning and recognizing the places we go. In *Ubicomp '05*, pages 159–176, 2005.
13. J. H. Kang, W. Welbourne, B. Stewart, and G. Borriello. Extracting places from traces of locations. In *WMASH '04*, pages 110–118, New York, NY, USA, 2004. ACM.
14. D. H. Kim, J. Hightower, R. Govindan, and D. Estrin. Discovering semantically meaningful places from pervasive rf-beacons. In *Ubicomp '09: Proceedings of the 11th international conference on Ubiquitous computing*, pages 21–30, New York, NY, USA, 2009. ACM.
15. D. H. Kim, Y. Kim, D. Estrin, and M. B. Srivastava. Sensloc: Sensing everyday places and paths using less energy. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 43–56, New York, NY, USA, 2010. ACM.
16. M. B. Kjaergaard, J. Langdal, T. Godskand, and T. Toftkjaer. Entracked: energy-efficient robust position tracking for mobile devices. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 221–234, New York, NY, USA, 2009. ACM.
17. K. Laasonen, M. Raento, and H. Toivonen. Adaptive on-device location recognition. In *Pervasive '04*, pages 287–304, 2004.
18. L. Liao, D. Fox, and H. Kautz. Extracting places and activities from gps traces using hierarchical conditional random fields. *Int. J. Rob. Res.*, 26(1):119–134, 2007.
19. J. Lin, G. Xiang, J. I. Hong, and N. Sadeh. Modeling people's place naming preferences in location sharing. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, Ubicomp '10, pages 75–84, New York, NY, USA, 2010. ACM.
20. K. Lin, A. Kansal, D. Lymberopoulos, and F. Zhao. Energy-accuracy aware localization for mobile devices. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys'10)*, 2010.
21. H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell. Soundsense: scalable sound sensing for people-centric applications on mobile phones. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 165–178, New York, NY, USA, 2009. ACM.
22. N. Marmasse and C. Schmandt. Location-aware information delivery with commotion. In *HUC '00*, pages 157–171. Springer-Verlag, 2000.
23. M. Monibi and D. J. Patterson. Getting places: Collaborative predictions from status. In *Proceedings of the European Conference on Ambient Intelligence*, AmI '09, pages 60–65, Berlin, Heidelberg, 2009. Springer-Verlag.
24. J. Paek, J. Kim, and R. Govindan. Energy-efficient rate-adaptive gps-based positioning for smartphones. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys'10)*, June 2010.
25. J.-g. Park, B. Charrow, D. Curtis, J. Battat, E. Minkov, J. Hicks, S. Teller, and J. Ledlie. Growing an organic indoor location system. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys '10, pages 271–284, New York, NY, USA, 2010. ACM.
26. D. J. Patterson, X. Ding, S. J. Kaufman, K. Liu, and A. Zaldivar. An ecosystem for learning and using sensor-driven im status messages. *IEEE Pervasive Computing*, 8:42–49, 2009.
27. T. Sohn, K. A. Li, G. Lee, I. E. Smith, J. Scott, and W. G. Griswold. Place-its: A study of location-based reminders on mobile phones. In M. Beigl, S. S. Intille, J. Rekimoto, and H. Tokuda, editors, *Ubicomp '05*, volume 3660 of *Lecture Notes in Computer Science*, pages 232–250. Springer, 2005.
28. N. Toyama, T. Ota, F. Kato, Y. Toyota, T. Hattori, and T. Hagino. Exploiting multiple radii to learn significant locations. In *LoCA '05*, pages 157–168, 2005.
29. Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma. Mining interesting locations and travel sequences from gps trajectories. In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 791–800, New York, NY, USA, 2009. ACM.
30. C. Zhou, D. Frankowski, P. Ludford, S. Shekhar, and L. Terveen. Discovering personally meaningful places: An interactive clustering approach. *ACM Trans. Inf. Syst.*, 25(3):12, 2007.
31. C. Zhou, P. Ludford, D. Frankowski, and L. Terveen. Talking about place: An experiment in how people describe places. 2005.
32. Z. Zhuang, K.-H. Kim, and J. P. Singh. Improving energy efficiency of location sensing on smartphones. In *MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 315–330, New York, NY, USA, 2010. ACM.