I'm feeling LoCo: A Location Based Context Aware Recommendation System

Norma Saiph Savage*, Maciej Baranski*, Norma Elva Chavez**, Tobias Höllerer *

* University of California, Santa Barbara, Computer Science Department
** Universidad Nacional Autonoma de Mexico, Facultad de Ingenieria, Departamento de Ingenieria en Computacion

Abstract. Research in ubiquitous location recommendation systems has focused on automatically inferring a user's preferences while little attention has been devoted to the recommendation algorithms. Location recommendation systems with a focus on recommendation algorithms generally require the user to complete complicated and time consuming surveys and rarely consider the user's current context. The purpose of this investigation is to design a more complete ubiquitous location based recommendation algorithm that by inferring user's preferences and considering time geography and similarity measurements automatically, betters the user experience. Our system learns user preferences by mining a person's social network profile. The physical constraints are delimited by a user's location, and form of transportation, which is automatically detected through the use of a decision tree followed by a discrete Hidden Markov Model. We defined a decision-making model, which considers the learned preferences, physical constraints and how the individual is currently feeling. Our recommendation algorithm is based on a text classification problem. The detection of the form of transportation and the user interface was implemented on the Nokia N900 phone, the recommendation algorithm was implemented on a server which communicates with the phone. The novelty of our approach relies on the fusion of information inferred from a user's social network profile and his/her mobile phone's sensors for place discovery. Our system is named: I'm feeling LoCo.

Keywords. Personalization, Recommendation Systems, Pervasive computing, Human Computer Interaction, Context Aware Recommendation Engines, Automatic Travel Guides.

1. Introduction

Personalization (see e.g., Mulvenna S. M and Buchner A., 2000) is a key component of modern location based services (LBS). To achieve personalization the system builds a model representing the user's needs and preferences. In the literature this process is generally called a user model (UM) (see e.g., Kobsa. A., 2001). Most LBS obtain this data by extended surveys. Yet responding to long questioners can be cumbersome and imposes a cognitive burden for the user.

Recommendation systems (see e.g., Resnick H. V. P. , 1997) are a particular type of personalized system which filters information and presents only what is relevant to an individual. There have been many attempts to build systems, which recommend places to visit and facilitate decision-making. For example, Rinner and Raubal (2004) designed a service named Hotel Finder which by considering a user's location, spatiotemporal constraints and preferences, recommended suitable hotels. Espeter and Raubal (2009) extended Hotel Finder and created a system that aided cohort decision-making. Albeit decision making in unfamiliar environments was improved, as these system required users to manually input the majority of their preferences, the user experience in both cases may have deteriorated

According to the studies done by Regula and Bachman (1981), when responding to long questionnaires individuals are more likely to give identical answers to most or all of the items. Therefore in Hotel Finder and in the work of Espeter and Raubal (2009), because the user had to manually provide a series of personal preferences, it is likely many of the user' responses did not truly reflect their interests. From the study of Regula and Bachman (1981) it is clear that a system, which could automatically and accurately infer an individual's preferences, would dramatically boost the user experience because the user would not require to spend time completing surveys.

The work done by Sugimoto Takeuchi Y (2006) took a first approach in place recommendation without needing the user to pass through any survey phase. By utilizing information from a user's location history, the system in their study automatically suggested stores to visit. The drawbacks were that for logging the history of visited places, the system assumed that all locations were stores. The system stored no information related to the nature of the place the user was visiting: grocery stores, Mexican restaurants and pawn shops were treated as being the same category. The system simply calculated the probabilities that existed for moving from one position to another. It was therefore impossible to query the system and ask for suggestions of where to eat or where to buy an inexpensive pair of jeans, which are some of the typical questions a user might have for a location-based recommendation system.

A number of previous systems automatically discovered the significant places related to a user and also stored information related to these sites. The work done by Marmasse, N. and Schmandt, C. (2006) is an example of this type of system, which discovered the places a user frequently visited and additionally stored information the user manually annotated, such as notes and to-do lists for a particular site. The work of Marmasse, N. and Schmandt, C. (2006) suffered a problem similar to that of Hotel Finder. It forced the user to spend time in selecting information from a small output device and entering data through uncomfortable input interfaces such as a thumb keyboard or a stylus. Another limitation of Marmasse, N. and Schmandt, C. (2006) was that its employed recommendation algorithm was very basic. The recommendations were based mostly on "reminders". For example if the system detected that the user was near a place for which they had set a to-do list, the system suggested visiting this site and completing the to-do list. It was not able to suggest new places for the user to visit and it provided little aid in the decision-making.

A system, which provided more assistance in the decision-making, while leveraging user input was TripTip by Kim et.al (2009). Given the places the user had visited and their characteristics, TripTip recommended sites. All the data was obtained by mining a person's 43places.com profile and crawling the 43place.com website. Despite TripTip's improvements over recommendation systems that automatically inferred user preferences, TripTip still suffered several limitations: it recommended places only within walking distance of where the user was last seen. The system disregarded spatial-temporal constraints. Additionally, because TripTip was not a mobile system it could not automatically detect the user's current location. Therefore to receive recommendations, the user had to actively update where they were. This evidently damaged the user experience. Furthermore, TripTip only suggested places that had a similarity with the sites in a user's 43places.com profile. Therefore an individual, whose 43places.com profile held only information about visits to educational institutes, would most likely have a difficult time receiving restaurant recommendations.

A system, which sought to suggest significant places, despite having little user data is foursquare's recommendation engine. foursquare is a location-based online social networking website, which permits users to "check-in" to places by either visiting a mobile version of their website, text messaging or by using a smart-phone specific application. In late May 2011, foursquare began offering place recommendations. foursquare's recommendation engine considers the user's location, check-in history (the places the user had visited) and "popular" sites near the user (foursquare developed a metric for inferring place popularity based on the number of user's that have visited the place, as well as the number of visits that all users have made). Due to the popularity metric, foursquare's recommendation engine can suggest relevant places to visit, without needing extensive user information. Albeit foursquare's recommendation engine can potentially suggest relevant places to visit without having the user provide extensive amounts of information, the system's recommendations are less user tailored and more generic. Furthermore, in cases where foursquare does have sufficient data, the recommendation engine fails to consider context for the recommendation: the engine does not acknowledge that the user's current transportation mode could affect the type of places a user would want to visit. Additionally because the engine only considers the category associated with the places the user visits, foursquare's recommendation engine disregards the place's contextual information, such as: Does the restaurant cater organic healthy food? What type of people do visit the restaurant: business men, students, surfers?

The research on place recommendation systems has paid little attention to the integration of contextual information for the recommendation algorithm. Systems, which do consider context, require the user to complete extensive surveys and constantly update their contextual information.

The aim of this investigation is to design a ubiquitous location based recommendation system, which by considering time geography and similarity measurements, presents a more complete recommendation algorithm. Our algorithm takes an approach similar to that of Rinner and Raubal (2004), but unlike Rinner and Raubal (2004) is not restricted to an extensive questionnaire phase or does it require the user to constantly update their contextual information.

Instead of including a survey phase, our system mines a person's social network profile and maps this information into user preferences. For inferring the user's preferences, our system, unlike TripTip and foursquare's recommendation engine, considers the contextual information related to the places the user has visited: tags and categories associated with a place are utilized for learning the user preferences. Our system additionally can also offer relevant place recommendations even under the circumstance that user data is lacking. Our system detects whether sufficient user data has been provided. If sufficient user content is not present, our system mines the information of the wikitravel page (http://wikitravel.org) of the city the user is in, and automatically finds the city's landmarks and adopts this data for the recommendation.

Furthermore, our system includes a mobile application, which automatically infers a user's current mode of transportation and utilizes this information to determine how far a person would be willing to travel to visit a location. In our approach the user is only required to input their mood. The user's mood is used to delimit even more the type of places, which will be recommended by the system. Figure 1 shows the features utilized for recommendation. Our system is named: I'm feeling LoCo. Where LoCo is short for Location and Context.

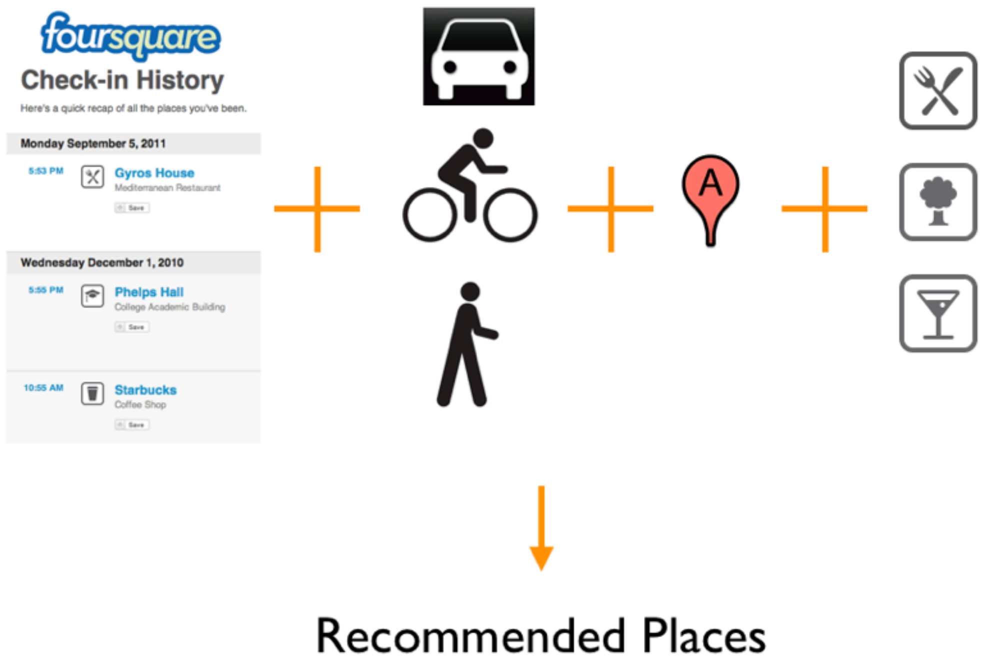

Figure 1. Features considered for the recommendation algorithm: the user preferences (based on the user's foursquare check-in history), the user's current transportation mode, the user's current location and the user's mood(the type of places the user is currently interested in visiting).

In the following sections we present in greater detail each component of our recommendation system. First, we explain how personal spatiotemporal constraints and preferences are automatically inferred and used as features for our place recommendation algorithm. Our place recommendation algorithm is presented afterwards. Insights gained from our study are given subsequently. In the final two sections we describe additional evaluations obtained via cognitive walkthrough methodology, and present our conclusions.

2. Automatic Integration Of A User's Spatiotemporal Constraints

Time geography considers that there are natural laws and social norms, which determine a person's ability to be present at a specific location and time. Communication and transportation services aid individuals in trading time for space and play an important role in allowing people to be physically present at a certain location and time. The work of Raubal et.al (2004) showed the importance of integrating time geography to personalized LBS, in particular personal spatiotemporal constraints.

Our study follows this guideline and incorporates in the personalization process capability constraints. Capability constraints are a particular type of spatiotemporal constraints, which confine human activities in time and space to available resources. Personal capability constraints can be delimited by the individual's mode of transportation, because the form of transportation bounds the places the person can visit. Our recommendation algorithm utilizes this constraint for delimiting the list of suggested places: only places "near" the user are analyzed and nearness is defined by the mode of transportation. For example, if an individual is riding a bicycle and requesting a restaurant recommendation, the system will not suggest places, which are an hour biking distance away. Whereas, if the person is driving a car, a restaurant that is an hour away by bike could still be recommended. Our study also considers that requiring the user to constantly update their current form of transportation is uncomfortable. Therefore a person's mode of transportation is automatically detected. The detection is done on the user's smartphone. In this case it was implemented on the Nokia N900 phone.

The method for automatically detecting a user's form of transportation is similar to the process proposed by Reddy et al. (2010). Their mobile system discriminated between a person who was stationary, walking, biking, or driving. The classification was realized by a decision tree (DT) followed by a first-order discrete Hidden Markov Model (DHMM). The combination of the decision tree with a discrete hidden Markov model improves classification, because the decision tree is tailored to differentiate between the boundaries of transportation modes, and the discrete hidden Markov model helps in reducing noise by utilizing temporal knowledge of the previous transportation mode that were detected. There are times for example, when the user is driving and due to traffic or stop lights their speed is decreased. This may cause the decision tree to classify the user's activity as biking, but because a transition from driving a car to biking is unlikely, the discrete hidden Markov model corrects the classification. As feature vectors, we used the variance of the accelerometer signal and the GPS speed data.

The approach of Reddy et al (2010) was selected because their system was capable of running without strict orientation or position requirements, which is fundamental when doing classification on smartphones, since the form in which individuals carry their phones varies widely. Some, for example, keep their mobile device in their backpacks, while others place them on their belts. This method was also selected, because the authors demonstrated that historical user pattern data is not required. It can therefore be immediately utilized by a person, without needing a prior training phrase. Figure 2 presents an outline of the algorithm.

Figure 2. Overview of the algorithm utilized for the detection of the transportation mode: The accelerometer variance and the GPS speed data from the n900 are the decision tree's inputs. The decision tree classifies the input to an activity: biking, walking or driving. A hidden Markov model is fed a series of activities encountered by the decision tree. Based on the presented pattern, the hidden Markov model determines the final activity classification.

3. Automatic Recollection Of User Preferences

For limiting the time spent on questionnaires, user preferences are automatically obtained by mining the user's social network profile. The utilized social network was foursquare. From a user's GPS coordinates, foursquare returns a list of possible places the person could be in. Each item on the list has an associated name, category and relevant tags which define the place. An item on the list might be for example "Ruben's Tacos," with the category of restaurant and tags such as: Mexican food, burritos, mariachi, margaritas etc. The user, with a simple click or tap, can select the place they are currently in, and foursquare will log to the user's profile, the place along with all the associated tags. Figure 3 shows the foursquare user interface with a list of places the person can check in to.

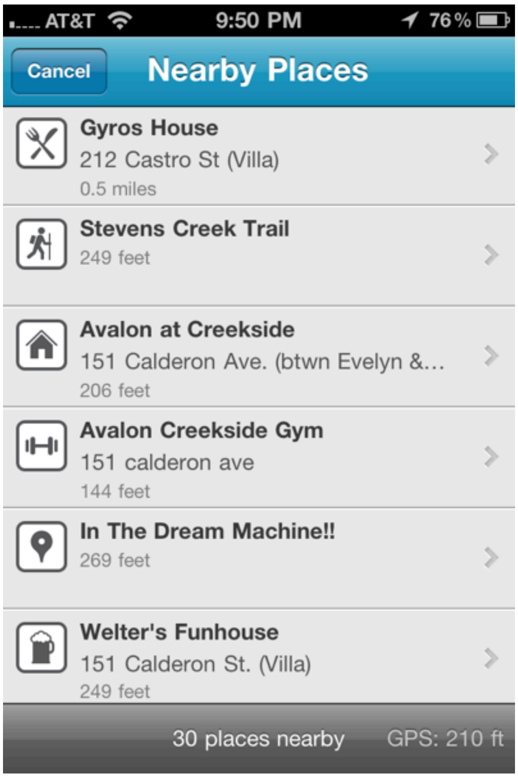

Figure 3. foursquare interface to check-in to a place.

The foursquare API permits the retrieval of all of the user check-ins along with the information conjoined with the places the user has visited, such as tags, name of place and category. We used the retrieved information to build a user model that holds contextual information related to the type of places the user visits. The user model is in essence a document, which holds a series of words. Each time a user visits a place, its name, category and tags are added to the end of the document. In our implementation, a server manages the creation of the user model and has a daemon, which periodically checks the user's foursquare profile for new check-ins, and performs an update if necessary.

4. Place Recommendation Algorithm

Recommendation algorithms are generally divided in two types: algorithm that utilize collaborative filtering and algorithms that utilize content based filtering (see e.g., Baudisch, 1999). Collaborative filtering establishes that personal recommendations can be computed by calculating the similarity between one user's preferences and the preferences of other individuals. In collaborative filtering the preferences of a large user group is registered. Given a user A that is seeking recommendations, similarity metrics are utilized to find a subgroup of people that present preferences similar to that of user A. An average of the preferences of that subgroup is computed. The returned preference function is what is utilized for user A's recommendations. By contrast, content-based filtering utilizes the information about an item itself for recommendations. The advantage of this method is that it is not limited to suggesting options that have previously been rated by users. Furthermore content based filtering can provide the user with a better explanation as to why option X was suggested. For example, the system can tell the user that 'Ruben's Tacos' was recommended, because the user had frequented restaurants before, which serve Mexican dishes. Content based filtering recommendation algorithms hold a set of items denoting the user's preferences. The task of the algorithm is to classify an unseen item (an option the user has not expressed any opinion about), as something relevant or irrelevant for the user.

Due to the nature of the data we were handling, it was decided to utilize a content based filtering approach for our recommendation algorithm. In this study, the items denoting the user preferences are the restaurants visited by the user along with their associated information: tags and assigned category. The unseen items to be classified are places "near" the user that they have never before visited. Because the space of places to analyze is immensely large, our system utilizes the user's capability constraints, preferences and mood, to delimit the search. Capability constraints influence the outcome of the algorithm as follows: given the current location of the user, the foursquare API is utilized to return all of the places within a certain radius to where the user is. The size of the radius depends on the user's mode

Figure 4. I'm feeling LoCo interface. The interface displays the detected transportation mode of the user and presents eight buttons from which the user can select their current mood.

of transportation: the faster the user moves, the larger the radius. A larger radius generally implies that more places will be considered for the recommendation. The radius size was empirically calculated. From this list of places, which are around the user, a second filtering step is performed. The filtering is now based on how the user is "feeling": foursquare labels every place with a category. As of December 2010 there were seven different categories: Arts & Entertainment, College & Education, Food, Work, Nightlife, Great Outdoors, Travel and Shops. Each of these categories was mapped to a particular feeling:

Arts & Entertainment= "feeling artsy"

College & Education="feeling nerdy"

Food="feeling hungry"

Home / Work / Other="feeling workaholic"

Nightlife="feeling like a party animal"

Great Outdoors="feeling outdoorsy"

Shops="feeling shopaholic"

Our system provides the user with an interface through which they can select with a tap one of the moods mentioned above and portray to the interface their feelings. Figure 4 shows the 'I'm feeling LoCo' interface. The selected mood delimits even more the places considered for recommendation: only places labeled with the category to which the chosen feeling is mapped to are selected. For example, if the user stated they were "feeling nerdy" only places labeled with the category of College & Education are acknowledged.

Aside from the "LoCo" mood, all the other user moods consider the contextual information associated with the places the user visited for the recommendation procedure: from each place in the newly filtered list of places, its associated tags are obtained. A set of words containing the intersection between the tags of the user and the tags of the particular place is created. For each term in this set of words its log frequency weight is obtained. The log frequency weight of a term t in a set d can be defined as a function F(t):

$$F(t) = \begin{cases} 1 + \log tf_{t,d} & tf_{t,d} > 0 \\ 0 & otherwise \end{cases}$$

where tf_{t,d} represents the number of times t occurs in d. In this case, d is the document, which holds all of the words associated with the places the user has visited, and t refers to one particular term or word present in the document. Once the weights for all of the tags of a particular place are calculated, a summation over all of the weights is done. This summation represents the log frequency weighting score of a particular place. The K places with the highest log frequency weighting are selected and are what is recommended to the user. K is a design parameter, which can be chosen arbitrarily. In our study, for visualization purposes, we set K to a value of 4. This list of K places represents the places that best match the user's personal preferences and spatiotemporal constraints. The list contains for each place, its name, the distance from the user's current location and the GPS coordinates of the place. The phone displays the recommended places on Google Maps, allowing the user to select the place he or she wishes to visit. For the particular case of when the LoCo mood is selected, the system retrieves all venues that are within a convenient distance to the user, nearness to the user being the only considered factor. The top K suggested places are picked randomly by the system.

Because our recommendation algorithm is content based, it depends on the user's foursquare check-ins to generate place suggestions. Relying on how active a user is on a social network can be problematic, especially in cases where the user rarely utilizes the social network site. For this reason, we developed a metric that will recommended meaningful venues to visit, regardless of the user provided content. The metric functions as follows: given the city the user is in, we mine the City's wikitravel page (wikitravel.org)

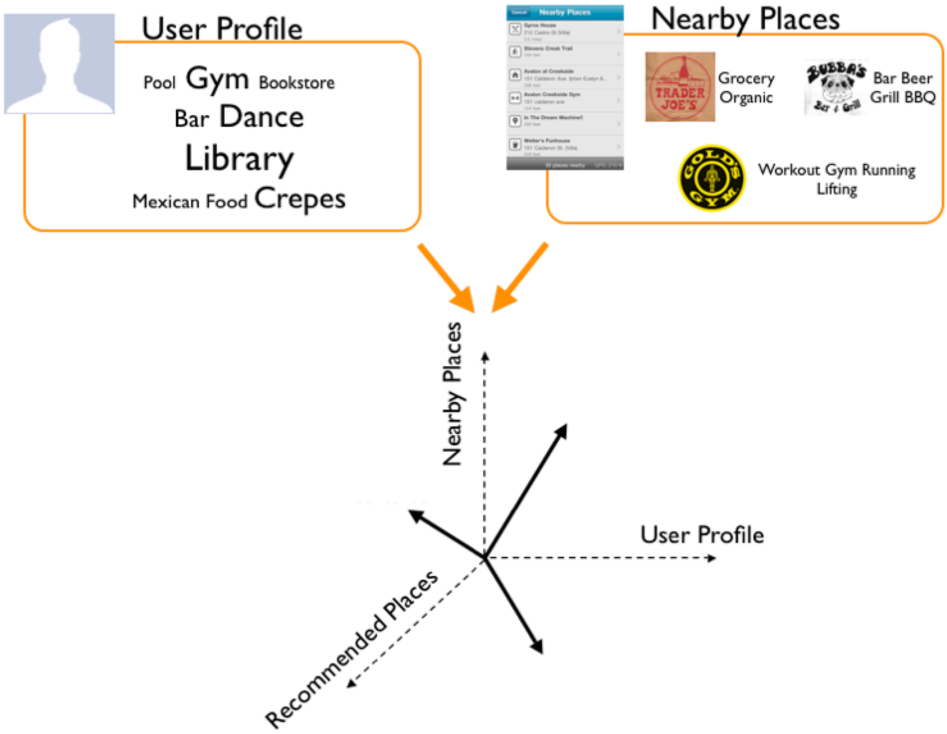

Figure 5. Outline of I'm feeling LoCo's recommendation algorithm. Depending on the user's profile and the places near the user, the algorithm recommends sites to visit. Only nearby places are considered for the recommendation. The definition of nearness changes according to the user's transportation mode. Places nearby are then further filtered based on the category they present and the similarity they have with the other places the user has previously visited. The K places with the highest similarity score are what is suggested to the user.

for the city's iconic places or landmarks. Each landmark is then searched on foursquare, where its address and associated category is retrieved. If the landmark is conveniently near the user and has the category the user requested, the landmark is suggested to the user. This metric permits the system to recommend significant venues without requiring excessive user generated content. Figure 5 presents an outline of the final recommendation algorithm and Figure 6 of the entire system.

5. Iterative Design

During the course of our work, which we tested throughout by frequent use by the authors and several volunteers in informal formative design evaluations, a number of limitations were encountered. In the following, we men-

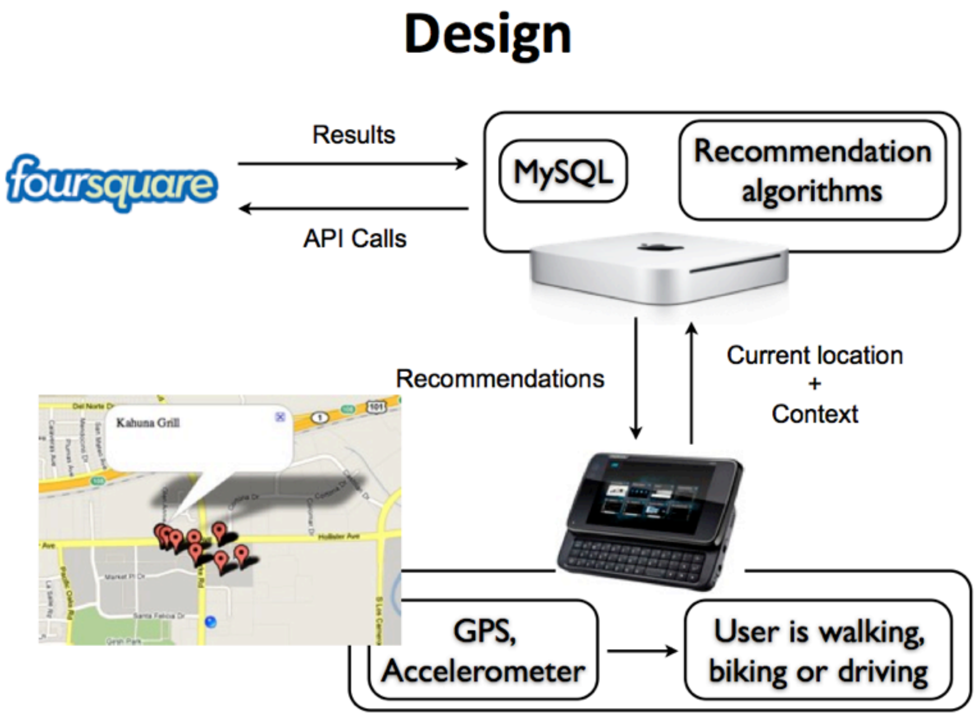

Figure 6. Overview of I'm feeling LoCo system operation. Through foursquare API calls, a server is constantly updating an individual's user model, and storing in a database this information. On the other hand, a mobile application is continuously detecting the user's transportation mode as well as offering an interface through which the user can request place recommendations. When the user queries the system for a recommendation, the mobile phone submits to the server, the user's location and contextual information: their transportation mode. The server inputs this data, along with the generated user model to the recommendation algorithm, that utilizes these features to decide what places are the most suitable to be recommended. The server returns to the mobile phone a list of the best K places for the user to visit. The mobile phone then displays these places on a virtual map.

tion a few of the most interesting ones: Foursquare returns only places within 400 meters radius from the user's current location. But our place retrieval algorithm requires a radius, which varies according to person's mode of transportation. An algorithm for automatically increasing the considered radius had therefore to be implemented. The functionality of this new feature is presented in Figure 7.

Another issue was that the places returned by the foursquare API only contained the name and categories associated with the place, its corresponding tags were not given. Therefore additional API calls were required per place. These extra API calls led to exceeding the permitted number of foursquare

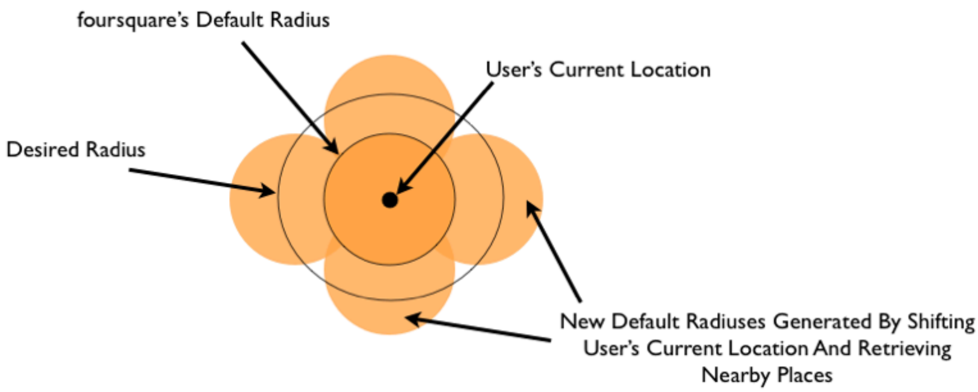

Figure 7: This diagram describes how the radius of places near the user is incremented beyond foursquare's limit.

API calls a user is allowed to make within an hour, which is 200. In order to overcome this limitation for the purpose of facilitating a meaningful evaluation (the cognitive walkthroughs described in Section 6), we resorted to the following temporary work-around, which is not adhering to foursquare's usage policies and which we don't advocate for any real use, but which was workable for the duration of our experiments: The API calls for attaining tags associated with a given place do not have to be authenticated by the user who is currently using the “I'm Feeling LoCo” application. Hence a number of foursquare “helper” accounts were created, for the purpose of making additional API calls. An automatic account switcher was implemented. This account switcher tracked per “helper account” the number of foursquare API calls which had been made in the last hour. When the number of calls was approximating the limit, it gracefully switched to another helper account: one that had not been utilized for the last hour. On top of account switching, a server-side caching strategy was also implemented to store the data associated with places that had previously been retrieved from foursquare. Before making an API call to retrieve the tags of a place, it was determined whether or not that specific place had already been added to the database.

We believe that an important lesson to be learned from our studies is that although integrating a social networking site into a system can allow for a quick profiling of an individual, there will be always certain shortcoming from the offered API, mainly because it is not tailored to the specifics of the system that is integrating it.

6. Usability Inspection of I'm feeling LoCo

In this section, we inspect the usability of I'm Feeling LoCo using cognitive walkthrough methodology. The cognitive walkthrough is a practical evaluation method based on the work of Polson, Lewis et al. (1992) in exploratory learning. In a cognitive walkthrough the user explores the interface and with the system tries to accomplish a series of assigned tasks. The user first searches the interface, analyzing what the system enables them to do. Next the user selects the specific actions that appear to aid them in achieving their final assigned goal. The cognitive walkthrough helps identify the ease of learning, use and usability of an application. The 'I'm Feeling LoCo' usability inspection was done in different US cities (Portland, OR; Beaverton, OR; Santa Barbara, CA; and Goleta, CA) by eight different users under diverse transportation modes. Each user utilized the system on at least two separate occasions.

6.1. Users

The people selected for the cognitive walkthrough were foursquare users that had at least 20 check-ins. All of the participants owned a Smartphone and had utilized the navigation assistant equipped in their phone. Only two of the users in the study had utilized a personalized travel guide. Five of the users obtained their place recommendations from casual conversations with acquaintances and reading online reviews in sites such as Yelp.

6.2. Tasks

The tasks we requested each user to perform were:

*Find a place to eat while walking in downtown Santa Barbara or Portland.

*Find a place for celebrating with friends while being a passenger and navigator in a car near Santa Barbara and Portland.

*Find a place for studying while biking in Goleta, CA.

Due to resource limitations, only three users did the task involving biking.

6.3. Results

The findings of the cognitive walkthrough were divided in two: the results obtained for the main menu and the results obtained for the mobile map in which the users visualized the recommended places.

1.1.1. Main Menu Usability

The main menu presents the user's detected transportation mode as well as a series of buttons denoting possible moods for the user to select. The observations and feedback received for the main menu was very positive. All of the users in the study were able to correctly select the mood associated with the type of place they were requested to find.

Additionally all of the users were inquired about what option they would select if they wanted to:

*Find a place for studying.

*Find a place for doing cultural activities.

*Find a park for walking.

*Find a place for shopping.

*Find a few of the corporations that had offices in the area.

All of the users selected the correct mood for each of the above tasks. It was straightforward for them to understand the mapping between their selected mood and the type of places returned by the system. All users appreciated the whimsical names selected for the moods. Furthermore, the users were surprised that I'm feeling Loco could flawlessly detect their current transportation mode. All users made positive comments about this feature. The users also made positive comments about the 'I'm feeling LoCo' button. They liked that this option allowed them to discover places that were not within their normal pattern. The users viewed this option as a fun component of the system that enabled them to explore their surroundings more. Three users suggested an option for querying the system for directions to a specific place. They mentioned that many times while shopping they had specific stores they wanted to visit. Therefore an alternative for directly searching for a particular place would be beneficial.

1.1.2. Mobile Map Usability

The participants were overall satisfied with the places returned by the recommendation system. For the first users participating in the study, the place recommendation requests were done in small US towns, such as Beaverton, OR; Hillsboro, OR; and Goleta, CA. Because these small towns offered a very limited selection of places, the system's suggestions were not very relevant for the user and did not fully portray the person's interest. It was due to this situation, that we opted to perform the cognitive walkthrough in downtown Portland and downtown Santa Barbara, where we could guarantee that a larger subset of places would be present, and better results would be obtained. When the study was done in downtown Portland and Santa Barbara, all the users expressed that the suggestions were places they would be interested in visiting. All the returned places were relevant to the user and were located at a distance reachable to the user.

An interesting pattern we observed was that on the second day of utilizing the system, the foursquare usage of all participants had incremented. We believe that after noticing that the application considered their check-ins for the recommendation, the users felt motivated to check-in to more places they visited and obtain therefore a far more tailored recommendation. On the second trial of the system, the users were pleased with how the recommendations better matched their personal preferences. We believe this shows that our system promotes the usage of location based social sites.

While visiting the downtown area many users were surprised that a few of the most popular and typical restaurants were not recommended by the system. This made us consider that in a location recommendation system, there are places that all users should be exposed to. Explicit serendipity should be enabled at all times and not only when there is not sufficient user content (Our current system explicitly recommends city landmarks, only when there is a lack of user data).

The participants enjoyed the idea that the recommendations changed accordingly to their transportation mode, but two of the users expressed difficulty in reaching the destination (These two users were the only ones that were not locals in the city where they tested our application). The difficulty arose because I'm feeling LoCo simply displays markers with the top recommendations, but offers no instructions as to how the destination can be reached. We conclude that this feature be integrated to the interface.

Additional suggestions that we derived from user feedback included modification of the type of presented map, providing detail based on the user's transportation mode; for example: show bike paths when biking. We also considered that eyes-free interaction with the system could be beneficial, especially while biking or driving. We plan on integrating an eyes-free approach similar to that of Savage et al. (2010). In summary, the cognitive walkthroughs demonstrated a good level of usability, yielded positive reactions from the participants, and generated ideas for further improvement.

7. Conclusions

In this study we presented a novel personalization system, which considers automatically inferred user preferences and spatiotemporal constraints for location recommendation.

This system can serve as an early research example, providing an outlook on future developments in personalized LBSs, in which the majority of the data utilized for generating the recommendations is automatically collected from different information sources, freeing the user from completing exhaustive surveys or manually updating their current state. Our system models personal spatiotemporal constraints by automatically discovering the user's mode of transportation. To infer personal preferences we proposed using a bag of words approach. We also presented the results from cognitive-walkthrough-style evaluations of our location recommendation system, comprising results from eight individuals, searching for places that best satisfied their personal priorities in the US cities of Goleta, CA; Beaverton, OR; Santa Barbara, CA; and Portland, OR. The cognitive walkthrough sessions demonstrated that our proposed system can be utilized to deliver useful location recommendations.

We believe that, the recommendation procedure could be improved by integrating other sources of user information, such as a person's Google calendar. The system would now consider the fact that the user has appointment X at time Y in location Z, creating a new spatiotemporal constraint. The inclusion of semantics may also provide additional useful information about the user. It would be interesting to analyze if meanings to certain visits could be inferred and effect the user model. For example, is there a meaning related to going to a nightclub every weekend or visiting a church every Sunday?

As this study concentrated on improving the user experience when utilizing a location recommendation system, little attention was paid to the energy efficiency. In our future work, we envision implementing a duty cycle on the mobile device, which could significantly extend the battery life. Instead of running the GPS and the accelerometer at all times, usage could be based on a duty cycle from the user's behaviors. We believe that if the user is actively using foursquare, it would be possible to construct a model of user behavior and predict when the user is at home, at work or sleeping. The system could predict user activity and therefore turn off sensors utilized for determining user activity.

Other areas of future research, which could provide a better user interaction with our system, are: changing the user interface accordingly to the person's form of transportation. For example, if the user is driving, an eyes-free interface could be presented and allow the user to keep their eyes on the road rather than on the recommendation system. Furthermore, if the user is biking, for example, the system could display routes of bike paths.

Our current system, due to foursquare's lack of user ratings to venues they visit, assumes that constant visits to a site means the user likes the place. But this is not always true, a user could frequent a place yet not enjoy it. In the future we hope to integrate to our study other social networks that do offer rating information (such as yelp). This would allow our system to return much better recommendations.

Despite its preliminary character, the research reported here indicates that it is possible to construct an adequate location recommendation system without requiring the completion of extended and complicated surveys. Furthermore, this study shows it is viable to integrate contextual information of the environment surrounding the user and user activities into the location recommendation engine.

Acknowledgments

This work was partially supported by a UC MEXUS-CONACYT fellowship, as well as NSF grants IIS-0747520 and IIS-1058132.

References

Bao L. and Intille S (2004) Activity recognition from user-annotated acceleration data. ACM Trans. Sen. Netw, pages 1–17.Bertin J (1967) Semiology of graphics: diagrams, networks, maps. University of Wisconsin Press, 1983 (first published in French in 1967, translated to English by Berg WJ in 1983)

P. Baudisch (1999) Joining Collaborative and Content-based Filtering. In: Proceedings of the ACM CHI Workshop on Interacting with Recommender Systems. ACM Press.

Espeter Martin and Raubal Martin (2009) Location-based decision support for user groups Journal of Location Based Services, Vol. 3, Iss. 3

Herzoh A. Regula and Jerald G. Bachman (1981) , Effects on Questionnaire Length on Response Quality, The Public Opinion Quarterly, Vol. 45, No. 4 (Winter, 1981), pp. 549-559

Kim Jinyoung, Kim , Hyungjin, and Ryu Jung-hee (2009). TripTip: a trip planning service with tag-based recommendation. In Proceedings of the 27th international conference extended abstracts on Human factors in computing systems (CHI EA '09). ACM, New York, NY, USA, 3467-3472

Kobsa. A. (2001) Generic user modeling systems. User Modeling and User-Adapted Interaction, 11(1-2):49–63.

Marmasse, N. and Schmandt, C (2000). Location-Aware Information Delivery with ComMotion. In: Proceedings of the 2nd international Symposium on Handheld and Ubiquitous Computing (Bristol, UK, September 25 - 27, 2000).

Mulvenna S. M and Buchner A. (2000), Personalization on the net using web mining. Communications of the ACM, 43:122–125.

Polson, P.G., Lewis, C., Rieman, J., and Wharton, C(1992). Cognitive walkthroughs: A method for theory- based evaluation of user interfaces. International Journal of Man-Machine Studies 36, 741-773.

Reddy S., Mun M, Burke J, Estrin D., Hansen M and Srivastava M. (2010) , Using mobile phones to determine transportation modes. ACM Trans. Sen. Netw, 69(2):13:1–13:27

Resnick H. V. P. (1997) ,Cross-representation mediation of user models. Communications of the ACM, 40(3):56–58.

Savage N.S., Ali S.R., Chavez N.E. and Savage R., (2010), Mmmmm: A multi-modal mobile music mixer. In Proceedings of the 2010 conference on New Interfaces for Musical Expression, NIME '10, pages 395-398, 2010

Seiler S, 2010, The impact of search costs on consumer behaviour : a dynamic approach. London School of Economics. Collaborative Management Talks.

Sugimoto Takeuchi Y, 2006, M. CityVoyager: An Outdoor. Recommendation System Based on User Location History. In Proc. UIC'2006, Springer Berlin

Rinner, C. and Raubal, M., 2004. Personalized multi-criteria decision strategies in location-based decision support. Journal of Geographic Information Sciences, 10, 149–156.