
Planning-Context Aware Mobile Recommendations

Chad A. Williams and Sean T. Doherty

Abstract

In the realm of mobile applications a significant effort has been made to develop recommender systems that customize results based off of one's current location and more recently even their inferred current activity. While this aspect of context has been shown to be quite successful, we suggest anticipating what they are currently planning for the future may help further improve the relevancy of the results as well. This work examines this problem as one of trying to predict the user's planning context, defined as what activities are currently being planned and how far in the future the event they are planning is going to be. An empirical analysis is made of the predictability of planning context and a discussion of the potential implications of this for mobile context aware recommenders.

Keywords

Mobile applications • Planning context • Locational context • Empirical analysis • Predictability

Introduction

Context aware recommender systems (CARS) have received significant focus in recent years as a way of increasing the relevance of results compared to traditional recommendation techniques. In this pursuit many different aspects of context have been examined regarding their ability to better understand what is of interest to the user at that particular point in time. As has been shown in previous work in the context of web recommendation, this type of insight can lead to superior predictions over basing recommendations on a general profile of the user alone. Within this work we study the

benefits a similar approach might provide in the realm of mobile recommendations.

One of the key differences between a traditional web site experience and a mobile application experience is the availability of additional information about the user beyond just their profile and/or click stream. As a result users typically have an increased expectation that mobile applications are more tailored based on their context. As such, with mobile applications passive observation such as location history take a much more central role in understanding the context of what is relevant to a person.

Identifying mobile user context has been a focus of the ubiquitous computing community for several years. While numerous studies have focused on current context awareness, using that context for recommendation and prediction of future context has received far less attention. Another aspect of this is that while predicting the next location has been well studied, other aspects of context may be just as relevant to a user. Some of these aspects have received little attention such as understanding when an activity and associated trip are planned.

C.A. Williams (✉)
Department of Computer Science, Central Connecticut State
University, New Britain, CT, USA
e-mail: cwilliams@ccsu.edu

S.T. Doherty
Department of Geography and Environmental Studies,
Wilfrid Laurier University, Waterloo, ON, Canada
e-mail: sdoherty@wlu.ca

We propose recognizing the context of a user in terms of their planning perspective to be a critical aspect in determining what is most relevant in mobile recommender system. For example, consider a user who is at their local coffee shop in the morning. Based on current models of context that focus on the immediate surroundings and/or time alone, the system might filter content to focus on what is nearby or related to their current activity. Consider, however, if the system knew the user was making plans in the near future of what they were going to be doing later that evening. In such a scenario, what the user considers to be most relevant extends beyond the immediate situation and also includes content that aids in making those plans.

This research focuses on including the planning behavior of mobile users as part of their recommender context. One of the key challenges addressed in this study is much of the data relevant to these aspects of context must be passively collected since user's typically do not explicitly identify this type of information. This work addresses how commonalities in planning behavior can be used to enhance what context is relevant to a mobile user outside of their immediate context despite limited information besides that that can be derived through passive data sources. The experiments conducted below attempt to identify both when plans are being made and the activity being planned. This work appears to be the first to address the integration of

these two goals. This is followed by a discussion of the implications of the findings and directions for future work in this area.

Mobile User Contextual Factors

Transportation planners have studied travel behavior extensively over the years. More recently, focus has shifted from looking at travel alone to understanding why a trip was made and when the decision to make the trip occurred. One of the approaches used for this has been examining the activity needs/desires of the person as the reason the travel is made [1–4]. These studies have shown that at an abstract generic level a person's context can help determine their future activity and planning behavior. This study, however, examines modeling a personalized context aware approach rather than using a generalized traveler prediction model.

Determining what is relevant in these terms requires an examination of several different types of contextual factors.

Below we describe these utilizing the terminology framework established in Adomavicius et al. [5].

With mobile devices, two types of dynamic fully observable data from the crux of determining the relevant context. Specifically the two categories are passive and active collection. As the architecture in Fig. 1, adapted from Williams

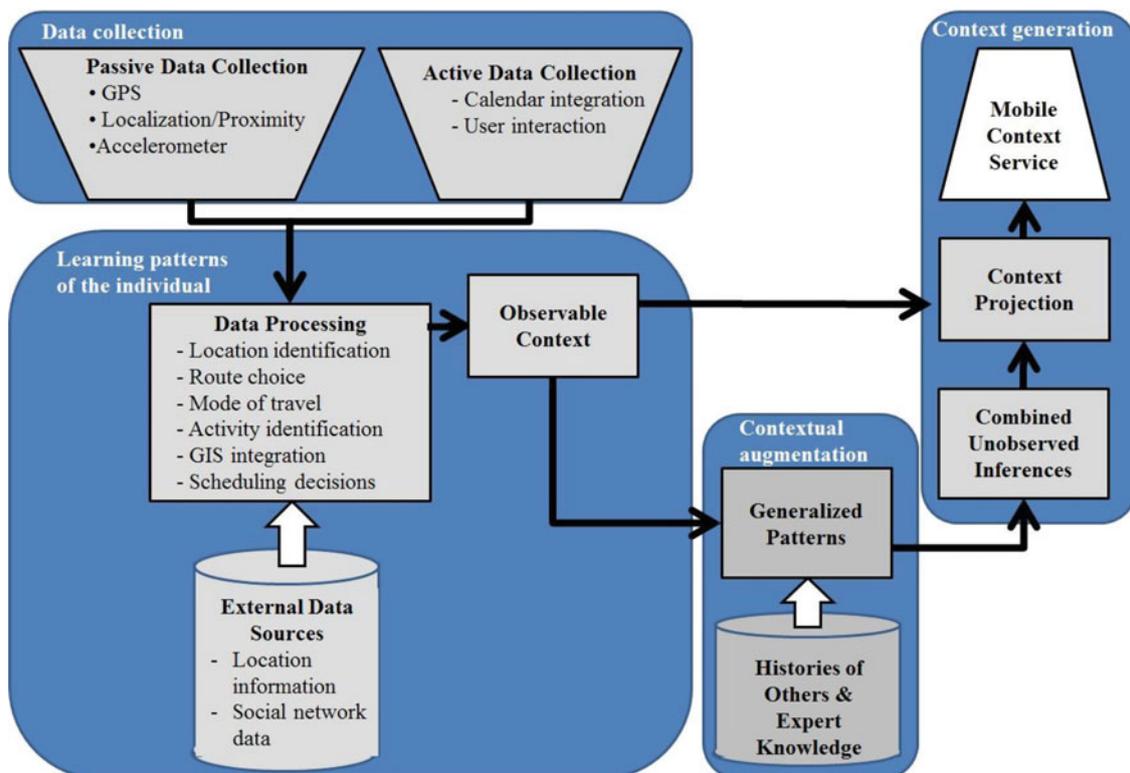


Fig. 1 Architecture of mobile context generation

and Mathew, shows, these two sources of data serve as the input to generating meaningful mobile user context.

Passive collection methods such as Global Positioning System (GPS) and accelerometer tracking are perhaps the most critical methods in mobile applications, because they provide location and movement context without requiring any active data entry by the user. Beyond the obvious location and route observations that can be collected, this data can often be processed to reveal partial context knowledge as well. For example by analyzing movement rate, stop frequency, and the accelerometer data, often the mode of travel such as walking, biking, train, bus, or car can be reliably inferred [6]. Combining location stop information with GIS map matching can sometimes be used to infer trip purpose [7]. Other sensors such as the device's microphones could be used to record ambient noise (e.g. voices, wind, and music) or determine whether the user is typing or talking providing additional environmental or use context. In theory, having tracked GPS data from all people would also allow passengers and social context to be inferred through examination of proximity and duration of contact [8]. Further processing of accelerometer data, especially from multiple body parts, can also be used to infer many specific human behaviors/context, such as sitting, standing, and specific exercise types [9–11].

In contrast, active collection methods involve manual data entry by the user. An example of active collection would be integration with calendar functionality, where calendar events that were entered manually could be used to ascertain event plans. There are obvious advantages to data entered explicitly by the user. One of these would be the potential for more accurate information, such as the specific activity that is actually taking place compared to just having the likelihood of what activity is occurring. Other advantages include the ability to collect information that simply cannot be captured completely through passive means such as the reasoning behind the planning decisions that are made. However, there are also significant limitations to active collection. Most notably any manual data entry puts a burden on the user, which in practice means users are less likely to provide the entry without incentive, particularly on an ongoing basis. However, one-time entry such as registration is often acceptable.

As noted in previous studies of transportation behavior, a demographic profile of a traveler can provide some significant insight in the types of activity patterns observed [12]. While some of this information might be explicitly given as part of a profile, such as age, overall it is partially observable as many demographic factors must be inferred to categorize the individual. Other types of information in this category such as work, school and home location, while not explicitly given, can often be reliably inferred through repeated past patterns.

Finally, within this work, what we are particularly interested in, and the goal of this study, is the largely unobservable dynamic context related to activity and travel planning. Towards this goal, this work examines the “Contextual augmentation” component of the architecture shown in Fig. 1, with the aim to output unobserved inferences of what the user is currently planning. We propose being able to predict when different aspects of plans are made is critical information to what is relevant to a user. As a result, recognizing that a plan is being made for a specific activity at a specific time window in the future may provide context that results in a significant improvement of determining relevant recommendations and thus improving the overall mobile user's experience.

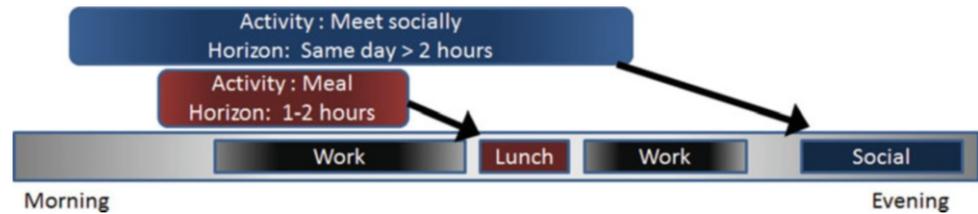
Planning Context

“Planning horizon” is a term used to describe the length of time before a trip or activity took place for which its plans were made [13]. For example if it was 10 AM and John decided to go out with friends at 6 PM that evening, the planning horizon for that activity would be 8 h. Planning context can range from far in advance, such as a visit to the doctor planned several weeks ahead, to spur of the moment, such as an urgent stop to get gas. For many activities, the planning may have become rather routine over time, such as a child regularly going to school in the morning, or regularly going to the same coffee shop every morning, wherein little active planning or thought was made. A typical person's day has a mix of routine and actively planned activities with a more finite planning horizon.

The focus of this work is on scheduling behavior and time horizons. Studies have confirmed that most people do not make their plans for all activities at a single point in time [14]. Instead, plans are continuously made and finalized throughout the day for varying planning horizons [15]. For instance, a person does not go throughout the day continually making plans two hours in advance of each activity, nor would they plan everything completely impulsively. In reality, scheduling is more fluid where a person might be making reservations for dinner the next morning, followed by an impulsive decision to grab a snack, followed by making plans for meeting a friend for lunch in an hour. The main point being it is not as simple as the person is just considering the next activity or looking any fixed period in the future.

We propose that determining what is relevant to a user at a given moment in time is in part dictated by their planning context. This planning context is a combination of the type of activity they are looking to plan and the planning horizon dictating when that activity will be carried out. The combination of these factors plays a critical role in what is relevant at a given point in time. For example, if the application knows when John is interacting with the system at 10 AM

Fig. 2 Planning context throughout day



that he is looking to make plans for that evening not just considering what to do next; that may greatly help in recommending the most relevant information.

If we were to envision what is relevant to a user from a planning perspective it is likely to fluctuate throughout the day. Take for example the planning context John wants to make plans to get together socially with friends in the evening at least 2 h before the event. John also wants to make plans to meet a family member for lunch an hour or two before the event. From a modeling perspective, we can potentially have more than one relevant planning context at a single time as shown in Fig. 2. In a portion of the morning both contexts would be relevant, but as lunch approaches the planning context for the meal would no longer be relevant. As this example illustrates when we are considering what context(s) is relevant it may not be a single answer.

Methodology

The purpose of this work is to assess how well a person's planning context can be predicted at a point in time during a person's day, given information that could be obtained either through an initial profile or by passive sources. Data for this paper is derived from a Computerized Household Activity Scheduling Elicitor (CHASE) survey conducted in Toronto in 2002–2003 [16]. The CHASE survey captured a detailed accounting of the activity scheduling process of 271 households over a 1-week period. The data was recorded such that for each activity that took place, both observed attributes (start time, end time, location, involved persons, category of the activity) and a time frame for when it was planned were all captured, via a scheduling program completed regularly throughout the week. The activities were broken down into 11 categories: active recreation; drop-off/pick-up; entertainment; household obligations; meals; night sleep, other needs; other; services; shopping; social; and work/school. The planning time frame included choices such as routine, X number of days ago, more than 2 h before, 1–2 h before, less than 1 h before, and just prior.

To model the planning context for a participant the activities recorded and their planning data were used to construct a time line for each day. The day was discretized into time segments early morning [12 AM, 9 AM), late morning [9 AM, 11 AM), midday [11 AM, 1 PM), early

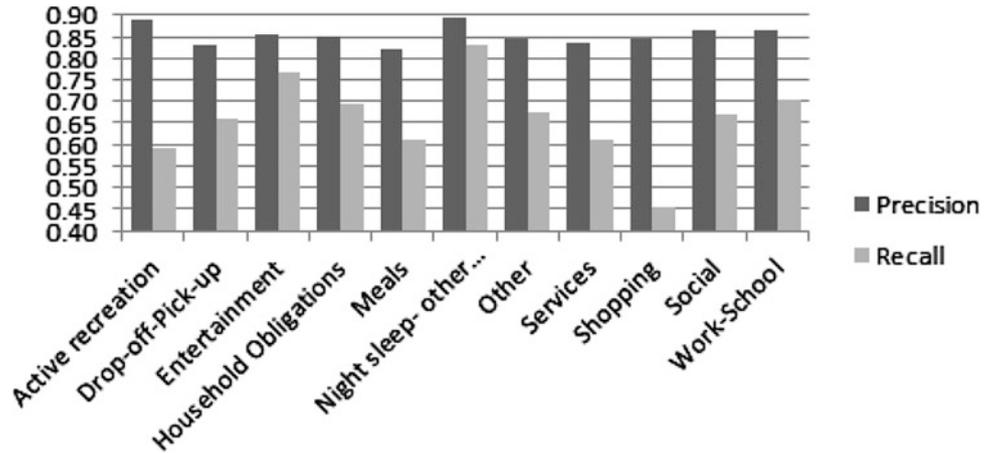
afternoon [1 PM, 3 PM), late afternoon [3 PM, 5 PM), early evening [5 PM, 7 PM), and late evening [7 PM, 12 AM). For each of these time segments a series of possible planning context entries were created that consisted of the combination of a specific activity type and one of three planning horizons (1–2 h prior, more than 2 h prior, or 1–2 days before) for a total of 33 possible planning contexts for each time segment.

To identify which planning context(s) were active, each of the planning context entries was defaulted as being not active. Next, each activity entry was then iterated through and based on the planning time frame the appropriate planning context time segments would be marked as active. For example if a 'Shopping' activity took place that was noted as having been planned '1–2 h before' the planning context would be noted as active for the period(s) that corresponded between 60–120 min before the *start* of the activity. If the planning time frame was noted as being 'more than 2 h before' the planning context was marked as active from 121 min prior through to the beginning of the day. For the '1–2 days before', no time on the activity's actual day was marked but the two full previous days would have the planning context ['Entertainment', '1–2 days before'] active for example. The result of this effort was a planning context schedule that indicates at a given time segment, what planning context(s) were active for the participant.

The goal was then to build a classifier that given a time of day, day of week, information that could likely be obtained through passive means, and a user profile; it could predict which planning context(s) were active. For the information that could be obtained through passive means, we are approximating that to be the current high-level activity type. While in current practice this cannot always be determined exactly, recent research advancements discussed above continue to make what can be inferred more reliable. In addition to these, a selection of personal data was selected based on its information gain: age, employment status, and gender.

The training data was then created by creating a record for each activity entry with the participant's data and for each activity type and planning horizon whether that planning context was active during that activity. Based on this training data a C4.5 classifier was built.

Fig. 3 Precision and recall for 1–2 h planning horizon



Evaluation Metrics

For measuring prediction performance, we use the information retrieval metrics of precision and recall [17]. The basic definition of recall and precision can be written as:

$$\text{precision} = \frac{\# \text{ true positives}}{(\# \text{ true positives} + \# \text{ false positives})}$$

$$\text{recall} = \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false negatives}}$$

For the purpose of this study, *# true positives* is the number of times a specific activity and planning horizon (planning context) is predicted correctly; *# false positives* is the number of planning context values incorrectly predicted in the time step, and *# false negatives* is the number of planning context values predicted as not in context that were in fact in context.

Experimental Results

For our experiments, we examine a comparison of precision and recall for the variety of different activity types captured in CHASE across the various planning horizons. All results reported are based on a tenfold cross validation methodology. Figure 3 captures how well the prediction model performs at identifying, for a specific time window, which activity is currently being planned that will take place in 1–2 h. In a practical sense, this is essentially taking the profile of the individual, their current activity, and given a time of day inferring that in the current hour they are making plans to go shopping an hour or 2 in the future. This information could potentially be very useful in adapting recommendations for what is currently of interest to a user. As the precision results in Fig. 3 demonstrate, the C4.5 classifier performs reasonably well with results

		Actual class	
		Planning	Not planning
Predicted class	Planning	3290	555
	Not planning	999	14561
1-2 hours before - Entertainment			

Fig. 4 Confusion matrix for 1–2 h planning horizon for activities entertainment related

		Actual class	
		Planning	Not planning
Predicted class	Planning	1473	235
	Not planning	725	16972
1-2 hours before - Social			

Fig. 5 Confusion matrix for 1–2 h planning horizon for activities social related

ranging from on the low end .82 for “meals” to .895 for “night sleep—other needs.” From a context recommendation perspective, this indicates that if a specific planning context is predicted as being relevant there is a high degree of likelihood that it is correct. From the perspective of tailoring recommendations this particularly important because it means there is a low percentage chance that the system would be making use of planning context that is irrelevant. The recall, on the other hand, varies significantly more from .452 for “shopping” to .831 for “night sleep—other needs.” From the perspective of recommendation, this would mean that the system did not recognize that some activity was being planned, and thus the missed

planning context would not be utilized in identifying the most relevant information.

A more detailed analysis of the activity types is illustrated in the confusion matrices. For “entertainment” depicted in Fig. 4, the results show that during 22 % of the time windows examined “entertainment” activities were being planned for 1–2 h in the future. This activity type was the second most observed for this planning horizon. A similar look at “social” activity planning context, which was present during 11 % of the time windows, is displayed in Fig. 5. Both of these planning contexts show similar precision, but recall is 10 % lower for “social” activities.

Figures 6 and 7 show the confusion matrices for the planning contexts of “services” and “shopping” 1–2 h in the future. As these results show, while the precision remains high, recall drops considerably for these contexts. In terms of the reasons why the planning context for these two activity types was more difficult to predict, may be due to the nature of these activities as compared to “entertainment” and “social” activity types. “Entertainment” and “social” activities tend to be more discretionary, whereas “shopping” and “services” are more task oriented. Thus, while additional study is needed, this may indicate identifying when optional

		Actual class	
		Planning	Not planning
Predicted class	Planning	923	180
	Not planning	583	17719

1-2 hours before - Services

Fig. 6 Confusion matrix for 1–2 h planning horizon for activities services related

		Actual class	
		Planning	Not planning
Predicted class	Planning	934	169
	Not planning	1132	17170

1-2 hours before - Shopping

Fig. 7 Confusion matrix for 1–2 h planning horizon for activities shopping related

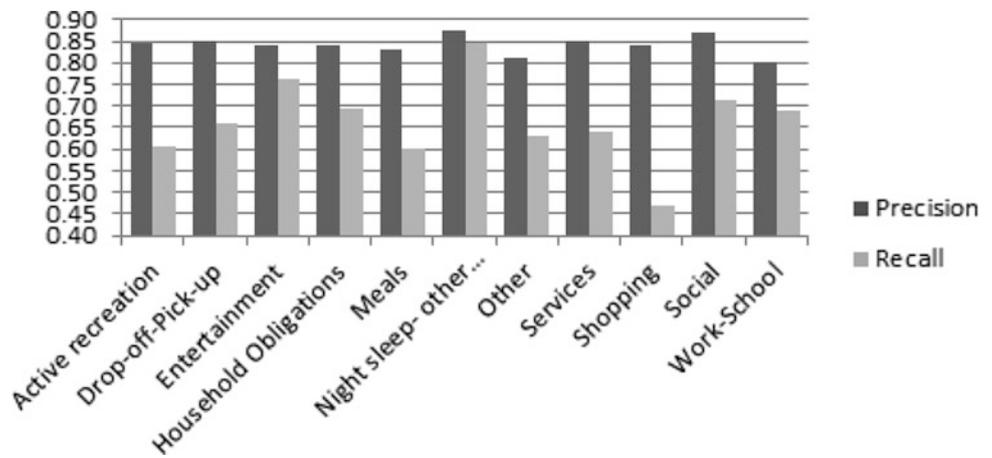


Fig. 8 Precision and recall for same day greater than 2 h planning horizon

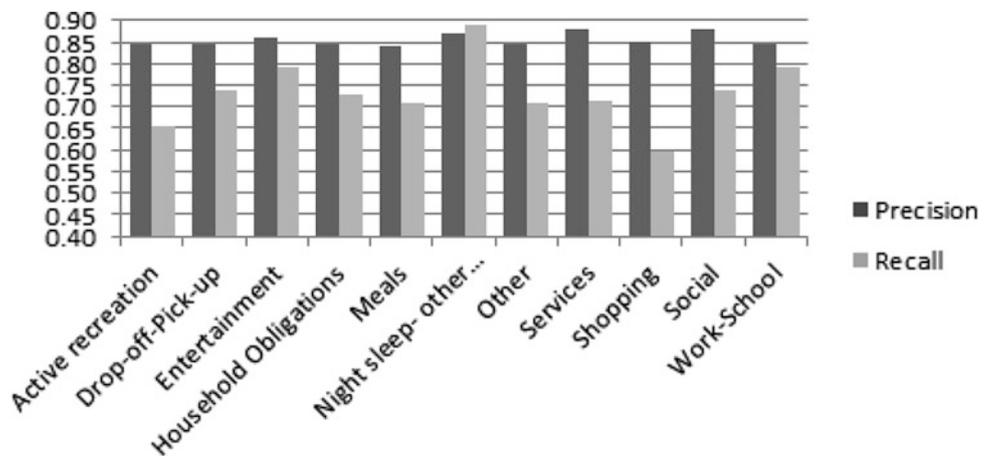


Fig. 9 Precision and recall for 1–2 day planning horizon

activities are being planned may prove to be an easier task than identifying when plans are made for activities related to maintaining the household.

Figures 8 and 9 depict the precision and recall for the activity types for the same day greater than 2 h and 1–2 days planning horizons respectively. The average precision results across the activity types for these two additional planning horizons is nearly identical to that of the 1–2 h planning horizon discussed above, approximately 85 %. For the recall results, the average across activity types for the same day greater than 2 h planning horizon (66.5 %) remains similar to that of the 1–2 h planning horizon (66 %). However, the average recall for the 1–2 day planning horizon rose to 73.2 %. Together these results indicate that, while there is room for improvement, there is a significant opportunity to take advantage of knowing what a person is currently planning. Furthermore since if a prediction is made it is likely correct, this increases the likelihood that if recommendations are made based on this predicted information there will be a have a lesser chance of decreasing the quality of recommendations.

Related Work

Several works have suggested other factors besides location of a mobile user that may help in determining what is relevant to users, but the vast majority has focused on context related to the immediate next activity. Some works have suggested that in addition to current location, current context refers to data such as time availability, real-time weather or traffic updates and real-time events such as accidents [18–20]. In addition, factors like financial situation and the group of people traveling along with the mobile user have been shown to affect the individual’s real-time travel decision-making [21]. Other works have demonstrated that the makeup of the group involved influences travel decisions resulting from different groups of people having different preferences depending on their backgrounds [22, 23]. Other studies have examined what aspects of context and planning horizon affect the selection of destination/activity for tourists [24, 25].

We see these works as complimentary. Once the type of activity and timing are identified as part of context, as addressed in this work, that enhanced context can be combined with the findings of these related works to better determine how that information should be used to filter the relevance of the results.

Conclusions

This paper presents a novel approach to enhancing user context. While the majority of prior work in context-based recommendation for mobile applications has focused on the user’s immediate situation, this study introduces the idea of extending the user’s context to include what plans are being made beyond just the next location as an aspect of the user’s context. An empirical analysis is conducted of how well predictions can be made of what activity type and how far in the future a user is currently planning with positive results. As this work demonstrates, this can be done with high precision potentially offering a significant opportunity to improve recommendations to mobile users through incorporating the aspect of what is being planning to context-based recommenders.

The findings of this study bring up many additional areas for future study. While using locational context on mobile applications can sometimes be intuitive such as tailoring results to what is nearby; how best to take advantage of a user’s future activity plans that are not necessarily nearby is far less understood. Using mobile planning context to determine what is most relevant poses some very interesting challenges as the most relevant information is likely not just related to the activity being planned, but also must take into account location constraints of the user’s upcoming schedule. Another area for future study is the adjustments in schedule a person makes throughout the day. For instance adding activities to a person’s schedule can often impact other activities such as changing their timing or even the location. Recognizing when these points occur and identifying what information may be useful in adjusting existing activity plans may be beneficial as well.

References

1. O. Ashiru, J. W. Polak, and R. B. Noland, “The Utility of Schedules: Theoretical Model of Departure-Time Choice and Activity-Time Allocation with Application to Individual Activity Schedules,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1894, pp. 84–98, 2004.
2. E. Miller and M. Roorda, “Prototype Model of Household Activity-Travel Scheduling,” *Transportation Research Record*, vol. 1831, pp. 114–121, 2003.
3. M. S. Lee and M. G. McNally, “On the structure of weekly activity/travel patterns,” in *Transportation Research Part A: Policy and Practice*, vol. 37, 2003, pp. 823–839.
4. M. Lee and M. McNally, “An empirical investigation on the dynamic processes of activity scheduling and trip chaining,” *Transportation*, vol. 33, no. 6, pp. 553–565, November 2006.
5. G. Adomavicius, B. Mobasher, F. Ricci, and A. Tuzhilin, “Context-aware recommender systems,” *AI Magazine*, vol. 32, no. 3, pp. 67–80, 2011.
6. L. Stenneth, O. Wolfson, P. S. Yu, and B. Xu, “Transportation mode detection using mobile phones and GIS information,” in *Proceedings of the 19th ACM SIGSPATIAL International*

- Conference on Advances in Geographic Information Systems*, ser. GIS'11. New York, NY, USA: ACM, 2011, pp. 54–63.
7. W. Bohte and K. Maat, "Deriving and validating trip purposes and travel modes for multi-day GPS-based travel surveys: A large-scale application in the Netherlands," *Transportation Research Part C: Emerging Technologies*, vol. 17, no. 3, pp. 285 – 297, 2009.
 8. S. T. Doherty and P. Oh, "A multi-sensor monitoring system of human physiology and daily activities," *Telemedicine and e-Health*, vol. 18, no. 3, pp. 185–192, 2012.
 9. S. T. Doherty, *Transport Survey Methods: Keeping up with a Changing World*. Emerald Group Publishing Limited, 2009, Ch. Emerging methods and technologies for tracking physical activity in the built environment, pp. 153–190.
 10. G. J. Welk, "Use of accelerometry-based activity monitors to assess physical activity," *Physical activity assessments for health-related re- search*, pp. 125–141, 2002.
 11. K. Y. Chen and D. R. Bassett, "The technology of accelerometry-based activity monitors: current and future," *Medicine and science in sports and exercise*, vol. 37, no. 11, pp. S490–S500, 2005.
 12. D. Ettema, T. Schwanen, and H. Timmermans, "The effect of Location, Mobility and Socio-Demographic Factors on Task and Time Allocation of Households," *Transportation: Planning, Policy, Research, Practice*, vol. 34, no. 1, 2007.
 13. A. Mohammadian and S. T. Doherty, "Modeling activity scheduling time horizon: Duration of time between planning and execution of pre- planned activities," *Transportation Research Part A: Policy and Practice*, vol. 40, no. 6, pp. 475–490, Jul. 2006.
 14. S. T. Doherty and E. J. Miller, "A computerized household activity scheduling survey," *Transportation*, vol. 27, no. 1, pp. 75–97, Feb. 2000.
 15. S. T. Doherty, "How far in advance are activities planned? Measurement challenges and analysis," *Journal of the Transportation Research Board*, vol. 1926, pp. 41–49, 2005.
 16. S. Doherty, E. Nemeth, M. Roorda, and E. Miller, "Design and Assessment of the Toronto Area Computerized Household Activity Scheduling Survey," *Journal of the Transportation Research Board*, vol. 1894, pp. 140–149, 2004.
 17. C. Cleverdon, "Evaluation of Tests of Information Retrieval Systems," *Journal of Documentation*, vol. 26, pp. 55–67, 1970.
 18. B. Tatomir, L. J. M. Rothkrantz, and A. C. Suson, "Travel time prediction for dynamic routing using ant based control," in *Winter Simulation Conference (WSC'09)*, ser. WSC'09. Winter Simulation Conference, 2009, pp. 1069–1078.
 19. A. Hinze and S. Junmanee, "Travel recommendations in a mobile tourist information system," in *Proceedings Fourth International Conference on Information Systems Technology and its Applications (ISTA'05)*, R. Kaschek, H. C. Mayr, and S. Liddle, Eds., 2005, pp. 86–100.
 20. L. Vanajakshi, S. Subramanian, and R. Sivanandan, "Travel time prediction under heterogeneous traffic conditions using global positioning system data from buses," *Intelligent Transport Systems, IET*, vol. 3, no. 1, pp. 1 –9, March 2009.
 21. H. Timmermans, *Progress in Activity-Based Analysis*, H. Timmermans, Ed. Elsevier, 2005.
 22. J. Anable, "Complacent car addicts' or'aspiring environmentalists'? Identifying travel behaviour segments using attitude theory," *Transport Policy*, vol. 12, no. 1, pp. 65–78, January 2005.
 23. H. J. Timmermans and J. Zhang, "Modeling household activity travel behavior: Examples of state of the art modeling approaches and research agenda," *Transportation Research Part B: Methodological*, vol. 43, no. 2, pp. 187–190, 2009.
 24. L. Baltrunas, B. Ludwig, and F. Ricci, "Context relevance assessment for recommender systems," in *Proceedings of the 16th International Conference on Intelligent User Interfaces (IUI'11)*. New York, NY, USA: ACM, 2011, pp. 287–290.
 25. S. Choi, X. Y. Lehto, A. M. Morrison, and S. S. Jang, "Structure of travel planning processes and information use patterns," *Journal of Travel Research*, vol. 51, no. 1, pp. 26–40, 2012.
 26. C. Williams and J. Mathew, "An architecture for mobile context services," in *Proceedings of the 8th International Joint Conferences on Computer, Information, Systems Sciences, & Engineering*, 2012.



<http://www.springer.com/978-3-319-06763-6>

New Trends in Networking, Computing, E-learning,
Systems Sciences, and Engineering

Elleithy, K.; Sobh, T. (Eds.)

2015, XIV, 635 p. 471 illus., 315 illus. in color.,

Hardcover

ISBN: 978-3-319-06763-6