Contexto: Lessons Learned from Mobile Context Inference

Moshe Unger
Information Systems Engineering and Telekom Innovation Laboratories at BGU
mosheun@post.bgu.ac.il

Ariel Bar
Information Systems Engineering and Telekom Innovation Laboratories at BGU
arielba@bgu.ac.il

Bracha Shapira
Information Systems Engineering and Telekom Innovation Laboratories at BGU
bshapira@bgu.ac.il

Lior Rokach
Information Systems Engineering and Telekom Innovation Laboratories at BGU
liorrk@bgu.ac.il

Ehud Gudes
Computer Science Department at Ben-Gurion University of the Negev
ehud@cs.bgu.ac.il

UbiComp'14 Adjunct, September 13-17, 2014, Seattle, WA, USA
ACM 978-1-4503-3047-3/14/09.
http://dx.doi.org/10.1145/2638728.2638781

Abstract
Context-aware computing aims at tailoring services to the user's circumstances and surroundings. Our study examines how data collected from mobile devices can be utilized to infer users' behavior and environment. We present the results and the lessons learned from a two-week user study of 40 students. The data collection was performed using Contexto, a framework for collecting data from a rich set of sensors installed on mobile devices, which was developed for this purpose. We studied various new and fine-grained user contexts which are relevant to students' daily activities, such as "in class and interested in the learned materials" and "on my way to campus". These contexts might later be utilized for various purposes such as recommending relevant items to the students' context. We compare various machine learning methods and report their effectiveness for the purposes of inferring the users' context from the collected data. In addition, we present our findings on how to evaluate context inference systems, on the importance of explicit and latent labeling for context inference and on the effect of new users on the results' accuracy.

Author Keywords
Context-aware; Inference; Mobile sensors; Machine learning

Introduction

The emergence and penetration of smart mobile devices gave rise to the development of context aware systems that utilize sensors to collect available data about users in order to improve services. In our research, "context" refers to aspects of the users' immediate location, activity or circumstances. Early context-aware applications used mainly GPS in order to infer users' location [4,5]. While most existing studies attempt to detect location or mobility state [1,2,3], we present a more fine grained approach that expresses specific and relevant contexts to students' daily activities which are more informative than traditional activity recognition contexts [6]. Reasoning of such contexts was inferred only by wearable sensors [7] (e.g. "smoking"). Our contribution is twofold: we present a set of new reported contexts that we were able to infer successfully by means of mobile sensors alone. Those contexts can be used for enhancing the user experience on a wide variety of applications (e.g. context aware recommender systems). Another contribution is the knowledge acquired from the inference process. We present findings about the identification of relevant features for each context and the best machine learning algorithms for the task. Finally, we suggest unsupervised approach to infer latent contexts to deal with the cold start problem.

Contexto Framework and Experiments

We developed Contexto, an application that records 10 different sensors from mobile devices. Our experiment involves measuring various sensors from users' mobile devices, which reported their contexts during their daily activities during the week. Each student was asked to report explicitly on any change of these pre-defined contexts. We asked the students to report on the following events: on campus, in class and interested (or not) in the learned materials, on a break between classes, on my way to…, at home, listening to music, eating, and smoking. We collected 13,484,489 records of sensor information from a group of 40 students over a period of two weeks. Records were grouped according to sessions consisting of 6 seconds of sensor recording. We applied various pre-processing on the raw data such as aggregation, selection, and mappings and generated 22 features such as: location (longitude, latitude), location address, speed, lighting condition, noise level, battery (plugged, level, temperature) and 3D axis accelerometer statistics (average, standard deviation and correlation). We applied and evaluated results of several classification algorithms, including: C4.5 decision tree, Naïve Bayes, k-Nearest Neighbors, Decision Stump, Logistic Regression and Ensemble learning (Bagging and Random Forest). We used the AUC (area under curve) as our accuracy measure.

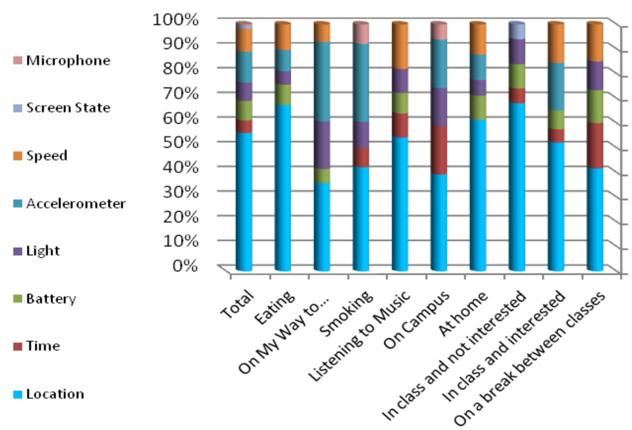

Figure 1. Significant sensors for context inference

Analysis and Results

In order to discover the most indicative features that affected the inference, we ranked them according to the info gain measure as shown in Figure 1. We can notice that in order to infer different contexts, it is required to apply a different set of sensors. For example: the most indicative sensor for "on my way to" is accelerometer.

We used two methods for splitting the dataset into training and test sets: traditional 10 fold cross validation, and time based splitting which considered the data collected over the last week of the experiment as the test set, while the rest considered as training set. The time based splitting represents real-life scenarios in a better way, whereas prediction models are trained based on historical data in order to predict future samples. There are noticeable differences in the accuracy between the two approaches as the time-based splitting approach produces lower accuracy results (as shown in Figure 2). This can be explained by the fact that many following instances were very similar to each other being sampled in small time intervals. However, the time split between weeks included less dependent instances, since they were sampled in different times, thus prediction is more challenging. Moreover, we noticed that the decision tree models (single C4.5 model or Bagging of C4.5 models) produced better accuracy in predicting the users' immediate context and were more resilient to the shift from the to the time-based evaluation.
Real-life context inference solutions might require handling new users whose records are not included in the training set. We examined the feasibility to predict context for new users. We split the data in the user perspective in a "leave one out" fashion, i.e. the test set contained records of a single user while the training set contained all other users' records. Accuracy results of this analysis were naturally much lower, as shown in Figure 2. Since these results were obtained from a relatively small dataset, we believe that adding data of more users will improve the feasibility to infer contexts of new users.

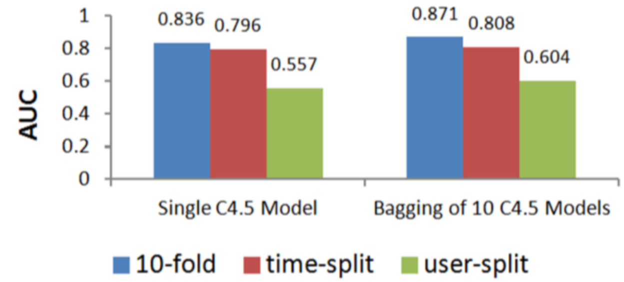

Figure 2. Context prediction accuracy among different models

Additional experiment that examines personal models and included the user id feature improved the accuracy result in 12.35% compared to the baseline model (time based splitting with C4.5 decision model). We believe that the personal models cannot express generic context inference patterns and thus our models did not include any personal features of the users.

Latent Contexts

The models learned in previous analyses require some labeled sensor records from the user in order to produce satisfactory accuracy results. Labeling the records of all users may become impractical in large scale environments. Discovering latent context may address this issue. The main idea is to extract a set of unknown patterns $P = \{P_1, P_2, \ldots, P_k\}$ from unlabeled sensor records S. We suggest using a clustering algorithm (we used k-Means, k=5) on the unlabeled sensor data to cluster the data into groups and treat each cluster-centroid as a latent context. We repeat this process for each user. Although there is no easy way to label these latent contexts, in many scenarios the actual label is not required (e.g. context aware recommendation systems). Figure 3 demonstrates the correlation between user's clusters and the distribution of the actual contexts (e.g. 78% of the contexts in cluster1 were "on my way to”). This encourages us to believe that it is possible to derive latent contexts from unlabeled sensor data. Although this method removes the necessity to label each data record, it still requires generating a separate model for each user.

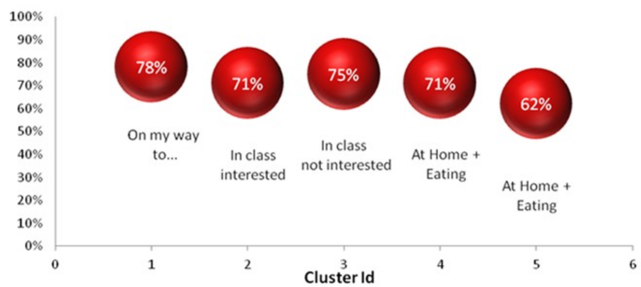

Figure 3. Dominate context distribution by cluster

Conclusion and Future Work

Evaluation of context inference systems is a challenging task especially when contexts are unique and captures daily activities. We showed that various contexts may require different sensors and features; supervised learning (e.g. decision tress) can be applied for this task in different scenarios, however it requires labeled data and a careful attention to the new user problem; finally, we showed that latent context analysis can be applied to address these limitations. Future work may include improving the prediction accuracy by extending the user profile, designing and evaluating additional latent context models (e.g. by deep learning) and finally, evaluating the effect of different context inference models and approaches on services such as context aware recommender systems.

References

[1] Huynh, Tâm, Mario Fritz, and Bernt Schiele. "Discovery of activity patterns using topic models." Ubiquitous computing. ACM, 2008.

[2] Oh, Keunhyun, Han-Saem Park, and Sung-Bae Cho. "A mobile context sharing system using activity and emotion recognition with Bayesian networks." Ubiquitous Intelligence & Computing and 7th ATC, 2010

[3] Borazio, Marko, and Kristof Van Laerhoven. "Using time use with mobile sensor data: a road to practical mobile activity recognition?" 12th International Conference on Mobile and Ubiquitous Multimedia. ACM, 2013.

[4] Cheverst, Keith, et al. "Experiences of developing and deploying a context-aware tourist guide: the GUIDE project." international conference on Mobile computing and networking. ACM, 2000.

[5] Chen, Zhenyu, et al. "Inferring social contextual behavior from bluetooth traces."Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication. ACM, 2013.

[6] Blanke, Ulf, and Bernt Schiele. "Daily routine recognition through activity spotting." Location and Context Awareness. Springer Berlin Heidelberg, 2009. 192-206.

[7] Scholl, Philipp M., Nagihan Kücükyildiz, and Kristof Van Laerhoven. "When do you light a fire?: capturing tobacco use with situated, wearable sensors." Pervasive and ubiquitous computing adjunct publication. ACM, 2013.