Evaluating Trust-Based Fusion Models for Participatory Sensing Applications

(Extended Abstract)

Erfan Davami
Department of EECS
University of Central Florida
Orlando, FL, USA
erfand@knights.ucf.edu

Gita Sukthankar
Department of EECS
University of Central Florida
Orlando, FL, USA
gitars@eecs.ucf.edu

ABSTRACT

Participatory sensing is a specialized form of crowdsourcing for mobile devices in which the users act as sensors to report on local environmental conditions, such as traffic, pollution, and wireless signal strength. This computing framework has great promise as a tool for urban planners, but deploying new applications is a challenge since the overall performance can be sensitive to the specific user population. This paper describes the process of prototyping a mobile phone crowdsourcing app for monitoring parking availability on a large university campus.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

Keywords

participatory sensing; crowdsourcing; trust-based fusion

1. INTRODUCTION

The term "participatory sensing" was originally coined by Burke et al. [1] to describe the use of human-carried mobile devices to create interactive sensor networks. This computing paradigm can be viewed as a special case of crowdsourcing in which the users' current locations are being leveraged, rather than their problem-solving ability. Unlike many software applications, the performance of a participatory sensing framework improves with user enrollment, which poses interesting challenges to the quality assurance and deployment process. Rather than concentrating on stress testing the application with large numbers of users, the question becomes how the system will perform during the initial recruitment phase when the number of participants is low. If the initial rollout is disappointing, a vicious dropout cycle can ensue as users stop using the application which further sabotages the experience of the remaining participants.

This paper presents a case study on the use of agent-based modeling to evaluate participatory sensing applications, specifically a mobile phone crowdsourcing app for monitoring parking lot usage on a large university campus.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Our evaluation focuses on a comparison of trust-based fusion techniques for modeling the users' reliability and aggregating individual reports. For participatory sensing, it is particularly important to make efficient use of every user report. Compared to workers on crowdsourcing platforms such as Mechanical Turk, participants are much less fungible. It is unlikely that another user will be able to report on the same location at precisely the same time, rendering simple majority voting mechanisms less useful. Conversely, accurate user modeling assumes a greater importance since it is more feasible to get multiple reports from a single worker.

In this paper, we treat the user modeling issue as a problem of estimating participant trustworthiness from a limited number of reports. Based on our deployment experiences, there are three main causes for inaccurate reporting: 1) problems with the user interface, 2) spatial confusion about the parking lot location, 3) deceptive behavior from users concealing their secret parking strategies. The experiments in this paper evaluate the use of a simplified trust model that collapses these three issues into a single dimension. The trustworthiness of the reporting user is then factored into the fusion process to aggregate the reports into a cohesive overview of parking lot occupancy.

2. PROPOSED METHOD

Our smart phone app (KPark) allows community members to view the occupancy levels of all parking garages on the University of Central Florida campus; it is freely available on the Apple App store and Google Play.[1] The app (shown in Figure 1) can be used in a passive viewing mode or in an active reporting mode in which the user can optionally provide information about the current occupancy status of parking lot sections through a menu system.

To model the transportation habits of our user population, we re-purposed a Netlogo agent-based urban transportation simulation [2] and supplemented it with our own specialized models for user parking and app usage behavior. The transportation simulation is an activity-oriented microsimulation that generates a population with realistic transportation schedules. We initialized the simulation with data collected from 1000 community members who opted to participate in an electronic survey on their transportation, dining, parking, and scheduling preferences.

The aim of our case study was to evaluate several design decisions for our mobile crowdsourcing app pre-deployment.

[1]https://www.facebook.com/kparkapp

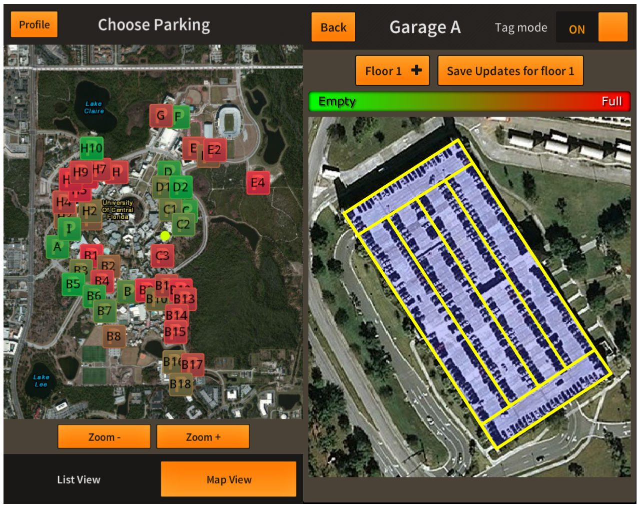

(a) (b)

Figure 1: Potential parking spaces are displayed on an interactive map shown in (a). Red denotes parking sections that are full or close to full; green sections have a higher probability of vacancy. (b) shows a more detailed view of the parking lot. Users have the option of saving updates to the occupancy level for a specific parking section.

One of the most important questions was the choice of trust prediction model. For privacy-preserving reasons, we opted not to use a strategy where we verify the user's location with GPS data. Instead the user's trustworthiness is inferred during a calibration period, when we compare the deviation of an individual user's reports against the average parking lot occupancy based on aggregated data indexed by the day and time. This data serves as a reasonable approximation to doing an actual majority vote across many user reports. If multiple reports from the same user deviate from the aggregated parking services data, it is likely to be the result of user error. This paper evaluates four different strategies for predicting trust: maximum likelihood estimation (MLE), Bayesian, beta reputation, and Gompertz functions. We also evaluate the performance of two data fusion methods for occupancy prediction: 1) weighted average (weight user reports by trust value) 2) max trust (select the report from the user with the maximum trust and ignore the rest).

3. RESULTS

The focus of this case study was to determine which methods would be most useful during the sensitive early deployment phases of the mobile phone app. Hence we specifically conducted experiments to test the robustness of the proposed algorithms against a set of commonly seen early deployment conditions, such as having a low number of total users and a high proportion of untrustworthy users who are unfamiliar with the system. The data used in the evaluation was generated by simulating one semester in the agent-based model, with 21,000 total drivers parking on campus using schedules generated by the activity-oriented microsimulation according to an electronic survey of the transportation habits of 1000 students.

Based on our analysis, we make the following observations.

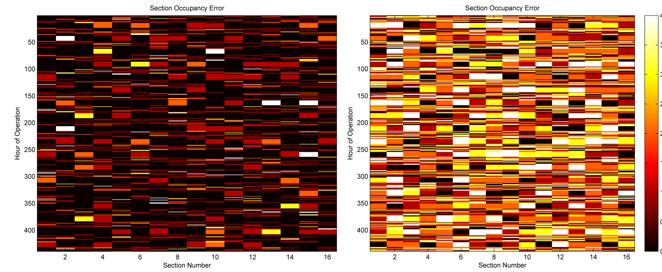

Figure 2: Heatmaps of occupancy prediction errors. Darker areas have less prediction error. The figure on the left (MLE trust prediction, max trust occupancy prediction) shows substantial improvement over the random baseline (right).

1. Having a high percentage of untrustworthy users who are reporting erroneous tags is the most problematic case for occupancy prediction. This indicates that it is important to make sure that the user interface is designed to make the reporting process as error free as possible. The beta reputation system is the most robust system for performing trust prediction with untrustworthy users.
2. Surprisingly, low user adoption is less problematic for occupancy prediction than expected. Applying a discount factor to stale data can improve the trust prediction results when there are few users submitting reports.
3. The rate of improvement for different trust prediction method vs. the number of submitted reports differs. The MLE and Bayesian methods leverage a high number of reports most effectively for increased trust prediction accuracy.

4. CONCLUSION

The aim of this case study was to evaluate the theoretical performance of a set of trust-based fusion algorithms under a range of hostile scenarios. Based on these experiments, we plan to use Bayesian updates for trust prediction, the max trust method for occupancy prediction, and apply discounting to reduce the effects of stale data. Our target recruitment is to have 1% of the campus drivers using the app, making at least one report per week. We plan to add additional functionality to the app, such as car finding, to incentivize the reporting process.

5. ACKNOWLEDGMENTS

This research was supported by NSF IIS-08451 and DARPA D13AP00002 awards.

6. REFERENCES

[1] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. Srivastava, "Participatory sensing," in *ACM Sensys Workshop on World Sensor Web*, 2006.

[2] R. Beheshti and G. Sukthankar, "Extracting agent-based models of human transportation patterns," in *Proceedings of the ASE/IEEE International Conference on Social Informatics, Washington, DC*, 2012, pp. 157–164.