

Spatiotemporal Behavior Profiling: A Treasure Hunt Case Study

Victor de Graaff¹, Dieter Pfoser², Maurice van Keulen¹, and Rolf A. De By³

¹ Department of Computer Science
University of Twente, The Netherlands
v.degraaff@utwente.nl, m.vankeulen@utwente.nl

² George Mason University, Fairfax, VA, USA
dpfoser@gmu.edu

³ Fac. of Geo-Information Science & Earth Observation (ITC)
University of Twente, The Netherlands
r.a.deby@utwente.nl

Abstract. Trajectories have been providing us with a wealth of derived information such as traffic conditions and road network updates. This work focuses on deriving user profiles through spatiotemporal analysis of trajectory data to provide insight into the quality of information provided by users. The presented behavior profiling method assesses user participation characteristics in a treasure-hunt type event. Consisting of an analysis and a profiling phase, analysis involves a timeline and a stay-point analysis, as well as a semantic trajectory inspection relating actual and expected paths. The analysis results are then grouped around profiles that can be used to estimate the user performance in the activity. The proposed profiling method is evaluated by means of a student orientation treasure-hunt activity at the University of Twente, The Netherlands. The profiling method is used to predict the students' gaming behavior by means of a simple team type classification, and a feature-based answer type classification.

Keywords: behavior analysis, behavior prediction, gps data, user generated content

1 Introduction

Trajectories are more than a simple collection of geographical coordinates with timestamps added to the mix. They represent user actions and can, when interpreted properly, lead to an in-depth analysis of behavior and, consequently, user profiling. As a matter of fact, *you are what you “where”!*

Research involving trajectories has so far focussed on data management and data mining aspects at the geometrical levels. Results have led, for example, to improved telematics services using live traffic assessment by means of vehicle tracking and, more recently, map construction algorithms resulting into automatic road network generation and updates (e.g. [1]).

Focusing on the qualitative aspects, semantic trajectory compression methods have been introduced to reduce the size of the actual trajectory data. Essentially these methods rely on decision points, or, landmarks, in combination with movement vectors to

reduce the number of position samples recorded for each trajectory. Based on such landmarks, one would be able to characterize movement at a high level based on traversed landmarks.

The scope of this work is now to identify an even higher level of abstraction and to abstract trajectories into user profiles based on their behavior as derived from movement. Specifically, we derive user profiles through spatiotemporal analysis of trajectory data to provide insight into the quality of information provided by users. The context of this work is the creation of a trajectory analysis component inside a bigger architecture where user profiles and location profiles play an important role, such as the one discussed in [2]. The presented behavior profiling method assesses user participation characteristics in a treasure hunt type event. We propose a method that allows us to map trajectories collected during a treasure hunt to a certain user typology. Our overall method consists of an analysis, and a profiling/prediction phase. The trajectory data is analyzed using a timeline and a stay point analysis, as well as a semantic trajectory inspection relating actual and expected paths. Timeline analysis detects the differences between the users and their changing behavior over time. Stay point analysis determines where teams spent a significant amount of time and helps us in assessing the impact of the environment on user behavior. Trajectory inspection is then used to distinguish between engaged teams and indifferent teams. The specific dataset used captures the answers and spatiotemporal characteristics of 100+ students using the dedicated smartphone application developed for making students acquainted with the city of Enschede. The outcome of the analysis phase is used as input to the profiling and prediction phase. This work proposes two methods, team type and answer type classification, to predict the students' gaming behavior. Team type classification uses two distinguishing features, correctness and distance, to create a total of four profile types. Answer type classification is based on a total of 10 features overall describing the teams and their answers. The generated profiles and their features are used as input to the automated generation of a decision tree to predict answer types. This is especially useful in scenarios where answer correctness is not as sharply defined, such as user generated reviews of products or places. To validate our method, we use a treasure hunt contest as a case study. The data relating to this event was collected using a mobile application. The application was used by several hundred students. This treasure hunt was one of the program elements of the welcome week for new students at the University of Twente, The Netherlands. The students used 132 devices to participate in the *Kick-In Quest* using a specific app on Android and iOS platforms. During the game, GPS data and answers were collected in real real-time. Using this data and applying our profiling methods, we can clearly distinguish characteristic user types and reason about the performance of specific profiles in the game. This final aspect demonstrates that trajectory data can be used to successfully reason about the behavior of users and that *we really are what we where!*

The outline of this work is as follows. related work is discussed in Section 2. Section 3 describes the data collection and treasure hunt game in the process. How this data is then analyzed and how profiles are derived is described in Section 4. Section 5 discusses the team type classification and the answer type prediction mechanism. Finally, Section 6 gives conclusions and directions for future work.

2 Related work

Work on behavior profiling based on GPS data often describes the analysis and prediction of travel patterns as for example described in [3–5]. In this paper however, we focus on the *behavioral patterns* that follow from this movement, but not the actual movement patterns themselves. This was also the focus of Giannotti et al. [6] who use visited regions of interest (such as a railway station, bridge, or a museum) and sequential patterns thereof to describe peoples’ movements. Zheng et al. [7] used a similar, but more formalized approach. Spaccapietra et al. [8] introduced a conceptual model using movement types and visits to specific points of interest to describe the behavior of people in sequential patterns including the movement. Yan et al. [9] introduce a framework for semantic annotation of trajectories using several abstraction layers, and discuss typical challenges when dealing with trajectory data from mobile devices.

Our timeline analysis resembles the approach taken by Guc et al. [10], but contrary to their manual annotation approach, the annotation in our work is done automatically. Our stay point analysis is based on the work by Zheng et al. discussed in [7].

Several other interesting approaches for computing with spatial trajectories are discussed in [11]. Recently, two in-depth overviews of the state of the art in the field of modeling and semantic enhancement of trajectory data were presented by Parent et al. [12] and Jiang et al. [13]. Behavior, as we discuss it in this paper, is defined as *semantic behavior* in [12]: “trajectory behavior whose predicate bears on some contextual data and possibly on some spatial and/or temporal data.”

3 Case Study and Data Collection

The trajectory data used in this work was collected as part of a case study that involved 54 teams comprised of first-year university students, using a total of 132 mobile devices participating in a treasure hunt as part of the new student orientation. In the following, we describe the event during which the data was collected, the technology that was used, we provide a description of the collected data, and describe the pre-processing of the data.

3.1 Event

Every year, new students are welcomed to the University of Enschede with a voluntary, but popular welcome week called the *Kick-In*. As part of the welcome week for new Bachelor students in 2013, the *Kick-In Quest*, as seen in Figure 1, took place on a Saturday morning, from 10:00am until noon. The *Kick-In Quest* manifested itself as a treasure hunt, in which students had to answer questions at certain locations using a mobile app. Points were awarded for (i) answering questions correctly, and for (ii) collected GPS location data. This was explained to the students using information screens in the app. The students worked together in teams, and were motivated to use multiple devices per team to obtain more location points (resulting in more collected tracking data). The awarded amount of points for correct questions was also based on the proximity to the question location, which forced students to move around the town center

of Enschede, even if the answer was already known, or could be found online. Each of the 54 teams that participated had its own designated question sequence, guiding them to 20 locations. These sequences were put together from a total of 24 questions and respective locations. It was not possible to look ahead or change a previously entered answer in the app. In case the app was closed, the app continued where it had previously stopped.

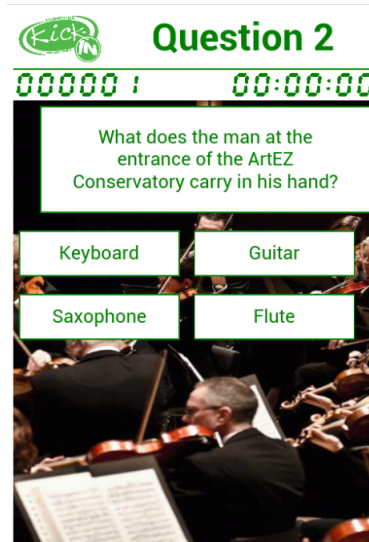


Fig. 1. Screenshot of the mobile application *Kick-In Quest*.

3.2 Technology

The mobile application was developed using the PhoneGap platform [14]. This platform allows app developers to build an application as if it were a web page using HTML5, CSS, and JavaScript. The features of the mobile device, such as the GPS sensor, can be accessed through asynchronous JavaScript calls. PhoneGap supports many different platforms, but for this app only Android and iOS were used.

3.3 Collected Data

During the game, trajectory data was uploaded from the mobile devices to a server as tuples containing the *trajectory id* and all fields of PhoneGap's *Position object*: latitude, longitude, altitude, accuracy, altitude accuracy, heading, speed, and timestamp. Furthermore, *answer data* was uploaded containing the following information: question ID, trajectory ID, answer, latitude, longitude, accuracy, and operating system (i.e., Android, or iOS).

3.4 Data pre-processing

Besides several minor data cleaning tasks, such as resolving differences between the date formats of the different platforms, data pre-processing involved the removal of location references based on cell phone tower locations and trajectory point outliers. Android offers two types of locations: *fine-grained locations* and *coarse-grained locations*. PhoneGap uses the coarse location data when the fine-grained location is unavailable. This results in trajectories as illustrated in 2. Coarse-grained locations can be detected by setting a threshold on the *accuracy value*. Removing all data points for which the accuracy value exceeds 50m has proven to filter out all such coarse points. In addition, by using a realistic threshold for the speed of the participants, extreme outliers were removed from the GPS trajectories as well.




Fig. 2. Trajectory before removal of location references based on cell phone tower locations. The signal keeps jumping back and forth between relatively accurate GPS positioning and less accurate cell phone tower locations.

4 Spatiotemporal Behavior Analysis

Our first goal in this work is to analyze the trajectory data. Here we use three methods to get an understanding of the data and the respective patterns in the participants' behavior. First, we perform a *timeline analysis* to observe the differences between users, and changing behavior over time. Secondly, we use a *stay point analysis* to determine where teams spent a significant amount of time, and to inspect the impact of the respective environment on the participants' behavior. Lastly, we inspect the respective trajectories with respect to the expected trajectories using the question locations.

4.1 Timeline Analysis

The spatiotemporal question-answer behavior results in a specific signature for each group. Our hypothesis is that by quantifying this behavior, we can easily identify more or less successful participants. A timeline analysis is used to relate the behavior of users by means of analyzing the spatiotemporal answering behavior.

The questions in the treasure hunt were location-bound. This means that the students were instructed to answer the question at a stated location, even if the answer would be known or found online. The students were motivated to visit the location by awarding an increasing amount of points based on the inverted distance to that location at the moment of answering the question. Therefore, to score the maximum number of points, the answers needed to be (i) correct, and (ii) answered at the respective location. Combining right and wrong answers with close or distance answer locations resulted in the four different answer types shown in Table 1.

Table 1. Spatial Question/Answer categories.

| | Close | Distant |
|-----------|----------------------|------------------------|
| Correct | Correct Close (CC) | Correct Distant (CD) |
| Incorrect | Incorrect Close (IC) | Incorrect Distant (ID) |

In Figure 3, we provide the timeline analysis graph for our case study. The x -axis represents the time, the y -axis represents the team number. The teams are ordered on the time of their first answer, starting at the bottom. The bar of each team consists of a sequence of answers and the time it took to answer each. The longer each bar, the longer it took to answer. The four colors represent the four answer types. CC answers are illustrated in green, CD in blue, IC in yellow, and ID in red. For each answer, the start of the bar represents the moment at which the question is presented on the screen, and the end of the bar the moment the answer is entered in the application (with a small degree of freedom for visualization purposes). For teams that used multiple devices, only the first given answer is considered, because the mobile application often hints at the correct answer after an answer is provided. This influenced the answering behavior on the other devices.




Fig. 3. Timelines of teams. Colors indicate answer types. Highlighted teams are discussed in detail below.

The green (CC) answers, which are correct in both actual answer and location, are more prominent towards the beginning of the game. CC answers take a substantial amount of time to obtain, since the students have to move to the right location. Therefore, a large part of the graph is green. Therefore, this visualization type is very suitable to get insight into the amount of time spent on certain behavior. If additionally it is of interest to see the ratio of behavior occurrences, other techniques can be used, such as the pie chart of answer types in Figure 4. A closer look at Team 33, for example, shows a very green timeline, while only 9 out of the 20 answers are actually CC answers. The team with the highest percentage of CC answers was Team 16, with 12 out of the 13 answers correct and at the right location. Not only did this team have the highest percentage, the score of 12 was a tie for the highest number of CC answers. A deeper analysis of the data for this specific team gave us insight in their behavior. The reason that this team could achieve this high number of CC answers without skipping questions, by providing a CD or ID answer in between, was that it participated with four devices, and split up into a walking group with two devices, and a cycling group with two more devices. Judging by the perfect timing of which devices were used to answer the question first, we can conclude that they must have communicated about this throughout the game.




Fig. 4. Answer type frequencies. Only the first answer for each team has been taken into account, in case a team participated with multiple devices, since the explanation in the next screen often revealed the correct answer.

From this timeline, we can detect several *team behavior types*. Some teams were indifferent about the outcome of the entire game, which is reflected by a late starting time, short answering times, and many incorrect answers (e.g., Teams 51 and 52 in Figure 3). Other teams were indifferent about answering the questions at the right location, while still trying to answer questions correctly (e.g., Team 2). Teams like Team 1 just wanted to finish the game. Finally, teams like Team 16 were dedicated throughout the entire game. From those teams that played the game all the way until the end, there is another behavior pattern that can be detected. They rushed towards the end of the game. This can be seen by the many CD and ID answers towards the end (e.g., Team 33). Another important observation is that there were barely any teams that were consistent in their answer types, except for team 16 that has a nearly entirely green timeline. A deeper analysis of this observation was done using the bar chart in Figure 5. This chart,

combined with the knowledge of question locations from Figure 6, shows that remotely located questions were far less popular to be answered near the stated location.

In addition to the timeline analysis, different visualizations are useful to compensate for the coloring bias that is induced by the fact that some behavior requires more time than other behavior. In the treasure hunt case for example, CC answers require more traveling time than CD answers. This results in an overall “greener” graph, while the pie chart in Figure 4 shows that CD answers (blue) are actually more common.




Fig. 5. Answer type frequencies by question ID. Questions 21 and 22 were located remotely, and therefore had a large amount of CD (blue) and ID (red) answers.

4.2 Stay Point Analysis

Stay point analysis is another form of a spatiotemporal analysis that provides insight into the most significant locations of the people’s stop-and-go behavior. Stay points are defined as points where the speed was below a certain threshold for a respective time period. This analysis is uncoupled from that what the students were asked to do (answering questions), and gives us more insight into what they actually did besides playing the game.

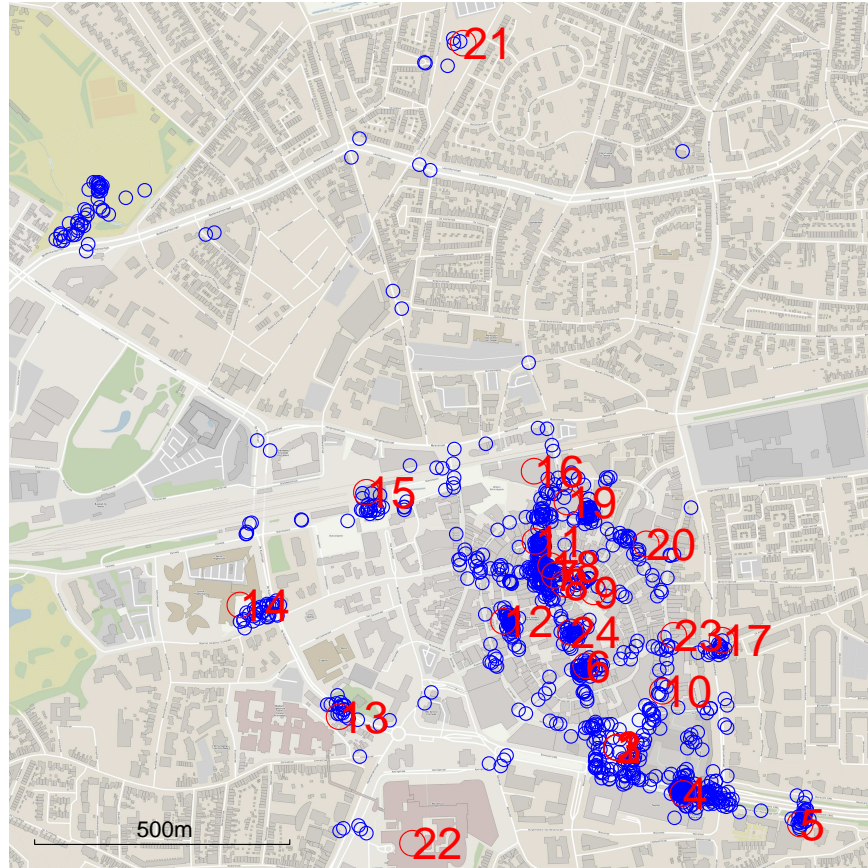


Fig. 6. Stay points derived purely from GPS data. Blue circles indicate detected stay points, the red circles and numbers indicate the question locations. The park in the northwest corner is the location of the event that took place afterwards. Question 4 (in the southeast) was in a supermarket.

In our case study, we use a speed threshold of $1m/s$ and a stay threshold of $30s$. To overcome the problem of GPS signals bouncing around when no, or little movement is observed, we did not use the typical sampling rate of $1Hz$, but compared the position with to a GPS signal sample of $5s$ ago (under-sampling).

The results of this analysis are shown in Figure 6. The red circles are the actual question locations. Blue circles are used to visualize the stay points after clustering them on a trajectory basis. Multiple blue circles located close to each other therefore represent multiple trajectories slowing down significantly in that area. For visualization purposes, the radius size of the question location circles has been reduced to $25m$, as opposed to the $100m$ used to distinguish between “close” vs. “distant” answers in the timeline analysis.

Several groups of stay points can be detected. The cluster of points in the *upper left corner* of the map reveals the location of the next program element for the students. Several teams went to this location early, or left the app running while the game was already over. The game was initiated near the center of the map. This is where the students were supposed to gather, and obtained a team code to start the app. Question no. 4 was to be answered inside a supermarket. The cluster of stay points in the lower right corner of the map surrounds that location. Judging by the time spent near this supermarket, several students went into the supermarket for more than just the answer of the question. The slow movement afterwards is probably caused by the consumption of snacks and drinks afterwards.

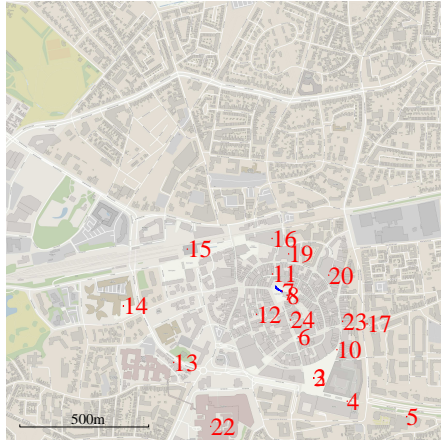
4.3 Trajectory Inspection

To gain further insight into the behavior of individual teams, we carried out a visual inspection of trajectories. This revealed some challenges as several teams had turned off their GPS tracking between answering questions. Also, we could cluster the trajectories into four types, (i) barely moving, (ii) not leaving the center of the town, (iii) leaving center of town once, and (iv) moving around.


Figure 7 shows examples of different trajectory types. However, no clear correlation could be found between the correctness of answers and the movement patterns of the teams. For example, Team 25 for example answered 10 questions, all correctly, without moving substantially.

5 Behavior Prediction


In many scenarios of behavior analysis, the correct behavior is unknown. An example scenario is the analysis of reviews and ratings of geo-referenced objects, such as restaurants. The ability to distinguish serious users from the less serious ones, allows us to predict the quality of their content. In the case study, we have this information on answer correctness, and used this as a ground truth to validate two prediction mechanisms built using our approach. The first one, a simple team type classification, is mainly based on the outcome of the timeline analysis. For the second one, a machine-learning approach to predict the answer types, we have used the outcome of all three analyses.



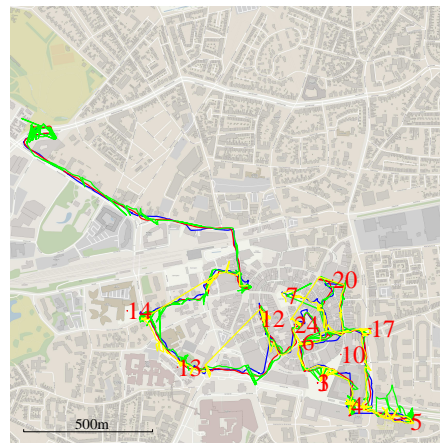
(a) Type 1 - Team 51 barely participated and answered the questions from a cafe near the start of the event.



(b) Type 2 - Team 1 participated, but did not leave the center.



(c) Type 3 - Team 14 actively participated and moved somewhat around to answer questions.



(d) Type 4 - Team 16 actively participated, and visited even remote question locations.

Fig. 7. Raw trajectories of several teams. Each device (per team) has an own color for the trajectory. The red numbers indicate the question locations for that specific team. Teams were unable to look ahead to the next questions.

5.1 Team Type Classification

A simple team type classification categorizes teams into four groups using (i) start time and (ii) average speed of each team. Table 2 shows the resulting four types. Teams that started within 2000s of the first starting team were classified as early starters. The cut-off speed between the *slow and fast teams* was set at $1.2859m/s$, the median speed of all teams. For teams that participated with multiple devices, the entire travelled distance was divided by the total amount of time the devices submitted a GPS signal. Teams that answered less than 10 out of the 20 questions were not taken into account for this classification.

Table 2. Team characteristics

| | High average speed | Low average speed |
|------------------|--------------------|-------------------|
| Early start time | Serious | Get-It-Over-With |
| Late start time | Rushed | Indifferent |

The radar charts in Figure 8 show the distribution of answer types for each team type. Each of the four axes represents the corresponding answer type.

Again, it can be seen that IC (incorrect + close) answers are very uncommon (none of the teams had more than 20% IC answers), which also makes them hard to predict. CC answers (correct + close) are much easier to predict and are especially common among the teams classified as Serious. CD (correct + distant) answers are especially common among Get-It-Over-With teams. ID (incorrect + distant) answers can be found primarily among Indifferent teams, and to a lesser degree among Rushing teams.

5.2 Feature-based Answer Type Classification


Rather than classifying team behavior, we, in the following, examine answer type behavior. We created nine features to describe the team's behavior in general, and one feature to compare the answer to the answer of other teams:

On a team basis we record the following features.


1. start time,
2. average speed of all the team's trajectories (total distance divided by the total time),
3. maximum covered distance,
4. time spent using the application,
5. number of devices,
6. median distances travelled between answers,
7. median time elapsed between answers, and
8. number of stay points

To relate the response to other teams, we also record


9. whether or not the answer is prevalent answer of all teams.




(a) Serious teams have a strong tendency towards CC answers.



(b) Get-it-over-with teams have a strong tendency towards CD answers.



(c) Rushed teams typically have a high ratio of CD and ID answers, and very little CC answers.



(d) Indifferent teams have a strong tendency towards ID answers.

Fig. 8. Radar charts of answer types for each team type. Each of the four axes represents the corresponding answer type.

We use a brute force approach to find the combination of features that leads to the best F -measure (the harmonic mean of precision and recall as commonly used in Information Retrieval) for answer correctness prediction. The *best combination* of features to predict CC answers turned out to be a combination of the (i) start time, (ii) the duration, (iii) the median distance between answers, and (iv) whether or not the answer is the prevalent answer of all the teams.

To validate the results of this method, we carried out a 10-fold cross validation. Compared to the baseline method, which assumes all answers are CC answers, this prediction mechanism performs well in precision (0.5369), decent in recall (0.5321) and also performs slightly better in F -measure (0.5274) than the baseline method. Since the majority vote is a very strong feature for answer correctness prediction, results are even better when we predict ID answers, with an F -measure of 0.8686.

This section has shown two methods of how to use aspects (features) of “rich” trajectory data to predict user behavior. While the results are encouraging, we will in subsequent work experiment with larger datasets to verify these findings and improve on our methods.

6 Conclusions and Future Work

Trajectory data is so more than a temporal sequence of position samples. Combined with the right metadata, as in this case question/answer behavior in a treasure hunt, it represents rich spatiotemporal data that can be used to analyze user behavior, which is a first step towards the creation of user profiles. This work has shown how behavior patterns can be detected from trajectory data using several types of analyses. Although these methods have been derived in the context of a treasure hunt scenario, they are generally applicable for the analysis of “rich” trajectory data. The stay point analysis can be used to detect station locations in commuter data, taxi stops, traffic lights, and traffic bottle necks. The timeline analysis is useful for contrasting the behavior of people. In the specific visualization, the colors of the bars can represent other classification technique, such as for example sentiment analysis. Using the various analysis methods, this work demonstrates user behavior prediction based on various feature sets of the semantic treasure hunt trajectory data. Experimental evaluation has shown that indeed user behavior is codified in this trajectory data.

We can give the following directions for future work. Recently, we collected a new trajectory set using a similar setup, but with less potential for game-breaking behavior. This new data set will be used to detect Point-of-Interest visits from trajectory data, which can be used for building user profiles at another abstraction level. Similar to the virtual disc around the question locations in this paper, induced by the 100 meter vicinity criterium, *Polygons-of-Interest*, as discussed in [15], can then be used to detect the presence in a pre-defined area. Our goal will be to derive automatic methods for the profiling of users based on their rich trajectory data. Such profiling approaches will be especially helpful when trying to assess movements of large groups of people, e.g., urban commuting behavior, tourism, and large scale events.

References

1. M. Ahmad, S. Karagiorgou, D. Pfoser, and C. Wenk, "A comparison and evaluation of map construction algorithms using vehicle tracking data," *GeoInformatica Journal*, 2015. in press.
2. V. de Graaff, M. van Keulen, and R. A. de By, "Towards Geosocial Recommender Systems," in *4th Intern. Workshop on Web Intelligence & Communities (WI&C 2012)*, Lyon, France, ACM, April 2012.
3. L. Liao, D. J. Patterson, D. Fox, and H. Kautz, "Learning and inferring transportation routines," *Artificial Intelligence*, vol. 171, no. 5, pp. 311–331, 2007.
4. Y. Zheng, Q. Li, Y. Chen, X. Xie, and W.-Y. Ma, "Understanding mobility based on gps data," in *Proceedings of the 10th international conference on Ubiquitous computing*, pp. 312–321, ACM, 2008.
5. A. Monreale, F. Pinelli, R. Trasarti, and F. Giannotti, "Wherenext: a location predictor on trajectory pattern mining," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 637–646, ACM, 2009.
6. F. Giannotti, M. Nanni, F. Pinelli, and D. Pedreschi, "Trajectory pattern mining," in *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Jose, California, USA, August 12-15, 2007*, pp. 330–339, 2007.
7. Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, "Mining interesting locations and travel sequences from GPS trajectories," in *Proceedings of the 18th international conference on World wide web*, pp. 791–800, ACM, 2009.
8. S. Spaccapietra, C. Parent, M. L. Damiani, J. A. de Macedo, F. Porto, and C. Vangenot, "A conceptual view on trajectories," *Data & knowledge engineering*, vol. 65, no. 1, pp. 126–146, 2008.
9. Z. Yan, D. Chakraborty, C. Parent, S. Spaccapietra, and K. Aberer, "SeMiTri: a framework for semantic annotation of heterogeneous trajectories," in *Proceedings of the 14th International Conference on Extending Database Technology*, pp. 259–270, ACM, 2011.
10. B. Guc, M. May, Y. Saygin, and C. Körner, "Semantic annotation of gps trajectories," in *11th AGILE International Conference on Geographic Information Science*, 2008.
11. Y. Zheng and X. Zhou, *Computing with spatial trajectories*. Springer, 2011.
12. C. Parent, S. Spaccapietra, C. Renso, G. Andrienko, N. Andrienko, V. Bogorny, M. L. Damiani, A. Gkoulalas-Divanis, J. Macedo, N. Pelekis, *et al.*, "Semantic trajectories modeling and analysis," *ACM Computing Surveys (CSUR)*, vol. 45, no. 4, p. 42, 2013.
13. S. Jiang, G. A. Fiore, Y. Yang, J. Ferreira Jr, E. Frazzoli, and M. C. González, "A review of urban computing for mobile phone traces: current methods, challenges and opportunities," in *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing*, p. 2, ACM, 2013.
14. S. Allen, V. Graupera, and L. Lundrigan, "PhoneGap," in *Pro Smartphone Cross-Platform Development*, pp. 131–152, Springer, 2010.
15. V. de Graaff, R. A. de By, M. van Keulen, and J. Flokstra, "Point of interest to region of interest conversion," in *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL GIS 2013)*, Orlando, FL, USA, (New York), pp. 378–381, ACM, November 2013.