Semantic Annotation of GPS Trajectories

Baris Guc*, Michael May**, Yucel Saygin*, Christine Körner**

*Sabanci University, Orhanlı-Tuzla, 34956 Istanbul, Turkey, {barisguc@su.sabanciuniv.edu, ysaygin@sabanciuniv.edu}

**Fraunhofer Institut Intelligente Analyse- und Informationssysteme, Schloss Birlinghoven, 53754 Sankt Augustin, Germany, {christine.koerner, michael.may@iais.fraunhofer.de}

Abstract. The analysis of mobile behavior is elementary for applications such as city planning, optimization of public transport or mobile communications. However, many applications in the mobility domain require a semantic interpretation of movement information. While physical trajectories can readily be recorded using Global Positioning System (GPS) technology, the semantic interpretation of the data is still a great research challenge. However, GPS trajectories can be utilized to facilitate the manual semantic annotation process, and thus render tedious interviews and manual mobility records unnecessary. In order to realize this schema in practice, two requirements are necessary. First, a conceptual annotation model, and second a graphical user interface which takes advantage of the rich information hidden in physical trajectory data. In this paper, we introduce an extensible trajectory annotation model, which is oriented on the notion of episodes and allows a clear separation of semantic and physical trajectory information. We implemented our model and developed a program to support trajectory annotation independent of the recorded location. We present our program architecture and various features that facilitate the annotation process.

1 MOTIVATION

The analysis of mobile behavior leads to many instructive insights about the habits of a city's or a country's population. In Germany, for example, the Federal Ministry of Transport, Building and Urban Affairs commissions in regular intervals of several years a nationwide survey concerning the day-to-day travel behavior of Germans [BMVBS, 2007]. The study evaluates, among others, mobility data with respect to travel distance, the means of transportation and the purpose of traveling. This kind of analysis requires a physical as well as a semantic description of a person's movements.

Physical and semantic trajectories can be obtained in different ways. A first approach is interrogation, where people recall their daily movements and relate them to geographic space. However, this approach is error-prone as people may forget parts of their movements, and need to estimate the respective times. Also, concurrent logging does not solve these problems, as people often forget to take notes. The second approach is benefited by the rapid development of Global Positioning System (GPS) technology. GPS devices can be utilized to record the movements of a person with high accuracy. The data consists of tuples of the form (time, coordinates) and usually follows a sampling interval of one second. After data cleansing, high quality information about a person's physical movements (e.g. location, velocity, distance of travelling) can be obtained. However, GPS trajectories lack semantic information. The home or working place of a person cannot simply be extracted from the data. Furthermore, the purpose of a given movement is unknown as well as the engaged means of transportation. Thus, GPS trajectories provide rich geographic information, but lack semantic information for analysis of mobile behavior.

Previous work in transportation research has addressed the task to derive semantic information from GPS trajectories using the trajectories themselves as well as further background knowledge. In particular, Wolf and Wolf et al. [Wolf, 2000; Wolf, 2001] evaluate the possibility to replace traditional travel diaries with GPS recordings, focusing on the derivation of trip purposes. In a first step, they detect separate trips and assign land use information to the ends of the respective trips. A previously defined relationship between land use and trip purpose along with trip destination points and further information, such as arrival time, are then used to derive the purpose of a trip. A similar approach is followed by Axhausen et al. [Axhausen, 2003]. They assign purposes to trips using land use information as well as personal information about a traveller's home or work address. However, annotation extraction is a complex task and the presented approaches are limited in several aspects. First, the data collections are restricted to periods of private vehicular movement. As a result, the heuristics used for trip segmentation cannot be applied in general. Second, only few categories for trip purposes are distinguished, which do not suffice for detailed studies of mobile behavior. Third, both researches utilize land uses, which may lead to ambiguous results for mixed parcels, as for example shopping malls containing shops, food stores and various service providers. In addition, in some cases a trip can have a purpose irrelevant to its destination. For instance, skating or walking for pleasure might not have a destination, or may return to the start location as the activity is not done in order to reach some destination.

We believe that the automatic annotation of GPS trajectories can further be advanced by methods from the field of data mining and machine learning. However, this requires the collection of an adequate set of training data. Many studies have shown the inconsistency between GPS trajectories and travel diaries [Stopher, 2007; Zmund, 2003]. However, provided with a physical trajectory in geographic and temporal context, a person will be well able to recall all daily activity and to complete the semantic information of the trajectory in a more accurate way. We therefore developed a software system to visualize, annotate, and store GPS trajectory data. The system follows a flexible conceptual trajectory model and utilizes Google Maps [Google Inc, 2007] to visualize trajectories in detailed geographic context at nearly all locations on the Earth.

In Section 2 we introduce the underlying conceptual trajectory model. Section 3 describes the system architecture and annotation workflow. In Section 0 we present various functionalities of our program interface that facilitate the annotation process. We conclude the paper with a short summary and future work.

2 TRAJECTORY ANNOTATION MODEL

When asked about our movements of the previous day, we will probably state to which places we have been on what purpose, and maybe reveal whether the trip was made by car or public transport. The information we give is aggregated and semantic. Geographic locations are interpreted with names, for example, on a shopping trip we are more likely to refer to the type or name of a visited shop than to its address. Similarly, the concept of an annotation model must allow aggregating the movement information and accounting for the underlying purpose of a trip. Aggregation can be achieved by partitioning the movement into a sequence of homogeneous sections with respect to the type of visited location (e.g. home, work, shopping) or mode of transportation (e.g. by car, bus, foot). The purpose of most trips corresponds to the destination of the trip. However, a clear distinction between movement undertaken to reach some destination and movement where the purpose of a trip lies in the trip itself needs to be made. Before we present our annotation model in more detail and show how these several aspects are handled, we will review related approaches in literature.

Current approaches analyze physical trajectories in order to extract meaningful information from the data itself. Therefore, trajectories are split into subsections according to physical characteristics. Mountain and Raper [Mountain, 2001] introduce the concept of episodes. They denote an episode as "a discreet time period for which the user's spatio-temporal behavior was relatively homogeneous". Homogeneity in movement thereby refers to speed, sinuosity, spatial range and freedom of movement

[Dykes, 2003]. Additional characteristics for the description and extraction of episodes are summarized in [Laube, 2006]. The underlying idea of this concept is that homogeneity in physical movement implies homogeneity in semantics. Therefore, sudden changes in movement characteristics identify possible break points for the interpretation of trajectory data. A second line of research follows the concept of stops and moves [Spaccapietra, 2007; Alvares, 2007]. Basically, a stop is defined as the visit of an (application specific) area for at least some minimum period of time. The remaining parts of the trajectory form the moves. This approach is advantageous when the application focuses on location-based activities. For example, the trajectory of a tourist could be analyzed with regard to all points of interest within a city. In addition, stops of various duration and extent can be identified algorithmically from the physical trajectory without further background knowledge [Palma, 2008]. In this way, waiting periods in a traffic jam or at an intersection can be identified. However, this approach concentrates on physical characteristics of a trajectory. On one side, semantic stops can occur without appearing in the data. Imagine, for example, posting a letter on the way to work. The moment of action is very brief and likely to drown in background noise. On the other side, not every stop in the physical trajectory possesses an (application dependent) interpretation. Transferring the concept of stops and moves to the semantic level is thus likely to cause problems, as it mixes physical activity and inactivity within these concepts.

For the annotation of trajectories, we developed a semantic model which follows the concept of episodes. Our model provides two kinds of annotation elements: episodes and trips. Episodes partition a trajectory and describe semantically homogeneous sections of the trajectory. Trips are defined as groups of a sequence of episodes on a higher semantic level.

We regard homogeneity of episodes with respect to the purpose of an action and the mode of transportation. However, these characteristics can be adapted to fit application needs. The mode of transportation is an inherent characteristic of human motion. Regardless of whether movement occurs or not, each moment in time of a geospatial lifeline can be classified according to the means of transportation, as for example by foot, bike, car or plane. Most of our daily journeys are undertaken to reach some destination where we have a specific action in mind. Our model therefore provides different categories of activities, as for example “work”, “shopping” or “education”, and the category “transition” for means of conveyance. However, the purpose of a journey may also lie in the journey itself, for instance when hiking or roller skating. In this case, the purpose of the episode is not mere transition, but may fall into a category named “recreation”. Thus, the semantic characteristics of an episode are clearly separated from the underlying physical movement. We will underline this aspect in another example. Imagine that a taxi driver picks up a passenger from the airport. From the passenger’s point of view, the purpose of the drive will be “transition”. From the taxi driver’s point of view, it will be “work”.

In order to bundle episodes on a higher semantic level, we provide the concept of a trip. A trip consists of a sequence of episodes which pertain to a common aim. For instance all episodes on the way to work can be grouped into a common trip (Figure 1). So far, we considered trips only at one level of hierarchy. However, depending on the application, grouping of episodes over several hierarchical layers may be desirable.

We applied our model in practice to annotate GPS trajectories that were collected over several days. In our model, we assume that episodes partition the surveying interval and leave no gaps in the annotation. However, gaps in the physical trajectory can easily occur due to the obstruction of GPS signals, malfunctioning devices or simply the forgetfulness of users. In these cases, the annotation process does not depend on physical evidence, but can be continued based on a time bar. We prefer this approach to the interruption of annotation, because although there is no GPS data, the user might still be able to annotate that period. For instance, a user could see in the map that she entered her place of work. Next, there is no data until she exits the building as the building blocks the GPS signals. In this case, the user can identify the period which she spent at work, and can annotate the data accordingly. On the other hand, a user might not always recall the purpose or exact times of a period without data. For such cases, the episode can be annotated as “undefined”.

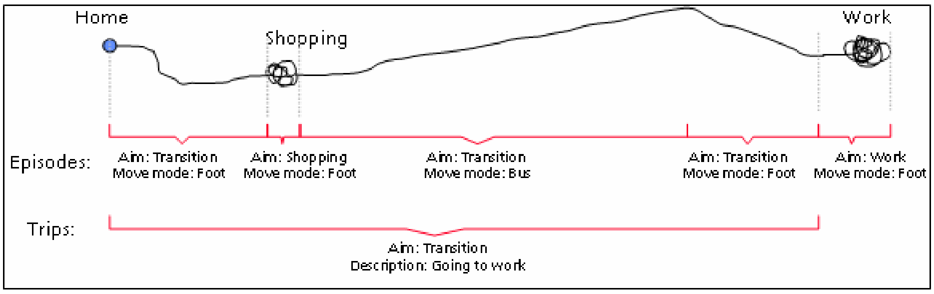

Figure 1: Example annotation

3 WORKFLOW AND SYSTEM ARCHITECTURE

Workflow

The general workflow of the annotation process is shown in Figure 2. After the data has been downloaded from the GPS logging device, it can be fed into the trajectory annotation program. The physical trajectories are displayed in geographic context through an interface to Google Maps. The advantage is that the annotation process becomes independent of location. The user does not need to provide a map of her city or country, and she can display trajectories of almost any location on Earth.

After the data has been annotated, the physical trajectories and their annotations can either be stored in a file or an Oracle database. The database allows a convenient collection of the data from all users and ensures the preservation of privacy. Trajectory and annotation data, once stored in a file or the database, can be re-imported and re-annotated anytime.

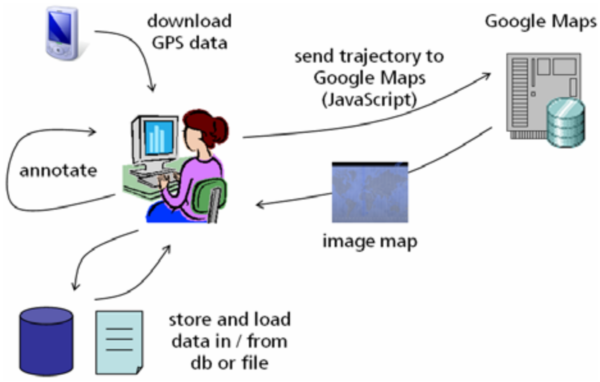

Figure 2: Annotation workflow

Architecture

The trajectory annotation tool (TrajAnn) has been developed as a desktop application using the Java programming language [Sun Microsystems, 2007]. Its architecture consists of three layers as shown in Figure 3. The first layer handles the storage of the trajectory and annotation data. A file handler component reads trajectory data from the original .dbf format which our GPS logger records. It can also export trajectory and annotation data in different file formats. The other component handles the database connection to store and retrieve user information, trajectory and annotation data from a database. This feature is developed for Oracle 10g.

The second layer controls the program flow. This layer receives the raw trajectory and annotation data from the first layer. It prepares the raw trajectory data for annotation and controls the creation and modification of annotation data. The animation component accesses trajectory data and controls the time shift with which the data will be displayed on the interface.

The third layer is the interface layer which consists of graphical user interface (GUI) components. The purpose of this layer is the visualization of physical trajectories in order to facilitate the annotation process. The interface layer receives the raw trajectory and annotation data from the second layer. The first component is a browser which displays Google Maps for the embedding of trajectories in a geographic context. The overlay of trajectory and map data can be controlled by executing JavaScript codes through an application programming interface (API) provided by Google [Google Maps API, 2007]. We developed a class which generates the required JavaScript code dynamically according to visualization needs, and execute these scripts on the browser component. The second component is a timeline which is used to annotate the data and to display various information about the physical trajectory. The timeline is developed on a component of the Jaret project [Jaret timebars component, 2007], which basically shows intervals in a Gantt chart. This component has been modified for our application by changing the behavior and control of the intervals, and by adding an animation feature. Several other Java Standard Widget Tool (SWT) components are used for displaying information about GPS data points and controlling the annotation and animation.

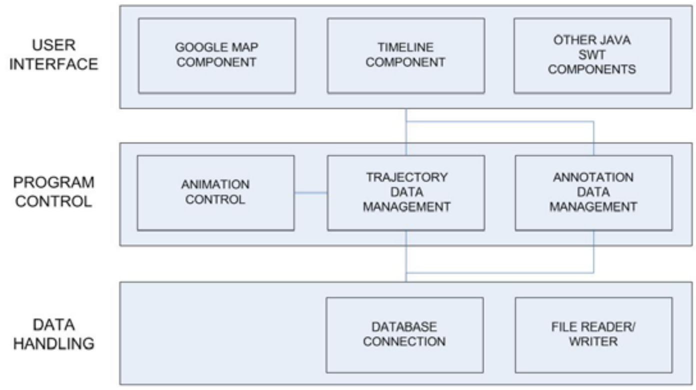

Figure 3: System architecture

4 INTERFACE FUNCTIONALITY

In this section we describe various features of our program which facilitate the annotation process of trajectories. These features include the display of trajectories in their spatial and temporal context, trajectory animation and the exploitation of a user's favourite places.

Visualization of GPS Trajectories

The fundamental requirement of the interface is to display raw GPS data in its geographical context. It helps the users to recall their trajectories and the semantics associated. In our application, the map's level of detail is very important. For proper orientation, it should display the street network with names of streets and important places. Another requirement is that the map is not limited to a specific area. On one hand, users may travel to another country during vacation or work. On the other hand, we intend our tool to be applicable on any collection of trajectory data, independent of location.

According to these requirements, we have decided to use Google Maps [Google Inc, 2007] in our application. Google Maps provides a map service with a high level of detail for many countries. It offers both satellite images and maps of the street network, and operates worldwide.

Google Maps also provides a zoom feature, which is necessary for the annotation process. The user can zoom out to take a broad view of the map and the trajectory. Thus, it is possible to view the start and end points of the trajectory and get an overview of the route the user has taken. When the map is zoomed in, details of the map and trajectory can be observed. For instance, a user can see the point where she left her home, which implies a semantic change in the trajectory. Also, the user can observe the streets she traveled and the places she visited (e.g. shops). In this way, a very accurate annotation can be performed by the user.

Another useful feature that Google Maps provides is the addition of overlays to the map. Several different types of geometric objects can be added to Google Maps as overlays with different settings such as width, color or opacity. These objects include points, polylines and polygons. In our application, we use points to mark the positions of the start and end of a trajectory as well as the current position selected on the timeline as explained in Section 4.2. The trajectories are displayed using polyline overlay. Finally, we use polygon overlay to mark a user's favorite places as explained in Section 4.4.

Display of Temporal Trajectory Aspects

Although the map visualizes the geographical positions of a user's movement, it is insufficient to display the temporal aspect of the movement. The annotation process is performed by segmenting the trajectory into several time periods with different semantics. Hence it is necessary to present the temporal aspect of the data to the user.

For this purpose, we implemented a timeline bar in our interface as shown in Figure 4. While the timeline bar displays various aspects of the data including the dates, it handles also annotation operations. The timeline covers the whole time period of the GPS log data. Similar to the map in our application, the timeline has a zoom function. It is possible to zoom out to get an overview of the data, for example, observing the start and end time of the data recorded the previous day, or the distribution of records recorded during the last week. The timeline can also be zoomed in to view details of the data such as the accuracy of a point at an exact date.

The timeline bar is divided into five rows which have different functions. The first three rows give information about different aspects of the data while the last two rows are used for annotation. The first row displays the horizontal accuracy of the data. In the application, the user can choose an accuracy threshold. The accuracy of the GPS data is compared with this threshold, and individual data points are marked in different colors as valid or invalid position data. In the second row, the mobility of a user is displayed. The GPS devices contain an additional mobility sensor, which records whether the device is actually moved or stays at a fixed position [MGE DATA, 2007]. This information helps the user to differentiate traveling episodes from episodes where she has reached some destination. The third row is used to display intervals in which the user visited her placemarks, which are explained in Section 4.4.

The last two rows in the timeline are used for the annotation process. In these rows, the annotation elements are visualized as intervals. The forth row displays the episodes, and the fifth row displays the trips. Unlike the first three rows, these intervals can be edited by the user. The annotation is done by resizing, merging and dividing episodes, and by forming trips over a sequence of episodes. Different attributes of annotation elements, such as aim and movement mode, can be set for every individual annotation element by selecting it on the timeline. The values for aim and movement mode are provided by a predefined list, which can be adapted to the need of application. This supports the consistent annotation among several users. Annotation elements are colored according to an attribute the user is free to select. For instance, episode intervals which have the same aim can be colored alike.

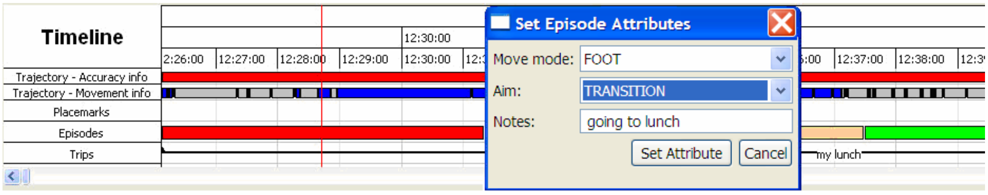

Figure 4*:* Timeline component of the interface

Trajectory Animation

Animation is another feature of TrajAnn that facilitates the annotation of trajectories. By viewing the trajectory on a map, it is hard for a user to understand the direction of the movement or the speed at a single point in the trajectory. However, if the trajectory is displayed and the movement animated, the user can easily include the temporal aspect of the data in her considerations. She can identify periods with slow or fast motion, or without motion, as well as the direction of her movement. The latter aspect is especially helpful, if parts of the trajectory overlap. This happens regularly, for example, on the user's way to and from work.

In our application, the map and the timeline bar are synchronized. During the animation of the movement, the indicator on the timeline bar is advanced with respect to the user's position on the map. We implemented the animation features similar to a video playback where the user can play, pause or seek the data. It is possible to skip to another date by selecting it on the timeline. In this way, the user can easily reach any segment of trajectory she is interested in. In addition, we implemented a variable playback speed. This is essential, as the annotation in real-time motion would be too time-consuming. It has the further advantage that the user can adapt the recorded speed to her needs.

Placemarks

Every person has some places that she visits frequently, such as her home or place of work. When a trajectory intersects these places, it is likely that all points inside this area form an episode of their own. Therefore, the annotation process can be facilitated by an automatic identification of those trajectory parts. In our application, we implemented a placemark feature which enables the user to specify her favorite places, and calculates all intersections along the trajectory.

In order to set her placemarks, the user can simply mark the boundary of some area in the map. The placemarks are saved in the database along with a name and place type. The area itself is stored as a spatial object of type polygon. When the user retrieves her placemarks from the database, the physical trajectory is checked for any intersections with the placemarks using an is-point-in-polygon function. All time intervals in which she visited any of her placemarks are indicated on the timeline. The user can further choose to display the placemarks along with the trajectory data in the map view.

5 SUMMARY AND FUTURE WORK

Many applications in the mobility domain require a semantic interpretation of trajectory data. Today, this information is mostly obtained by tedious interviews and manual records. However, the ready availability of GPS technology, which provides an accurate measure of the physical movement, allows us to expedite the process. Being able to view one's own trajectory in geographic and temporal context revives that person's memory and provides for precise movement information.

In this paper, we introduce a conceptual annotation model, which is based on the notion of episodes. Our model is extensible and can easily be adapted to different application needs.

Furthermore, it allows for a clear separation of semantic and physical information. We implemented our model and developed a software tool to support trajectory annotation. The annotation process is facilitated by features such as a zoom-enabled timeline, animation and placemarks.

As future work, we plan to extend our software to take an increasing share in the annotation process. In this vision, the annotation tool provides the user with a possible annotation, and the user needs only to verify or correct the annotation. A first step in this direction is the automatic extraction of placemarks. Some research has already been conducted to extract frequently visited locations and places of interest [Laasonen, 2004; Palma 2008; Giannotti, 2007]. In combination with background knowledge, a suggestion for the type of place can be made. A second step will be to automatically derive the means of transportation. The velocity, acceleration and direction information of the physical trajectory provide a sound basis for such a classification, but they need to be further enhanced with background knowledge, such as information about public transportation.

BIBLIOGRAPHY

Alvares, L. O.; Bogorny, V.; Kuijpers, B.; Macedo, J. A. F.; Moelans, B. and Vaisman, A.: A Model for Enriching Trajectories with Semantic Geographical Information. In: Proc. of the 15th ACM International Symposium on Advances in Geographic Information Systems (ACM-GIS'07), 2007.

Axhausen, K.W., S. Schönfelder, J. Wolf, M. Oliveira and U. Samaga: 80 weeks of GPS-traces: Approaches to enriching the trip information, Arbeitsbericht Verkehrs- und Raumplanung, 178, Institut für Verkehrsplanung und Transportsysteme, ETH Zürich, Zürich, 2003.

BMVBS - Federal Ministry of Transport, Building and Urban Affairs, Mobility in Germany, http://www.mobilitaet-in-deutschland.de, accessed 2007.

Dykes, J. A. and Mountain, D. M.: Seeking Structure in Records of Spatio-Temporal Behaviour: Visualization Issues, Efforts and Applications. In: Computational Statistics and Data Analysis, 43 (4) 581-603, 2003.

Giannotti, F.; Nanni, M.; Pedreschi, D. and Pinelli, F: Trajectory Pattern Mining. In: Proc. of the 13th ACM International Conference on Knowledge Discovery and Data Mining (KDD 2007), 2007.

Google Inc: Google Maps, http:// maps.google.com, accessed 2007.

Google Maps API: http://code.google.com/apis/maps/index.html, accessed 2007.

Jaret timebars component: http://jaret.de/timebars/, accessed 2007.

Laasonen, K.; Raento, M. and Toivonen, H: Adaptive On-Device Location Recognition. Proc. of the 2nd International Conference on Pervasive Computing, 2004.

Laube, P., Dennis, T., Walker, M. and Forer, P.; Movement Beyond the Snapshot – Dynamic Analysis of Geospatial Lifelines. In: Proc. of the ICA Workshop on Geospatial Analysis and Modelling, 2007.

MGE DATA: Mobitest Product information. http://www.mobitest.de, accessed 2007.

Mountain, D. M. and Raper J. F.: Modelling Human Spatio-Temporal Behaviour: A Challenge for Location-based Services. In: Proc. of the 6th International Conference on GeoComputation, 2001.

Palma, a. T., Bogorny, V., Kuijpers, B. and Alvares, L. O.: A Clustering-based Approach for Discovering Interesting Places in Trajectories. In: Proc. of the 23rd Annual ACM Symposium on Applied Computing (SAC 2008). ACM, 2008 (to appear).

Spaccapietra, S., Parent, C., Damiani M. L., de Macedo, J. A., Porto, F. and Vangenot, C.: A Conceptual View on Trajectories. Technical Report LBD-REPORT-2007-001. Ecole Polytechnique Fédérale de Lausanne, 2007.

Stopher, P., Xu, M. and FitzGerald, C.: Assessing the Accuracy of the Sydney Household Travel Survey with GPS, Transportation, 34 (6), 723-741, 2007.

Sun Microsystems: Java Programming Language. http://www.java.sun.com, accessed 2007.

Wolf, J.: Using GPS data loggers to replace travel diaries in the collection of travel data, Dissertation, Georgia Institute of Technology, School of Civil and Environmental Engineering, Atlanta, Georgia, 2000.

Wolf, J., Guensler R. and Bachman, W.: Elimination of the Travel Diary: An Experiment to Derive Trip Purpose from GPS Travel Data, Transportation Research Record, 1768, 125-134, 2001.

Zmud, J. and Wolf, J.: Identifying the Correlations of Trip Misreporting – Results fro the California Statewide Household Travel Survey GPS Study. In: Proc. of the 10th International Conference on Travel Behaviour Research, 2003.