

Travel History: Reconstructing Travelers' Trajectories Using Digital Footprints

Amon Veiga Santana¹ and Jorge Campos^{1,2}

¹Graduate Program in Systems and Computation – UNIFACS

²Bahia State University – UNEB

Salvador, Bahia, Brazil

amoncaldas@yahoo.com.br, jorge@unifacs.br

Abstract: *This work proposes a solution for reconstructing travel histories using heterogeneous social track posts in social networks, GPS positioning data, location history data generated by cloud services or any digital footprint with an associated geographic position. The solution encompasses a conceptual model; a methodology to reconstruct travel histories based on heterogeneous social tracks sources; and an application to present the reconstructed travel itinerary in a graphical and interactive fashion. An experiment conducted with real travelers showed that the proposed solution is a reasonable way to reconstruct semantic-rich travel histories in an automatic fashion.*

1. Introduction

The popularization of Online Social Network (OSN) and User Generated Content (UGC) have modified the way people search, find, read, access, and share information on the Internet (Ye et al. 2011). OSNs have an important role in the production and search for information. OSN users' are frequently involved in activities to find relevant contents, advices, opinions, or to simply interact with their mates to have fun (Lange-faria and Elliot 2012). UGCs (e.g., posts in social networks and comments in websites and forums) have become an important and recognized source of information in the tourism domain (Akehurst 2009). Travel specialized websites, for instance, have increased its sociability and usage by adopting mechanisms that facilitates content sharing in real time between users. A 2011 PhocusWright report¹ shows that nine of ten cyber travelers read and trust online reviews in touristic related sites. Unfortunately, there are far more people willing to consume this kind of content, than people disposed to generate them (Pan and Crotts 2012). It is because most people see UGC as a time consuming and boring task, but they will not mind to contribute if there exists some kind of application or service that captures their contribution in an automatic fashion.

A special kind of information incorporated by most OSNs that has attracted the attention of the travel and tourism community is the users' position while they are moving. The increasing number of location-enabled devices opens the possibility of making the position of the user a mandatory piece of information to virtually any kind of social interaction or user generated content. Moreover, the capability of keeping track of the position of a user at high detailed levels opens the possibility to combine traveler's

¹http://www.researchandmarkets.com/reports/1866967/phocuswrights_social_media_in_travel_2011

trajectory data and georeferenced social interactions to produce, in an automatic fashion, a structured and semantic rich dataset of traveler's preferences and behaviors.

This paper introduces Travel History, a conceptual model and a methodology to reconstruct the trajectory of travelers based on records of their position and their interactions posted on social networks. Position information may vary from the usual detailed GPS logs to any evidence of places visited by the traveler and recovered from the traveler social network repository. Thus, Travel History model supports the representation of the trajectory with different levels of granularities mixed and interleaved with travelers' social interactions.

The remainder of this paper is structured as follows: Section 2 discusses related work. Section 3 presents the Travel History Model. Section 4 discusses the techniques used to instantiate entities of the model. Section 5 introduces a prototype tool and presents some results of an experimental evaluation with real travelers' volunteers. Section 6 presents conclusions and indicates future work.

2. Related work

The analysis of trajectories of moving objects has been intensively discussed by the GIS community over the last decade. Fed initially by the profusion of data captured from sensors and location devices, studies in this field have evolved from the generation of trajectories using GPS raw trajectory data to the use of novel means to enrich trajectories semantically. One salient new source of information comes from the growing habit among users to interact in social networks, posting, commenting, or sharing contents that contain geographic references. This source of information has proven its value for many different fields and purposes. It is of special interest of this work the combination of trajectory semantic enrichment techniques and georeferenced post in social networks to produce semantic rich set of information about travelers and their visits.

Semantic enrichment and annotation in trajectory data are very active research topics. (Spaccapietra et al. 2008) proposed the first model that treats trajectories of moving objects as a spatiotemporal concept. The *Stops* and *Moves* model is one of the most accepted model to represent trajectory of moving objects. (Andrienko, Andrienko, and Wrobel 2007) used the time for adding semantic annotation to the stationary part of a trajectory and argued that the more time is spent in a place more important it is the place to a person. In the (Zheng et al. 2009)'s work, it was exposed a technique that considered, beyond the spending time, the geographic coincidence with Points of Interest (POI) defined in the application. (Zheng et al. 2010) proposed a technique based on speed, acceleration and the orientation of the user to detect the transportation mode used to move from one place to another. A comprehensive set of solutions for semantic annotation of heterogeneous trajectories can be found in (Yan et al. 2011) and for semantic trajectories modeling and analysis in (Parent et al. 2013).

Researches in the trajectory domain provide a solid base for the development of effective solutions to extract information from raw trajectory data. In another branch, several initiatives focus in pattern and knowledge discovery from User Generated Georeferenced Content (UGGC). In our context, UGGC is defined as a UGC that carries some kind of information that allows the identification of the geographic location of the related content, not necessarily the location of the user. A georeferenced picture of

Copacabana beach posted in Instagram and a Web review made by someone in New York about the Copacabana Palace Hotel, for instance, are both examples of UGCC of the same geographic region. UGCC do not have the same spatial granularity of positioning devices, such as GPS, but allows a more refined semantic extraction.

Related with initiatives that deal with UGCC for semantic enrichment, (Ji et al. 2009)(Hao et al. 2010) proposed a solution for mining city attractions from touristic blogs posts, (Rattenbury, Good, and Naaman 2007) proposed an approach to extract semantic from georeferenced picture posted in the Flickr social network, and (Gao et al. 2010) proposed a method to identify touristic attractions from Flickr's georeferenced pictures and to enrich the description of such attractions with information extracted from collaborative websites like Yahoo Travel Guide² and WikiTravel³. (Lu et al. 2010) proposed a picture-based customized trip planning. This system allows trip planners to specify personal preferences and generates travel routes from geo-tagged photos. The proposed solution is limited to surrounding attractions in a given city or region and does not support travel plans lasting more than one day and involving multiple destinations. (Yoon et al. 2012) proposed a framework for itinerary social recommendation using trajectories generated by local residents and expert travelers.

Despite the enormous potential of aforementioned initiatives, few works have combined the use of trajectory data and UGCC in the process of trajectories reconstruction and semantic enrichment. (Gil et al. 2014) proposed a method for trajectory annotation based on the spatiotemporal compatibility of Twitter posts. Although the spatial component of the post is mentioned in the work of (Gil et al. 2014), the proposed methodology takes into account only the temporal compatibility between trajectory data and Twitter posts, that is, they do not use the posts' content or location to enrich the trajectory semantically.

Analyzing early related work, it is noticeable that most solutions used detailed logs of position devices to analyze people's movements. This tendency is switching to incorporate location information embedded in social interactions and stored in the cloud. In a social post, for instance, the position comes with some kind of information or even a personal opinion about the place visited. Thus, trajectory reconstruction using georeferenced social interactions can be the strategy to recover context elements and semantics of trajectory. As the process of trip planning involves information search and retrieval, it is natural that travelers also look for this kind of information among their friends and people from their social circle. At the best of our knowledge, however, there is no service that, considering previous users' experience registered as social tracks, offers efficient means for travelers to access information about structured travel itineraries, including attractions and transportation means. Next section introduces an attempt to represent this kind of information.

3. Travel History

Travel History conceptual model encompasses all information needed to represent relevant actions and movements of a traveler. The central entity of the model is *Travel History*, which is

² <https://www.yahoo.com/travel/guides>

³ <http://wikitravel.org>

an entity that aggregates *Stays* and *Trails* traveled by an individual during a given time interval. Figure 1 shows the relationship among models entities using UML notation.

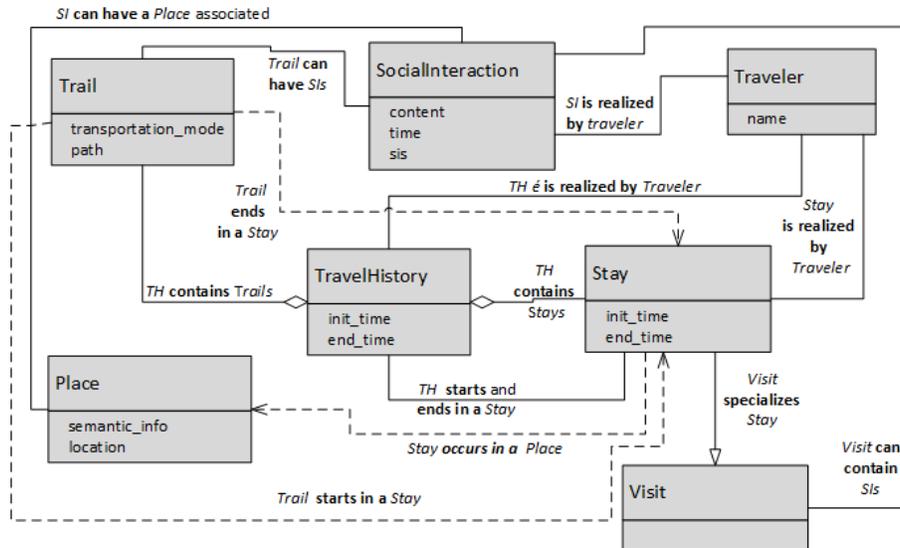


Figure 1: Travel History Conceptual Model.

A *Trail* is an entity that captures the traveler movement. Each *Trail* has an associated path and a transportation mode. The path is a collection of geographic points that represents the geometry of the movement. The path may vary from a pair of points indicating only the endpoints of the movement up to a collection of points representing the detailed path fulfilled by the traveler. The transportation mode indicates how the *Traveler* goes from one place to another (walking, by train, etc.).

Stays represent the places where the *Traveler* remained for a while or change the transportation mode. Each *Stay* occurs at a *Place*. A *Place* is a geographic location plus some semantic incorporated, like a description, political categorization or a combination of them. Consider, for instance, the place where the traveler stops walking and takes a cab. This place is represented as a *Stay* only because the change in the transportation mode. If the traveler remains around a place for a while, these place also becomes a *Stay* in our model.

A *Stay* may or may not have a special meaning for the trip. When the *Traveler's* permanency at some place is relevant for the trip, the *Stay* is specialized and becomes a *Visit*. The relevance of the *Stay* takes into account the amount of time spent at the place or the amount of social interactions related to the place. Thus, a *Visit* represents a place where the *Traveler* has made and registered some social interaction or stayed for certain amount of time above a given threshold.

Visits and *Trails* may have one or more *Social Interactions*. These interactions are contents that help to understand *Traveler's* intentions or activities. The association of *Visits* and *Trails* with *Social Interactions* considers both temporal and geographical matches. Temporal matches consider the time of the realization of the *Social Interaction*, (i.e. the interaction is associated with a *Stay* or a *Trail* that is going on at the time of the occurrence of the post). Geographical matches occur when a *Social Interaction* has some geographic information associated to its content (i.e., it is an UGGC). In this case, the location of the interaction is determinant to establish the association between the *Social*

Interaction and the respective *Stay* or *Trail*. The amount of *Social Interactions* related to a *Stay* or *Trail* is also an indicator of the place's relevance for the trip.

Travel History model was conceived aiming at handling multiple types of UGGCs retrieved from different OSNs and combined with any sort of positioning data about the user movement. Next section discusses the mechanism of converting these kinds of heterogeneous data into entities of the model.

4. Rebuilding Travel Histories

The *Travel History* reconstruction process is based on heterogeneous sources of data. Sometimes it is available a very fine set of registers of a traveler's movement captured by some kind of position device. Other times there is a less fine position records, but the data comes with some kind of semantics attached, and not rarely, there is a social interaction that can be used as a source of information about the travel. Even coming from different sources, these datasets share common concepts and structures. Thus, they were grouped in three main categories: 1) Raw Trajectory Data (RTD), 2) Semantic Trajectory Data (STD), and 3) Social Interactions. When a Social Interaction has an associated geographic location, it is called Georeferenced Social Interaction (GSI).

RTD and STD are sequences of spatiotemporal records. Although they share the same basic structure, they differ significantly considering both the spatial-temporal granularity and content. RTD are generated by positioning devices and contains only registers with the position and the timestamp of each reading. RTD is usually obtained from a single device and store the data as they are produced. STD, on the other hand, is a result of preprocessed data acquired from many different sources. This information comes from some kind of cloud service such as "Google Takeout", which allows users to recover information about their movement. Users have to authorize Google to keep track of their whereabouts and Google uses this information in lots of different services. Although RTD records are denser than STD records, the semantic of the later is much more relevant to the trajectory reconstruction than the former. The last category of source of information is GSI. GSI records are even sparser than RTD and STD records. Thus, GSI alone contributes little to reconstruct the detailed trajectory geometry, but they are very important to enrich the trajectory semantically.

The process of rebuilding *Travel Histories* goes from gathering all pieces of information to the semantic enrichment of models entities. This process can be split in three phases: data acquisition; data processing; and entities generation (Figure 2). In the first phase, data are acquired from different sources, like social networks, location history web services, and location's tracks recorded in the user device. In the second phase, the data are processed to identify *Stay* candidates and transportation modes. In the last phase, *Stays*, *Visits* and *Trails* are generated and semantically enriched.

The data acquisition phase starts with the definition of the time window when the travel happened. After defining the temporal window of the trip, the reconstruction process continues by gathering all relevant information about user movement and his/her *Social Interactions*. The data acquisition strategy depends on the category of the sources available. RTD and STD are collected as a single file. RTD records come from mobile applications that continuously record the position of the travelers over time. These records are stored in the device internal memory and can be imported at any time. STD records come from cloud services (like Google Takeout). These records are requested by the

owner of the data, the only person able to retrieve them. RTD and STD files are traversed and relevant information is extracted and stored in a local database. The process of gathering social interactions, on the other hand, can be fully automated. OSN's users authorize a computer application to search and retrieve all social interactions of a given period. These data are also stored in a local database.

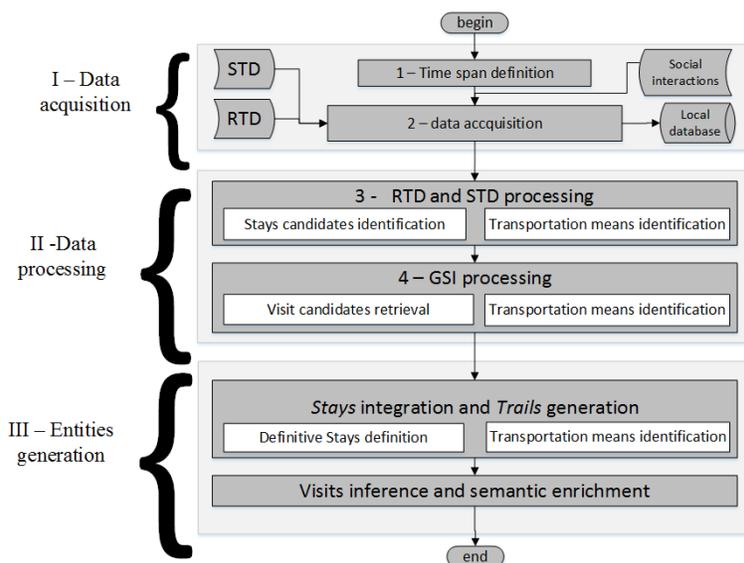


Figure 2: The three phases of the Travel History generation process.

In order to illustrate the reconstruction process, consider a hypothetical travel with samples of information gathered from the traveler's digital footprints. Figure 3 shows a travel timeline with samples of information where, in the first segment there is only intermittent RTD records. In the second segment, RTD and GSI (content posted in social networks like Facebook or Instagram) are available. In the last segment, a combination of RTD, STD and GSI is available and in some parts they overlap.

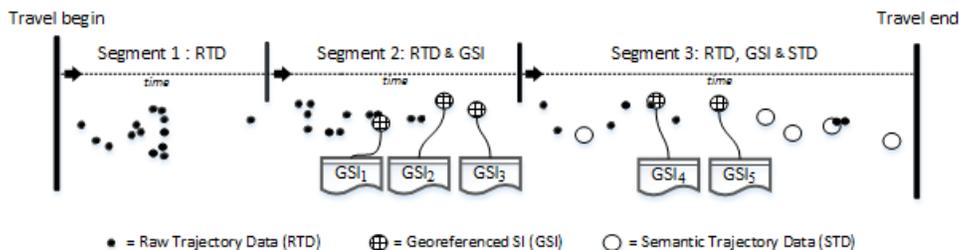


Figure 3: Graphical presentation of RTD, STD and GSI records.

With all pieces of information in place, the second phase of the reconstruction process starts by identifying *Stays* candidates and transportation modes between these candidates. *Stays* candidates represent the locations where the traveler hangs around for a while or change the transportation mode.

Stays candidates are generated considering the category of the data source to be processed. To detect *Stays* from RTD and STD sources, it was developed an algorithm

capable of recognizing these entities based on the geometric configurations of the track (clusters of points or isolated points) and based on some semantic information already present in the data. When a traveler stays in a place, a cluster of points (i.e., a dense formation of points) is formed. The algorithm identify this formation of points and group these points to form a *Stay* candidate (Figure 4 - case A). On the other extreme, isolated points also becomes a *Stay* candidate (Figure 4 - case B). This case occurs when there is a record distant from both the previous and the next point in the sequence and it is not considered an outlier. Outliers are treated during the pre-processing phase of the reconstruction process. Most outliers are disregarded based on the physical unviability for a traveler being at a certain place considering, for example, the maximum speed of known transportation means. Outliers are disregarded from the dataset and do not reach this phase of the reconstruction process.

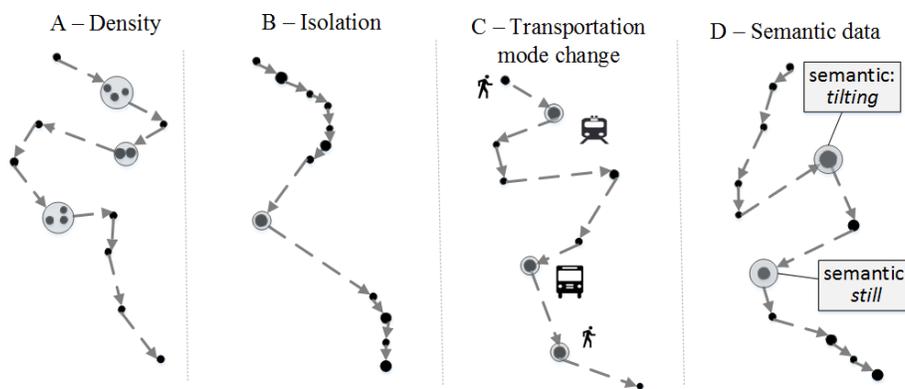


Figure 2: Stays identification techniques based in RTD and STD and the result when it is integrated with Social Interactions.

While processing RTD files, the transportation mode used between two *Stays* candidates is also computed. The definition of the transportation mode takes into account the following aspects: speed, speed variation, acceleration, orientation variation and continuity. Each transportation mode has a single combination of these factors. By taking them together, it is possible to infer how the *Traveler* moved between *Stays*. Due to the lack of space, it is abstracted away the details of determining transportation modes. *Stays* are also defined at every location where a transportation mode change occurs (Figure 4 - case C). Finally, a *Stay* can be inferred from the semantic already embedded in STD files. These data sometimes have semantic information like “still” or “tilting” associated with a place. (Figure 4 - case D). These places always become a *Stay* candidate.

Stay candidates are also generated considering GSI information. In this case, the rule is simple, that is, every GSI generate a *Stay* candidate. Later, some of these *Stays* will be grouped, becoming a single *Visit*, others will not be confirmed as a *Stay* and will become a social interaction of a *Trail*. All *Stays* candidates, no matter the source of information, are stored in a common persistence entity.

During the last phase of the reconstruction process, all high-level entities of the model representing parts of the travel history are generated, integrated and semantically enriched. Throughout the integration step, issues related to the duplicity and overlaps are solved. At this point, each *Stay* candidate is processed, confirmed as a definitive *Stay*, promoted to a *Visit*, or merged with others *Stays*. Since *Stays* candidates are generated

from different sources separately, it is possible that some *Stays* candidates refer the same event of the trip. The *Stays* merge process occurs when the distance between two *Stays* is less than a given threshold.

The next step in the reconstruction process is the *Trails* generation. *Trails* connect two *Stays* and describe the *Traveler* movement between them. During *Trails* generation, it is necessary to identify the transportation mode. If the *Stays* were generated from the same source, the transportation mode between them has been already defined in the second phase, but if the *Stays* were generated based on different sources, the same algorithm to detect transportation mode discussed earlier is used. The last step of the phase of the reconstructing process is the semantic enrichment of *Trails* and *Visits*. *Visits* are specialized versions of *Stays*. For a *Stay* to become a *Visit*, it is considered the amount of time spent on the site and the number of *Social Interactions* carried out by the traveler. Once all model entities are instantiated, an application using geovisualization techniques can easily depicts the graphical realization of the reconstruct travel history.

5. Experiment and results

To evaluate the *Travel History* model, it was developed a prototype tool that employs all techniques presented in this paper (available at <http://th.fazendoasmalas.com>). The prototype allows the acquisition, processing and generation of *Travel Histories*. At the end, the tool shows the user travel history in an interactive map. *Stays*, *Visits*, and *Trails* are presented in a graphical and user-friendly web application.

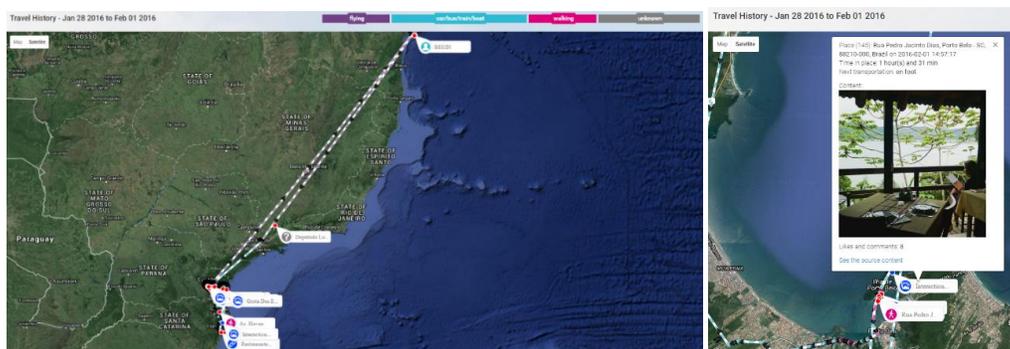


Figure 5: Graphical presentation of a Travel History generated using RTD, STD and GSI data sources: an overview of the entire trip and a detailed view of part of the trip.

Figure 5 shows the reconstruction of a Travel History generated using different and rich sets of information. This travel occurred between January 28 and February 1, 2016. It was a five days' trip in the south part of Brazil, including a visit to the capital of Paraná, a train trip between the capital and some place along the state coastline, a visit to Mel Island in the Paranaguá Bay, and a visit to the state of Santa Catarina, including the capital Florianópolis and other small towns nearby. To reconstruct this travel history, it was used as source of information GPS log files, location history files generated by Google and the online social networks Facebook, Instagram and Twitter. When RTD, STD, and GSI are all available to reconstruct the travel, it is possible to zoom in the map presentation to analyze the detailed path of the traveler and to visualize comments and social interactions posted along the path. Figure 5 shows on the left side an overview of the entire trip and on the right side a detailed view of part of the trip to Parana and Santa

Catarina and a box with the one of the GSI posted about the trip. We used this application to support the realization of an experimental evaluation of the proposed model and methodology. Next section discusses the result of such an experiment.

Volunteers from RBBV (acronym, in Portuguese, for Brazilian Travel Bloggers Network) have been invited to use the tool, submit their data, reconstruct their travels and evaluate the travel history generated by the application. A total of 58 volunteers started the experiment, but only 23 travelers completed the entire process successfully. Some volunteers did not submitted data, others submitted inconsistent data, and some did not perform the evaluation. The volunteers were oriented to answer a questionnaire after analyzing and exploring their reconstructed travel. The questions presented to the volunteers aimed at verifying the level of satisfaction with the accuracy and similarity of the Travel History created based on their digital footprints when compared with the events and destinations of the real trip they have made. An interactive map allowed volunteers to check visited *Places*, analyze the performed *Trails*, verified transportation means used, and examine associated semantics. The tool used to present the graphical realization of the trip was not evaluated.

In the process of evaluating the travel reconstruction process, the volunteers have answered five questions. For each question, volunteers were asked to give a grade ranging from 0 to 10. The results of the experiment are presented in the Table 1.

The first question aims to evaluate the accuracy of *Visits* identification. The result for this question indicates that most of the evaluators (~79%) considered that the identification of *Visits* was completely accurate or precise in most cases, while ~13% felt it was accurate in some places, and 8% found the process of *Visit* identification was slightly or completely inaccurate.

The second question measures the satisfaction with temporal order of visits and trails. This question is related to the integration process of *Stays* and *Visits*, which is responsible for identifying overlaps, to perform merges, and to sort these entities. The result indicates that most evaluators (82.6%) considered that the order of visits and trail order has been completely precise or precise in most cases, 13.05% found that the identification of the order was accurate in some places and only one evaluator (4.31%) considered that the order of *Visits* and *Trails* was imprecise.

Table 1: Travel History reconstruction evaluation results

| Aspect analyzed | Avg. grade | Standard Deviation | Totally accurate or in most of the cases |
|---|-------------|--------------------|--|
| <i>Visits</i> identification accuracy | 7,71 | 2,7 | 78,25% |
| <i>Visits</i> and <i>Trails</i> order accuracy | 7,82 | 2,03 | 82,6% |
| Transportation means identification accuracy | 7,39 | 2,19 | 69,56% |
| Activities and semantic identification accuracy | 7,06 | 1,79 | 74% |
| Travel History rebuilt represents the real travel made. | 8,69 | 2,14 | 95,65% |
| <i>Averages considering all aspects</i> | 7,73 | 2,17 | 80,01% |

The third question evaluates the accuracy of identifying the transportation mode used in each *Trail*. Although the level of satisfaction with the identification of the transportation mean is close to 70%, this aspect has the worst evaluation on the survey. The identification of the transportation mode is directly linked to the existence and granularity of RTD and STD sources and the accuracy of the location of GSI. It is noticeable that the identification of transportation mean improves when interactions in OSNs are made in real time during the trip.

The fourth question is more subjective and it is related to the accuracy of the semantic enrichment process. In this regard, 73.9% of the evaluators answered that the identification of activities and interactions was completely accurate or accurate in most cases. The semantic enrichment process can be improved by incorporating the capability of including textual content of Social Interactions and with the ability to access structured information about users' activities. Facebook, for example, has such kind of information, but, at the time of writing, it is not possible to access such kind of information using third-party applications.

The fifth question evaluated the overall perception of the reconstruction process. It is by far the best-rated item of the survey. Almost 96% of the evaluators considered that the Travel History reconstructed represents, totally or in the major part, the travels they have made. Taking all aspects together, about 80% are satisfied with the proposal of reconstructing semantic trajectories based on heterogeneous social tracks sources.

6. Conclusion

Considering the digital socialization growth and the search for online social recognition, travelers begin to demand ways to share their travel experiences in a systematic and intuitive way. This paper proposes conceptual and data models and a methodology to reconstruct semantic-rich traveler's trajectories. The central entity of the model is *Travel History*, which is an entity that aggregates *Stays*, *Visits* and *Trails* traveled by an individual during a given time interval. A *Trail* is an entity that captures the traveler movement and the transportation mode used. A *Stay* represent places where the *Traveler* remained for a while or changes the transportation mode. A *Stay* becomes a *Visit* if it is a place of intense online social interaction or if it is a place where the traveler spent a considerable amount of time.

Model's entities are instantiated based on a myriad of sources of information, varying from detailed low-level GPS registries and going up to high-level georeferenced social interactions. Thus, the proposed methodology used to generate models entities and to identify transportation mode is based on techniques to process, analyze, and integrate data with different levels of semantic and spatial-temporal granularity. The ability to reconstruct the trip successfully is directly related to the quality and quantity of the sources of information available. Different, reliable, and abundant sources of information will produce rich and accurate travel histories. On one hand, RTD and STD are good sources of information for detailed analysis of the trip. On the other hand, GSI generates semantic richer entities. As expected, the combination of all sources produces the best result.

In order to evaluate the proposed model and methodology, an experiment with travelers from a social traveler's network was designed and run. The results of the experiment show an overall level of satisfaction of 80%, considering the identification of

the model's entities (i.e., *Stays*, *Visits* and *Trails*), the temporal order of these entities, the identifications of the transportation mode used by the traveler, the identification of activities performed by the traveler during the trip, the semantic enrichment of travelers' activities, and the level of adherence of the modeled travel history with the real trip.

As future work, there are several aspects that can be investigated. Semantic enrichment, for instance, can be improved by incorporating text-mining algorithms. Moreover, a travel social media ontology can be developed to improve semantic identification. To improve the data accuracy and granularity, mobile applications can be used to collect other kinds of social interaction, like offline media capture or any other type of interactions on the device. Algorithms for transportation means identification can be improved to become more accurate and to support the identification of other kinds of transportation.

The use of the proposed model and methodology in Web application in the tourism domain will allow the reconstruction of large number of Travel Histories, which in turn, can be a way to generate a knowledge base for travel itineraries, preferences, attractions and other aspects and events inherent to travels. This knowledge can be used as the base of a travel recommendation system or other initiatives such as urban planning, demographic and behavioral studies, intelligent transportation systems, social recognizing research, among others. Despite the fact that the model is generic and that it can be, in principle, used in several domains to describe semantic trajectories, the usage for other domains requires further investigations.

7. Acknowledgements

This work was partially supported by CAPES.

8. References

- Akehurst, Gary. 2009. "User Generated Content: The Use of Blogs for Tourism Organisations and Tourism Consumers." *Service Business* 3 (1): 51–61.
doi:10.1007/s11628-008-0054-2.
- Andrienko, Gennady, Natalia Andrienko, and Stefan Wrobel. 2007. "Visual Analytics Tools for Analysis of Movement Data." *ACM SIGKDD Explorations Newsletter*.
doi:10.1145/1345448.1345455.
- Gao, Yue, Jinhui Tang, Richang Hong, Qionghai Dai, Tat-Seng Chua, and Ramesh Jain. 2010. "W2Go: A Travel Guidance System by Automatic Landmark Ranking." *Proceedings of the International Conference on Multimedia - MM '10*: 123.
doi:10.1145/1873951.1873970.
- Gil, Ricardo, Belther Nabo, Renato Fileto, Mirco Nanni, and Chiara Renso. 2014. "Annotating Trajectories by Fusing Them with Social Media Users' Posts." In *Geoinfo Synposiun 2014*.
http://www.geoinfo.info/proceedings_geoinfo2014.split/Paper04-F-p21.pdf.
- Hao, Qiang, Rui Cai, Changhu Wang, Rong Xiao, Jiang-Ming Yang, Yanwei Pang, and Lei Zhang. 2010. "Equip Tourists with Knowledge Mined from Travelogues." *Proc. of the 19th International World Wide Web Conference*: 1–10.
doi:10.1145/1772690.1772732.

- Ji, Rongrong, Xing Xie, Hongxun Yao, and Wei-Ying Ma. 2009. "Mining City Landmarks from Blogs by Graph Modeling." In *ACM International Conference on Multimedia*, 105–114. doi:10.1145/1631272.1631289.
- Lange-faria, Wendy, and Statia Elliot. 2012. "Understanding the Role of Social Media in Destination Marketing." *Tourismos: An International Multidisciplinary Journal of Tourism* 7 (1): 193–211.
- Lu, Xin, Changhu Wang, Jiang-Ming Yang, Yanwei Pang, and Lei Zhang. 2010. "Photo2Trip: Generating Travel Routes from Geo-Tagged Photos for Trip Planning." *Proceedings of the International Conference on Multimedia - MM '10*: 143–152. doi:10.1145/1873951.1873972.
- Pan, Bing, and John C. Crotts. 2012. "Theoretical Models of Social Media, Marketing Implications, and Future." *Social Media in Travel, Tourism and Hospitality: Theory, Practice and Cases* (1965): 1–19.
- Parent, Christine, Stefano Spaccapietra, Chiara Renso, Gennady Andrienko, Natalia Andrienko, Vania Bogorny, Maria Luisa Damiani, et al. 2013. "Semantic Trajectories Modeling and Analysis." *ACM Computing Surveys* 45 (4): 42:1–42:32. doi:10.1145/2501654.2501656.
- Rattenbury, Tye, Nathaniel Good, and Mor Naaman. 2007. "Towards Automatic Extraction of Event and Place Semantics from Flickr Tags." *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval SIGIR 07* pages: 103. doi:10.1145/1277741.1277762.
- Spaccapietra, Stefano, Christine Parent, Maria Luisa Damiani, Jose Antonio, José Antônio De Macedo, Fábio Porto, and Christelle Vangenot. 2008. "A Conceptual View on Trajectories." *Data & Knowledge Engineering* 65.1 (May 2007): 126–146.
- Yan, Zhixian, Dipanjan Chakraborty, Christine Parent, Stefano Spaccapietra, and Karl Aberer. 2011. "SeMiTri: A Framework for Semantic Annotation of Heterogeneous Trajectories." In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT/ICDT '11)*, 259–270. doi:10.1145/1951365.1951398.
- Ye, Qiang, Rob Law, Bin Gu, and Wei Chen. 2011. "The Influence of User-Generated Content on Traveler Behavior: An Empirical Investigation on the Effects of E-Word-of-Mouth to Hotel Online Bookings." *Computers in Human Behavior* 27 (2): 634–639. doi:10.1016/j.chb.2010.04.014.
- Yoon, Hyoseok, Yu Zheng, Xing Xie, and Woontack Woo. 2012. "Social Itinerary Recommendation from User-Generated Digital Trails." *Personal and Ubiquitous Computing* 16: 469–484. doi:10.1007/s00779-011-0419-8.
- Zheng, Yu, Yukun Chen, Quannan Li, Xing Xie, and Wei-Ying Ma. 2010. "Understanding Transportation Modes Based on GPS Data for Web Applications." *ACM Transactions on the Web* 4 (1): 1–36. doi:10.1145/1658373.1658374.
- Zheng, Yu, Lizhu Zhang, Xing Xie, and Wei-Ying Ma. 2009. "Mining Interesting Locations and Travel Sequences from GPS Trajectories." *Proceedings of the 18th International Conference on World Wide Web - WWW '09*: 791. doi:10.1145/1526709.1526816.