Inferring Ratings for Custom Trips from Rich GPS Traces

Theodoros Chondrogiannis
Department of Computer and Information Science
University of Konstanz
Konstanz, Germany
theodoros.chondrogiannis@uni.kn

Mouzhi Ge
Faculty of Informatics
Masaryk University
Brno, Czech Republic
mouzhi.ge@muni.cz

ABSTRACT

Trip planning services are employed extensively by users to compute paths between locations and navigate within a road network. In some real-world scenarios such as planning for a hiking trip or running training, users usually require personalized trip planning. Although some existing systems can recommend trips that other users have posted, along with a set of ratings w.r.t. the difficulty of the route, conditions, or the enjoyment it provides. Very often though users want to define a custom trip that fits their personal needs, for which existing systems are unable to provide any rating. In this paper we therefore define the problem of inferring ratings for custom trips. We also outline a solution to infer ratings by utilizing the ratings of trips previously posted by users and their similarity with a given custom trip. Finally, we present the results of preliminary experiments were we evaluate the efficiency of our proposed approach on inferring ratings for trips related to hiking and other similar activities.

CCS CONCEPTS

• **Information systems** → *Geographic information systems*; *Social recommendation.*

KEYWORDS

route planning,route recommendation,mapping services

1 INTRODUCTION

Nowadays, outdoor activities such as jogging, hiking or cycling have become an important part of daily life. One key component in these outdoor sports is trip selection that depends on individual preferences and purposes. For example, someone prefers to enjoy nice views while hiking even if she has to take a longer trail, while someone else may care only about how challenging a hiking trails is.
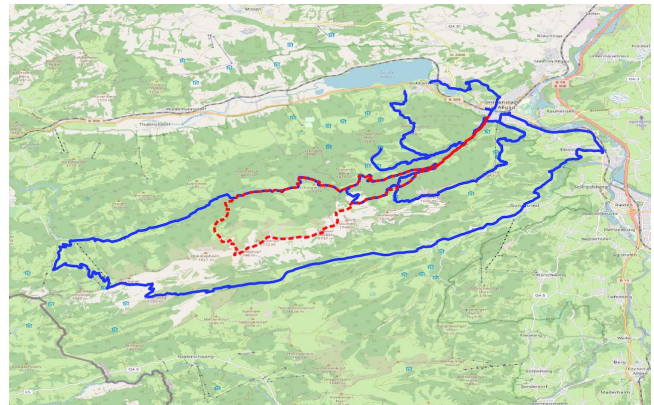
Figure 1: Motivational example.

The quality of the trip selection process can affect the enjoyment of the activity significantly, e.g., training result or user mood. Towards this end, many applications allow users to design their own trips for conducting their preferred sport activity. Consider that in a system different hiking trips with GPS traces have been uploaded and rated by users. These trips can be visualized as the solid blue lines in Figure 1. One user proposes a custom trip that is displayed as the dashed red line. This newly designed trip may overlap with some existing trips. In most cases, the rating of the designed trip is unknown and, to the best of our knowledge, there is currently no research or applications that directly tackle how to infer the rating for the designed trip.

In this paper, we study the problem of rating inference for custom trips. Rating inference can be considered as predicting a rating for a non rated item and has been widely used in recommender systems research. A recommender system can be formally presented as a set of users $U = \{u_1, u_2, ..., u_i, ..., u_m\}$ and a set of objects $O = \{o_1, o_2, ..., o_j, ..., o_n\}$. The relations between U and O can be further presented as a rating matrix. The size of U and O can be very large, such as thousands or even millions of users and items. Many ratings in the user-item matrix are empty. Thus, a recommender is to compute a score $r_{i,j}$ that matches the expected interest or utility of the user u_i to object o_j. More recently, multi-criteria recommender systems are also developed to consider different attribute ratings of an object [1]. In our context, the rating inference for custom trip can be considered as multi-criteria and cold-start recommendation problem. Each trip has different attributes such as difficulty or views, and each attribute can have a rating. Also, a trip is defined by a user for the first time and has never been used by others. Thus, the trip is a cold-start object to the system.

To tackle this problem, we propose a solution to predict the attribute rating of a custom trip by using the fact that the query trip shares segments with existing trips for which the ratings are known. We believe that the outcome of the rating inference can be of high practical value to many applications related to outdoor sport activities. Our preliminary evaluation on real-world data shows that the proposed solution is promising and it can be considered as baseline for future research.

2 RELATED WORK

A number of route recommendation algorithms are related to the rating inference for custom trips. In these works, the routes are either driving routes [4] or running routes [7]. The route recommendation can usually be dealt with as a route ranking problem based on personalization a user's location and preferences [5]. Different aspects can be further taken into account when ranking the relevant routes. For example, Wang et al. [9] consider real-time traffic data and historical taxi driving data to offer real-time route recommendations. Wang et al. [10] study the recommendation of driving routes based on the mood, fatigue and social information of the user. Su et al. [8] and Li et al. [6] utilize the knowledge of the crowd to improve the recommendation quality. Chen et al. [3] extend the problem by adding a set of places to visit. Throughout this paper, the term trip and route may be used interchangeably.

To the best of our knowledge, the existing route recommendation algorithms cannot be directly applied for inferring custom trip ratings. In contrast to existing works where routes are automatically generated [7], in our use-case scenario trips are designed by the users themselves. We argue that in the motivating scenarios we presented in the previous section w.r.t. outdoor sport activities, the user may prefer a trip she designs herself rather that a ranked list of already defined routes. Such a trip may capture preferences of the user that are not modelled in the system. Even in such cases, the user would still like to know the evaluation of the trip w.r.t. attributes supported by the system.

3 BACKGROUND & PROBLEM DEFINITION

Let $G = (N, E)$ be a directed weighted graph representing a spatial network with nodes N and edges $E \subseteq N \times N$. Each edge $e = (u, v) \in E$ has an assigned positive weight $l(e)$ which captures the network distance from u to v. A GPS trace L is a sequence $\langle l_1, \ldots, l_m \rangle$ where each l_i is a position represented by geographical coordinates, i.e., longitude and latitude. A trip is represented by a tuple $t = \{L, r\}$, where L is the GPS trace associated with trip t, and r is a rating assigned to t by the user who created it. If the value of r is set, then we call t a *rated* trip. Otherwise t is called an *unrated* trip. Also, the GPS trace of a trip t can be mapped to some path of the spatial network. We denote this path by $p(t)$.

In this paper, given a set of rated trips $T = \{t_1, \ldots, t_{|T|}\}$ and an unrated query trip t_q, our aim is to infer a rating for t_q by taking into account the similarity of t_q with the trips in T. Consider our example in Figure 1. The blue/solid lines represent trips that are stored in the system and have already been rated by the users who created them. The red/dotted line represents a trip that a user has defined and which does not exist in the database, and is therefore unrated. However, notice that the user-defined trip is overlapping,

i.e., shares segments, with several rated trips. We argue that it is possible to infer a rating for the user-defined/red trip, by taking into account the ratings of the rated/blue trips. Furthermore, the higher the overlap between two trips the higher the chances that their rating is similar.

4 SOLUTION OVERVIEW

The main idea behind our approach is to first map GPS traces for both the rated trips in the input set T and the unrated query trip t_q to the spatial network, thus replacing the GPS trace of each trip with a path in the spatial network G. Considering paths instead of GPS traces facilitates the computation of the shared length/overlap between trips. Next, by considering the ratings of trips in T the paths of which overlap with the path of the query trip t_q, we infer a rating for t_q. In brief, our approach consists of two phases: a) the *data preparation* phase, which involves the mapping of the GPS traces of rated trips in the input set T to paths in the spatial network G, and the construction of an index to optimize the retrieval of overlapping trips, and b) the *query processing* phase, which involves the mapping of the GPS trace of the query trip t_q to a path $p(t_q)$ in the spatial network G, the retrieval of the overlapping trips $t_i \in T$ with t_q, and the computation/inference of the rating r_q for t_q.

4.1 Data Preparation

The fist step of the data preparation is the map matching of the GPS traces of trips t_i in the inpute set T to paths $p(t_i)$ in the spatial network G. In this step the list of coordinates of the input trip t_q is mapped to a path, i.e., a list of edges, of the spatial network. As the map matching of GPS traces to a spatial network is a well studied problem, we use the *Viterbi map matching* algorithm proposed by Wei et al. [11] which, according to the results of the ACM SIGSPATIAL Cup 2012 [2], offers both good performance and accuracy.

Next, to optimize the retrieval of trips the paths of which overlap with the path of a given query trip t_q, we build an inverted index on the edges of each trip. More specifically, for each edge e of the spatial network, we store the list of trips for which e is part of. By storing this information, given a query trip t_q, we avoid computing the overlap of $p(t_q)$ with all the paths of trips in the database, i.e., we need to compute the overlap only with the paths of trips $t_i \in T$ that contains at least one edge of $p(t_q)$.

4.2 Rating Inference

The second phase of our approach begins with the mapping of the GPS trace of the unrated trip t_q to the underlying spatial network G. Similar to the data preparation phase, we employ the Viterbi map matching algorithm proposed by Wei et al. [11], and we compute the path $p(t_q)$ in the underlying spatial network that the GPS trace of t_q is mapped to.

Having completed the map matching step, we need to retrieve all the trips $t_i \in T$ such that every associated path $p(t_i)$ overlaps with $p(t_q)$. The overlap ratio is the fraction of $p(t_q)$ that is shared with $p(t_i)$ and is given by given by

$$Ol(t_i, t_j) = \frac{\sum_{\forall e \in p(t_i) \cup p(t_j)} \ell(e)}{\ell(p(t_j))}.$$
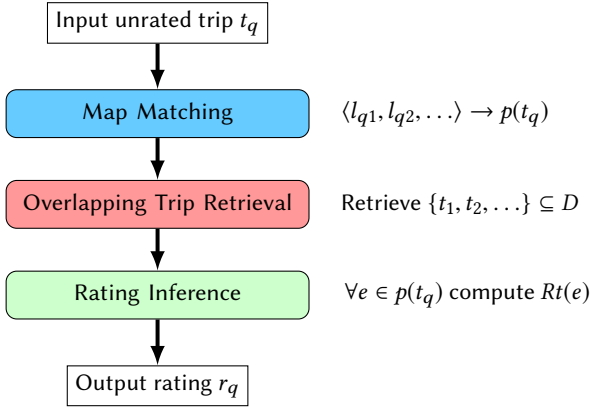
Figure 2: Rating inference overview.

We first retrieve from the edge inverted index the set that contains all the trips $t_i \in T$ the paths of which have at least some overlap with $p(t_q)$. Then, we compute the overlaps $Ol(t_i, t_q)$ for all the paths of the retrieved trips from the edge inverted index.

Subsequently, for each edge $e \in p(t_q)$ we define the set T_e as

$$T_e = \{t_i \mid t_i \in D \wedge e \in p(t_i)\}$$

which is the set of trips the paths of which contain e. To define the rating of each edge $e \in p(t_q)$ we have to consider not only the rating of each trip $t_i \in T$ the path $p(t_i)$ of which shares e with $p(t_q)$, but also the fraction of $p(t_q)$ that is shared with $p(t_i)$. As such, the rating of e is given by

$$Rt(e) = \sum_{\forall t_i \in T_e} r_i \cdot \frac{Ol(t_i, t_q)}{\sum_{\forall t_j \in T_e} Ol(t_j, t_q)}.$$

Note that in order to ensure the proper computation of $Rt(e)$ we normalize the overlap of t_i with t_q.

Finally, the rating of trip t_q is defined as the weighted average of the ratings of its edges, i.e.,

$$r_q = \sum_{\forall e \in p(t_q)} Rt(e) \cdot \frac{\ell(e)}{\ell(E_q)}$$

where $E_q \subseteq p(t_q)$ is the set of edges of $p(t_q)$ that are also crossed by some path $p(t_i) : t_i \in T_e$, and $\ell(E_q)$ is the sum of the lengths of all edges $e \in E_q$. Figure 2 outlines all the steps of our proposed rating inference approach.

5 PRELIMINARY EVALUATION

In this section, we present a case study and preliminary evaluation using data obtained from Outdooractive[1]. The dataset contains a large list of trips. For each trip, multiple fields are provided, e.g., trip type, user that submitted the trip, information about how to access the starting point etc. Also, the dataset provides five different rated attributes, i.e., **C**ondition, **D**ifficulty, **T**echnique, **Q**uality of experience, and **L**andscape, the ratings of which are integers within the interval $[1, 6]$. We also obtained data for four different spatial networks from OpenStreetMap[2]. For each experimental scenario

[1]https://www.outdooractive.com/
[2]https://www.openstreetmap.org/

Table 1: Spatial network and trip data.

network	nodes	edges	trips (hiking)	trips (all)
Swabia	491213	630094	544	353
Austria	2484861	3033885	516	260
NE Italy	1467754	1884450	696	419
Bavaria	3045179	3928652	1346	754

we used trips from the Outdooractive dataset that could be matched with each spatial network of OpenStreetMap. For the matching of the GPS traces with the spatial networks, we obtained the implementation of the Viterbi maps algorithm [11] from Graphhopper[3]. The details of the data we used are shown in Table 1.

In our experiments, we use for each dataset 90% of the trips as training data, i.e., rated trips, and 10% of the trips as test data, i.e., trips the rating of which we infer. We report the results for five different runs/combinations of training-test data. For each set of test trips, we measure the *Mean Absolute Error* (MAE) of our inferred ratings for all five attributes. Moreover, we round our inferred rating to the closest integer value and we measure the accuracy of our prediction, i.e., for how many test trips our approach was able to infer the correct rating. Furthermore, we distinguish two different cases for the trips: a) we consider **all** trips regardless of their type, an b) we restrict the experiments to trips that are related to **hiking**. We do so in order to get insight on how the trip type influences the quality of the inferred rating for each attribute.

To infer trip ratings, our approach takes into account only the part of an unrated trip that overlaps with some already rated trip. The average overlap of each unrated trip with already rated trips was 48.6% for Schwabia, 14.1% for Austria, 22.4% for NE Italy, and 15.5% for Bavaria. In case where a given query trip had no overlapping segments with any other rated trip, then our approach was unable to determine a rating. More specifically, our approach was able to determine a rating for 91.8% on average of the test trips in Schwabia, 63.5% of the test trips in Austria, 68.4% of the test trips in NE Italy, and 72.2% of the test trips in Bavaria.

Figure 3 shows the mean absolute error of our proposed approach. We observe that the MAE is clearly influenced by the rating attribute. More specifically, our approach demonstrates the lowest MAE for the difficulty attribute and the highest MAE for the condition attribute, in all scenarios. Nevertheless, the MAE of our approach for all rating attributes is relatively low in all cases, i.e., below 1.2 with the exception of condition in Austria. With regard to the difference between the scenarios that consider all trips and those that consider only hiking related trips, we do not observe an indication that the type restriction reduces the MAE, although it happens in half of all cases. There are various reasons why this might be the case, e.g. certain attributes are more universal, similar types of activities may interest very diverse users etc., which indicates that a more thorough investigation is necessary.

Figure 4 shows the accuracy of our proposed approach. Similar to the MAE, our approach demonstrates the highest accuracy for the difficulty attribute and the lowest accuracy for the condition attribute in all scenarios. Our approach was able to infer the correct
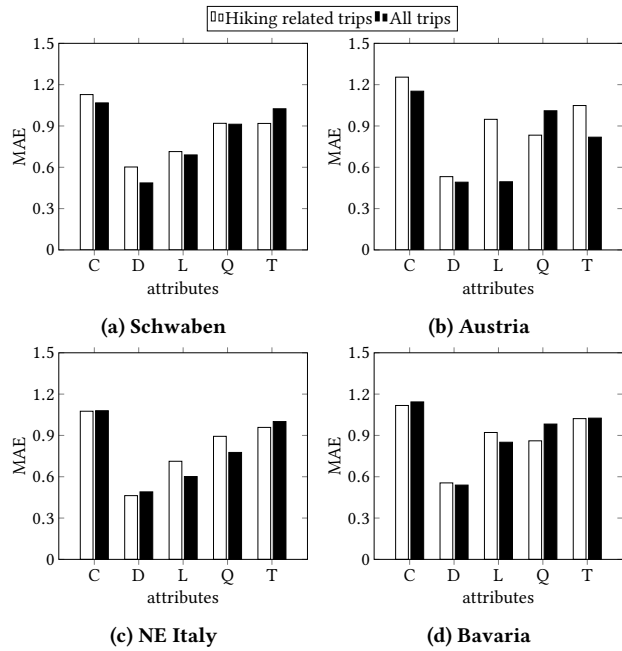
[3]https://www.graphhopper.com

Figure 3: Inference Mean Absolute Error for all four datasets for all trips and hiking related trips only.

Figure 4: Inference accuracy for all four datasets for all trips and hiking related trips only.

difficulty rating for more than half of the test trips in all cases. Also similar to MAE, the trip type restriction does not indicate any significant improvement in the accuracy of the inferred ratings.

6 CONCLUSIONS

In this paper, we studied the problem of inferring ratings for user-defined trips by considering the overlap of such trips with already rated trips. More specifically, we first map the GPS traces of already rated trips and the GPS of a given unrated query trip to paths in the underlying spatial network. Then, we compute for each edge of the path associated with the query trip a rating, by considering the rating of all the trips the paths of which contain that particular edge. Then we define the inferred rating as the weighted average of the ratings of all rated edges on the path of the query trip. Our preliminary evaluation shows the efficiency of our approach in inferring trips ratings related to hiking, thus demonstrating the promise of our approach.

In the future, we plan to investigate different directions in order to improve the efficiency of our approach. First, we need to expand our approach to take into account the edges for which a rating cannot be inferred, as we are currently considering only edges which lie in the path of at least one rated trip. Second, we need to develop a solution for determining the rating of trips the paths of which do not overlap with any path of an already rated trip. As we are also planning to evaluate our approach on different use-cases, we believe it is possible to infer ratings using scenario specific properties. Finally, we intend to study the effect of the overlap among the trips on the accuracy of rating inference. This will allow us to gain more insight and improve the inference algorithm.
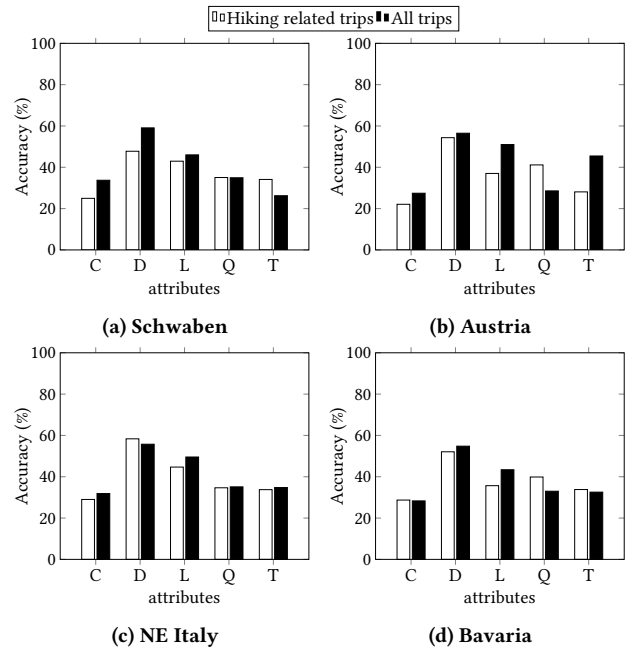
REFERENCES

[1] Gediminas Adomavicius and YoungOk Kwon. 2015. *Multi-Criteria Recommender Systems.* Springer US, Boston, MA, 847–880.
[2] Mohamed Ali, John Krumm, Travis Rautman, and Ankur Teredesai. 2012. ACM SIGSPATIAL GIS Cup. In *Proc. of the 20th ACM SIGSPATIAL GIS Conf.* 597–600.
[3] Lisi Chen, Shuo Shang, Christian S. Jensen, Bin Yao, Zhiwei Zhang, and Ling Shao. 2019. Effective and Efficient Reuse of Past Travel Behavior for Route Recommendation. In *Proc. of the 25th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining.* 488–498.
[4] Jian Dai, Bin Yang, Chenjuan Guo, and Zhiming Ding. 2015. Personalized route recommendation using big trajectory data. In *Proc. of the 31st IEEE ICDE.* 543–554.
[5] Takeshi Kurashima, Tomoharu Iwata, Go Irie, and Ko Fujimura. 2010. Travel Route Recommendation Using Geotags in Photo Sharing Sites. In *Proc. of the 19th ACM CIKM Conf.* ACM, 579–588.
[6] Yu Li, Man Lung Yiu, and Wenjian Xu. 2015. Oriented Online Route Recommendation for Spatial Crowdsourcing Task Workers. In *Proc. of the 14th Int. Symp. on Spatial and Temporal Databases.* 137–156.
[7] Benedikt Loepp and Jürgen Ziegler. 2018. Recommending Running Routes: Framework and Demonstrator. In *Proc. of the 2nd Workshop on Recommendation in Complex Scenarios.* 543–554.
[8] H. Su, K. Zheng, J. Huang, H. Jeung, L. Chen, and X. Zhou. 2014. CrowdPlanner: A crowd-based route recommendation system. In *In Proc. of the 30th IEEE ICDE.* 1144–1155.
[9] Henan Wang, Guoliang Li, Huiqi Hu, Shuo Chen, Bingwen Shen, Hao Wu, Wen-Syan Li, and Kian-Lee Tan. 2014. R3: A Real-time Route Recommendation System. *Proc. VLDB Endowment* 7, 13 (Aug. 2014), 1549–1552.
[10] Jiaqi Wang, Yunyao Lu, Xiaojie Wang, Jing Dong, and Xiping Hu. 2018. SAR: A Social-Aware Route Recommendation System for Intelligent Transportation. *Comput. J.* 61, 7 (2018), 987–997.
[11] Hong Wei, Yin Wang, George Forman, Yanmin Zhu, and Haibing Guan. 2012. Fast Viterbi Map Matching with Tunable Weight Functions. In *Proc. of the 20th SIGSPATIAL GIS Conf.* 613–616.