Towards the Recommendation of Personalised Activity Sequences in the Tourism Domain

Gunjan Kumar, Houssem Jerbi, and Michael P. O'Mahony
Insight Centre for Data Analytics,
School of Computer Science, University College Dublin, Ireland
{firstname.lastname}@insight-centre.org

ABSTRACT

In this paper we consider the problem of recommending sequences of activities to a user. The proposed approach leverages the order as well as the context associated with the user's past activity patterns to make recommendations. This work extends the general activity recommendation framework proposed in [16] to iteratively recommend the next sequence of activities to perform. We demonstrate the efficacy of our recommendation framework by applying it to the tourism domain and evaluations are performed using a real-world (checkin) dataset.

CCS CONCEPTS

• **Information systems → Recommender systems**; *Decision support systems*; *Spatial-temporal systems*;

KEYWORDS

Sequence Recommendation, Recommender Systems, Activity Recommendation, Activity Timeline Matching

1 INTRODUCTION

Internet and digital technologies have significantly influenced the tourism sector in the last decade resulting in a steady growth in e-tourism [7]. Users now have easy access to vast amounts of information on the web which assists them to plan trips, make reservations, and purchase products etc. However, the number of available choices have increased so rapidly that it has become difficult to find the right information at the right time. Thus, recommender systems, which have found immense success in e-commerce, have the potential to play a crucial role in e-tourism by providing personalised and relevant content to users [5, 14, 24, 27].

To provide useful recommendations, it is essential to capture the behaviour and needs of users, which has been particularly challenging in e-tourism [26]. However, as digital technologies have now permeated our daily lives to a great extent, many aspects of our lives can now be easily recorded in digital format. For example, physical activities performed, locations visited and media consumed by users can be recorded using mobile devices [12]. Moreover, mobile personal assistants, such as Google Now and Microsoft Cortana, are capable of passively recording the digital activities of users. These recordings, which contain the activity patterns and preferences of users, can facilitate the development of personalised recommender systems capable of generating recommendations at the right time and in the right way for a given user and context [11, 31, 32].

In our previous work [15, 16], we proposed a generic activity recommendation framework to recommend the *next activity* to perform to a user. Our approach was applied successfully in the lifelogging and urban computing domains, where activities included socialising, eating, etc. and modes of transport, respectively. In this paper, we extend the activity recommendation framework to address the task of recommending a *sequence of activities* to the user. Moreover, we apply our framework to the tourism domain, where a recommended sequence of activities might be, for example, visiting a zoo, eating Italian food, and then listening to live music.

Our work is motivated by the assumption that people tend to repeat similar patterns of activities under similar circumstances [29]. Hence, in order to infer the next activities for a user, it is important to consider the activity patterns performed in the past. At the same time, the context surrounding these activities significantly affects the next activities the user performs. The importance of modelling context has been recognised in both tourism [6, 17, 18] as well as recommender systems research [1]. Context is particularly important in tourism as the user is predominantly mobile [10]. For example, features such as the time of day, location and weather can determine whether a user visits a particular amusement park in the city or not.

In recommender systems research, the task of recommending sequences is comparatively under-explored [13, 27]. However, there exists works, particularly for points of interest/itinerary (LBSN) [20, 30, 34] and music playlists [2, 4, 8, 22, 23, 27] recommendation, which address this task. A popular approach for modeling sequences has been Markov-based models [4] and all-k^{th}-order Markov models [3, 9, 25, 28]. However, in general, these approaches are not suitable for modelling sequences of activities with multiple features or context and are limited to the Markov assumption which does not apply in all cases [4]. An alternative hierarchical graph-based approach to capture sequences and geographical hierarchies in location trajectories is presented in [19]. This is further enhanced in [35] by modeling location popularity and user experiences to mine popular travel sequences across users in a non-personalised manner. Similarly, graph-based models have been used for collaborative itinerary recommendation [33]. However, these approaches do not capture the context information associated with user activities.

The key distinguishing characteristic of our work is that the model captures both the past activities of users, together with the context associated with these activities, in order to recommend the next sequence of activities for users to perform. The main contributions of this work can be summarised as follows:

- The extension of the generic activity recommendation framework in [15, 16] to recommend the next sequence of activities that should be performed by users. For this, an iterative, content-based recommendation approach is proposed, which takes the sequence as well as the features associated with

previous activity occurrences into consideration to build the recommendation model (Section 2);
- The application of our proposed algorithm to the tourism domain. Experiments using a location checkin dataset [21] demonstrate the efficacy of our approach in recommending sequences given a diverse variety of activities and user activity patterns (Section 3).

2 RECOMMENDATION APPROACH

In this section, we formulate the problem of recommending the next sequence of activities to a user. These activities can be, for example, eating Italian food, shopping at a bookstore, listening to live music, etc. The proposed content-based sequence recommendation algorithm leverages sequential patterns in a user's past activities as well as the contextual information (for example, time of day, location, weather, etc.) associated with each activity occurrence.

2.1 Problem Formulation

We introduced the concept of an *activity object* and an *activity timeline* in [15]. An *activity object*, ao_i, refers to a single occurrence of an activity and consists of a set of features, $ao_i = \{v_i^1, v_i^2, ..., v_i^m\}$, which describe the context surrounding that particular occurrence of the activity. For example, an activity object can refer to an instance of 'a visit to a zoo' (i.e. the *activity name*) with associated contextual features, such as *time of day, geo-location, weather, popularity of the location*, etc. An *activity timeline* (or *timeline* for short) for a user is then a chronological sequence of all activity objects performed by that user, $\mathcal{T} =< ao_1, ao_2, ..., ao_n >$.

2.2 Recommendation Algorithm

The proposed recommender is based on previous work [16], in which the past activities performed by a user were modelled as a timeline, $\mathcal{T}$, and the objective was to recommend the next activity to a user to perform. Here, we extend this approach to recommend the next *sequence of activities* for users to perform, $\mathcal{T}_{rec} =< ao_{rec_1}, ao_{rec_2}, ..., ao_{rec_L} >$.

Referring to Algorithm 1, a sequence of activities at a given recommendation time (*RT*) are generated as follows. The most recent activity object performed by the user, referred to as the *current activity object*, ao_c, is initialised as the activity object occurring at time *RT* in the user's timeline. The *current timeline*, $\mathcal{T}_c$, is then extracted from the user's timeline; it consists of the subsequence of the N activity objects occurring prior to ao_c and ends with ao_c (Step 1).

The recommendation of each activity object ao_{rec_i} in $\mathcal{T}_{rec}$ is performed iteratively (Step 4) as follows (see [16] for details). For each previous occurrence in the user's timeline of an activity with the same name as ao_c (e.g. 'Italian Food'), a *candidate timeline* ($\mathcal{T}_j$) is extracted (Step 5). Let $\mathscr{T}$ be the set of all candidate timelines in a given iteration. A two-level edit distance $\big(d(.,.)\big)$ between each candidate and the current timeline is computed [15]; based on these distances, a score (Eqn. 1) is assigned to the activity that occurs immediately after each candidate timeline $\mathcal{T}_j$ in $\mathscr{T}$ (Steps 7–8).

Algorithm 1: *SeqNCSeqRec*

Input: User, u; user's past timeline, $\mathcal{T}$; recommendation time, *RT*; current activity object, ao_c; N-count value, N
Output: a recommended timeline (sequence) $\mathcal{T}_{rec}$ of L activity objects, $\mathcal{T}_{rec} =< ao_{rec_1}, ao_{rec_2}, ...ao_{rec_i} ..., ao_{rec_L} >$

1. Extract the current timeline $\mathcal{T}_c$ from $\mathcal{T}$; the final element of $\mathcal{T}_c$ is ao_c
2. $\mathcal{T}_{rec} \leftarrow < >$
3. $i \leftarrow 1$
4. **while** $i \leq L$ **do**
5. Extract candidate timelines $\mathscr{T}$ from $\mathcal{T}$ (each $\mathcal{T}_j \in \mathscr{T}$ ends with an activity object ao_f^j such that ao_f^j.name = ao_c.name)
6. $\mathcal{R} \leftarrow \{\ \}$
7. **for** each $\mathcal{T}_j \in \mathscr{T}$ **do**
   $\mathcal{R} \leftarrow \mathcal{R} \cup ao_{f+1}^j$
8. **for** each $ao \in \mathcal{R}$ **do**
   Compute $Score(ao)$
9. ao_{rec_i}.name $\leftarrow$ top-1(ao.name : $ao \in \mathcal{R}$)
10. Compute and assign features to ao_{rec_i}
11. $\mathcal{T}_{rec} \leftarrow append(\mathcal{T}_{rec}, ao_{rec_i})$
12. $\mathcal{T}_c \leftarrow append(\mathcal{T}_c, ao_{rec_i})$
13. $RT \leftarrow ao_{rec_i}$.time
14. $i \leftarrow i + 1$
15. **return** $\mathcal{T}_{rec}$

From this set of scored activity objects, the top-1 activity name with the highest score is returned as the name for ao_{rec_i} in $\mathcal{T}_{rec}$ (Step 9). The values for the other features of ao_{rec_i} are then computed (Step 10) based on the average values for each feature from the user's past timeline. For example, if the recommended activity name is eating 'Italian Food', the time at which this activity should occur (ao_{rec_i}.time) is calculated as follows. The median difference between all occurrences of 'Italian Food' and the immediately preceeding activity in the user's past timeline is calculated; ao_{rec_i}.time is then given by the current recommendation time (*RT*) plus this difference.

Before the next iteration of the algorithm, ao_{rec_i} is appended to the current timeline $\mathcal{T}_c$ (and becomes the current activity object in the next iteration) (Step 12) and the recommendation time (*RT*) is set to ao_{rec_i}.time (Step 13). Thus, the L activity objects in the recommended timeline $\mathcal{T}_{rec}$ are generated in L iterations.

$$Score(ao) = 1 - \frac{d(\mathcal{T}_j, \mathcal{T}_c) - \min\limits_{\mathcal{T}_p \in \mathscr{T}} d(\mathcal{T}_p, \mathcal{T}_c)}{\max\limits_{\mathcal{T}_p \in \mathscr{T}} d(\mathcal{T}_p, \mathcal{T}_c) - \min\limits_{\mathcal{T}_p \in \mathscr{T}} d(\mathcal{T}_p, \mathcal{T}_c)} \ . \qquad (1)$$

2.2.1 Distance between Timelines. For the purpose of determining the similarity between two timelines $\mathcal{T}_1$ and $\mathcal{T}_2$, the two-level similarity algorithm proposed in our earlier work [15] is used. This algorithm first computes the minimum cost of rearranging the activities to achieve the same activity sequence and then aligns the values of the features of the corresponding activity objects. See [15] for further details on this approach.

2.2.2 N-count matching. The matching unit determines the length of the subsequences to be considered when calculating the distances between timelines. The *SeqNCSeqRec* algorithm uses the N-count matching approach as proposed in [16]. Thus, the N activity objects in the timeline preceding the current activity object form the current timeline (and likewise for candidate timelines). Note that the optimal value of N for each user will differ, depending on the degree of repetition and regularity of activities performed by each.

3 EVALUATION

We first describe the dataset used to construct activity timelines for users and the experimental methodology employed. This is followed by an evaluation of the proposed N-count based sequence recommender.

3.1 Dataset

For our experiments, we used a subset of the Gowalla checkins dataset [21]. The complete dataset obtained contains around 36 million checkins, 2.8 million locations and 0.3 million users. Every checkin is bound to a specific location and timestamp. A subset of these locations have categories assigned to them, such as, 'Italian Food', 'Bookstore', 'City Park', etc. These locations also have contextual features such as latitude, longitude, number of users checking in to it, number of photos taken at the location, etc. In relation to our recommendation framework, each of the location categories is considered as an 'activity name' and the recommendations made will be sequences of these categories. Hence, for evaluation, we select only those checkins locations which have assigned categories.

Further, categories are organised in a three-level hierarchy, consisting of 7, 134 and 151 level 1, 2, and 3 categories, respectively. For example, the level 1 category 'Food' has child categories 'African', 'American', 'Asian', 'Coffee Shop', etc. at level 2, while 'Coffee Shop' has child categories 'Starbucks' and 'Dunkin Donuts' at level 3. Given our objective is to recommend activities (categories) to users, we consider level 2 categories as the most suitable level of granularity, and hence any checkin locations with level 3 categories are assigned the parent category at level 2. As such, the names of activity objects in user timelines are given by the level 2 categories of the locations checked in to by users.

Since the characteristics of the timelines on weekdays and weekends are different, here we considered data corresponding to weekdays only. To address multiple consecutive checkins by users at the same location, we merged such checkins for a given user if they had the same category, were less than 600 meters apart and occurred within an interval of 10 minutes. Further, we selected only those users which have checkin data for at least 50 days with a minimum of 10 checkins per day. The sampled dataset had 916 users with 2.7 million checkins in total. The median number of checkins per day for users varied from 11–134, while the median number of distinct categories of checkins per day for users varied from 4–58.

3.2 Methodology

An offline evaluation was conducted for the proposed recommendation approach. Each user's complete timeline was split into training and test timelines, where the test timeline contained data for the most recent 20% of available days. For each user, a recommended sequence of categories of length 3 was generated at different *recommendation times* (RTs), which corresponded to the end time of each activity object in the test timeline. Recommendation performance is evaluated using *agreement @ k* ($k = 1, 2, 3$) which is the percentage of RTs for a user where the first k categories in the recommended sequence and the actual sequence are an exact match.

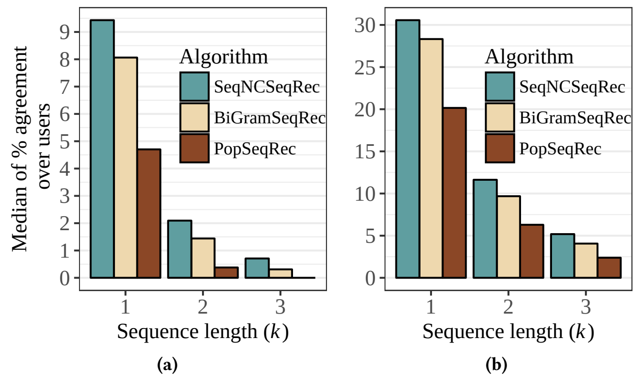

Figure 1: Median percentage agreements for recommended sequences for *SeqNCSeqRec* and baseline algorithms using timelines constructed from categories at (a) level 2 and (b) level 1 in the hierarchy.

For the computation of two-level edit distances between timelines, the following operation costs and feature weights were used: $c_{ins} = c_{del} = 1$, and $c_{sub} = 2$; $w_{category} = 2$, $w_{start-time} = 1$, $w_{popularity} = 1$, $w_{location} = 1$. These weights were set according to their hypothesised importance from the perspective of comparing timelines; for example, the weight associated with updating the category was set to the highest value since this is clearly a key consideration when computing distances between timelines. See [15] for details on the two-level edit distance approach.

3.3 Recommendation Performance

The performance of our proposed sequence-based N-count sequence recommendation algorithm (*SeqNCSeqRec*) is compared to the following baselines:

- The bi-gram-based sequence recommender (*BiGramSeqRec*) is based on the Markov assumption that the next activity depends only on the current activity. For each user, the frequency of occurrence of all activity name (category) bi-grams in the user's timeline are computed. For a given RT, a sequence of activity objects is recommended iteratively as per *SeqNCSeqRec* except that, at each iteration, the most frequently occurring bi-gram beginning with the current activity name is identified, and the recommended activity is simply that of the second element of this bi-gram. Such Markov-based approaches have proved to be quite successful in modelling sequences in previous studies [9].
- For a given user and RT, at each iteration of the algorithm, the popularity-based sequence approach (*PopSeqRec*) recommends the activity that the user performed most frequently at that time in the past.

3.3.1 Algorithm Performance. Figure 1(a) shows the median percentage agreements (k = 1, 2, 3) over all users for the proposed *SeqNCSeqRec* recommender and the two baselines. For *SeqNCSeqRec*, the results shown correspond to the optimal value of N-count for each user. It is clear from these results that the proposed approach significantly outperforms the baseline approaches. For example, *SeqNCSeqRec* improves upon *BiGramSeqRec* by 16.98%, 45.38%, and 129.3% for recommended sequences of length 1, 2, and 3, respectively, and improves upon *PopSeqRec* by more than 100% in all cases. Differences in results between the proposed and baselines algorithms are statistically significant (Wilcoxon-Mann-Whitney rank sum test) at the $p<.05$ level. The results also indicate that performance declines when larger sequences are recommended, which is to be expected, given the increased challenges involved in making such recommendations.

While the above findings are promising, it can be seen that the percentage agreements achieved by all algorithms are relatively low; for example, the percentage agreement is 9.5% for sequence lengths of 1 using *SeqNCountSeqRec*. We make the following observations in this regard. Firstly, as described in the previous section, in order to generate a sequence of recommendations, only the top-1 recommended activity is considered at each iteration of the *SeqNCSeqRec* algorithm. In addition, the evaluation is based on only a single recommended sequence being made to users, which clearly represents a strict approach.

Secondly, while many (although not all) level 2 categories are semantically similar, they are not considered a match according to the evaluation metric. For example, consider the level 2 categories 'Mexican' and 'South American/Latin' which relate to dining and are children of the level 1 category 'Food'. From a recommendation perspective, these different types of dining experiences are clearly related and (arguably) should represent a match. Thus, we also evaluate our recommender when all checkin locations are mapped to level 1 categories in the hierarchy – i.e. user timelines are constructed from activity objects with names given by the level 1 categories of locations checked in to by users. The results are shown in Figure 1(b). While similar trends as before are seen, the percentage agreements achieved are much greater; for example, over 30% for *SeqNCSeqRec* compared to the previous 9.5% for sequences of length 1. The 'true' performance of the recommender lies somewhere in between these values (since not all level 2 categories are semantically related); a further analysis of this matter is left to future work.

3.3.2 Performance across Users. A key intuition behind our approach is that the next activities performed by individual users depends, to a lesser or greater extent, on their past activity patterns. In the proposed *SeqNCSeqRec* recommendation algorithm, the number of past activities to be considered when generating recommendations is determined by the N-count value (see Section 2.2). In previous work [16], where the task was to recommend a single activity to users, it was seen that the optimal N-count value varied across users. In this section, we investigate whether a similar affect is seen when recommending sequences of activities to users.

As per [16], we hypothesise three distinct groups of users to capture the degree to which past activity patterns reflect future activity performance – Group 1: next activities are based on the current activity only (N-count = 0); Group 2: next activities are based on the current activity and a small number of past activities (N-count lies in the interval [1,4]); and Group 3: next activities are based on the current activity and a larger number of past activities (N-count = 5+).

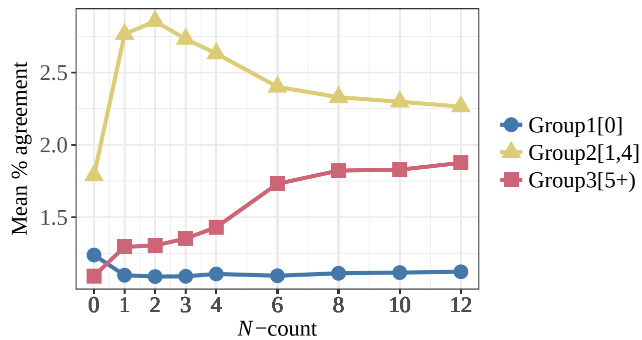

Figure 2: Mean percentage agreement for recommended sequences over users in each group.

In this experiment, users were assigned to one of the above groups based on the range in which their optimal N-count value appears (optimal in the sense that best percentage agreement was seen for sequences of length 3). Overall, 421, 374 and 121 users were assigned to Groups 1, 2 and 3, respectively. Results are shown in Figure 2. It can be seen that the mean recommendation performance for Group 1 users (46% of all users) was significantly lower than that seen for users in the other groups. This finding is to be expected, since it indicates that it is easier to recommend sequences of activities to users which are more consistent in their activity patterns. Thus, it can be concluded that adopting a personalised approach for users, by selecting the optimal N-count value for each user, is important. While it is not feasible to determine this value by experiment for large user bases, an approach to automatically learn a suitable value for individual users such as proposed in previous work [16] can be applied.

4 CONCLUSIONS AND FUTURE WORK

In this paper, we have expanded on our previous work to suggest sequences of activities for users based on past activity patterns. Notwithstanding the strict evaluation metric used in this work, the proposed approach shows promising performance and outperforms the baseline algorithms considered. In future work, we will investigate collaborative approaches in which candidate timelines will be drawn from the activities of other users in the system. Further, we will consider new approaches to suggest sequences of activities (for example, using RNNs) and investigate the recommendation of context (for example, where, when, with whom etc.) associated with each of the suggested sequence of activities.

5 ACKNOWLEDGMENTS

The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289.

REFERENCES

[1] Gediminas Adomavicius, Bamshad Mobasher, Francesco Ricci, and Alex Tuzhilin. 2011. Context-Aware Recommender Systems. *AI Magazine* 32, 3 (2011).

[2] Jie Bao, Yu Zheng, David Wilkie, and Mohamed Mokbel. 2015. Recommendations in Location-based Social Networks: A Survey. *Geoinformatica* 19, 3 (July 2015), 525–565. https://doi.org/10.1007/s10707-014-0220-8

[3] Thorsten Bohnenberger and Anthony Jameson. 2001. When Policies Are Better Than Plans: Decision-theoretic Planning of Recommendation Sequences. In *Proceedings of the 6th International Conference on Intelligent User Interfaces (IUI '01)*. ACM, New York, NY, USA, 21–24. https://doi.org/10.1145/359784.359829

[4] Geoffray Bonnin and Dietmar Jannach. 2014. Automated Generation of Music Playlists: Survey and Experiments. *Comput. Surveys* 47, 2, Article 26 (Nov. 2014), 35 pages. https://doi.org/10.1145/2652481

[5] Joan Borràs, Antonio Moreno, and Aida Valls. 2014. Intelligent Tourism Recommender Systems: A survey. *Expert Systems with Applications* 41, 16 (2014), 7370 – 7389. https://doi.org/10.1016/j.eswa.2014.06.007

[6] Matthias Braunhofer, Mehdi Elahi, Francesco Ricci, and Thomas Schievenin. 2013. *Context-Aware Points of Interest Suggestion with Dynamic Weather Data Management*. Springer-Verlag, Cham, 87–100. https://doi.org/10.1007/978-3-319-03973-2_7

[7] Dimitrios Buhalis. 2003. *eTourism: Information technology for strategic tourism management*. Pearson Education.

[8] Shuo Chen, Josh L. Moore, Douglas Turnbull, and Thorsten Joachims. 2012. Playlist Prediction via Metric Embedding. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*. ACM, New York, NY, USA, 714–722. https://doi.org/10.1145/2339530.2339643

[9] Mukund Deshpande and George Karypis. 2004. Selective Markov Models for Predicting Web Page Accesses. *ACM Transactions on Internet Technology* 4, 2 (May 2004), 163–184. https://doi.org/10.1145/990301.990304

[10] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati Pantziou. 2014. Mobile recommender systems in tourism. *Journal of Network and Computer Applications* 39 (2014), 319 – 333. https://doi.org/10.1016/j.jnca.2013.04.006

[11] Ramanathan Guha, Vineet Gupta, Vivek Raghunathan, and Ramakrishnan Srikant. 2015. User Modeling for a Personal Assistant. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining (WSDM '15)*. ACM, New York, NY, USA, 275–284. https://doi.org/10.1145/2684822.2685309

[12] Cathal Gurrin, Alan F. Smeaton, and Aiden R. Doherty. 2014. LifeLogging: Personal Big Data. *Foundations and Trends in Information Retrieva* 8, 1 (June 2014), 1–125. https://doi.org/10.1561/1500000033

[13] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems* 22, 1 (Jan. 2004), 5–53. https://doi.org/10.1145/963770.963772

[14] Katerina Kabassi. 2010. Personalizing Recommendations for Tourists. *Telematics and Informatics* 27, 1 (2010), 51 – 66. https://doi.org/10.1016/j.tele.2009.05.003

[15] Gunjan Kumar, Houssem Jerbi, Cathal Gurrin, and Michael P. O'Mahony. 2014. Towards Activity Recommendation from Lifelogs. In *Proceedings of the 16th International Conference on Information Integration and Web-based Applications & Services (iiWAS '14)*. ACM, New York, NY, USA, 87–96. https://doi.org/10.1145/2684200.2684298

[16] Gunjan Kumar, Houssem Jerbi, and Michael P. O'Mahony. 2016. Personalised Recommendations for Modes of Transport: A Sequence-based Approach. *The 5th ACM SIGKDD International Workshop on Urban Computing (UrbComp 2016)* (2016).

[17] Carlos Lamsfus, David Martin, Zigor Salvador, Alex Usandizaga, and Aurkene Alzua-Sorzabal. 2009. Human-Centric Ontology-Based Context Modelling In Tourism. *Mediterranean Conference on Information Systems* (2009).

[18] Carlos Lamsfus, Dan Wang, Aurkene Alzua-Sorzabal, and Zheng Xiang. 2015. Going Mobile. *Journal of Travel Research* 54, 6 (2015), 691–701. https://doi.org/10.1177/0047287514538839

[19] Quannan Li, Yu Zheng, Xing Xie, Yukun Chen, Wenyu Liu, and Wei-Ying Ma. 2008. Mining User Similarity Based on Location History. In *Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '08)*. ACM, New York, NY, USA, Article 34, 10 pages. https://doi.org/10.1145/1463434.1463477

[20] H. Liu, L. Y. Wei, Y. Zheng, M. Schneider, and W. C. Peng. 2011. Route Discovery from Mining Uncertain Trajectories. In *11th IEEE International Conference on Data Mining Workshops*. 1239–1242. https://doi.org/10.1109/ICDMW.2011.149

[21] Xin Liu, Yong Liu, Karl Aberer, and Chunyan Miao. 2013. Personalized Point-of-interest Recommendation by Mining Users' Preference Transition. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management (CIKM '13)*. ACM, New York, NY, USA, 733–738. https://doi.org/10.1145/2505515.2505639

[22] Brian McFee and Gert RG Lanckriet. 2011. The Natural Language of Playlists. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR)*.

[23] Brian McFee and Gert RG Lanckriet. 2012. Hypergraph Models of Playlist Dialects. In *Proceedings of the 13th International Conference on Music Information Retrieval (ISMIR)*.

[24] Antonio Moreno, Aida Valls, David Isern, Lucas Marin, and Joan Borràs. 2013. SigTur/E-Destination: Ontology-based Personalized Recommendation of Tourism and Leisure Activities. *Engineering Applications of Artificial Intelligence* 26, 1 (Jan. 2013), 633–651. https://doi.org/10.1016/j.engappai.2012.02.014

[25] James Pitkow and Peter Pirolli. 1999. Mining Longest Repeating Subsequences to Predict World Wide Web Surfing. In *Proceedings of the 2nd Conference on USENIX Symposium on Internet Technologies and Systems - Volume 2 (USITS'99)*. USENIX Association, Berkeley, CA, USA, 13–13. http://dl.acm.org/citation.cfm?id=1251480.1251493

[26] Francesco Ricci. 2002. Travel Recommender Systems. *IEEE Intelligent Systems* (2002).

[27] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor. 2010. *Recommender Systems Handbook* (1st ed.). Springer-Verlag New York, Inc., New York, NY, USA.

[28] Guy Shani, David Heckerman, and Ronen I. Brafman. 2005. An MDP-Based Recommender System. *Journal of Machine Learning Research* 6 (Dec. 2005), 1265–1295. http://dl.acm.org/citation.cfm?id=1046920.1088715

[29] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of Predictability in Human Mobility. *Science* 327, 5968 (2010), 1018–1021. https://doi.org/10.1126/science.1177170

[30] Chih-Hua Tai, De-Nian Yang, Lung-Tsai Lin, and Ming-Syan Chen. 2008. Recommending Personalized Scenic Itinerary with Geo-tagged Photos. In *IEEE International Conference on Multimedia and Expo*. 1209–1212. https://doi.org/10.1109/ICME.2008.4607658

[31] Iis P. Tussyadiah and Dan Wang. 2016. Tourists' Attitudes toward Proactive Smartphone Systems. *Journal of Travel Research* 55, 4 (2016), 493–508. https://doi.org/10.1177/0047287514563168

[32] Artem Umanets, Artur Ferreira, and Nuno Leite. 2014. GuideMe – A Tourist Guide with a Recommender System and Social Interaction. *Procedia Technology* 17 (2014), 407 – 414. https://doi.org/10.1016/j.protcy.2014.10.248 Conference on Electronics, Telecommunications and Computers – CETC 2013.

[33] Hyoseok Yoon, Yu Zheng, Xing Xie, and Woontack Woo. 2010. Smart Itinerary Recommendation Based on User-generated GPS Trajectories. In *Proceedings of the 7th International Conference on Ubiquitous Intelligence and Computing (UIC'10)*. Springer-Verlag, Berlin, Heidelberg, 19–34. http://dl.acm.org/citation.cfm?id=1929661.1929669

[34] Hyoseok Yoon, Yu Zheng, Xing Xie, and Woontack Woo. 2012. Social Itinerary Recommendation from User-generated Digital Trails. *Personal Ubiquitous Comput.* 16, 5 (June 2012), 469–484. https://doi.org/10.1007/s00779-011-0419-8

[35] Yu Zheng, Lizhu Zhang, Xing Xie, and Wei-Ying Ma. 2009. Mining Interesting Locations and Travel Sequences from GPS Trajectories. In *Proceedings of the 18th International Conference on World Wide Web (WWW '09)*. ACM, New York, NY, USA, 791–800. https://doi.org/10.1145/1526709.1526816

