Exploiting Contextual Information for Recommender Systems Oriented to Tourism

Pablo Sánchez
pablo.sanchezp@uam.es
Information Retrieval Group
Universidad Autónoma de Madrid
Madrid, Spain

ABSTRACT

The use of contextual information like geographic, temporal (including sequential), and item features in Recommender Systems has favored their development in several different domains such as music, news, or tourism, together with new ways of evaluating the generated suggestions. This paper presents the underlying research in a PhD thesis introducing some of the fundamental considerations of the current tourism-based models, emphasizing the Point-Of-Interest (POI) problem, while proposing solutions using some of these additional contexts to analyze how the recommendations are made and how to enrich them. At the same time, we also intend to redefine some of the traditional evaluation metrics using contextual information to take into consideration other complementary aspects beyond item relevance. Our preliminary results show that there is a noticeable popularity bias in the POI recommendation domain that has not been studied in detail so far; moreover, the use of contextual information (such as temporal or geographical) help us both to improve the performance of recommenders and to get better insights of the quality of provided suggestions.

CCS CONCEPTS

• **Information systems → Recommender systems**; **Retrieval effectiveness**.

KEYWORDS

Recommender Systems; Point-Of-Interest recommendation; Contextual Information; Evaluation

ACM Reference Format:
Pablo Sánchez. 2019. Exploiting Contextual Information for Recommender Systems Oriented to Tourism. In *Thirteenth ACM Conference on Recommender Systems (RecSys '19), September 16–20, 2019, Copenhagen, Denmark.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3298689.3347062

1 INTRODUCTION

Since the emergence of the first Recommender Systems (RS) in the early 1990s, their main objective remains the same: to provide hypothetically interesting recommendations to users by analyzing their tastes and needs [36]. Many recommendation strategies have been proposed in the literature to achieve this goal, with collaborative filtering, content-based, and hybrid approaches being the most widespread ones [2], normally oriented at predicting the missing values of the user-item matrix. However, in recent years it has been observed that incorporating additional contextual information, in some cases, is key to improve the performance of recommenders [3]. Some of these additional contexts include temporal and geographical information, item features or even the weather, and although they have been applied to various recommendation domains, some of these contexts are especially useful in the Point-of-Interest (POI) area, that is, in the tourism domain.

In fact, the development of POI recommender systems has been recently motivated by the rise of Location-Based Social Networks (LBSN), mostly due to the global increase on tourism, prevalent use of mobile phones, and the pervasiveness of social networks in general. In these LBSN the users record check-ins they make to certain venues and share information with other users in the system [49]. This information can be exploited and enriched by a POI recommender in order to make suggestions of new places to visit for the users.

Even if the general goal is the same as in traditional recommendation, the POI recommendation task has some particularities that differ from other, more widespread domains such as movies or books, and hence, there are specific details that require further analysis. Some of the most important ones are shown hereunder [26, 29, 46]:

- Sparsity: while the user-item matrix in RS is normally very sparse, in POI recommendation the data sparsity is even more severe. For example, the densities for the Movielens20M and Netflix datasets are 0.539% and 1.77% respectively, while the densities of the check-ins dataset from Foursquare [48] and Gowalla [13] are 0.0034% and 0.0047% respectively.
- Implicit and repeated information: in many typical datasets the interactions between the users and the system is encoded through ratings, whereas in POI recommendation the only available information is the moment at which a user visited a specific venue. In addition, since users may check-in several times on the same POI (something that in traditional recommendation is not usually taken into account [32]), some researchers build frequency matrices in order to model those repeated preferences.
- External influences: besides the collaborative information that can be obtained from the user-item matrix, POI recommendation is highly affected by temporal, social (user friends), and geographical influence. The latter is the most important effect to consider in this kind of recommendation as there is a clear correlation between the places the user

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
RecSys '19, September 16–20, 2019, Copenhagen, Denmark
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6243-6/19/09.
https://doi.org/10.1145/3298689.3347062

visits and the distance between them [29, 30]. However, temporal information is also critical in POI recommendation: it is not only necessary to know the schedules or the availability of the venues but it is also useful to analyze seasonal trends and establish sequential patterns from the items.

At the same time, the evaluation of both POI and traditional RS has typically been oriented towards the analysis of the relevance of recommendations. That is, by establishing a comparison between the ground truth and the users' recommendations using classic ranking metrics from Information Retrieval such as Precision, Recall, or NDCG. However, even though there have been several analyses of alternative evaluation perspectives such as novelty, diversity, and serendipity [11] in traditional recommendation, it is still necessary to delve deeper into these aspects for the tourism domain, as well as considering and analyzing other possible biases inherent in this specific type of recommendation since these special characteristics are normally ignored in the experiments reported in the area.

Therefore, motivated by the previously mentioned open issues, the research of this PhD thesis is focused on analyzing new context-aware perspectives of the POI recommendation domain (especially by exploiting the temporal and geographical component in offline experiments), including the adaptation of these perspectives towards how should the quality of the recommendations be assessed (i.e., also adapting the evaluation step). Firstly, for POI recommendation we intend to exploit some of its intrinsic characteristics and see the effects and biases that different POI approaches may have in their recommendations to be able to explain better their performance. At the same time, we also plan to develop new POI recommendation strategies by taking advantage of the above-mentioned contexts to propose new models and improve the performance of existing strategies. Finally, regarding the evaluation aspect, we aim to incorporate these additional contexts (such as the temporal, sequential, or geographical information) into the evaluation metrics to analyze other dimensions beyond accuracy, if possible, under a unified evaluation framework.

2 RELATED WORK

2.1 Sequential and POI Recommendation

Traditionally, the recommendation task has been formulated as an optimization problem [2]:

$$i^*(u) = \underset{i \in I}{\arg\max}\ g(u, i) \tag{1}$$

where the goal is to find the optimal item i^* which maximizes the interest of user u on any item $i \in I$, measured by function g. Usually, this utility function has been modeled by exploiting the user-item matrix of the system, leaving aside contextual information that may be important for improving the recommendations. Temporal information is one of the most important contexts to take into account – and one of the easiest to capture too as we only need to store the moment when the user consumed the item. It has been successfully included in some approaches like Matrix Factorization ([25]) or nearest neighbors ([9, 15]). Additionally, it has also favored the emergence of recommender systems specially designed to model sequences of interactions in a more natural way [34].

There is an important number of works proposing recommendation models that somehow exploit the sequential information. For example, the authors of [35] proposed the FPMC (Factorizing Personalized Markov Chain) model, by combining Matrix Factorization (MF) techniques and Markov Chains (MC) into a single algorithm. Another hybrid approach is presented in [21] where the authors combined an L-order MC with a personalization approach and Factored Item Similarity Models (FISM) into an algorithm named Fossil (Factorized Sequential Prediction with Item Similarity Models), outperforming many state-of-the-art algorithms, including the FISM techniques, FPMC, and Bayesian Personalized Ranking (BPR) models, although the researchers did not analyze the performance of other traditional recommenders such as standard MF approaches or neighborhood-based algorithms. However, one the most fruitful strategies nowadays in the area to model sequences are Neural Networks. Some examples include the GRU4Rec model, proposed in [22] to capture the sequential dependencies between the user preferences in order to make recommendations, and Caser, a Convolutional Neural Network method proposed in [44], that exploits both the sequential patterns of the data and the general preferences of the users in order to make the predictions. There are other approaches including Multi Layer Perceptrons, Autoencoders, Recurrent and Convolutional Neural Networks; for a complete survey of these and other approaches, we refer the reader to [52].

Besides the temporal information, the geographical influence – as discussed before – is another important context in the POI recommendation domain. Several POI-specific approaches model this aspect specifically, including Matrix Factorization schemes (such as [26, 27, 30] and in particular [16] that also models the temporal context) or hybrid models combining both geographical and social influence (user friends) [49–51]. It should be mentioned that the research on POI recommendation is also related to the so-called Tourist-Trip Design problem (TTDP) [17], where the objective is no longer to recommend independent POIs but to be able to generate meaningful routes taking into account additional considerations such as the travel time, POI availability, or the distance of the total route. However, we have detected a gap between route recommendation and some of the aforementioned models as sometimes the personalization strategies or the user satisfaction are left in the background ([5, 14, 47]) or they do not take into account some of the important features of the TTDP as POI availability or travel time ([12, 41]). In this sense, we believe that route recommendation proposals can benefit from both areas of User Modeling and Recommender Systems to improve the personalization of recommendations by relaxing some of their original restrictions.

2.2 Recommender Systems Evaluation

The research advances made in the area of RS has also affected the way of measuring the usefulness of the recommendations produced. Although error metrics such as MAE or RMSE were initially used in the area to analyze the performance of the algorithms (mainly motivated by the Netflix prize from 2009), their use was displaced by other metrics used in the Information Retrieval area (like Precision, Recall, or NDCG) in order to measure the quality of item rankings, not simply by individual predictions. However, even if these metrics allow us to perform a better analysis of the quality of the recommendations, there are some authors who have warned about the necessity to assess other dimensions beyond relevance.

Two well-known examples of these complementary dimensions are novelty and diversity, aimed at measuring whether the recommended items differ from what was consumed in the past (novelty) or whether the items are sufficiently different from each other (diversity). In [11] the authors make an exhaustive analysis of both novelty and diversity summarizing some important metrics that are used in this context; then, they propose in [45] a probabilistic framework to derive some of these metrics using a common formulation. Moreover, dimensions like robustness, confidence, or serendipity have also been proposed and studied in other works [19]; however, due to the large number of dimensions that can be evaluated, it is widely accepted that finding a model outperforming every other recommender in all dimensions is a difficult problem to solve [18].

In the area of POI recommendation, some works have proposed complementary metrics using additional information to take advantage of the available contexts mentioned in Section 1. For instance, in [8] the authors compute the popularity score and visiting time of the POIs as well as the Recall at the category level instead of the item level (something also proposed in [20] for Precision), whereas in both [28] and [33], the authors proposed metrics specifically designed to exploit the categories of the POIs. These metrics allow, even though through an ad-hoc formulation, to perform a deeper and more complete analysis of the recommendations produced by POI recommenders. However, additional contexts remain unexplored – to the best of our knowledge – in terms of evaluation dimensions. The temporal information and the sequences followed by the users, although exploited when performing recommendations (e.g., in Neural Networks and Markov Chains models), they have been mostly ignored when evaluating the quality of the recommendations (except in few works such as [12], where the authors propose an adaptation of the F1 metric to take into account the pairwise order of users visiting the POIs). Besides, exploiting the sequential component in the evaluation step could be useful in other domains, such as music, where sequences (in this case, in the form of playlists) are also a natural way of representing the user consumptions [43].

Besides, considering recent works where popularity and other biases are showcased [6, 23], we believe it is necessary to perform a more in-depth study on the tourism domain to understand the behavior of the algorithms under its specific and intrinsic properties, such as repeated consumption, high sparsity, and the tendency to prefer venues that are closer to each other [31].

3 RESEARCH GOALS

The purpose of this thesis is twofold: on the one hand, we aim to advance in the research of the POI recommendation problem by improving the performance of the models through incorporating other contexts and exploiting sequences in the data; on the other hand, we expect to provide an exhaustive analysis on the evaluation of the recommendations by defining new metrics that explore other dimensions besides relevance. The whole thesis will be oriented towards an offline experimental setup.

Based on this, first of all, we want to **make an in-depth study of the POI recommendation problem by analyzing its special features and considerations** and see the effects that they have in different algorithms. As a first step, we propose to compare classical recommendation algorithms against specific POI recommendation proposals (something normally ignored in the area) in order to analyze their performance in different datasets together with an exhaustive analysis of the different biases that may exist in this domain, as discussed in the previous section.

Secondly, we aim to **propose novel mechanisms to generate recommendations by exploring new ways of exploiting user preferences and incorporating sequential components oriented to the POI recommendation domain**. Although there are specific approaches to model both the temporal and sequential components (see Section 2.1), this part of the thesis would be more focused on algorithms that are computationally less expensive and more easily interpretable, such as nearest neighbors. Besides, since the proposed algorithms would somehow exploit the temporal information available in the datasets, time-aware evaluation settings should be used [9], a practice not so common in the area, which would require to perform a thorough comparison of the state-of-the-art algorithms and the effect on performance of their parameters, as well as how temporal splits affect the typical process of parameter tuning. At the same time, these novel mechanisms should be able to benefit from the available geographical information.

Finally, and to cover all the necessary steps to define a full recommender system, we plan to **incorporate contextual features into offline evaluation metrics** to redefine the concept of "relevance". By doing so, we would obtain a set of complementary metrics that will allow us to analyze the specific contexts available in our problem. The candidate features that could potentially be included into classical offline evaluation metrics are the temporal information (both from the perspective of the users – as the sequences followed by them or how repeated interactions may affect the recommenders – or of the items – as their freshness or temporal novelty), content (such as the categories), and the geographical context (by measuring the distances of the produced recommendations). In general, we are interested in incorporating any information source that could help us understanding how and why some recommenders may perform better than others. As a result, we plan to use this information to propose novel algorithms that would be considered in the previous research goal, reinforcing the whole recommendation cycle – being careful that the metrics are as independent as possible from the context features used by the recommenders, otherwise, the evaluation methodology could be tuned towards those features used in the algorithms. Although some of the metrics derived will be specific to the POI recommendation area, we intend to generalize as much as possible the developed framework so that the recommendation analysis of any domain can benefit from it.

4 PRELIMINARY RESULTS

In this section, we present a brief summary of some of the results obtained so far when addressing the research goals described before.

As a starting point to investigate in the area of POI recommendation, we decided to analyze some of the particularities of this specific recommendation problem in more detail [37]. Firstly, we observed that the results obtained using classic ranking metrics were much lower than in other domains like movies, mostly due to the high sparsity and the high number of POIs to recommend. Secondly, we observed that in POI recommendation there is a strong popularity bias (also motivated by the great sparsity in typical datasets), which means that the basic popularity recommender obtains a similar performance (and sometimes better) than other more complex models like nearest neighbors or Matrix Factorization. Besides, we found that repeated interactions had a huge impact in the recommenders performance: through different temporal splits (either allowing the recommenders to return POIs already seen by the user or not) we observed that the results were remarkably higher when in the test set we allowed repeated preferences, where instead of popularity, a simple recommendation algorithm that returns just the POIs already visited by the user would outperform most of the tested methods. In this work (and in others that we are currently developing), we used the global check-in Foursquare dataset from [48], which includes more than 33M check-ins on 415 cities covering 18 months of users interactions; for our experiments, we performed a temporal split and decided to work with isolated cities since there is no item overlap between them.

Furthermore, and as a preliminary work prior to the thesis, we investigated some mechanisms to add sequentiality in traditional Recommender Systems. In this regard, we found that the Longest Common Subsequence (LCS) algorithm [4] could be used in recommendation as a similarity metric (instead of cosine or Pearson) by creating a method that transforms user preferences into sequences. This proposal was presented in [7] (for movies and music) and [40] (only movies but extending the framework to also exploit the item features), although the sequences were not generated using temporal information, and this was left as a future work and to be exploited in other domains with more realistic interactions. We have also explored specific techniques for the tourism domain by studying cross-domain strategies, where each city could be seen as a different domain under common definitions in the area [10], since items from different cities do not overlap with each other. Inspired by the works where diversity is injected into retrieval and recommendation lists [11, 42], and more recently, popularity and novelty [1, 24], we have analyzed the feasibility of reranking techniques to increase the sequentiality of the recommendations, by starting with pure sequence-agnostic recommendations, and building sequences according to different reranking strategies.

Finally, regarding the research goal related to evaluation, we have advanced on different but complementary directions. Firstly, we extended the novelty and diversity framework of [45] to integrate the temporal novelty (*freshness*) of the items, in such a way that items with more recent preferences will be more temporary novel or *fresh* [39]. We observed some biases in the results obtained, like temporal activity bursts and, in some datasets, a correlation between the item ids and their age in the system. Besides, we also observed that, usually, the recommenders that were able to retrieve more fresh items, were also the best in terms of relevance, evidencing that in some cases the users are more attracted to newer items instead than to older ones. Furthermore, we also observed that the user coverage of recommenders was an aspect that needs a detailed analysis when performing a temporal split as sometimes there is a large number of new users in the ground truth. Additionally, in [38] we proposed a new and generic (anti) relevance model to measure the recommendations that the user specifically indicated she did not like, this refers to those preferences in the test set with low ratings. The main conclusion in our study was that the mere fact of learning the tastes of users with personalized models (like BPR, nearest neighbors, or Matrix Factorization) does not avoid to recommend items the user scored with a low rating in the ground truth; on the other hand, non-personalized algorithms like random or popularity, even though their performance was worse in terms of relevance, they also recommended fewer anti-relevant items in the test set, which was a surprising and interesting result. Finally, and as a first attempt on incorporating relevance and sequentiality in evaluation metrics, in [37] we reformulated the LCS algorithm to be integrated in traditional evaluation metrics and to produce, at least, results comparable to Precision and Recall but considering the sequence followed by the user in the test set. The results were promising, although not many changes were found between the sequential and classical variations of these metrics, mostly due to the very small number of relevant items included in the test set. In the future, we aim to continue studying this problem and extending the experiments with further datasets and evaluation settings. In fact, we believe these analyses can be applied to other recommendation domains besides tourism as long as they share some characteristics; for instance, the music domain could be a good candidate since it also suffers from low sparsity, repeated consumptions, and the sequential information is very important [43].

5 CONCLUSIONS

In this paper we have presented the current status of a thesis project aimed at developing new perspectives for Recommender Systems oriented to the tourism domain (with special emphasis on POI recommender systems) where additional contexts like geographical and temporal information are used to improve and evaluate recommendations. Firstly, we have analyzed some of the differences between the POI recommendation area and traditional recommendation, emphasizing the importance of contextual information (highlighting the geographical influence in particular) and data sparsity, and we have also described some of the strategies used for both POI and sequential recommendation. At the same time, we have also proposed the extension of traditional Information Retrieval metrics used to measure the quality of the recommendations to take into account other dimensions besides relevance, for instance by considering the temporal or geographical context and information associated with users and items when assessing the usefulness of the recommendations for the user.

Although we have a preliminary basis of published papers validating some of our hypotheses, we believe that further research is needed on some of the presented aspects. Specifically, we consider that the next steps must involve the use of contextual information to analyze additional biases in the datasets for generating novel POI recommenders and the analysis of the different recommenders under the same evaluation methodology using the different, but complementary, metrics proposed in this thesis.

ACKNOWLEDGMENTS

This work has been funded by the Ministerio de Ciencia, Innovación y Universidades (reference: TIN2016-80630-P) and by the European Social Fund (ESF), within the 2017 call for predoctoral contracts.

REFERENCES

[1] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing Popularity Bias in Recommender Systems with Personalized Re-ranking. *CoRR* abs/1901.07555 (2019).

[2] Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Trans. Knowl. Data Eng.* 17, 6 (2005), 734–749.

[3] Gediminas Adomavicius and Alexander Tuzhilin. 2015. Context-Aware Recommender Systems. In *Recommender Systems Handbook*. Springer, 191–226.

[4] Alberto Apostolico. 1997. String Editing and Longest Common Subsequences. In *Handbook of Formal Languages 2*, Grzegorz Rozenberg and Arto Salomaa (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 361–398.

[5] Victor Anthony Arrascue Ayala, Kemal Cagin Gülsen, Anas Alzogbi, Michael Färber, Marco Muñiz, and Georg Lausen. 2017. A Delay-Robust Touristic Plan Recommendation Using Real-World Public Transportation Information. In *RecTour@RecSys (CEUR Workshop Proceedings)*, Vol. 1906. CEUR-WS.org, 9–17.

[6] Alejandro Bellogín, Pablo Castells, and Iván Cantador. 2017. Statistical biases in Information Retrieval metrics for recommender systems. *Inf. Retr. Journal* 20, 6 (2017), 606–634.

[7] Alejandro Bellogín and Pablo Sánchez. 2017. Collaborative filtering based on subsequence matching: A new approach. *Inf. Sci.* 418 (2017), 432–446.

[8] Igo Ramalho Brilhante, José Antônio Fernandes de Macêdo, Franco Maria Nardini, Raffaele Perego, and Chiara Renso. 2013. Where shall we go today?: planning touristic tours with tripbuilder. In *CIKM*. ACM, 757–762.

[9] Pedro G. Campos, Fernando Díez, and Iván Cantador. 2014. Time-aware recommender systems: a comprehensive survey and analysis of existing evaluation protocols. *User Model. User-Adapt. Interact.* 24, 1-2 (2014), 67–119.

[10] Iván Cantador, Ignacio Fernández-Tobías, Shlomo Berkovsky, and Paolo Cremonesi. 2015. Cross-Domain Recommender Systems. In *Recommender Systems Handbook*. Springer, 919–959.

[11] Pablo Castells, Neil J. Hurley, and Saul Vargas. 2015. Novelty and Diversity in Recommender Systems. In *Recommender Systems Handbook*. Springer, 881–918.

[12] Dawei Chen, Cheng Soon Ong, and Lexing Xie. 2016. Learning Points and Routes to Recommend Trajectories. In *CIKM*. ACM, 2227–2232.

[13] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *KDD*. ACM, 1082–1090.

[14] Munmun De Choudhury, Moran Feldman, Sihem Amer-Yahia, Nadav Golbandi, Ronny Lempel, and Cong Yu. 2010. Automatic construction of travel itineraries using social breadcrumbs. In *HT*. ACM, 35–44.

[15] Yi Ding and Xue Li. 2005. Time weight collaborative filtering. In *CIKM*. ACM, 485–492.

[16] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2013. Exploring temporal effects for location recommendation on location-based social networks. In *RecSys*. ACM, 93–100.

[17] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati E. Pantziou. 2014. A survey on algorithmic approaches for solving tourist trip design problems. *J. Heuristics* 20, 3 (2014), 291–328.

[18] Asela Gunawardana and Guy Shani. 2009. A Survey of Accuracy Evaluation Metrics of Recommendation Tasks. *Journal of Machine Learning Research* 10 (2009), 2935–2962.

[19] Asela Gunawardana and Guy Shani. 2015. Evaluating Recommender Systems. In *Recommender Systems Handbook*. Springer, 265–308.

[20] Jing He, Xin Li, and Lejian Liao. 2017. Category-aware Next Point-of-Interest Recommendation via Listwise Bayesian Personalized Ranking. In *IJCAI*. ijcai.org, 1837–1843.

[21] Ruining He and Julian McAuley. 2016. Fusing Similarity Models with Markov Chains for Sparse Sequential Recommendation. In *ICDM*. IEEE, 191–200.

[22] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *International Conference on Learning Representations ICLR 2016, May 2-4, 2016, Caribe Hilton, San Juan, Puerto Rico*.

[23] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What recommenders recommend: an analysis of recommendation biases and possible countermeasures. *User Model. User-Adapt. Interact.* 25, 5 (2015), 427–491.

[24] Marius Kaminskas and Derek Bridge. 2017. Diversity, Serendipity, Novelty, and Coverage: A Survey and Empirical Analysis of Beyond-Accuracy Objectives in Recommender Systems. *TiiS* 7, 1 (2017), 2:1–2:42.

[25] Yehuda Koren. 2009. Collaborative filtering with temporal dynamics. In *KDD*. ACM, 447–456.

[26] Xutao Li, Gao Cong, Xiaoli Li, Tuan-Anh Nguyen Pham, and Shonali Krishnaswamy. 2015. Rank-GeoFM: A Ranking based Geographical Factorization Method for Point of Interest Recommendation. In *SIGIR*. ACM, 433–442.

[27] Defu Lian, Cong Zhao, Xing Xie, Guangzhong Sun, Enhong Chen, and Yong Rui. 2014. GeoMF: joint geographical modeling and matrix factorization for point-of-interest recommendation. In *KDD*. ACM, 831–840.

[28] Kwan Hui Lim, Jeffrey Chan, Christopher Leckie, and Shanika Karunasekera. 2015. Personalized Tour Recommendation Based on User Interests and Points of Interest Visit Durations. In *IJCAI*. AAAI Press, 1778–1784.

[29] Yiding Liu, Tuan-Anh Pham, Gao Cong, and Quan Yuan. 2017. An Experimental Evaluation of Point-of-interest Recommendation in Location-based Social Networks. *PVLDB* 10, 10 (2017), 1010–1021.

[30] Yong Liu, Wei Wei, Aixin Sun, and Chunyan Miao. 2014. Exploiting Geographical Neighborhood Characteristics for Location Recommendation. In *CIKM*. ACM, 739–748.

[31] Harvey J. Miller. 2004. Tobler's First Law and Spatial Analysis. *Annals of the Association of American Geographers* 94, 2 (2004), 284–289. https://doi.org/10.1111/j.1467-8306.2004.09402005.x

[32] Xia Ning, Christian Desrosiers, and George Karypis. 2015. A Comprehensive Survey of Neighborhood-Based Recommendation Methods. In *Recommender Systems Handbook*. Springer, 37–76.

[33] Enrico Palumbo, Giuseppe Rizzo, Raphaël Troncy, and Elena Baralis. 2017. Predicting Your Next Stop-over from Location-based Social Network Data with Recurrent Neural Networks. In *RecTour@RecSys (CEUR Workshop Proceedings)*, Vol. 1906. CEUR-WS.org, 1–8.

[34] Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-Aware Recommender Systems. *ACM Comput. Surv.* 51, 4 (2018), 66:1–66:36.

[35] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *WWW*. ACM, 811–820.

[36] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender Systems: Introduction and Challenges. In *Recommender Systems Handbook*. Springer, 1–34.

[37] Pablo Sánchez and Alejandro Bellogín. 2018. Challenges on Evaluating Venue Recommendation Approaches: Position Paper. In *RecTour@RecSys (CEUR Workshop Proceedings)*, Vol. 2222. CEUR-WS.org, 37–40.

[38] Pablo Sánchez and Alejandro Bellogín. 2018. Measuring anti-relevance: a study on when recommendation algorithms produce bad suggestions. In *RecSys*. ACM, 367–371.

[39] Pablo Sánchez and Alejandro Bellogín. 2018. Time-Aware Novelty Metrics for Recommender Systems. In *ECIR (Lecture Notes in Computer Science)*, Vol. 10772. Springer, 357–370.

[40] Pablo Sánchez and Alejandro Bellogín. 2019. Building user profiles based on sequences for content and collaborative filtering. *Inf. Process. Manage.* 56, 1 (2019), 192–211.

[41] Jitao Sang, Tao Mei, Jian-Tao Sun, Changsheng Xu, and Shipeng Li. 2012. Probabilistic sequential POIs recommendation via check-in data. In *SIGSPATIAL/GIS*. ACM, 402–405.

[42] Rodrygo L. T. Santos, Craig MacDonald, and Iadh Ounis. 2015. Search Result Diversification. *Foundations and Trends in Information Retrieval* 9, 1 (2015), 1–90.

[43] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current challenges and visions in music recommender systems research. *IJMIR* 7, 2 (2018), 95–116.

[44] Jiaxi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. In *WSDM*. ACM, 565–573.

[45] Saul Vargas and Pablo Castells. 2011. Rank and relevance in novelty and diversity metrics for recommender systems. In *RecSys*. ACM, 109–116.

[46] Hao Wang, Manolis Terrovitis, and Nikos Mamoulis. 2013. Location recommendation in location-based social networks using user check-in data. In *SIGSPATIAL/GIS*. ACM, 364–373.

[47] Wolfgang Wörndl, Alexander Hefele, and Daniel Herzog. 2017. Recommending a sequence of interesting places for tourist trips. *J. of IT & Tourism* 17, 1 (2017), 31–54.

[48] Dingqi Yang, Daqing Zhang, and Bingqing Qu. 2016. Participatory Cultural Mapping Based on Collective Behavior Data in Location-Based Social Networks. *ACM TIST* 7, 3 (2016), 30:1–30:23.

[49] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik Lun Lee. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *SIGIR*. ACM, 325–334.

[50] Jia-Dong Zhang and Chi-Yin Chow. 2013. iGSLR: personalized geo-social location recommendation: a kernel density estimation approach. In *SIGSPATIAL/GIS*. ACM, 324–333.

[51] Jia-Dong Zhang, Chi-Yin Chow, and Yanhua Li. 2014. LORE: exploiting sequential influence for location recommendations. In *SIGSPATIAL/GIS*. ACM, 103–112.

[52] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep Learning Based Recommender System: A Survey and New Perspectives. *ACM Comput. Surv.* 52, 1 (2019), 5:1–5:38.