Discovering Related Users in Location-Based Social Networks

Sergio Torrijos
Universidad Autónoma de Madrid
Madrid, Spain
sergio.torrijos@estudiante.uam.es

Alejandro Bellogín
Universidad Autónoma de Madrid
Madrid, Spain
alejandro.bellogin@uam.es

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

ABSTRACT

Users from Location-Based Social Networks can be characterised by how and where they move. However, most of the works that exploit this type of information neglect either its sequential or its geographical properties. In this article, we focus on a specific family of recommender systems, those based on nearest neighbours; we define related users based on common check-ins and similar trajectories and analyse their effects on the recommendations. For this purpose, we use a real-world dataset and compare the performance on different dimensions against several state-of-the-art algorithms. The results show that better neighbours could be discovered with these approaches if we want to promote novel and diverse recommendations.

KEYWORDS

location-based social networks, neighbours, trajectory similarity

ACM Reference Format:
Sergio Torrijos, Alejandro Bellogín, and Pablo Sánchez. 2020. Discovering Related Users in Location-Based Social Networks. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20), July 14–17, 2020, Genoa, Italy*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3340631.3394882

1 INTRODUCTION

The aim of Recommender Systems (RS) is to help users in finding relevant items, usually by filtering large catalogues and taking into account the users' preferences. Collaborative Filtering (CF) systems can be considered as the earliest and most widely deployed recommendation approach [16], suggesting interesting items to users based on the preferences from "similar" or related people [23]. Other types of recommendation algorithms include content-based systems – that suggest items similar to those the user preferred or liked in the past and based on content features of the items in the system [22] –, demographic systems – where users are categorised based on their demographic attributes [2] –, social filtering systems – exploiting contacts, interactions, and trust between users [13] –, and hybrid recommenders – where different techniques are combined in order to palliate individual drawbacks of the algorithms [4]. In this work, we focus on Location-Based Social Networks (LBSNs), where users share the venues, places, or Points-of-Interest (POI) they visit, establish connections with other users, and check venue properties, such as their opening times, opinions, and pictures.

Because of the increasing number of users registered in LBSNs and similar systems, POI recommendation approaches have become particularly useful and several specific models have been proposed in recent years. In particular, such approaches tend to incorporate inherent properties of these systems, such as social, geographical, or temporal information [20, 21]. However, nearest neighbour techniques have been, in general, neglected in most of these studies, in favour of matrix factorisation or neural networks models [21, 24]. Nonetheless, we believe that algorithms based on similarities have a huge potential, since they may provide efficient computation, easy implementation, and explainable recommendations [23], but also because it has been demonstrated recently that these techniques can also achieve quite competitive results, even when compared against neural network approaches [9].

More specifically, in this work we propose to integrate user mobility patterns in the computation of the similarity between users. Hence, in order to decide which users should be considered when producing recommendations, we present two different techniques: one based on exploiting common check-ins and another that accounts for how close two trajectories (defined as the trail left by each user in the system) are. We show results on one city (Tokyo) extracted from a real-world dataset of Foursquare check-ins. We demonstrate that the proposed similarities are not better than classical methods in terms of performance, but they allow to better predict the category of the next POI to visit, while producing more diverse and novel recommendations, especially for those users identified as tourists (in contrast with those identified as locals). We also show that the neighbours found by our approaches are more similar to those in the social network of the users in comparison with the neighbours found by pure CF user similarities.

2 BACKGROUND

Collaborative Filtering Recommender Systems are usually classified in two categories: model-based and memory-based. Model-based approaches build statistical models of user/item interaction patterns to provide automatic predictions, through, e.g., matrix factorisation (MF) or probabilistic models [16]; memory-based algorithms, on the other hand, make predictions based on the entire collection of interactions, usually by computing similarities between users or items and taking those similarities into account when producing the recommendations [23].

In this paper, our focus is on the second type of CF algorithms, also known as *nearest-neighbour recommender systems* since they exploit those similarities to rank the users/items and use the closest ones (neighbours) to generate recommendations. While there are several variations of the formulations for these methods, those tailored for ranking purposes, instead of rating prediction, are more

UMAP '20, July 14–17, 2020, Genoa, Italy

ACM ISBN 978-1-4503-6861-2/20/07…$15.00
https://doi.org/10.1145/3340631.3394882

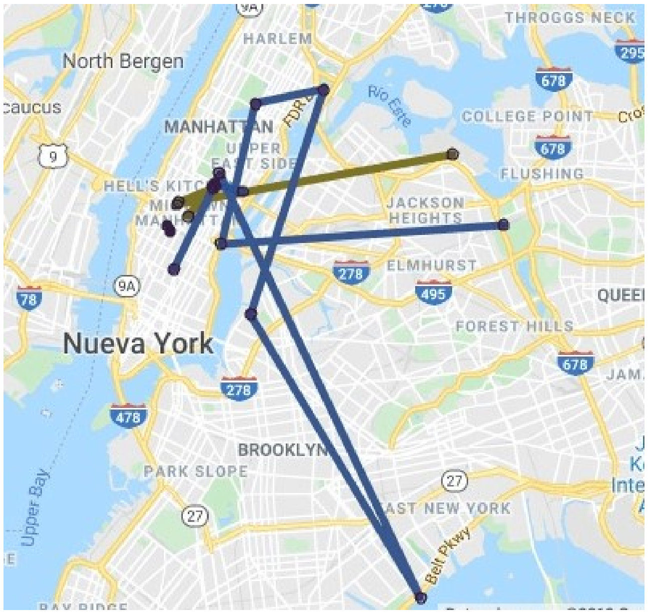

Figure 1: Check-ins collected from two users (one in blue and the other in green) in the city of New York. Note that there are many more places in common in the center of Manhattan.

popular nowadays [1, 8], where the predicted score is not normalised to fall in a pre-defined range:

$$\hat{r}(u,i) = \sum_{v \in N_k(u)} r(v,i)\mathrm{sim}(u,v) \tag{1}$$

where $r(u,i)$ and $\hat{r}(u,i)$ denote the recorded rating and predicted score by user u on item i respectively, $\mathrm{sim}(u,v)$ is a user similarity function, and $N_k(u)$ provides a neighbourhood of size k for user u.

3 DISCOVERING RELATED USERS

In this paper, we propose different strategies to select users in a Location-Based Social Network. Hence, we shall use the same formulation as in the classical user-based nearest neighbour algorithm (i.e., Equation 1) but changing the computation of the similarity between users, hence impacting in the neighbourhood computation and in the weight used to estimate the preferences. Thus, in the next sections we present our user similarity approaches tailored for LBSNs.

3.1 Exploiting common check-ins within a temporal window

The first approach we propose is based on the assumption that those users who visit the same venue *around the same time* are more similar. In this sense, this strategy – that we name `Ad-hoc` – is an adaptation of the frequently used overlap measure in classical scenarios but tailored for a domain where the temporal dimension is very important. It is formulated as follows:

$$\mathrm{sim}_\delta(u,v) = \|\{i \in I : |t(u,i) - t(v,i)| < \delta\}\| \tag{2}$$

where $\|S\|$ denotes the number of elements in S, I is the set of all items in the system, $t(u,i)$ returns the timestamp when user u checked-in in item i (if this did not happen, then it is infinity), and δ is a parameter to control the temporal window we allow to consider two users as similar if they have an item in common.

Under this perspective, and considering the toy example shown in Figure 1, those two users may be found as very similar or not depending on whether they tend to visit the same places at similar times. Hence, with this similarity function we impose a harder constraint on the time dimension than other strategies.

3.2 Exploiting common trajectories

Our second approach is based on the assumption that users are similar if their trajectories in the system are similar. For this, we use standard methods to compute trajectory similarities such as Dynamic Time Warping and Hausdorff distance [31]. We need first to transform the user check-ins into trajectories; as in previous works [24], we split the check-ins in such a way that if two consecutive user interactions are too distant in time (in our case, 8 hours) they are assigned to different trajectories of that user.

Thus, once a user u has been split in several trajectories $(x_1^u, \cdots, x_n^u)$, we compute the following similarity by averaging the trajectory similarity values over all pairs for both users:

$$\mathrm{sim}(u,v) = \frac{1}{n \cdot m} \sum_{j=1}^{n} \sum_{k=1}^{m} \mathrm{tsim}(x_j^u, x_k^v) \tag{3}$$

where n and m correspond to the number of trajectories of users u and v, respectively, and $\mathrm{tsim}(\cdot,\cdot)$ is a trajectory similarity function. In this work we shall use either a fast implementation of the Dynamic Time Warping function described in [29] and implemented in Python[1] (that we denote as `TS-DTW`), or an efficient version of the Hausdorff distance described in [30] and implemented in Python's SciPy (named as `TS-Haus`).

If we revisit the example shown in Figure 1 with this type of similarity, we argue that some parts of each user trajectory are quite close to each other and, hence, similar, whereas other parts are very distant to each other. Thus, we conclude that the proposed similarity is more flexible in both temporal and spatial dimensions, since all the points visited within a time frame (8 hours) will be considered in a trajectory basis.

In the future, we would like to explore other, more complex algorithms tailored at finding objects that move together, such as Flock, Convoy, or ST-DBSCAN [6, 11, 28]. These methods are computationally very expensive and, to the best of our knowledge, they have not been applied to check-in data, but to spatio-temporal data at different granularity levels. However, they could help us discriminate check-ins that are geographically similar but temporally different from those that are simply geographically or temporally similar, in contrast to the ones presented in this work.

4 EXPERIMENTS

We performed experiments on the Foursquare global check-in dataset[2] used in [32]. This dataset is formed by 33M check-ins in different cities around the world, although we decided to focus

[1] https://github.com/slaypni/fastdtw
[2] https://sites.google.com/site/yangdingqi/home/foursquare-dataset

Table 1: Performance comparison for all users. Best result on each metric is identified by a †, whereas the best result among those methods using neighbours (Ad-hoc, TS-DTW, TS-Haus, and UB) is highlighted in bold.

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.054	0.206	**3,314**	0.741
TS-DTW	0.036	0.203	408	0.871
TS-Haus	0.036	**0.215**	442	0.870
UB	**0.062**	0.201	1,899	**0.912**
IB	0.046	†0.264	†13,685	†0.956
BPR	0.066	0.229	28	0.872
IRenMF	†0.069	0.225	1,475	0.888

Table 2: Performance comparison on local users.

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.050	0.181	**2,829**	0.735
TS-DTW	0.032	0.179	358	**0.869**
TS-Haus	0.032	**0.191**	405	0.869
UB	**0.057**	0.177	1,440	0.704
IB	0.042	†0.250	†11,929	†0.872
BPR	0.059	0.204	26	0.608
IRenMF	†0.063	0.200	1,199	0.647

on one city (Tokyo) since it is among those with more interactions in the dataset, consistent with its cultural importance and tourism impact. It is also a frequent city used in related works to test this type of recommenders [10, 15, 33].

We created a temporal split, with 6 months of data for training (May - Oct '12) and one month for test (Nov '12); this period of the year was selected to capture seasonal trends related to summer, while including enough data to obtain significant results. We performed a 2-core before splitting the data, so that every user and item have at least two preferences [26]. This resulted in a dataset with 328K check-ins made by 8.6K users on 28.6K items.

In terms of the algorithms being tested, although our main baseline is the user-based nearest neighbour (since our proposals use the same formulation where we change the similarity function) [23], we also include in our comparison other state-of-the-art recommenders that have demonstrated good performance for POI recommendation [20]: IB [23] provides item-based nearest neighbour recommendations; BPR [25] is the Bayesian Personalised Ranking method using a MF technique; and IRenMF [21] is a weighted MF method that also exploits the geographical influence between neighbour venues.

4.1 Performance comparison

In this section, we analyse the performance of the proposed user similarities. For this, we report the following evaluation metrics (always at top-10): NDCG captures the classical ranking-based accuracy in terms of Normalised Discounted Cumulative Gain [12], FA measures precision at the category level, i.e., a match is found whenever a category in the recommendation list appears in the test set of the user (named FA from feature agreement, as in [17]),

Table 3: Performance comparison on tourist users.

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.071	†**0.271**	**1,217**	0.770
TS-DTW	0.069	0.246	320	**0.847**
TS-Haus	0.068	0.252	333	0.839
UB	**0.095**	0.245	1,133	0.710
IB	0.064	0.262	†3,208	†0.858
BPR	0.108	0.266	18	0.566
IRenMF	†0.111	0.262	106	0.580

and AD and EPC measure recommendation diversity and novelty, respectively [5]: AD (Aggregate Diversity) considers how many different items are being recommended whereas EPC (Expected Popularity Complement) measures how unpopular (as a measurement of novelty) the returned items are. These four dimensions are significant in location-based recommendation since accuracy, diversity, and novelty are important for users of these systems, whereas the feature agreement provides a necessary tradeoff to balance the difficulty of predicting the actual next POI that the user will visit; because of that, in some works researchers measure precision in terms of the item features (usually categories), so that a correct recommendation is considered whenever the features of the recommended items are found in the user test set [3, 14, 17, 19, 24].

Table 1 shows the results of the previously described metrics for the recommenders optimised according to their NDCG value. In this table we observe that the methods based on neighbours (either those proposed in this work or the baselines UB and IB) are not competitive when compared against BPR or IRenMF in terms of accuracy. They obtain, however, better results in the other dimensions, especially IB, where the proposed approaches outperform UB, more specifically, in FA and AD.

Nonetheless, it is well-known that users from LBSNs are varied in nature, in particular, these systems encompass local and tourist users with different motivations and preferences when moving around a city [18]. Because of this, in Tables 2 and 3 we restrict the computations of the metrics to those users categorised as locals or as tourists. We do this by discriminating their check-in behaviour, as in previous works [7]: if the difference in timestamp between the first and the last check-in is larger than 21 days, the user is classified as a local, otherwise as a tourist – this period is taken from the literature where it is assumed that most tourists concentrate their visits within a short time period, whereas local users will perform the check-ins on a city over a much longer period of time. It should be noted that other approaches exist where, instead of considering a temporal period between the check-ins, the average distance between all the check-ins of the user are considered to estimate her home location (and decide whenever the user is abroad) [7, 18, 20, 32]. In the same process, we also remove those bogus users which show unrealistic behaviour, such as too many check-ins in a short time or travels that would require moving too fast for a person, as in [24, 32].

Under this perspective, we observe that neighbour recommenders improve their performance: even though IRenMF is still the best algorithm, UB shows NDCG values closer to that method. However, it is for the rest of the dimensions – especially, FA – where these methods, and our proposed approaches in particular, stand out. We

Table 4: Overlap between top neighbours and explicit social network. The first column denotes the number of neighbours selected: whether those that optimise performance (Best, matching what is reported in Table 3) or the same number as the baseline (As UB). Columns T-NDCG and T-FA show the value of these metrics on tourist users.

Neigh. Sel.	Method	Avg.	Total	T-NDCG	T-FA
Best	Ad-hoc	2	1,008	0.071	0.271
	TS-DTW	9	4,833	0.069	0.246
	TS-Haus	9	4,833	0.068	0.252
As UB	Ad-hoc	2	1,008	0.071	0.273
	TS-DTW	2	1,074	0.056	0.271
	TS-Haus	5	2,685	0.050	0.270
	UB	4	2,140	0.095	0.245

observe that for tourist users the Ad-hoc approach is the best one in terms of FA, and all the proposed strategies outperform classical UB in this scenario, not only in FA, but also in terms of diversity (AD) and novelty (EPC). Our hypothesis is that these strategies work better for tourists because their trajectories are more meaningful than those for local users, which might be shorter or too monotonous.

In summary, we conclude that obtaining related users (neighbours) by exploiting common check-ins or trajectories is more beneficial for tourists than for other users, especially when deciding the category of the next place to visit (measured by FA) and to enhance the recommendation with novel and diverse items.

4.2 Social network analysis

In this section we explore whether the neighbours found by the proposed approaches are more similar to the social connections explicitly followed by the users. For this, we use an additional dataset provided by the same authors as the global check-in dataset used previously, where an external social network (Twitter) is extracted for the same users whose Foursquare check-ins were collected. Table 4 reports the average and total overlap between the groundtruth connections (actual social network) and the top-N neighbours depending on different strategies to discover related users. For this, we explore two different number of neighbours for our approaches: either the one that optimises performance (as in the previous section) or, for the sake of comparison, the same number of neighbours as the baseline.

Based on these results, we observe that the overlap with respect to the explicit social network of the user does change depending on the similarity employed to find users when optimising for performance. Whereas classical UB is somewhat in the middle, both strategies based on trajectory similarity are clearly better than the rest, evidencing that in Location-Based Social Networks there is a connection between the social network and the places visited by the users.

Furthermore, when using the same number of neighbours as the baseline, TS-Haus is the only approach that outperforms (by a slight margin) the overlap with the explicit social network with respect to the baseline. However, it should be noted that even for this (smaller) amount of neighbours, the obtained FA for tourist users remains competitive.

It is, however, not obvious which part is the cause and which the effect: whether users visit some places because their social network instigates it or if they follow certain users in the network because they observe that they share some places in common. Investigating an explanation for this phenomenon is left for future work, together with an exhaustive comparison against recommender systems approaches based on social networks, such as those using semantic associations or trust and reputation techniques [13, 27].

5 CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed two novel similarity metrics between users for a recommendation system that exploits user check-ins in Location-Based Social Networks. One metric captures the amount of common visited places within a predefined temporal window, so as to consider users as similar whenever they tend to visit the same places around the same time. The other metric is the first adaptation, to the best of our knowledge, of trajectory similarity metrics for recommendation; we make use of previous methods to transform check-ins into trajectories, and, after that, two popular trajectory metrics (Dynamic Time Warping and Hausdorff distance) are integrated.

The empirical evaluation of our proposal shows competitive results in terms of beyond-accuracy metrics (diversity and novelty). Especially positive results have been obtained when we isolate the evaluation only for tourists users (separated from users classified as locals), evidencing the inherently different behaviour of users in these systems, which should entail different strategies for each type of users. We have also found that our approaches were very accurate on predicting the category of the next item to visit, although not so good on predicting the actual items, as measured by classical accuracy metrics. Our analysis shows that one possible explanation for this effect is that the proposed similarity functions – in particular, those that exploit common trajectories – are able to predict better than classical similarity metrics the users that belong to a user's social network.

Several directions open up from this point to explore the potential of the proposed similarity metrics, some of them have already been mentioned in the paper. We are particularly interested in exploring algorithms from the co-movement pattern mining community, even though they tend to be particularly expensive. We would also like to analyse in detail how the advances done in POI recommendation based on check-ins could be translated into scenarios where GPS trajectories with high granularity are available, such as data generated from mobility and smart city applications.

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. The first author also acknowledges the Ministerio de Educación y Formación Profesional, within the 2019-2020 call for collaboration grants.

REFERENCES

[1] Fabio Aiolli. 2013. Efficient Top-n Recommendation for Very Large Scale Binary Rated Datasets. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, New York, NY, USA, 273–280. http://doi.acm.org/10.1145/2507157.2507189

[2] Mohammad Yahya H. Al-Shamri. 2016. User profiling approaches for demographic recommender systems. *Knowl.-Based Syst.* 100 (2016), 175–187.

[3] 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 *22nd ACM International Conference on Information and Knowledge Management, CIKM'13, San Francisco, CA, USA, October 27 - November 1, 2013*, Qi He, Arun Iyengar, Wolfgang Nejdl, Jian Pei, and Rajeev Rastogi (Eds.). ACM, 757–762. https://doi.org/10.1145/2505515.2505643

[4] Robin D. Burke. 2002. Hybrid Recommender Systems: Survey and Experiments. *User Model. User-Adapt. Interact.* 12, 4 (2002), 331–370.

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

[6] Wen Chen, Minhe Ji, and Jianmei Wang. 2014. T-DBSCAN: A Spatiotemporal Density Clustering for GPS Trajectory Segmentation. *Int. J. Online Eng.* 10, 6 (2014), 19–24. http://www.online-journals.org/index.php/i-joe/article/view/3881

[7] 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'10, Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, Toronto, Ontario, Canada, June 13-16, 2010*, Mark H. Chignell and Elaine G. Toms (Eds.). ACM, 35–44. https://doi.org/10.1145/1810617.1810626

[8] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of Recommender Algorithms on Top-n Recommendation Tasks. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. ACM, New York, NY, USA, 39–46. http://doi.acm.org/10.1145/1864708.1864721

[9] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). ACM, 101–109. https://doi.org/10.1145/3298689.3347058

[10] Ruifeng Ding, Zhenzhong Chen, and Xiaolei Li. 2018. Spatial-Temporal Distance Metric Embedding for Time-Specific POI Recommendation. *IEEE Access* 6 (2018), 67035–67045. https://doi.org/10.1109/ACCESS.2018.2869994

[11] Qi Fan, Dongxiang Zhang, Huayu Wu, and Kian-Lee Tan. 2016. A General and Parallel Platform for Mining Co-Movement Patterns over Large-scale Trajectories. *PVLDB* 10, 4 (2016), 313–324. https://doi.org/10.14778/3025111.3025114

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

[13] Ido Guy. 2015. Social Recommender Systems. In *Recommender Systems Handbook*. Springer, 511–543.

[14] Jing He, Xin Li, and Lejian Liao. 2017. Category-aware Next Point-of-Interest Recommendation via Listwise Bayesian Personalized Ranking. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 1837–1843. https://doi.org/10.24963/ijcai.2017/255

[15] Xu Jiao, Yingyuan Xiao, Wenguang Zheng, Hongya Wang, and Ching-Hsien Hsu. 2019. A novel next new point-of-interest recommendation system based on simulated user travel decision-making process. *Future Generation Comp. Syst.* 100 (2019), 982–993. https://doi.org/10.1016/j.future.2019.05.065

[16] Yehuda Koren and Robert M. Bell. 2015. Advances in Collaborative Filtering. In *Recommender Systems Handbook*. Springer, 77–118.

[17] Gunjan Kumar, Houssem Jerbi, and Michael P. O'Mahony. 2017. Towards the Recommendation of Personalised Activity Sequences in the Tourism Domain. In *Proceedings of the 2nd Workshop on Recommenders in Tourism*, Vol. 1906. CEUR-WS.org, 26–30.

[18] Anh Le, Konstantinos Pelechrinis, and Prashant Krishnamurthy. 2014. Country-level spatial dynamics of user activity: a case study in location-based social networks. In *ACM Web Science Conference, WebSci '14, Bloomington, IN, USA, June 23-26, 2014*, Filippo Menczer, Jim Hendler, William H. Dutton, Markus Strohmaier, Ciro Cattuto, and Eric T. Meyer (Eds.). ACM, 71–80. https://doi.org/10.1145/2615569.2615689

[19] 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 *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 1778–1784. http://ijcai.org/Abstract/15/253

[20] 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.

[21] Yong Liu, Wei Wei, Aixin Sun, and Chunyan Miao. 2014. Exploiting Geographical Neighborhood Characteristics for Location Recommendation. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014*. ACM, 739–748.

[22] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2011. Content-based Recommender Systems: State of the Art and Trends. In *Recommender Systems Handbook*. Springer, 73–105.

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

[24] 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 *Proceedings of the 2nd Workshop on Recommenders in Tourism co-located with 11th ACM Conference on Recommender Systems (RecSys 2017), Como, Italy, August 27, 2017 (CEUR Workshop Proceedings)*, Julia Neidhardt, Daniel R. Fesenmaier, Tsvi Kuflik, and Wolfgang Wörndl (Eds.), Vol. 1906. CEUR-WS.org, 1–8. http://ceur-ws.org/Vol-1906/paper1.pdf

[25] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 452–461.

[26] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 811–820.

[27] Domenico Rosaci. 2015. Finding semantic associations in hierarchically structured groups of Web data. *Formal Asp. Comput.* 27, 5-6 (2015), 867–884. https://doi.org/10.1007/s00165-015-0337-z

[28] Omar Ernesto Cabrera Rosero and Andres Oswaldo Calderon Romero. 2018. FP-Flock: An algorithm to find frequent flock patterns in spatio-temporal databases.

[29] Stan Salvador and Philip Chan. 2007. Toward accurate dynamic time warping in linear time and space. *Intell. Data Anal.* 11, 5 (2007), 561–580. http://content.iospress.com/articles/intelligent-data-analysis/ida00303

[30] Abdel Aziz Taha and Allan Hanbury. 2015. An Efficient Algorithm for Calculating the Exact Hausdorff Distance. *IEEE Trans. Pattern Anal. Mach. Intell.* 37, 11 (2015), 2153–2163. https://doi.org/10.1109/TPAMI.2015.2408351

[31] Dong Xie, Feifei Li, and Jeff M. Phillips. 2017. Distributed Trajectory Similarity Search. *PVLDB* 10, 11 (2017), 1478–1489. https://doi.org/10.14778/3137628.3137655

[32] 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.

[33] Lina Yao, Quan Z. Sheng, Yongrui Qin, Xianzhi Wang, Ali Shemshadi, and Qi He. 2015. Context-aware Point-of-Interest Recommendation Using Tensor Factorization with Social Regularization. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 1007–1010. https://doi.org/10.1145/2766462.2767794