

A Joint Two-Phase Time-Sensitive Regularized Collaborative Ranking Model for Point of Interest Recommendation

Mohammad Aliannejadi, Dimitrios Rafailidis, and Fabio Crestani

Abstract—The popularity of location-based social networks (LBSNs) has led to a tremendous amount of user check-in data. Recommending points of interest (POIs) plays a key role in satisfying users needs in LBSNs. While recent work has explored the idea of adopting collaborative ranking (CR) for recommendation, there have been few attempts to incorporate temporal information for POI recommendation using CR. In this article, we propose a two-phase CR algorithm that incorporates the geographical influence of POIs and is regularized based on the variance of POIs popularity and users activities over time. The time-sensitive regularizer penalizes user and POIs that have been more time-sensitive in the past, helping the model to account for their long-term behavioral patterns while learning from user-POI interactions. Moreover, in the first phase, it attempts to rank visited POIs higher than the unvisited ones, and at the same time, apply the geographical influence. In the second phase, our algorithm tries to rank users favorite POIs higher on the recommendation list. Both phases employ a collaborative learning strategy that enables the model to capture complex latent associations from two different perspectives. Experiments on real-world datasets show that our proposed time-sensitive collaborative ranking model beats state-of-the-art POI recommendation methods.

Index Terms—point-of-interest recommendation, time-aware recommendation, collaborative ranking, location-based social networks



1 INTRODUCTION

WITH the availability of *location-based social networks* (LBSNs), such as Yelp, TripAdvisor, and Foursquare, users can share check-in data using their mobile devices. LBSNs collect valuable information about users' mobility records with check-in data. Generating *points-of-interest* (POIs) recommendations plays a crucial role in satisfying the user needs, for example when exploring a new POI or visiting a city [1]. In fact, every town has numerous POIs and a user may have visited only a few in her hometown as well as when she is out-of-town [2]. POI recommendation tries to ensure user's satisfaction by suggesting her the most interesting locations in her vicinity, taking into account her preferences and contextual constraints [3], [4].

Many challenges limit the accuracy of POI recommendation. For instance, a major challenge in POI recommendation is data sparsity [5], [6]. Despite the fact that LBSNs feature a huge number of locations with a large variety, in practice users visit a very limited number of locations, making the user-item matrix sparse [7]. Moreover, since users spend most of their time in their hometown [8], the data sparsity problem is aggravated when a user visits a new city where she has no history of visited locations [9]. Several studies seek to address the data sparsity problem by incorporating additional information into the model, such as geographical and temporal information [10], [11], [12], [13]. In fact, the data sparsity problem is even worse when recommending

POIs as opposed to other items such as movies or songs. This mainly happens because the check-in data provides an implicit feedback [14], whereas users usually express their opinion about movies or songs with different ratings. [11] argued that check-ins offer only positive examples whereas POIs with no check-in remain undiscovered which can potentially be interesting to the user. However, we argue that if a user has visited a place only once, we cannot simply infer a positive feedback but we can also infer what types of locations the user is interested in. In the case of places with higher check-in frequency, we assume that they are more preferred than those with fewer check-ins [11].

In relevant literature, item recommendation is often treated as a rating prediction or matrix completion task [15]. However, [16] argued that considering the square loss as a measure of prediction effectiveness is not accurate. In other words, being able to produce a more accurate ranked list to a user should be rewarded. *Collaborative Ranking* (CR) is based on this idea and focuses on the accuracy of recommendation at the top of the recommendation list for each user [17]. Much work has adopted CR using explicit feedback from users such as item ratings. For example, [18] and [16] optimized a ranking loss to recommend movies to users. [19] showed that ranking-based learning of recommendation is also effective when one is dealing with implicit feedback from users. However, exploring CR for POI recommendation using implicit feedback is challenging because an effective strategy for sampling unvisited venues based on users' mobility and POIs proximity is required as part of the learning process.

Several studies started to incorporate temporal information to improve POI recommendation [20], [21], [22]. For example, temporal information has been incorporated in an

- M. Aliannejadi and F. Crestani are with the Faculty of Informatics, Università della Svizzera italiana (USI), Lugano, Switzerland.
E-mail: {mohammad.ali.anejadi, fabio.crestani}@usi.ch.
- D. Rafailidis is with the Maastricht University, Maastricht, the Netherlands.
E-mail: dimitrios.rafailidis@maastrichtuniversity.nl

hourly basis in [20] and in a sequential manner in [23] to recommend the “next” POI. Also, more advanced models considered both time and the sequential order of check-ins [22]. However, the long-term dynamics of users check-in behavior and venues popularity still need to be analyzed. For example, an open-air bar is popular mainly during summertime, or a student is supposedly more active during her holidays.

In this article, we propose a two-phase CR algorithm for POI recommendation. Our model is inspired by successful results of CR in other domains with explicit user feedback [18], [24] and of ranking methods using implicit user feedback for POI recommendation [11], [19]. We use a two-phase implicit feedback inference in our algorithm. In fact, we assume that single check-in means that a user “likes” the POI and multiple visits mean that user prefers the POI more. Based on this assumption, in the first phase we push POIs with single or multiple check-ins at the top of the recommendation list, taking into account the geographical influence of POIs in the same neighborhood. In the second phase, we push POIs with multiple check-ins over the ones with a single visit. As argued in [11], considering both visited and unvisited POIs in the learning, alleviates the sparsity problem. Therefore, the first phase of our algorithm addresses the sparsity problem, whereas the second stage boosts the accuracy of our model by pushing more relevant POIs at the top of the list. To take into account the dynamics of user and place check-in, we introduce a *time-sensitive regularizer* in the ranking loss. The regularizer models the activity pattern of every user and venue over time. Adding this regularizing parameter to the objective function fuses the activity patterns into the ranking function in a collaborative way.

In summary, our contribution in this article can be summarized as follows:

- We perform an extensive analysis to demonstrate the underlying patterns of preference and popularity over time.
- We propose a general time-sensitive regularizer, taking into account the variance of users activities and venues popularity over time.
- We propose a novel two-phase CR-based POI recommendation algorithm incorporating users implicit check-in feedback with a focus on the top of the list.
- We propose a geographical similarity measure and add its influence to the model’s objective function.

The experiments on two benchmark datasets show that the proposed approach outperforms state-of-the-art POI recommendation and CR methods. In particular, we show that the joint learning strategy enables the model not only to address the sparsity problem but also to rank relevant POIs higher. The first phase mainly addresses the sparsity problem by adding the geographical influence as well as considering both visited and unvisited venues for training. The second phase improves the accuracy of the model, pushing to the top of the recommendation list the POIs that users prefer more. Moreover, we demonstrate the effectiveness of the time-sensitive regularizer, that is applied to both phases of the algorithm taking into consideration the long-term behavior of users and the popularity of POIs.

The remainder of the article is organized as follows, Section 2 briefly reviews the related work while a deep analysis on the datasets is performed in Section 3. We describe our method in Section 4 and in Section 5, evaluate the performance of our model against competitive models. Finally, Section 6 concludes the article.

2 RELATED WORK

The related work can be divided to the following topics: POI recommendation, collaborative ranking, and time-aware recommendation.

2.1 POI Recommendation

POI Recommendation plays an important role in satisfying users’ expectations on LBSNs. Much work has been carried out in this area based on the core idea that users with similar behavioral history tend to act similarly [25]. This is the underlying idea of *collaborative-filtering-based* (CF-based) approaches [9], [10], [26]. CF can be divided into two categories: memory-based and model-based. Memory-based approaches consider user rating as a similarity measure between users or items [27]. Model-based approaches, on the other hand, employ techniques like matrix factorization [28]. However, CF-based approaches often suffer from data sparsity since there are a lot of available locations, and a single user can visit only a few of them. As a consequence, the user-item matrix of CF becomes very sparse, leading to poor performance in cases where there is no significant association between users and items. Many studies have tried to address the data sparsity problem of CF by incorporating additional information into the model [20], [29]. More specifically, [29] argued that users check-in behavior is affected by the spatial influence of locations and proposed a unified location recommender system incorporating spatial and social influence to address the data sparsity problem. [26] proposed a model that captures user interests as well as local preferences to recommend locations or events to users when they are visiting a new city. [2] proposed to consider both geographical and temporal influences while recommending POIs to the users via a geographical-temporal influences aware graph. They proposed to propagate these influences using a breadth-first strategy. [9] took into consideration user preference, geographical proximity, and social influences for POI recommendation. [30] proposed a multi-center Gaussian model to capture users’ movement pattern as they assumed users’ movements consist of several centers. In a more recent work, [31] considered three travel-related constraints (i.e., uncertain traveling time, diversity of the venues, and venue availability) and use them to prune the search space. [10] also proposed an approach integrating temporal and geographic influences into matrix factorization. [32] proposed a probabilistic mapping approach to determine the most salient information from a venue’s content to reduce dimensionality of data. More recently, [33] considered the appropriateness of a POI given a user’s context while ranking the venues. [34] fused virtual ratings derived from online reviews into CF. [35] investigated how geotagged photos can be linked to venues to study users’ tastes. Finally, [36] addressed the data sparsity problem

assuming that users tend to rank higher the POIs that are geographically closer to the one that they have already visited. Differently from these studies, our work addresses the data sparsity problem by including the unvisited venues in the learning process. We consider unvisited venues as negative training examples based on the assumption that a user “likes” and prefers venues that she has visited over the ones that she has not [22].

2.2 Collaborative Ranking

Another line of research lies in combining the ideas of CF and *Learning to Rank* (LTR). LTR methods have been proved to be effective in *Information Retrieval* (IR) [37]. LTR learns a ranking function which can predict a relevance score given a query and document. There are three categories of LTR, namely, point-wise [38], pair-wise [39], and list-wise [40]. In short, point-wise approaches predict ranking scores for individual items. Pair-wise approaches, on the other hand, learn the order of the items. List-wise approaches consider an entire list of items as individual training example. CR takes the idea of predicting preference order of items from LTR and combines it with the idea of learning the loss function in a collaborative way. [18] used a surrogate convex upper bound of *Normalized Discounted Cumulative Gain* (nDCG) error together with matrix factorization as the basic rating predictor. [41] explored optimizing a surrogate lower bound for *Expected Reciprocal Rank* (ERR) for data with multiple levels of relevance. [16] followed the idea of pair-wise LTR approaches. In particular, the authors base their work on LTR methods with an emphasis on the top of the recommendation list. This approach, however, is limited to explicit user feedback such as user ratings for movies. [42] presented an LTR model, taking into account the relevant items of users and their friends, pushing these items at the top of the list. [19] presented a generic optimization criterion as well as a learning algorithm for incorporating implicit feedback while learning personal ranking for users and they demonstrated its effectiveness on approaches such as matrix factorization. In [43], authors combined various LTR methods into a joint model aiming to enhance the recommendation accuracy with trust relationships. In a more recent work, [44] proposed a model considering not only relevant items of the user and her trusted friends, but also the items of her distrusted foes. Unlike the above works, our method uses a two-phase strategy to use the implicit feedback, i.e. check-ins. Also, we regularize the ranking loss based on users activity and venues popularity variances on a monthly basis. The social influence on POI recommendation, however, is left for future work.

2.3 Time-Aware Recommendation

Many researchers have studied the temporal influence on users’ preferences. A group of studies conducts time-aware recommendation learning of users’ temporal preference for specific time slots and for recommending POIs for a given time slot, like an hour of a day [45]. [20] computed the similarity between users by finding the same POIs at the same time slots in their check-in history to train a user-based CF model. [46] matched the temporal regularity of users with the popularity of POIs to improve a factorization-based

algorithm. [47] proposed a time-aware personalized model adopting a fourth-order tensor factorization-based ranking which enables the model to capture short-term and long-term preferences. [48] proposed topic-region model that discovers the semantic, temporal, and spatial patterns of users’ check-ins and uses the additional information to address the data-sparsity problem. [49] defined the temporal context as the public’s attention at a certain time and proposed a temporal context-aware mixture model, modeling the topics related to users’ interests and temporal context in a unified way. This work was later extended in [50] to a dynamic temporal context-aware mixture model, capturing users evolving interests. [21] preserved the similarity of personal preference in consecutive time slots by considering different latent variables at each time slot for each user. However, the long-term behavior of users still needs to be studied. For example, a high school student visits more places during summer, or a beach bar could be even shut down during winter. Therefore, it is crucial to consider the time-dependent activity patterns of both users and venues while training a recommender system. There also exists another category of approaches which tries to recommend the next POI to visit, also known as successive POI recommendation. For example, [51] captured sequential check-in behavior of users by training personalized Markov chain. [22] combined the ideas of both categories by recommending POIs for a particular time, exploiting sequential patterns of users. Our work distinguishes itself from these studies by considering a long-term activity shift of both users and locations in the course of time.

3 DATA ANALYSIS

In this section, we conduct an extensive analysis of two real-world datasets to explore user preferences’ dynamics over time.

3.1 Data

We use two real-world check-in datasets from Foursquare and Gowalla provided by [20]¹. The Foursquare’s dataset originally consists of 342,850 check-ins of users in Singapore in the period of Aug. 2010 and Jul. 2011 [20]. The Gowalla’s dataset, on the other hand, includes 736,148 check-ins made by users in California and Nevada in the period of Feb. 2009 and Oct. 2010 [20]. Every check-in contains a user id, POI id, time, and geographical coordinates. For a fair comparison, we also use the preprocessed data as in [20]. Users who have less than 5 check-ins, as well as POIs with less than 5 check-ins are removed from the datasets. Finally, we have 194,108 check-ins on Foursquare’s and 456,988 check-ins on Gowalla’s. We used the geographical coordinates of POIs to retrieve their corresponding categories such as bar, pizza place. More details are listed in Table 1. Notice that multiple check-ins refer to the check-ins that were made by a particular user to a particular POI more than once. As shown in Table 1, 45.51% of check-ins on the Foursquare’s dataset refer to users visiting POIs more than once. In Gowalla’s, we observe fewer multiple check-ins, namely, 32.69% of all check-ins.

1. <http://www.ntu.edu.sg/home/gaocong/data/poidata.zip>

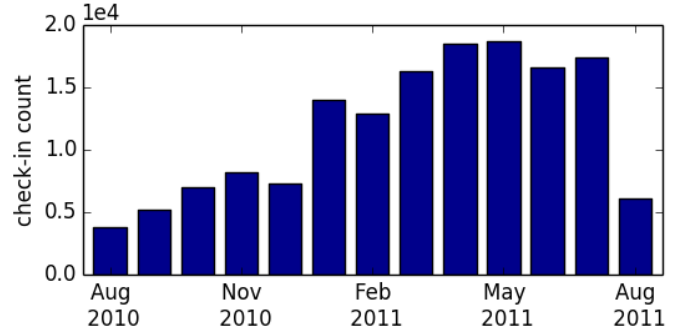
TABLE 1: General statistics of the datasets

	Foursquare's	Gowalla's
# of users	2,231	10,162
# of POIs	5,596	24,250
# of check-ins	194,108	456,988
Avg. POIs per user	45.57	30.356
Avg. users per POI	18.90	12.69
Multiple check-ins	45.51%	32.69%
Density	0.0081	0.0012

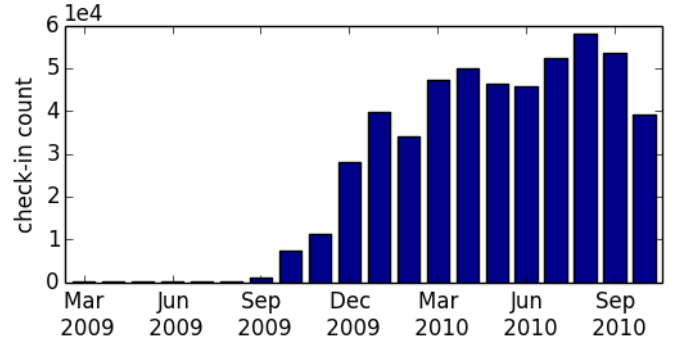
3.2 Time-Dependency of User Activities and Interests

Many studies analyzed users activity patterns over time [2], [20], [21], [46]. However, user activities have been mostly carried out at an hourly basis. We conduct a long-term user activity analysis, namely we study the user behavior over several months. This monthly analysis allows us to realize the long-term shift of users interest. Users' interests could evolve over time while sustaining their personal patterns. For instance, a teenager who reaches the legal age of drinking often starts going to bars (long-term interest shift), whereas an adult user often chooses to travel on Easter (personal pattern). Figure 1 depicts the distribution of check-ins of both datasets every month. There is a significant shift of activity on each month. For instance, we observe that the number of check-ins increases from Feb. to Jul. in both datasets. Next, we report further analysis only on Gowalla's, as we observed similar attributes when analyzing the Foursquare's dataset. Figure 2 shows the popularity of the top-8 venue categories. As we can see, before Oct. 2009 we observe dramatic shift of popularity. However, after Oct. 2009, the top categories exhibit a rather more stable popularity pattern as the dataset grows. Interestingly, we still can observe popularity shifts between "Coffee Shop", "American", "Mexican", "City Park", and "Asian". This suggests that users also exhibit a time-dependent pattern in visiting popular categories of POIs.

Figure 3 shows the activity of two samples of users in a long period. It also shows how the popularity of two samples of POI categories change over time. Note that, all the plots of Figure 3 illustrate the distribution of each user or category over time, that is, for each user or category, summing all the values of Y over time (X) equals to 1. We compute the variance of check-ins for both users and categories. Then, we pick the ones with the lowest in Figure 3a & c and the ones with the highest variances in Figure 3b & d. Our aim is to pick some representative user and categories from either end of the spectrum. As we can see, even users who are supposed to be the most stable users (Figure 3a) exhibit very time-dependent check-in patterns. It is worth noting that we ignored users with less than 50 check-ins in plotting Figure 3b, as those users mostly appear once in the dataset and hence are listed among the top variant users. However, we still observe that the users with high check-in variance are those who appear a few times and have less regular check-in behavior. This suggests that the more active a user is, the less variant her check-in pattern is. We also observed a *negative correlation between the users' check-in variance and quantity* (Spearman: -0.5583)



(a) Foursquare's



(b) Gowalla's

Fig. 1: Number of check-ins per month on Foursquare's and Gowalla's.

which supports our assumption. As for the POI categories, Figure 3c depicts the least variant categories. First, all of the categories relate to users' daily activities, for example going to a coffee shop or sandwich store. These categories are more time-independent since they are less affected by seasonal changes or weather conditions. Similar to what we observed for users, more popular categories are also less time-variant. In fact, we found a *negative correlation between POI categories variance and popularity* (Spearman: -0.8411). On the other hand, the more time-variant categories (Figure 3d) are less popular and they depend mainly on weather and political events. For example, "Ski Shop" only appears during winter and "Democratic Event" appears as a political campaign takes place.

All our observations suggest that there is an underlying temporal pattern in check-in activities from the perspective of both users and place categories. As we argued, check-in variance reveals meaningful information and thus we design our model focusing on these aspects.

3.3 Users' Multiple Check-ins

As we discussed earlier in Section 1, a good indicator of a user's preference is the fact that she visits the same POI multiple times. A multiple check-in between a user ρ and a venue l occurs if ρ visits l more than once. Other studies assumed multiple check-ins as an implicit feedback and demonstrated its effectiveness [11]. We also observed that a considerable amount of check-ins are multiple check-ins (Table 1). We randomly picked 100 users from the 500 most active users of Gowalla's. Figure 4 shows the quantity

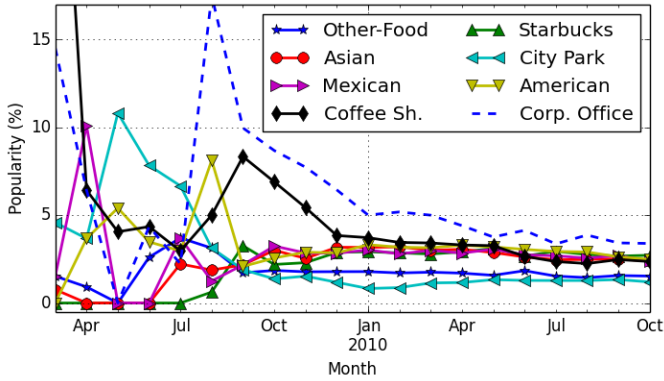


Fig. 2: Popularity of the top-8 categories over time on Gowalla’s (best viewed in color).

of single check-ins as well as multiple check-ins of the sample users. As we can see, users spend a considerable amount of their time visiting POIs they have visited before. In fact, multiple check-ins constitute 20% of check-ins for an average user on Gowalla’s (with a standard deviation of 21). According to Figure 4, the more active a user is, the more multiple check-ins she has. We calculated the correlation between the number of check-ins and the percentage of multiple check-ins per user. We observe that there is a positive correlation (Spearman: +0.4257) between the two variables, supporting our assumption. Based on these observations, we confirm that multiple check-ins provide valuable information regarding users’ preference and since these account for a considerable amount of the datasets, we design our model with a focus on them.

3.4 Remarks

Here, we present the main findings of our analysis. As observed in Section 3.2, users exhibit a long-term shift in their check-in behavior. As a consequence, POIs witness a shift in their popularity. However, many of these shifts are natural because users grow old and their habits evolve. Also, even though some types of POIs exhibit less time-sensitivity, many POIs are highly dependent on temporal phenomena such as seasonal changes or political campaigns. In addition, over a long period of time, less active users exhibit more variance in their check-in behavior. The same behavior is also observed for POIs, that is, less popular POIs suffer from more variance in terms of users’ check-ins. Thus, it is crucial to take into consideration how variant a user or POI has been in the past while learning the model.

According to 3.3, there is a positive correlation between the activity of a user and the number multiple check-ins she has had. This also means that a user who has visited the same POI several times in the past is more likely to visit the same POI in the future. Moreover, single check-ins are already crucial and should be considered in the model to account for users’ preferences as they choose to visit a particular POI when they could have chosen other venues in the same neighborhood. This gives us a rough estimate of their interest and preferences while multiple check-ins are more accurate indicators of users’ interest. Thus, both single

and multiple check-ins need to be involved while learning the model.

4 PROPOSED METHOD

Following the remarks based on our analysis in Section 3, in this section, we present our model. We first introduce the notations of our model. Let $\mathcal{P} = \{\rho_1, \dots, \rho_n\}$ be the set of n users and $\mathcal{L} = \{l_1, \dots, l_m\}$ be the set of m POIs. As an implicit feedback, we consider POIs that each user ρ_i has visited, \mathcal{L}_i , as “relevant” and all other unvisited POIs as “irrelevant” in the neighborhood where the user has visited all relevant items. Also, we consider POIs with multiple check-ins as “more relevant”. Therefore, we define \mathcal{L}_i^+ as the set of relevant POIs, \mathcal{L}_i^- as the set of irrelevant POIs, and \mathcal{L}_i^* the set of “more-relevant” POIs for user ρ_i . We also define $n_i^+ = |\mathcal{L}_i^+|$, $n_i^- = |\mathcal{L}_i^-|$, and $n_i^* = |\mathcal{L}_i^*|$.

Our aim is to define a ranking function $f_i(l)$ for each user ρ_i to rank more-relevant POIs higher than relevant POIs, and relevant POIs higher than irrelevant POIs. Moreover, we intend to incorporate the influence of POIs that are located close to each other. Let $U \in \mathbb{R}^{d \times n}$ be the latent factor for users and $V \in \mathbb{R}^{d \times m}$ be the latent factor for POIs. It is worth noting that \mathbf{u}_i corresponds to ρ_i and \mathbf{v}_j corresponds to l_j . The ranking function for the i^{th} user ρ_i and the j^{th} POI l_j would be $f_i(l_j) = \mathbf{u}_i^T \mathbf{v}_j$.

We design a two-phase objective function, where the first phase constructs ranking functions $f_i(l)$ such that relevant POIs are ranked higher than irrelevant POIs. In the second phase, functions $f_i(l)$ are updated to rank more-relevant POIs higher than relevant POIs. In Section 4.1, we first describe how we compute the distance between two POIs and in Section 4.2 we explain the first phase of the objective function in which the geographical influence is also incorporated. In Section 4.3, we give details on the second phase of the objective function. In Section 4.4, we describe the proposed time-sensitive regularizer which takes into consideration the users’ long-term behavioral patterns, and in Section 4.5 we present an overview of the proposed algorithm and how both phases of the objective function are optimized jointly.

4.1 Geographical Similarity

We compute the geographical similarity between two POI to incorporate the geographical context while characterizing the user’s geographical preferences. The similarity is inversely proportional to the distance between two POIs. Inspired by the relevant studies [11], [14], [52], [53] where a simple geographical measure improved the models significantly, we use the Haversine formula to compute the angular distance between l_i and l_j :

$$h_{ij} = 2 \times \arcsin \left(\sqrt{\sin^2(\Delta\phi_{ij}/2) + \cos\phi_i \times \cos\phi_j \times \sin^2(\Delta\eta_{ij}/2)} \right),$$

where ϕ_i and ϕ_j are latitudes of l_i and l_j in radian, respectively. Accordingly, η_i and η_j are longitudes of l_i and l_j in radian. Then we calculate the geographical similarity between l_i and l_j as follows:

$$g_{ij} = \frac{1}{1 + (h_{ij} \times R)},$$

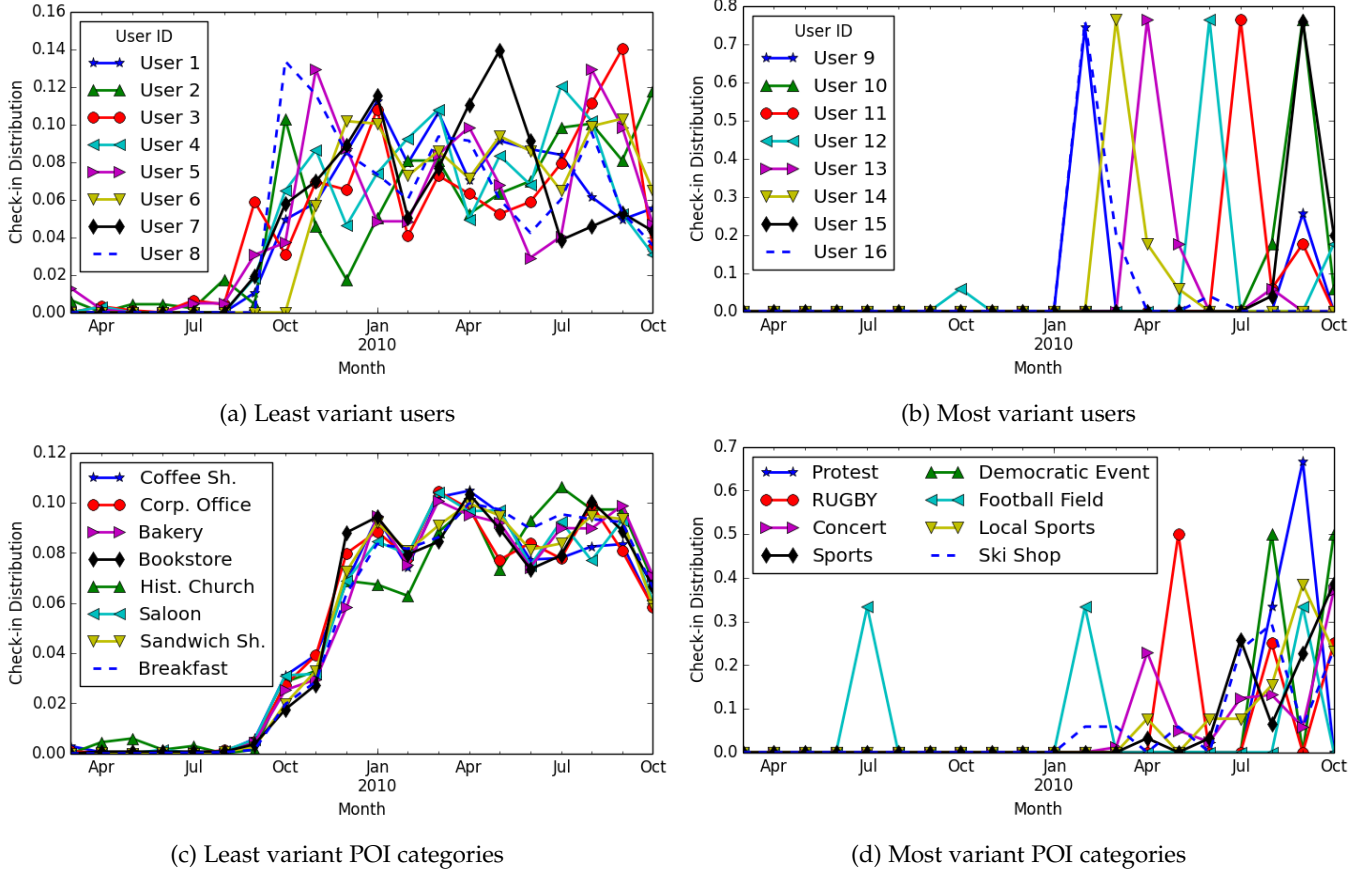


Fig. 3: Check-in distribution of users and POI categories over time in Gowalla’s. Figures a & c depict the least time-variant users and POI categories, respectively. Figures b & d, in contrast, show the distribution of the most time-variant users and POI categories, respectively (best viewed in color).

where R is the earth’s radius ($R = 6,371\text{KM}$). In the following section, we demonstrate how d_{ij} is incorporated in the our proposed method.

4.2 Phase 1: Visited vs. Unvisited POIs

In the first phase, we focus on ranking higher the POIs that a user has visited (no matter how many times) than the ones she has not visited. Formally, we aim at ranking \mathcal{L}_i^+ higher than \mathcal{L}_i^- . More specifically, our goal is to rank the POIs with emphasis on the top of the list. Moreover, with take into consideration the geographical distance between POIs. Building ranking functions that incorporate the distance between POIs also allows us to model latent associations between users living in the same neighborhood, who would not be associated with each other in a traditional CR setting. This happens because our method takes into account venues’ distances as it updates the user and item latent matrices. .

Let $H_i(l_j^-)$ be the “height” of an irrelevant venue:

$$H_i(l_j^-) = \sum_{k \in \mathcal{L}_i^+} \left(\mathbf{1}_{[f_i(l_k^+) \leq f_i(l_j^-)]} / (1 + \alpha \exp(g_{kj})) \right), \quad (1)$$

where α is the weight of geographical influence and $\mathbf{1}_{[\cdot]}$ is an indicator function. Note that α controls the model’s bias towards POIs in the same neighborhood and can be used to prevent the “Harry Potter” problem [54]. Dividing the

indicator function by G allows the model to incorporate the geographical distances into the model while constructing the height for irrelevant items. For example, if an irrelevant item is ranked higher than a relevant item, but they are very close then the denominator will be higher, reducing the height of the irrelevant POI. The objective function should aim at minimizing H_i for all given POIs. A lower value of H_i also means that there are fewer irrelevant POIs ranked higher than relevant ones. From an optimization perspective, indicator functions are not convex. Therefore, we use the logistic loss of the difference between the two functions as a convex upper bound surrogate. For reading simplicity we define

$$G_\alpha(i, j) = 1 + \alpha \exp(g_{ij}) .$$

We define the difference between the k^{th} POI and the j^{th} as follows:

$$\delta_i(k, j) = \mathbf{u}_i^T ((\mathbf{v}_k - \mathbf{v}_j) / G_\alpha(k, j)) . \quad (2)$$

Therefore, the surrogate height function $H'_i(l_j^-)$ becomes:

$$H'_i(l_j^-) = \sum_{k \in \mathcal{L}_i^+} \log [1 + \exp(-\delta_i(k, j))] , \quad (3)$$

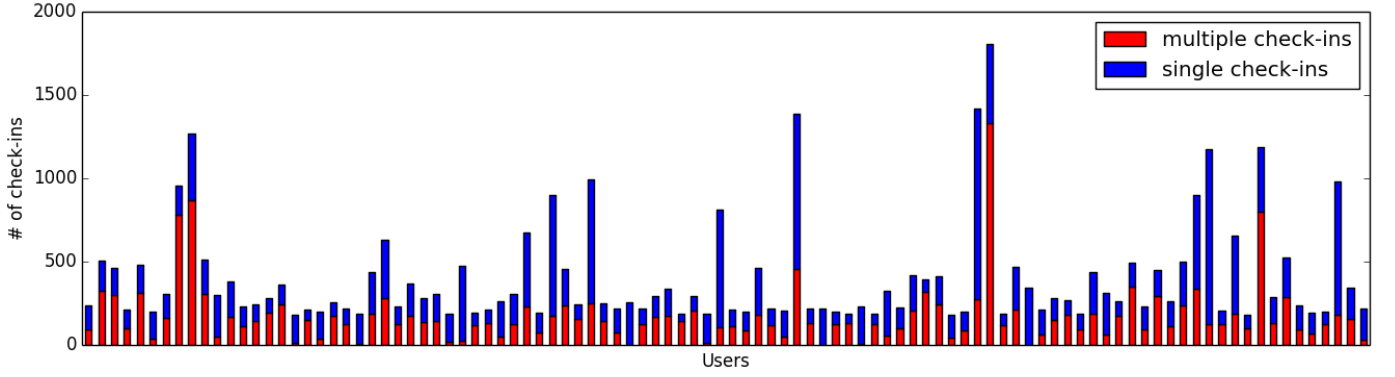


Fig. 4: Check-in histogram of 100 randomly-sampled users from the 500 most active users of Gowalla's. For each user, the red bar denotes the number of multiple check-ins, while the blue bar denotes the number of single check-ins.

where $\ell(\delta) = \log(1 + \exp(-\delta))$ is the logistic loss of δ . We consider ℓ_2 -norm of H'_i as the objective function, following [24]. Therefore, the objective function is:

$$R(U, V) = \sum_{i=1}^m \frac{1}{n_i} \sum_{j \in \mathcal{L}_i^-} (H'_i(l_j^-))^2 = \sum_{i=1}^m \frac{1}{n_i} \times \sum_{j \in \mathcal{L}_i^-} \left(\sum_{k \in \mathcal{L}_i^+} \log \left(1 + \exp \left(-\mathbf{u}_i^T [(\mathbf{v}_k - \mathbf{v}_j) / G_\alpha(k, j)] \right) \right) \right)^2. \quad (4)$$

For solving the above optimization problem, we use a gradient-descent-based alternating optimization algorithm. We first keep V fixed and update U , and then keep U and update V . Therefore, the update rules of the $t + 1$ iteration are:

$$\mathbf{u}_i^{t+1} = \mathbf{u}_i^t - \gamma \nabla_{\mathbf{u}_i} R(U^t, V^t), \forall i = 1, \dots, n, \quad (5)$$

$$\mathbf{v}_j^{t+1} = \mathbf{v}_j^t - \gamma \nabla_{\mathbf{v}_j} R(U^{t+1}, V^t), \forall j = 1, \dots, m. \quad (6)$$

For reading simplicity we define

$$\theta(k, j) = (1 + \exp(\delta(k, j))),$$

and the gradients of $R(U, V)$ with respect to \mathbf{u}_i and \mathbf{v}_j are computed as follows:

$$\begin{aligned} \nabla_{\mathbf{u}_i} R(U, V) &= \frac{2}{n_i} \sum_{j \in \mathcal{L}_i^-} \left(H'_i(l_j^-) \sum_{k \in \mathcal{L}_i^+} (\mathbf{v}_j - \mathbf{v}_k) / (G_\alpha(k, j) \theta(k, j)) \right), \end{aligned}$$

$$\begin{aligned} \nabla_{\mathbf{v}_j} R(U, V) &= \sum_{i \in \mathcal{P}_j^-} \frac{2}{n_i} \sum_{h \in \mathcal{L}_i^-} \left(H'_i(l_h^-) \sum_{k \in \mathcal{L}_i^+} \mathbf{u}_i / (G_\alpha(k, h) \theta(k, h)) \right) \\ &- \sum_{i \in \mathcal{P}_j^+} \frac{2}{n_i} \sum_{h \in \mathcal{L}_i^-} \left(H'_i(l_h^-) \sum_{k \in \mathcal{L}_i^+} \mathbf{u}_i / (G_\alpha(k, h) \theta(k, h)) \right), \end{aligned}$$

with \mathcal{P}_j^+ being the set of users who have visited l_j and \mathcal{P}_j^- the set of users who have not visited l_j .

4.3 Phase 2: Multiple vs. Single Check-ins

In the second phase, we focus on ranking POIs with multiple check-ins (i.e., more-relevant POIs) higher than the ones with single check-in (i.e., relevant POIs). As we observed in Section 3.3, there is a positive correlation between the number of multiple check-ins and total number check-ins. Therefore, it is crucial to take into consideration the fact that the POIs that users have visited more often in the past are more relevant. Formally, let $\mathcal{L}_i^{1+} = \mathcal{L}_i^+ - \mathcal{L}_i^*$ be the set of POIs that the i^{th} user has visited exactly once. Our goal is to rank \mathcal{L}_i^* higher than \mathcal{L}_i^{1+} . Let $\Pi_i(l_j^*)$ be the ‘‘reverse height’’ of a more-relevant venue, that is:

$$\Pi_i(l_j^*) = \sum_{k \in \mathcal{L}_i^{1+}} \mathbf{1}_{[f_i(l_k^*) \leq f_i(l_j^{1+})]}. \quad (7)$$

Lower values of Π_i mean that there are fewer relevant POIs ranked higher than the more-relevant ones. Similar to the first phase in Section 4.2, we use a logistic loss as the surrogate:

$$\Pi'_i(l_j^{1+}) = \sum_{k \in \mathcal{L}_i^*} \log(1 + \exp(-\mathbf{u}_i^T (\mathbf{v}_k - \mathbf{v}_j))). \quad (8)$$

However, Π'_i is not easy to be optimized using typical ranking loss, like DCG. Hence, we reformulate the objective functions as follows:

$$\sum_{k \in \mathcal{L}_i^*} \log(1 + \Pi_i(l_k^*)).$$

Then, the objective function $R_{\Pi}(U, V)$ of reverse height becomes:

$$\begin{aligned} R_{\Pi}(U, V) &= \sum_{i=1}^m \frac{1}{n_i} \sum_{j \in \mathcal{L}_i^*} \log(1 + \Pi'_i(l_j^*)) = \sum_{i=1}^m \frac{1}{n_i} \times \\ &\sum_{j \in \mathcal{L}_i^*} \log \left(1 + \sum_{k \in \mathcal{L}_i^{1+}} \log \left(1 + \exp(-\mathbf{u}_i^T (\mathbf{v}_k - \mathbf{v}_j)) \right) \right). \end{aligned} \quad (9)$$

We optimize $R_{\Pi}(U, V)$ similarly to Section 4.2, that is, we first keep V fixed and update U , then keep U updating V . Therefore, we consider the following update rules:

$$\mathbf{u}_i^{t+1} = \mathbf{u}_i^t - \gamma \nabla_{\mathbf{u}_i} R_{\Pi}(U^t, V^t), \forall i = 1, \dots, n, \quad (10)$$

$$\mathbf{v}_j^{t+1} = \mathbf{v}_j^t - \gamma \nabla_{\mathbf{v}_j} R_{\Pi}(U^{t+1}, V^t), \forall j = 1, \dots, m. \quad (11)$$

Similarly, the gradients are defined as follows:

$$\begin{aligned} & \nabla_{\mathbf{u}_i} R_{\Pi}(U, V) \\ &= \frac{1}{n_i} \sum_{j \in \mathcal{L}_i^*} \left(\frac{1}{1 + \Pi'_i(l_j^*)} \sum_{k \in \mathcal{L}_i^*} ((\mathbf{v}_j - \mathbf{v}_k) / (1 + \exp(\delta_i(k, j)))) \right), \\ & \nabla_{\mathbf{v}_j} R_{\Pi}(U, V) \\ &= \sum_{i \in \mathcal{P}_j^{1+}} \frac{1}{n_i} \sum_{h \in \mathcal{L}_i^*} \left(\frac{1}{1 + \Pi'_i(l_h^*)} \sum_{k \in \mathcal{L}_i^*} (\mathbf{u}_i / (1 + \exp(\delta_i(k, h)))) \right) \\ & - \sum_{i \in \mathcal{P}_j^*} \frac{1}{n_i} \sum_{h \in \mathcal{L}_i^*} \left(\frac{1}{1 + \Pi'_i(l_h^*)} \sum_{k \in \mathcal{L}_i^*} (\mathbf{u}_i / (1 + \exp(\delta_i(k, h)))) \right), \end{aligned}$$

where we define \mathcal{P}_j^{1+} the set of users who have visited l_j only once and \mathcal{P}_j^* the set of users who have visited l_j multiple times.

4.4 Time-Sensitive Regularizer

As we discussed in Section 3.2, the long-term temporal activity patterns of both users and POIs should be taken into account. One way to account for the activity patterns of users and the popularity of POIs is to consider how variant they are over time. For example, if a POI is a coffee shop and receives approximately the same number of people every month, it is more likely that it receives the same number of users in the next month. Whereas, POIs like ski shop are only popular during the ski season. In particular, we observed in Figure 3d that the popularity of certain POI categories are highly time-dependent. Here, we propose to incorporate a novel time-sensitive regularizer into the objective function of both of our objective phases in (4) and (9). Adding a regularizer that is calculated for each user and POI based on their past activities enables us to model the time-sensitivity of users and POIs. The main goal here is to penalize those users and POIs which are less stable. A more stable user or POI is one that exhibits less activity variance over time. This regularizer is defined based on our extensive analysis and observation in Section 3.2 where we observed that POIs that are less popular are more time-sensitive. We also had a similar observation for users, where we observed that less active users exhibit less stability in their check-in behavior.

Let $\sigma^{2,U} \in \mathbb{R}^{1 \times n}$ be the variance vector for users, where $\sigma_i^{2,U}$ is the activity variance of the i^{th} user ρ_i . For each user ρ_i , we count the number of check-ins per month and normalize the values. Then, calculating the variance of the monthly check-ins of ρ_i produces $\sigma_i^{2,U}$. Similarly, let $\sigma^{2,V} \in \mathbb{R}^{1 \times m}$ be the variance vector for POIs, with $\sigma_j^{2,V}$ being the popularity variance of the j^{th} POI l_j . Note that we calculate the variance of POIs based on their corresponding categories since we observed more meaningful popularity patterns with respect to the categories. The time-sensitive regularizer parameter for users and POIs are calculated as follows:

$$\Lambda^U = \lambda \log(1 + \exp(-\sigma^{2,U})), \quad (12)$$

$$\Lambda^V = \lambda \log(1 + \exp(-\sigma^{2,V})). \quad (13)$$

Algorithm 1: The Joint Two-Phase Collaborative Ranking Algorithm (JTCR).

Input: $\mathcal{P}, \mathcal{L}, \maxIter, \{\lambda, \gamma, \alpha, \epsilon, d\}$
Output: $U_{\text{out}}, V_{\text{out}}$

- 1 $t \leftarrow 0$
- 2 Initialize U^{t+1}, V^{t+1}
- 3 $\theta^{t+1} \leftarrow R(U^{t+1}, V^{t+1}) + R_{\Pi}(U^{t+1}, V^{t+1}), \theta^t = \frac{\theta^{t+1}}{2}$
- 4 **while** $(\text{abs}(\theta^{t+1} - \theta^t) > \epsilon) \wedge (t < \maxIter)$ **do**
- 5 $t \leftarrow t + 1$
- 6 // Phase 1
- 7 Update $\mathbf{u}_i^{t+1}, \forall i = 1, \dots, n$ Eq. (5)
- 8 Update $\mathbf{v}_j^{t+1}, \forall j = 1, \dots, m$ Eq. (6)
- 9 // Phase 2
- 10 Update $\mathbf{u}_i^{t+1}, \forall i = 1, \dots, n$ Eq. (10)
- 11 Update $\mathbf{v}_j^{t+1}, \forall j = 1, \dots, m$ Eq. (11)
- 12 $\theta^{t+1} \leftarrow R(U^{t+1}, V^{t+1}) + R_{\Pi}(U^{t+1}, V^{t+1})$
- 13 **end**
- 13 $U_{\text{out}} \leftarrow U^{t+1}, V_{\text{out}} \leftarrow V^{t+1}$

It is worth noting that we consider the logistic function of variances to prevent underflow and take the hyper parameter λ as a controlling parameter to prevent the model from overfitting. Ultimately, we add the time-sensitive regularizer to the objective functions of our two phases in (4) and (9). Thus, when updating \mathbf{u}_i , we add the regularizer term $\Lambda_i^U \mathbf{u}_i$, and for updating \mathbf{v}_j we add the regularizer term $\Lambda_j^V \mathbf{v}_j$.

4.5 Joint Two-Phase Collaborative Ranking Algorithm

Algorithm 1 presents the proposed joint two-phase collaborative ranking method. Line 2 initializes the factor matrices randomly. θ is initialized at line 3, summing up the values of the two phases of our objective function, namely, $R(U^{t+1}, V^{t+1})$ and $R_{\Pi}(U^{t+1}, V^{t+1})$. The joint optimization of the two phases is done between lines 4 and 12. As we see, in every iteration, \mathbf{u}_i and \mathbf{v}_j are first updated according to (5) and (6) (lines 7 and 8) to push visited POIs higher in the ranking. Each iteration is then followed by optimizing \mathbf{u}_i and \mathbf{v}_j according to (10) and (11) (lines 10 and 11), respectively. Therefore, U and V factor matrices are optimized jointly to push visited POIs higher than unvisited POIs and multiple visited POIs higher than single visited POIs simultaneously. After convergence, the final values of the latent factor matrices are stored at line 13. Note that the proposed time-sensitive regularizer is applied at lines 7, 8, 10, and 11. Also, the geographical influence is applied at lines 7 and 8. One can argue that employing a two-phase learning strategy might be computationally expensive. However, since in the second phase we only focus on the POIs that each user has checked in, the optimization algorithm does not add a substantial overhead to the whole system. In fact, for each user the complexity of one iteration is $\mathcal{O}(n_i^{1+} n_i^*)$, which is very small compared to the first phase, which is $\mathcal{O}(n_i^- n_i^+)$.

5 EXPERIMENTAL EVALUATION

In this section, we evaluate the performance of our model compared with state-of-the-art methods and study the impact of different parameters on the performance of our model.

5.1 Experimental Setup and Metrics

We evaluate our method on two real-world datasets, namely, Foursquare’s and Gowalla’s. Both datasets were provided by the authors of [20]. The statistical details of the datasets are listed in Table 1. We take the first 70% of the data for each user as the training set, 10% as the validation set, and the remaining 20% as the test set, following the evaluation protocol of [55].

We compare the performance of our model in terms of *Precision at k* (Prec@k) and *Normalized Discounted Cumulative Gain at k* (nDCG@k). Let $L_{ch}(\rho)$ be the set POIs that a user has visited in the test set and $L_{rec}^k(\rho)$ be the set of recommended POIs of size k . Prec@k(ρ) for a user ρ is defined as $\text{Prec@k}(\rho) = (|L_{ch}(\rho) \cap L_{rec}^k(\rho)|) / k$ and Prec@k for the whole dataset is the average Prec@k(ρ) for all the users in the test set.

To calculate nDCG@k, we need to define relevance values in the test set. Following the same strategy of Section 4, we define a three-level relevance for each POI based on the frequency of check-in for a particular user:

$$rel(l, \rho) = \begin{cases} 2 & \text{if } \rho \text{ visited } l \text{ multiple times} \\ 1 & \text{if } \rho \text{ visited } l \text{ only once} \\ 0 & \text{if } \rho \text{ did not visit } l \end{cases}.$$

Therefore, nDCG@k(ρ) for a given user ρ is defined as follows:

$$\text{DCG@k}(\rho) = \sum_{r=1}^k \frac{2^{rel(l_r, \rho)} - 1}{\log_2(r + 1)},$$

$$\text{nDCG@k}(\rho) = \frac{\text{DCG@k}(\rho)}{\text{IDCG@k}(\rho)},$$

where l_r is the POI at the r^{th} rank and IDCG@k(ρ) is the ideal DCG@k value for user ρ , that is, the highest possible value for DCG@k. The reported values of nDCG@k are the average of the nDCG@k(ρ) values for all the users in the test set. We report the values of nDCG@k and Prec@k for three values of k , namely 5, 10, 20.

5.2 Compared Methods

We compare our Joint Two-Phase Collaborative Ranking (JTCR) model with approaches that consider geographical influence for POI recommendation and approaches based on collaborative ranking with emphasis on ranking relevant items higher. Also, we include two variations of the proposed JTCR to demonstrate the effectiveness of different elements of our algorithm. Note that for each model, we find the optimum set of parameters using the validation set and report the mean and standard deviation of results of 5 different runs with the same parameters. We compare our JTCR model with the following methods:

- **JTCR-Phase1** reports the performance of the first phase of JTCR. We include this model as a baseline to demonstrate the effectiveness of the first phase of JTCR and the significance of the second phase of the algorithm.
- **JTCR-NoVar** reports the result of our proposed JTCR without using the time-sensitive regularizer. Instead, we use $\lambda/2(\|U\|^2 + \|V\|^2)$ as the regularizer. Our goal

is to demonstrate the effectiveness of the time-based regularizer.

- **JTCR-NoGeo** reports the result of our proposed JTCR without applying the geographical influence (i.e., $\alpha = 0$).
- **WRMF** [56] proposes an MF method for item prediction from implicit feedback. It is an adaptation of SVD, minimizing the square-loss.
- **IRenMF** [57] is based on weighted MF [58] exploiting two levels of geographical neighborhood characteristics: nearest neighboring locations share more similar user preferences, while locations in the same geographical region may share similar user preferences.
- **GeoMF** [14] augments users and venues latent factors in the factorization model with activity area vectors of users and influence area vectors of venues, respectively.
- **Rank-GeoFM** [11] is a ranking-based MF model that includes the geographical influence of neighboring venues while learning users preference rankings for venues.
- **Rank-GeoFM-NoGeo** reports the result of Rank-GeoFM without considering the geographical influence.
- **RH-Push / Inf-Push / P-Push** [16] are three push CR models based on reverse height, infinite, and p -norm. For each user, we considered the venues she visited as positive training samples and selected k venues randomly as negative training samples.

We aim to compare the performance of JTCR against state-of-the-art methods in POI recommendation which consider recommendation as a ranking problem and the ones that do not. Also, it is crucial to compare our method with approaches that incorporate geographical influence into the model. The other set of methods is based on CR. Our aim is to demonstrate the effectiveness of our two-phase regularized CR in comparison with other CR baselines.

5.3 Performance Comparison

Tables 2, 3, 4, and 5 report the performance of our method compared with 11 baselines in terms of nDCG@k and Prec@k for Foursquare’s and Gowalla’s. Based on the results we observe that our proposed JTCR significantly outperforms all the baseline methods on both datasets with respect to both nDCG@k and Prec@k. It is worth noting that the improvement is achieved for all values of k .

Moreover, Rank-GeoFM performs best among the geographical-based methods, as Rank-GeoFM propagates geographical influences using the constructed graph, which confirms that geographical neighborhood is a major factor for recommendation. Rank-GeoFM considers the implicit feedback while training the model similar to us, however, as we observe our two-phase collaborative ranking approach beats Rank-GeoFM indicating the effectiveness of our approach. Moreover, JTCR outperforms Rank-GeoFM by a large margin on both datasets. It is worth noting that JTCR beats all geographical-based methods in terms of both nDCG@k and Prec@k for all different values of k in both datasets. GeoMF and IRenMF do not consider POI recommendation as a ranking problem. Hence, they

TABLE 2: Performance evaluation on Foursquare’s in terms of nDCG@k. Statistically significant differences with JTCR are denoted by ∇ for $p < 0.05$ in paired t-test. Δ values express the relative difference, compared to JTCR. For each model we report the mean and standard deviation of 5 different runs.

	nDCG@5	Δ	nDCG@10	Δ	nDCG@20	Δ
JTCR	0.0639 \pm 0.0015	-	0.0529 \pm 0.0033	-	0.0394 \pm 0.0019	-
JTCR-Phase1	0.0534 \pm 0.0022 ∇	-16.43%	0.0436 \pm 0.0019 ∇	-17.58%	0.0339 \pm 0.0027 ∇	-13.96%
JTCR-NoVar	0.0605 \pm 0.0025 ∇	-5.32%	0.0497 \pm 0.0043 ∇	-6.05%	0.0375 \pm 0.0036 ∇	-4.82%
JTCR-NoGeo	0.0613 \pm 0.0017 ∇	-4.07%	0.0494 \pm 0.0022 ∇	-6.62%	0.0381 \pm 0.0019	-3.3%
WRMF	0.0248 \pm 0.0016 ∇	-61.19%	0.0210 \pm 0.0014 ∇	-60.3%	0.0178 \pm 0.0025 ∇	-54.82%
GeoMF	0.0422 \pm 0.0018 ∇	-33.96%	0.0336 \pm 0.0016 ∇	-36.48%	0.0250 \pm 0.0021 ∇	-36.55%
IRenMF	0.0430 \pm 0.0011 ∇	-32.71%	0.0348 \pm 0.0009 ∇	-34.22%	0.0286 \pm 0.0010 ∇	-27.41%
Rank-GeoFM	0.0438 \pm 0.0005 ∇	-31.46%	0.0359 \pm 0.0001 ∇	-32.14%	0.0277 \pm 0.0001 ∇	-29.70%
Rank-GeoFM-NoGeo	0.0418 \pm 0.0018 ∇	-34.59%	0.0314 \pm 0.0012 ∇	-40.64%	0.0234 \pm 0.0005 ∇	-40.61%
RH-Push	0.0251 \pm 0.0044 ∇	-60.72%	0.0187 \pm 0.0040 ∇	-64.65%	0.0137 \pm 0.0039 ∇	-65.23%
Inf-Push	0.0433 \pm 0.0053 ∇	-32.24%	0.0361 \pm 0.0051 ∇	-31.76%	0.0302 \pm 0.0054 ∇	-23.35%
P-Push	0.0423 \pm 0.0068 ∇	-33.8%	0.0309 \pm 0.0065 ∇	-41.59%	0.0218 \pm 0.0071 ∇	-44.67%

TABLE 3: Performance evaluation on Foursquare’s in terms of Prec@k. Statistically significant differences with JTCR are denoted by ∇ for $p < 0.05$ in paired t-test. Δ values express the relative difference, compared to JTCR. For each model we report the mean and standard deviation of 5 different runs.

	Prec@5	Δ	Prec@10	Δ	Prec@20	Δ
JTCR	0.0591 \pm 0.0025	-	0.0456 \pm 0.0041	-	0.0303 \pm 0.0046	-
JTCR-Phase1	0.0462 \pm 0.0013 ∇	-21.83%	0.0357 \pm 0.0022 ∇	-21.71%	0.0260 \pm 0.0021 ∇	-14.19%
JTCR-NoVar	0.0536 \pm 0.0018 ∇	-9.31%	0.0414 \pm 0.0015 ∇	-9.21%	0.0282 \pm 0.0025 ∇	-6.93%
JTCR-NoGeo	0.0551 \pm 0.0032 ∇	-6.77%	0.0411 \pm 0.0027 ∇	-9.87%	0.0292 \pm 0.0018	-3.63%
WRMF	0.0224 \pm 0.0020 ∇	-62.1%	0.0181 \pm 0.0012 ∇	-60.31%	0.0151 \pm 0.0017 ∇	-50.17%
GeoMF	0.0385 \pm 0.0015 ∇	-34.86%	0.0281 \pm 0.0016 ∇	-38.38%	0.0186 \pm 0.0012 ∇	-38.61%
IRenMF	0.0385 \pm 0.0009 ∇	-34.86%	0.0280 \pm 0.0011 ∇	-38.60%	0.0219 \pm 0.0015 ∇	-27.72%
Rank-GeoFM	0.0397 \pm 0.0004 ∇	-32.83%	0.0304 \pm 0.0002 ∇	-33.33%	0.0215 \pm 0.0001 ∇	-29.04%
Rank-GeoFM-NoGeo	0.0339 \pm 0.0007 ∇	-42.64%	0.0229 \pm 0.0003 ∇	-49.78%	0.0157 \pm 0.0003 ∇	-48.18%
RH-Push	0.0214 \pm 0.0042 ∇	-63.79%	0.0140 \pm 0.0051 ∇	-69.3%	0.0094 \pm 0.0046 ∇	-68.98%
Inf-Push	0.0397 \pm 0.0051 ∇	-32.83%	0.0313 \pm 0.0061 ∇	-31.36%	0.0217 \pm 0.0066 ∇	-28.38%
P-Push	0.0326 \pm 0.0072 ∇	-44.84%	0.0211 \pm 0.0101 ∇	-53.73%	0.0132 \pm 0.0094 ∇	-56.44%

attempt to optimize the overall error rate which proves to be less effective for POI recommendation mainly because the users are only interested in top k recommended POIs. Consequently, JTCR beats GeoMF and IRenMF with a large margin.

In addition, we observe that in most cases CR-based baseline methods, namely P-Push, Inf-Push, and RH-Push perform better than other baseline methods. This suggests that a CR-based approach leads to better performance for POI recommendation in general. However, we observe that JTCR outperforms all CR-based baselines. This indicates that all CR-based methods suffer from the sparsity problem while our two-phase CR strategy alleviates this problem by considering both visited and unvisited POIs in the same neighborhood. Also, none of the CR-based methods consider time in their ranking loss function. While it is important to consider POI recommendation as a ranking problem, it is also important to consider time to generate accurate recommendations (Section 3). Our proposed model beats the CR-based methods by a significant margin showing out time-sensitive regularizer (Section 4.4) based on the temporal behavior of users and the temporal popularity of POIs leads to a more accurate performance. Although

the CR-based baseline approaches focus on the top of the ranked list, they fail to rank more-relevant POIs higher in the ranking. In fact, these methods consider a binary relevance between users and POIs. Higher nDCG@k values indicate that our model ranks POIs with multiple check-ins more accurately, compared with the CR-based baseline approaches. This suggests that our two-phase model ranks the POIs with higher relevance more effectively than the CR approaches by considering multi-level implicit user feedback. It is worth noting that the variants of our model also outperform most of the baselines. In particular, JTCR-NoVar outperforms all the baselines and JTCR-Phase1 performs better than most of the baselines, including P-Push. This is important, since it indicates that incorporating the temporal information together with geographical similarities improves the performance of this model when the first phase is only considered. Also, we see that JTCR-NoGeo beats all the baselines. Specifically, it performs better than Rank-GeoFM-NoGeo and other baselines that do not consider geographical information. This indicates the effectiveness of the proposed model even when the geographical influence is not considered.

TABLE 4: Performance evaluation on Gowalla’s in terms of nDCG@k. Statistically significant differences with JTCR are denoted by ∇ for $p < 0.05$ in paired t-test. Δ values express the relative difference, compared to JTCR. For each model we report the mean and standard deviation of 5 different runs.

	nDCG@5	Δ	nDCG@10	Δ	nDCG@20	Δ
JTCR	0.1158 \pm 0.0022	-	0.0854 \pm 0.0007	-	0.0633 \pm 0.0008	-
JTCR-Phase1	0.1090 \pm 0.0021 ∇	-4.32%	0.0823 \pm 0.0015	-3.04%	0.0607 \pm 0.0012	-2.53%
JTCR-NoVar	0.1092 \pm 0.0029 ∇	-5.79%	0.0802 \pm 0.0022 ∇	-5.85%	0.0593 \pm 0.0009 ∇	-7.42%
JTCR-NoGeo	0.1099 \pm 0.0015 ∇	-5.09%	0.0823 \pm 0.0010	-3.63%	0.0608 \pm 0.0006	-3.95% ∇
WRMF	0.0620 \pm 0.0023 ∇	-46.46%	0.0523 \pm 0.0011 ∇	-38.76%	0.0425 \pm 0.0012 ∇	-32.86%
GeoMF	0.0604 \pm 0.0031 ∇	-47.84%	0.0495 \pm 0.0037 ∇	-42.04%	0.0374 \pm 0.0031 ∇	-40.92%
IRenMF	0.0606 \pm 0.0043 ∇	-47.67%	0.0499 \pm 0.0039 ∇	-41.57%	0.0389 \pm 0.0025 ∇	-38.55%
Rank-GeoFM	0.0593 \pm 0.0008 ∇	-48.79%	0.0525 \pm 0.0006 ∇	-38.52%	0.0451 \pm 0.0004 ∇	-28.75%
Rank-GeoFM-NoGeo	0.0675 \pm 0.0011 ∇	-41.71%	0.0514 \pm 0.0009 ∇	-39.81%	0.0387 \pm 0.0008 ∇	-38.86%
RH-Push	0.0985 \pm 0.0031 ∇	-14.94%	0.0765 \pm 0.0026 ∇	-10.42%	0.0569 \pm 0.0037 ∇	-10.11%
Inf-Push	0.1090 \pm 0.0035 ∇	-5.87%	0.0803 \pm 0.0028 ∇	-5.97%	0.0585 \pm 0.0018 ∇	-7.58%
P-Push	0.1026 \pm 0.0029 ∇	-11.40%	0.0805 \pm 0.0013 ∇	-5.74%	0.0596 \pm 0.0015 ∇	-5.85%

TABLE 5: Performance evaluation on Gowalla’s in terms of Prec@k. Statistically significant differences with JTCR are denoted by ∇ for $p < 0.05$ in paired t-test. Δ values express the relative difference, compared to JTCR. For each model we report the mean and standard deviation of 5 different runs.

	Prec@5	Δ	Prec@10	Δ	Prec@20	Δ
JTCR	0.0949 \pm 0.0021	-	0.0621 \pm 0.0007	-	0.0425 \pm 0.0007	-
JTCR-Phase1	0.0889 \pm 0.0018 ∇	-6.32%	0.0596 \pm 0.0012 ∇	-4.03%	0.0414 \pm 0.0011	-2.59%
JTCR-NoVar	0.0865 \pm 0.0032 ∇	-8.85%	0.0570 \pm 0.0015 ∇	-8.21%	0.0378 \pm 0.0012 ∇	-11.06%
JTCR-NoGeo	0.0866 \pm 0.0022 ∇	-8.75%	0.0590 \pm 0.0020 ∇	-4.99%	0.0401 \pm 0.0012 ∇	-5.65%
WRMF	0.0556 \pm 0.0018 ∇	-41.41%	0.0448 \pm 0.0012 ∇	-27.86%	0.0346 \pm 0.0010 ∇	-18.59%
GeoMF	0.0540 \pm 0.0028 ∇	-43.10%	0.0415 \pm 0.0042 ∇	-33.17%	0.0284 \pm 0.0032 ∇	-33.18%
IRenMF	0.0545 \pm 0.0051 ∇	-42.57%	0.0423 \pm 0.0043 ∇	-31.88%	0.0305 \pm 0.0023 ∇	-28.24%
Rank-GeoFM	0.0564 \pm 0.0009 ∇	-40.57%	0.0472 \pm 0.0008 ∇	-23.99%	0.0384 \pm 0.0004 ∇	-9.65%
Rank-GeoFM-NoGeo	0.0549 \pm 0.0008 ∇	-42.15%	0.0383 \pm 0.0008 ∇	-38.33%	0.0269 \pm 0.0007 ∇	-36.71%
RH-Push	0.0820 \pm 0.0039 ∇	-13.59%	0.0588 \pm 0.0030 ∇	-5.31%	0.0400 \pm 0.0042 ∇	-5.88%
Inf-Push	0.0864 \pm 0.0034 ∇	-8.96%	0.0569 \pm 0.0026 ∇	-8.37%	0.0377 \pm 0.0016 ∇	-11.29%
P-Push	0.0844 \pm 0.0029 ∇	-11.06%	0.0573 \pm 0.0012 ∇	-7.73%	0.0396 \pm 0.0018 ∇	-6.82%

5.4 Effect of the 2nd Phase

To study the effect of the second phase of JTCR, we compare the performance of JTCR when only the first phase is used (JTCR-Phase1) with the performance of JTCR when both phases are considered. As we can see in Tables 2 - 5, JTCR exhibits a significant improvement over JTCR-Phase1 in terms of all evaluation metrics for both datasets. This indicates that while JTCR-Phase1 is able to beat all other baselines, the second phase of the algorithm enables JTCR to model multiple check-ins more accurately. This validates the remark based on our analysis in Section 3.4 which states that a user who has visited a POI multiple times in the past is likely to visit the same venue in the future. Moreover, while this remark based on our analysis applies to a user, it is also valid with respect to similar users. Therefore, in the second phase, similar users and their corresponding collaborative associations are mainly determined based on how similar they are in terms of multiple check-ins. This helps the model rank “more relevant” items higher and hence improves the accuracy of the model. Moreover, we observe a higher relative difference on Foursquare’s. According to Table 1, Foursquare’s consists of more multiple check-ins than Gowalla’s (45.51% as opposed to 32.69%)

which explains why the second phase of JTCR achieves higher improvement on Foursquare’s.

5.5 Effect of the Time-Sensitive Regularizer

Next, we discuss the effect of the time-sensitive regularizer. To this end, we compare the performance of JTCR without using the time-sensitive regularizer (JTCR-NoVar). As seen in Tables 2 & 3, a statistically significant improvement of JTCR over JTCR-NoVar is observed in terms of all evaluation metrics for Foursquare’s, suggesting that using the time-sensitive regularizer enables JTCR to place more relevant venues higher in the ranking. As for Gowalla’s, we also see significant improvements in Tables 4 & 5 indicating that our proposed time-sensitive regularizer improves the performance of JTCR by penalizing those users and POIs that exhibit less stability in their check-in and popularity, respectively. This validates the remark based on our analysis that we had in Section 3.4 where we showed that there is a negative correlation between a POI’s popularity and its popularity variance. A similar observation was made for users. Based on this remark, we defined the time-sensitive regularizer to penalize those users and POIs that have been variant in the past. In other words, variant users are less

probable to visit variant POIs and the introduced regularizer enables the model to take this into account while training.

5.6 Effect of the Geographical Influence

Here, we discuss the effect of the geographical influence. Therefore, we compare the performance of JTTCR without applying the geographical influence (JTTCR-NoGeo). As seen in Tables 2 & 3, we observe a statistically significant improvement of JTTCR over JTTCR-NoGeo in terms of all evaluation metrics for Foursquare’s except for nDCG@20 and Prec@20. However, the significant improvement in terms of other evaluation metrics for Foursquare’s suggests that applying the geographical influence enables JTTCR to model users’ geographical behavior and activities more effectively. As for Gowalla’s, we see significant improvements in Tables 4 & 5 for all evaluation metrics except nDCG@10 and nDCG@20. This indicates that the geographical influence improves the performance of JTTCR by considering how users like POIs that are in the same neighborhood.

5.7 Effect of the Model Parameters

We demonstrate the effect of the model’s parameters. The results reported in the previous sections are achieved after the best parameter set was found on the validation set. We fixed the learning rate ($\gamma = 1e-4$) for both datasets to ensure the generalization of our model.

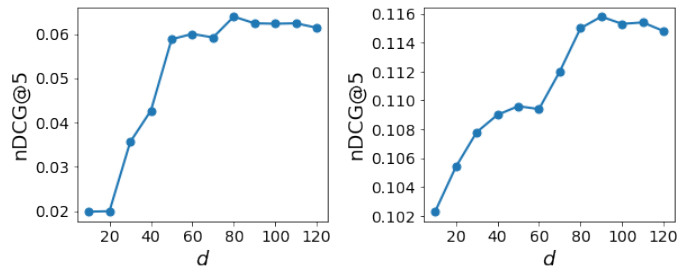
In Figure 5 we study the effect of latent factors d on the performance of our model and report nDCG@5 while keeping other parameters fixed. As we can see in Figure 5a, the optimal number of latent factors d for Foursquare’s is 80. For higher values of latent factors, nDCG@5 drops. Also, nDCG@5 drops when selecting lower values for d . We observe a similar behavior on Gowalla’s in Figure 5b with the difference that the optimal number of latent factors is 90. For all other values of d , we observe a drop in the performance.

Furthermore, in order to study the effect of the regularizing control parameter (λ), we varied λ while keeping d and α fixed. As shown in Figure 6a, the best λ for Foursquare’s is $1e-4$, while according to Figure 6b, for Gowalla’s is $1e-4$. The performance of our model drops using different values of λ for both datasets. While lower performance achieved with lower values of λ indicate that the introduced regularizer is essential to avoid overfitting, higher values of λ also hurt the performance.

Next, in Figure 7 we study the effect of geographical influence weight α on the performance of our model and report nDCG@5 while keeping other parameters fixed. As we can see in Figure 7a, the optimal value of α for Foursquare’s is 0.5 and the performance drops for all other values of α . We observe a similar behavior on Gowalla’s in Figure 7b where the best performance is achieved with $\alpha = 0.9$. For all other values of α , we see a drop in the performance.

5.8 Model’s Convergence

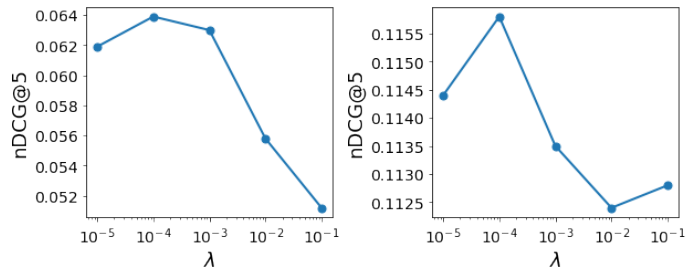
Finally, Figure 8 plots the value of the joint objective function, $\Theta^t = R(U^t, V^t) + R_{\Pi}(U^t, V^t)$ (Eq. (4) and (9)) when training the model in t iterations/epochs. We observe the value of Θ^t consistently decreases as the training epochs



(a) Foursquare’s

(b) Gowalla’s

Fig. 5: Effect of the number of latent factors.



(a) Foursquare’s

(b) Gowalla’s

Fig. 6: Effect of λ .

increase until the proposed JTTCR model converges. The behavior of the proposed objective function is as expected, since it is the summation of a logistic loss and a quasi convex function. R , that is a logistic loss, is convex and monotonic decreasing. Also, R_{Π} is strictly positive and monotonic decreasing. Hence, the summation of both loss functions converges as illustrated in Figure 8.

6 CONCLUSIONS AND FUTURE WORK

In this article, we presented an extensive data analysis of two POI recommendation datasets studying various attributes related to sparsity, time-sensitivity, and multiple check-ins. Based on the intuitions we got from data analysis, we proposed a two-phase CR model, called JTTCR. In addition, we showed how to incorporate the geographical influence into the objective functions and proposed a time-sensitive regularizer to capture the long-term user behavior and POI popularity patterns. The experimental results on the two benchmark datasets demonstrated that our proposed model outperforms other state-of-the-art methods. The results indicated that our model is able to address the data sparsity problem taking into account both visited and unvisited POIs in the training phase and their respective geographical distances. We also showed that the second phase is able to rank more-relevant POIs higher in the ranking, explaining the superiority of our two-phase model over the baselines as well as the first phase of the algorithm. This suggests that while single check-ins provide valuable information about the users’ preferences, multiple check-ins give us a more clear picture of their behavior and habits. Therefore, in the first phase of our algorithm, a CR model that focuses on ranking visited venues higher than unvisited

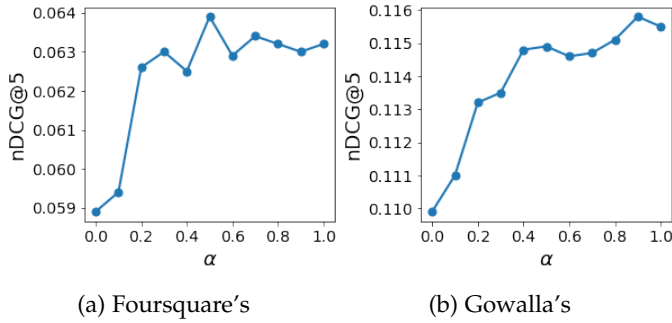
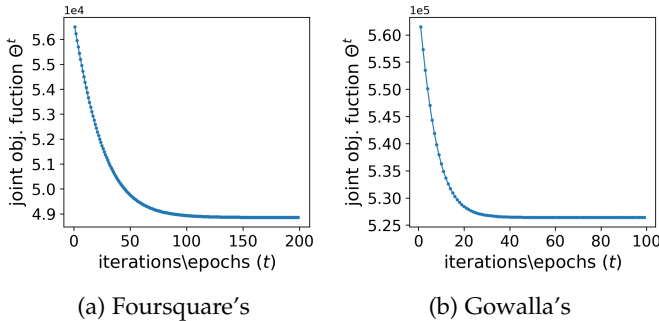
Fig. 7: Effect of α .

Fig. 8: Convergence of the joint objective function.

ones addresses the data sparsity by taking into account the unvisited venues in the training phase. In the second phase, a different CR approach is employed focusing on placing users' favorite POIs higher in the ranking. Throughout this process, we have regularized the learning procedure following the intuition that we had by analyzing time-sensitivity of users and POIs. Our aim was to penalize those POIs and users that have been more time-sensitive in the past.

A wealth of information about users and POIs are available on various LBSNs. As future work, we plan to extend our two-phase model to generate POI recommendations considering users' behavior in different LBSNs [59], [60]. Previous works have shown that considering a cross-platform and multimodal behavioral analysis improves the performance of a model dramatically. Hence, we are very interested in investigating how we can extend our current work to consider multimodal information such as POI category, user reviews, and opening hours. Also, with the recent advances in applying deep neural models for POI recommendation [61], [62], [63], [64] and their power to capture complex structures of user-POI interactions, we plan to combine our joint learning approach with the existing deep recurrent neural models to explore its potential benefits to a deep neural recommender model. Furthermore, since our training strategy requires negative training examples, we considered all unseen POIs as negative samples which increases the complexity of the model. Inspired by the relevant studies [29], [65], [66], we plan to explore various strategies for negative sampling and evaluate their effect on JTCR.

ACKNOWLEDGEMENT

This research was partially funded by the RelMobIR project of the Swiss National Science Foundation (SNSF).

REFERENCES

- [1] L. Chen, G. Chen, and F. Wang, "Recommender systems based on user reviews: the state of the art," *User Modeling and User-Adapted Interaction*, vol. 25, no. 2, pp. 99–154, 2015.
- [2] Q. Yuan, G. Cong, and A. Sun, "Graph-based point-of-interest recommendation with geographical and temporal influences," in *CIKM*, 2014, pp. 659–668.
- [3] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin, "Incorporating contextual information in recommender systems using a multidimensional approach," *ACM Trans. Inf. Syst.*, vol. 23, no. 1, pp. 103–145, 2005.
- [4] M. Aliannejadi, I. Mele, and F. Crestani, "Personalized ranking for context-aware venue suggestion," in *SAC*, 2017, pp. 960–962.
- [5] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, 2005.
- [6] J. Bao, Y. Zheng, D. Wilkie, and M. F. Mokbel, "Recommendations in location-based social networks: a survey," *GeoInformatica*, vol. 19, no. 3, pp. 525–565, 2015.
- [7] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang, "Collaborative location and activity recommendations with GPS history data," in *WWW*, 2010.
- [8] J. J. Levandoski, M. Sarwat, A. Eldawy, and M. F. Mokbel, "LARS: A location-aware recommender system," in *ICDE*, 2012, pp. 450–461.
- [9] G. Ference, M. Ye, and W. Lee, "Location recommendation for out-of-town users in location-based social networks," in *CIKM*, 2013, pp. 721–726.
- [10] J. Griesner, T. Abdesslem, and H. Naacke, "POI recommendation: Towards fused matrix factorization with geographical and temporal influences," in *RecSys*, 2015, pp. 301–304.
- [11] X. Li, G. Cong, X. Li, T. N. Pham, and S. Krishnaswamy, "Rank-geomf: A ranking based geographical factorization method for point of interest recommendation," in *SIGIR*, 2015, pp. 433–442.
- [12] M. Ye, D. Shou, W. Lee, P. Yin, and K. Janowicz, "On the semantic annotation of places in location-based social networks," in *KDD*, 2011, pp. 520–528.
- [13] M. Aliannejadi, I. Mele, and F. Crestani, "User model enrichment for venue recommendation," in *AIRS*, 2016, pp. 212–223.
- [14] D. Lian, C. Zhao, X. Xie, G. Sun, E. Chen, and Y. Rui, "Geomf: joint geographical modeling and matrix factorization for point-of-interest recommendation," in *KDD*, 2014, pp. 831–840.
- [15] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," in *NIPS*, 2007, pp. 1257–1264.
- [16] K. Christakopoulou and A. Banerjee, "Collaborative ranking with a push at the top," in *WWW*, 2015, pp. 205–215.
- [17] S. Balakrishnan and S. Chopra, "Collaborative ranking," in *WSDM*, 2012, pp. 143–152.
- [18] M. Weimer, A. Karatzoglou, Q. V. Le, and A. J. Smola, "COFI RANK - maximum margin matrix factorization for collaborative ranking," in *NIPS*, 2007, pp. 1593–1600.
- [19] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "BPR: bayesian personalized ranking from implicit feedback," in *UAI*, 2009, pp. 452–461.
- [20] Q. Yuan, G. Cong, Z. Ma, A. Sun, and N. Magnenat-Thalmann, "Time-aware point-of-interest recommendation," in *SIGIR*, 2013, pp. 363–372.
- [21] H. Gao, J. Tang, X. Hu, and H. Liu, "Exploring temporal effects for location recommendation on location-based social networks," in *RecSys*, 2013, pp. 93–100.
- [22] Y. Liu, C. Liu, B. Liu, M. Qu, and H. Xiong, "Unified point-of-interest recommendation with temporal interval assessment," in *KDD*, 2016, pp. 1015–1024.
- [23] S. Feng, X. Li, Y. Zeng, G. Cong, Y. M. Chee, and Q. Yuan, "Personalized ranking metric embedding for next new POI recommendation," in *IJCAI*, 2015, pp. 2069–2075.
- [24] C. Rudin, "The p-norm push: A simple convex ranking algorithm that concentrates at the top of the list," *Journal of Machine Learning Research*, vol. 10, pp. 2233–2271, 2009.
- [25] D. Goldberg, D. A. Nichols, B. M. Oki, and D. B. Terry, "Using collaborative filtering to weave an information tapestry," *Commun. ACM*, vol. 35, no. 12, pp. 61–70, 1992.
- [26] H. Yin, B. Cui, Y. Sun, Z. Hu, and L. Chen, "LCARS: A spatial item recommender system," *ACM Trans. Inf. Syst.*, vol. 32, no. 3, pp. 11:1–11:37, 2014.

- [27] B. M. Sarwar, G. Karypis, J. A. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *WWW*, 2001, pp. 285–295.
- [28] Y. Koren, "Factorization meets the neighborhood: a multifaceted collaborative filtering model," in *CIKM*, 2008, pp. 426–434.
- [29] M. Ye, P. Yin, W. Lee, and D. L. Lee, "Exploiting geographical influence for collaborative point-of-interest recommendation," in *SIGIR*, 2011, pp. 325–334.
- [30] C. Cheng, H. Yang, I. King, and M. R. Lyu, "A unified point-of-interest recommendation framework in location-based social networks," *ACM Trans. Intell. Sys. and Tech.*, vol. 8, no. 1, pp. 10:1–10:21, 2016.
- [31] C. Zhang, H. Liang, and K. Wang, "Trip recommendation meets real-world constraints: POI availability, diversity, and traveling time uncertainty," *ACM Trans. Inf. Syst.*, vol. 35, no. 1, pp. 5:1–5:28, 2016.
- [32] M. Aliannejadi, D. Rafailidis, and F. Crestani, "Personalized keyword boosting for venue suggestion based on multiple LBSNs," in *ECIR*, 2017, pp. 291–303.
- [33] M. Aliannejadi and F. Crestani, "Venue appropriateness prediction for personalized context-aware venue suggestion," in *SIGIR*, 2017, pp. 1177–1180.
- [34] W. Zhang, G. Ding, L. Chen, C. Li, and C. Zhang, "Generating virtual ratings from chinese reviews to augment online recommendations," *ACM Trans. Intell. Sys. and Tech.*, vol. 4, no. 1, pp. 9:1–9:17, 2013.
- [35] C. Cui, J. Shen, L. Nie, R. Hong, and J. Ma, "Augmented collaborative filtering for sparseness reduction in personalized POI recommendation," *ACM Trans. Intell. Sys. and Tech.*, vol. 8, no. 5, pp. 71:1–71:23, 2017.
- [36] F. Yuan, J. M. Jose, G. Guo, L. Chen, H. Yu, and R. S. Alkhalwaleh, "Joint geo-spatial preference and pairwise ranking for point-of-interest recommendation," in *ICTAI*, 2016, pp. 46–53.
- [37] T. Liu, "Learning to rank for information retrieval," *Foundations and Trends in Information Retrieval*, vol. 3, no. 3, pp. 225–331, 2009.
- [38] K. Crammer and Y. Singer, "Pranking with ranking," in *NIPS*, 2001, pp. 641–647.
- [39] C. J. C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. N. Hullender, "Learning to rank using gradient descent," in *ICML*, 2005, pp. 89–96.
- [40] Z. Cao, T. Qin, T. Liu, M. Tsai, and H. Li, "Learning to rank: from pairwise approach to listwise approach," in *ICML*, 2007, pp. 129–136.
- [41] Y. Shi, A. Karatzoglou, L. Baltrunas, M. Larson, and A. Hanjalic, "Gapfm: optimal top-n recommendations for graded relevance domains," in *CIKM*, 2013, pp. 2261–2266.
- [42] D. Rafailidis and F. Crestani, "Collaborative ranking with social relationships for top-n recommendations," in *SIGIR*, 2016, pp. 785–788.
- [43] —, "Joint collaborative ranking with social relationships in top-n recommendation," in *CIKM*, 2016, pp. 1393–1402.
- [44] —, "Learning to rank with trust and distrust in recommender systems," in *RecSys*, 2017, pp. 5–13.
- [45] Y. Ding and X. Li, "Time weight collaborative filtering," in *CIKM*, 2005, pp. 485–492.
- [46] Z. Yao, Y. Fu, B. Liu, Y. Liu, and H. Xiong, "POI recommendation: A temporal matching between POI popularity and user regularity," in *ICDM*, 2016, pp. 549–558.
- [47] X. Li, M. Jiang, H. Hong, and L. Liao, "A time-aware personalized point-of-interest recommendation via high-order tensor factorization," *ACM Trans. Inf. Syst.*, vol. 35, no. 4, pp. 31:1–31:23, 2017.
- [48] H. Yin, B. Cui, X. Zhou, W. Wang, Z. Huang, and S. W. Sadiq, "Joint modeling of user check-in behaviors for real-time point-of-interest recommendation," *ACM Trans. Inf. Syst.*, vol. 35, no. 2, pp. 11:1–11:44, 2016.
- [49] H. Yin, B. Cui, L. Chen, Z. Hu, and Z. Huang, "A temporal context-aware model for user behavior modeling in social media systems," in *SIGMOD*, 2014, pp. 1543–1554.
- [50] H. Yin, B. Cui, L. Chen, Z. Hu, and X. Zhou, "Dynamic user modeling in social media systems," *ACM Trans. Inf. Syst.*, vol. 33, no. 3, pp. 10:1–10:44, 2015.
- [51] C. Cheng, H. Yang, M. R. Lyu, and I. King, "Where you like to go next: Successive point-of-interest recommendation," in *IJCAI*, 2013, pp. 2605–2611.
- [52] M. Aliannejadi and F. Crestani, "A collaborative ranking model with contextual similarities for venue suggestion," in *IIR*, 2018.
- [53] M. Aliannejadi, D. Rafailidis, and F. Crestani, "A collaborative ranking model with multiple location-based similarities for venue suggestion," in *ICTIR*, 2018, pp. 19–26.
- [54] M. Koolen, T. Bogers, A. van den Bosch, and J. Kamps, "Looking for books in social media: An analysis of complex search requests," in *ECIR*, 2015, pp. 184–196.
- [55] S. Zhao, T. Zhao, I. King, and M. R. Lyu, "Geo-teaser: Geotemporal sequential embedding rank for point-of-interest recommendation," in *WWW*, 2017, pp. 153–162.
- [56] Y. Hu, Y. Koren, and C. Volinsky, "Collaborative filtering for implicit feedback datasets," in *ICDM*, 2008, pp. 263–272.
- [57] Y. Liu, W. Wei, A. Sun, and C. Miao, "Exploiting geographical neighborhood characteristics for location recommendation," in *CIKM*, 2014, pp. 739–748.
- [58] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. M. Lukose, M. Scholz, and Q. Yang, "One-class collaborative filtering," in *ICDM*, 2008, pp. 502–511.
- [59] D. Rafailidis and F. Crestani, "Top-n recommendation via joint cross-domain user clustering and similarity learning," in *ECML/PKDD*, 2016, pp. 426–441.
- [60] M. Aliannejadi and F. Crestani, "Personalized context-aware point of interest recommendation," *ACM Trans. Inf. Syst.*, vol. 36, no. 4, pp. 45:1–45:28, 2018.
- [61] Q. Liu, S. Wu, L. Wang, and T. Tan, "Predicting the next location: A recurrent model with spatial and temporal contexts," in *AAAI*, 2016, pp. 194–200.
- [62] N. Du, H. Dai, R. Trivedi, U. Upadhyay, M. Gomez-Rodriguez, and L. Song, "Recurrent marked temporal point processes: Embedding event history to vector," in *KDD*, 2016, pp. 1555–1564.
- [63] H. Yin, W. Wang, H. Wang, L. Chen, and X. Zhou, "Spatial-aware hierarchical collaborative deep learning for POI recommendation," *IEEE Trans. Knowl. Data Eng.*, vol. 29, no. 11, pp. 2537–2551, 2017.
- [64] J. Manotumruksa, C. Macdonald, and I. Ounis, "A deep recurrent collaborative filtering framework for venue recommendation," in *CIKM*, 2017, pp. 1429–1438.
- [65] —, "A personalised ranking framework with multiple sampling criteria for venue recommendation," in *CIKM*, 2017, pp. 1469–1478.
- [66] C. Cheng, H. Yang, I. King, and M. R. Lyu, "Fused matrix factorization with geographical and social influence in location-based social networks," in *AAAI*, 2012, pp. 17–23.



Mohammad Aliannejadi is a Ph.D. candidate at the Università della Svizzera italiana (USI) in Lugano, Switzerland, working on Mobile Information Retrieval as well as POI Recommendation. His main emphasis is on studying cross-app mobile search. Before that, he was an M.Sc. student at Amirkabir University of Technology (Tehran Polytechnic) in Tehran, Iran, where he worked on Spoken Language Understanding.



Dimitrios Rafailidis received the B.Sc., M.Sc., and Ph.D. degrees from the Department of Informatics of the Aristotle University of Thessaloniki, Greece, in 2005, 2007, and 2011, respectively. He is an Assistant Professor at Maastricht University. His research interests include Machine Learning, Information Retrieval, Recommender Systems and Social Media Mining.



Fabio Crestani is a Full Professor at the Faculty of Informatics of the Università della Svizzera italiana (USI) in Lugano, Switzerland, since 2007. Previously he was a Professor at the University of Strathclyde in Glasgow, United Kingdom. Prof Crestani is an internationally recognized researcher in Information Retrieval, Text Mining and Digital Libraries. In these areas he has published over two hundred refereed papers on both theoretical and experimental topics. He was the Editor-in-Chief of *Information Processing and Management* until 2015, has served in the organizing and program committee of several conferences and in the editorial board of several journals.