Personalized Landmark Recommendation Based on Geotags from Photo Sharing Sites

Yue Shi[1], Pavel Serdyukov[2], Alan Hanjalic[1], Martha Larson[1]
[1]Multimedia Information Retrieval Lab, Delft University of Technology, [2]Yandex
[1]{y.shi, a.hanjalic, m.a.larson}@tudelft.nl, [2]pserdukov@yandex.ru

Abstract

Geotagged photos of users on social media sites provide abundant location-based data, which can be exploited for various location-based services, such as travel recommendation. In this paper, we propose a novel approach to a new application, i.e., personalized landmark recommendation based on users' geotagged photos. We formulate the landmark recommendation task as a collaborative filtering problem, for which we propose a category-regularized matrix factorization approach that integrates both user-landmark preference and category-based landmark similarity. We collected geotagged photos from Flickr and landmark categories from Wikipedia for our experiments. Our experimental results demonstrate that the proposed approach outperforms popularity-based landmark recommendation and a basic matrix factorization approach in recommending personalized landmarks that are less visited by the population as a whole.

Introduction

Since GPS positioning has become a standard functionality of mobile digital devices, like cell phones or digital cameras, the history of user locations is readily available for a variety of location-based services (Zheng et al. 2010). As a result of such ubiquity, the volume of location data stored at social media sites has also greatly increased. The largest repositories of user location histories are, in fact, photo-sharing services, like Flickr (www.flickr.com), where locations (geotags) are attached to a significant number of photos uploaded by their users. As a result, for groups of users that regularly share not only photos, but also their locations, it is possible to provide additional location-based services by mining their personal location data (Clements et al. 2010; Serdyukov et al. 2009). In this paper, we focus on personalized landmark recommendation based on geotagged photos. The technique we propose is designed to be deployed in an application scenario that exploits knowledge from geotagged photos of users in an online community to recommend landmarks to a user (i.e., traveler) given a city that is new to that user.

Our motivation for this work is twofold: First, a traveler may benefit from recommendations of landmarks to visit, when she, is traveling or plans to travel to a city for the first time. Deriving the traveler's preference from her activity at social media sharing websites for landmark recommendations requires no effort from the traveler. Second, different travelers may probably have different preference for places to see in the same city. To meet the specific needs of individuals, landmark recommendations must be personalized. Therefore, a promising solution we propose in this paper is to provide landmark recommendations in a city for a traveler based on both the other travelers' preference for landmarks in the city and her preference on landmarks in the cities that she has visited before.

We propose to tackle the personalized landmark recommendation problem via the collaborative filtering (CF) paradigm (Adomavicius and Tuzhilin 2005), which recommends, for a given user, favored items of similar users. We reason that a user in a new city may like landmarks that are already favored by other users who have similar landmark visiting experience in other cities in the past. Note that we define a landmark in our work as a place with a significance for history, culture or contemporary society. As such, we consider locations that are landmarks to be broader than only officially designated landmarks (such as monuments). We emphasize that the personalized landmark recommendation studied in this paper differs from conventional CF scenarios in the following aspects. First, item ratings, i.e., graded preferences on certain landmarks are not directly available from the users' photos, as they would be in a conventional recommendation scenario such as movie recommendation, where users rate movies explicitly. Second, travelers often seek to enjoy a unique experience or to avoid tourist traps. For this reason, recommending popular and well-known landmarks is much less useful than providing information about landmarks that are less well-known, but are potentially interesting to a specific user. For example, a traveler visiting Paris for the first time would not need the recommendation of Eiffel Tower, since he would know about its existence and likely go there an-

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.
This work is funded from PetaMedia Network of Excellence EC FP7 grant n° 216444.

yway. However, he would appreciate recommendations of landmarks that are less visited by the general population, but rather fit his own interests well. Third, the number of visited landmarks for most of the users could be very limited, resulting in a very sparse user-landmark relation. Although data sparseness is a typical problem in common CF scenarios, we anticipate that it is more severe in landmark recommendation than in other scenarios, e.g., movie or music recommendation. Visiting landmarks is based on the availability of money, time and travel documents (e.g., valid visa). For these reasons, an individual user might have visited only relatively few of the very large total number of landmarks.

Our contributions in this paper can be summarized as follows. First, we put forward a new research topic, namely, personalized landmark recommendation using users' geotagged photos, and contribute a new data collection for this research. Second, we propose a novel approach, category-regularized matrix factorization, that exploits both user-landmark preference and category-based landmark similarity for personalized landmark recommendation, and demonstrate its effectiveness for recommending less-visited landmarks.

Related Work

Our work in this paper is related to location-based recommendation. Location data from user-uploaded photos, i.e., the geotags of photos, has been exploited to approach various tasks. A travel guidance system (Gao et al. 2010) was presented to automatically recognize and rank landmarks based on the tags and geotags of photos in Flickr and on the information extracted from online community travel sites, like Yahoo Travel Guide (travel.yahoo.com). The geotags of photos were also exploited for facilitating trip planning (Lu et al. 2010) in order to help travelers to discover attractive places/landmarks, find a proper path around a landmark, and find a proper route between landmarks. Arase et al. (2010) proposed to use geotagged photos to mine frequent trip patterns, i.e., frequently visited city sequences and typical visit duration, factors that are widely recognized for their potential for improving travel recommendation. De Choudhury et al. (2010) proposed to construct travel itineraries automatically by aggregating users' location data, e.g., staying time in a place, transit time between places, etc., which are extracted from geotagged photos. These approaches all address the problem of the general optimization of travel guides by relying on different available information resources. In contrast, the work in this paper addresses personalized landmark recommendation task, i.e., recommendations are generated in view of each user's individual preference.

To the best of our knowledge, there are only two studies that are closely related to this paper in that they also focus on personalized location-based recommendation. One of them proposed to personalize location prediction by exploiting geotagged photos to re-rank popular locations by integrating predictions based on similar users (Clements et al. 2010). This work discovered the dominant effect of popular locations in recommendation, which is also confirmed in our experiments, and also suggested that recommendations for less-frequently visited locations would be more interesting to travelers, the focus evaluated in this paper. The other instance of closely related work, (Kurashima et al. 2010), proposed to construct a generative photographer behavior model for personalized travel route recommendation by combining a Markov model for location dependence and a topic model for user interest dependence. Users' travel route histories were represented by geotags of their photos, and the ordering of locations was represented by the timestamps of users' photos. Compared to their work, our approach is substantially different in two aspects: First, our work focuses on a new application, i.e., landmark recommendation. Second, we specifically target the data sparseness problem and the challenge of recommending less frequently visited locations, both of which have not been investigated before.

Our work in this paper is also related to matrix factorization (MF) techniques in recommender systems. Generally, MF techniques (Koren et al. 2009) learn latent features of users and items from the observed preference in the user-item matrix and these features are then used to predict unobserved preferences. Recently, different regularizations have been proposed to extend the basic MF in order to make it suited for different purposes. For example, Ma et al. (2009) proposed to use the social relationships among users in a networked community to regularize the factorization of the user-item rating matrix for improved rating prediction. Furthermore, Zheng et al. (2010) proposed to deploy user activity correlation and location correlation in terms of location features to regularize the factorization of a location-activity matrix for improved location and activity recommendation. Also, Shi et al. (2010) proposed to use movie mood properties to regularize the factorization of a user-movie rating matrix for improved mood-specific movie recommendation. In this paper, we exploit landmark categories derived from Wikipedia to regularize the factorization of user-landmark preference matrix for improved landmark recommendation.

Category-regularized Matrix Factorization

In this section, we present the category-regularized matrix factorization (CRMF) approach for personalized landmark recommendation. Supposing there are M users and N landmarks, we integrate the category-based landmark similarity with the basic MF approach, which leads to the following objective function of CRMF:

$$L(\mathbf{U},\mathbf{V}) = \frac{1}{2}\sum_{i=1}^{M}\sum_{j=1}^{N} I_{ij}^{R}\left(R_{ij} - U_i^T V_j\right)^2 + \frac{\alpha}{2}\sum_{j=1}^{N}\sum_{n=1}^{N} I_{jn}^{S}\left(S_{jn} - V_j^T V_n\right)^2 + \frac{\lambda}{2}\left(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2\right) \quad (1)$$

where R_{ij} denotes the user i's preference on landmark j. Note that in this paper we use the normalized number of photos that user i took around landmark j to represent R_{ij}. U_i denotes the d-dimensional latent features (as a column vector) of user i, and V_j denotes the d-dimensional latent features (as a column vector) of landmark j. I_{ij}^X is an indicator function that is equal to 1 if X_{ij}>0, and 0 otherwise. $\|\mathbf{U}\|_F$ and $\|\mathbf{V}\|_F$ are the Frobenius norm of **U** and **V**, respectively, that serve to alleviate overfitting. λ is the norm regularization parameter. S_{jn} denotes cosine similarity between landmark j and landmark n in terms of their categories. Note that in this paper we exploit landmark categories from Wikipedia. A tradeoff parameter, i.e., α, serves to control the relative contribution from the category-based landmark similarity to the learned latent features. Note that CRMF becomes equivalent to the basic MF when α is equal to 0. Minimizing the objective function results in the latent features of users and landmarks learned from the user-landmark preference in **R** and from the category-based landmark similarity in **S**.

Since the objective function is not jointly convex over **U** and **V**, we choose to use alternated gradient descent to solve the minimization problem. With the learned latent features of **U** and **V**, we can predict user i's preference on landmark j as $U_i^T V_j$, which provides the basis for ranking landmarks for personalized recommendations. Note that in the recommendation list we exclude the landmarks that the user has already visited.

Data Description

The data for our experiments was collected using public Flickr API. First, we downloaded metadata for 42.9M geotagged photos and then filtered this data according to the following criteria. Considering that the overwhelming majority of Flickr users are from the US, we focused only on users with geotagged photos made outside the US in order to increase the share of "travelers" in our dataset. We also only considered the photos of city landmarks and continued with 40 cities in which the most photos were taken. The above steps resulted in 126,123 geotagged photos from 40,084 users. We extracted landmarks for each photo corresponding to geotagged Wikipedia articles that are less than one kilometer away from the geotag of that photo. In addition, we eliminated landmarks for which there was negligible overlap between the words in the title of the Wikipedia article and the tags assigned to the photo by the user. Finally, as mentioned above, we used the normalized number of a user's photos related to a landmark to quantify a user's preference for a landmark. The resulting user-landmark preference matrix consists of 260,362 scores from 40,084 users and 9,557 landmarks.

Fig. 1 indicates that both users and landmarks closely follow a power-law distribution, from which two conclusions could be drawn: 1) most of the users in our dataset visited only a limited number of landmarks (Fig. 1(a)), resulting in a sparse user-landmark preference matrix; 2) there are quite a few landmarks visited by lots of users (Fig. 1(b)), which makes the challenge and importance of recommending less-popular landmarks quite substantial.

In addition, we extracted categories for each landmark from Wikipedia, which resulted in a landmark-category binary matrix with 7,379 categories. The landmark category is used as another modality, independent of user-landmark preference, from which category-based landmark similarity can be obtained. The dataset is available at (http://homepage.tudelft.nl/q0v1y/yueshi/plr_data.zip).

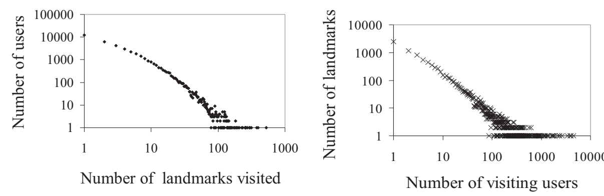

(a) (b)

Fig. 1. (a) Log-log plot relating the number of users to the numbers of visited landmarks, (b) Log-log plot relating the number of landmarks to the number of visiting users.

Experimental Evaluation

The most straightforward application from this study is to provide a user with landmark recommendations when she plans to visit a city. Our experiments are intended to evaluate the proposed CRMF approach for personalized landmark recommendation under a simulated setting that approximates this application. For this reason, in our experiments we use only the data of those users who have visited at least one landmark in each of at least two cities. In this way, we can define at least one "already visited" and at least one "target" city for each user. The user-landmark matrix used for the experiments contains 14,031 users and has a sparseness of 99.89%.

We randomly separate the dataset into a training set, a validation set and a test set. For the training set, we randomly select 60% users and their landmark preferences. We then randomly select another 20% users and their landmark preferences to form the validation set, which is used to tune parameters. The remaining 20% users and their preferences are used as the test set, which is used to evaluate the performance of CRMF and compare it to that of the baselines. We adopted one of the most widely used evaluation metrics, i.e., mean average precision (MAP) to evaluate the recommendation quality. The higher the MAP, the better is the recommendation performance. When evaluating the CRMF performance, for each user in the validation and test set, we randomly select one city that she has visited as that user's target city. Then we remove the user's landmark preferences for this selected target city and try to reproduce them. We compare landmark recommendation algorithms on the basis of their ability to predict the deleted landmark preferences in terms of MAP.

According to observations from the validation set, we set λ=1 and the tradeoff parameter α=0.001. With this fixed parameter setting, we compare CRMF with two baseline approaches on the test set:

- **PopRec**: Landmarks are recommended to users based on their popularity, which is defined in terms of the number of users who visited them in the training set. It is a non-personalized recommendation approach: every user in the test set will get the same recommendations.
- **MF**: We use the basic matrix factorization approach to represent a state-of-the-art CF approach.

By assuming that the numbers of most popular (and therefore irrelevant) landmarks in each city are different and by letting them vary between 0 and 10, we can observe the performance of the proposed CRMF and its relative improvement over the baseline approaches, as shown in Table 1. We use "0" (in the first column) to denote the case where no assumption of the relevance of most popular landmarks is made, and we measure the performance according to the ground truth in the test set.

As can be seen from Table 1, PopRec outperforms MF and CRMF under the condition that no popular landmarks are assumed irrelevant. However, as stated before, we are mainly concerned about the recommendation performance for landmarks less frequently visited by the general population, i.e., those a traveler may not know about beforehand. If the most popular landmark in each city is ignored, we can see that CRMF outperforms PopRec by ca. 7% in MAP. When different numbers of the most popular landmarks are ignored, CRMF achieves improvement over PopRec up to ca. 32% in MAP. Finally, we can observe that CRMF consistently outperforms MF under all conditions, achieving improvement up to ca. 12% in MAP.

Table 1. MAP comparison of CRMF and baseline approaches under removal of a different number (first column) of the most popular landmarks in each city.

	PopRec	MF	CRMF	Gain over PopRec	Gain over MF
0	0.402	0.377	0.384	-4.6%	1.8%
1	0.224	0.238	0.240	7.2%	0.8%
2	0.153	0.168	0.172	12.5%	2.4%
3	0.121	0.138	0.140	15.6%	1.3%
4	0.100	0.123	0.126	26.3%	2.3%
5	0.085	0.102	0.106	24.9%	3.6%
6	0.073	0.083	0.087	20.2%	5.7%
7	0.064	0.075	0.082	28.3%	8.9%
8	0.058	0.066	0.072	25.1%	8.9%
9	0.052	0.060	0.065	26.4%	8.6%
10	0.048	0.056	0.063	31.6%	11.7%

We also observe that, in general, the greater the number of popular landmarks that are ignored, the more relative improvement is achieved by CRMF over MF. This pattern indicates that category-based landmark similarity is not only helpful for improving general recommendation performance, but also for recommending less-visited landmarks. Note that all the improvements are statistically significant according to Wilcoxon signed rank significance test with $p<0.05$ measured across all users in the test set.

Conclusions and Future Work

In this paper, we put forward a new research challenge on personalized landmark recommendation based on geotagged photos from users on photo sharing sites. We propose a novel category-regularized matrix factorization approach to recommend landmarks to individual users based on both user-landmark preference information and category-based landmark similarity. We observe that personalized landmark recommendation must go beyond recommending the most popular, widely visited landmarks, and instead focus on less frequently visited landmarks that are well-fit to users' individual tastes. The proposed CRMF approach is able to exploit category-based landmark similarity to provide improved recommendation according to this criterion. Improvement is shown with respect to a popularity-based baseline and a MF approach that exploits only user-landmark preference information.

Our future work on this topic will involve investigation on integrating additional information resources besides landmark category into personalized landmark recommendation, and analyzing the influence of users' travel experience.

References

Adomavicius, G., and Tuzhilin, A., 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE TKDE*, 17, 6, 734-749.

Arase, Y., Xie, X., Hara, T., and Nishio, S., 2010. Mining people's trips from large scale geo-tagged photos. In *MM '10*, 133-142.

Clements, M., Serdyukov, P., de Vries, A. P., and Reinders, M. J., 2010. Using Flickr geotags to predict user travel behaviour. In *SIGIR '10*, 851-852.

De Choudhury, M., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R., and Yu, C., 2010. Automatic construction of travel itineraries using social breadcrumbs. In *HT '10*, 35-44.

Gao, Y., Tang, J., Hong, R., Dai, Q., Chua, T.-S., and Jain, R., 2010. W2Go: a travel guidance system by automatic landmark ranking. In *MM '10*, 123-132.

Koren, Y., Bell, R., and Volinsky, C., 2009. Matrix factorization techniques for recommender systems. *IEEE Computer*, 42, 8, 30-37.

Kurashima, T., Iwata, T., Irie, G., and Fujimura, K., 2010. Travel route recommendation using geotags in photo sharing sites. In *CIKM '10*, 579-588.

Lu, X., Wang, C., Yang, J.-M., Pang, Y., and Zhang, L., 2010. Photo2Trip: generating travel routes from geo-tagged photos for trip planning. In *MM '10*, 143-152.

Ma, H., Lyu, M. R., and King, I., 2009. Learning to recommend with trust and distrust relationships. In *RecSys '09*, 189-196.

Serdyukov, P., Murdock, V., and van Zwol, R, 2009. Placing Flickr photos on a map. In *SIGIR '09*, 484-491.

Shi, Y., Larson, M., and Hanjalic, A., 2010. Mining mood-specific movie similarity with matrix factorization for context-aware recommendation. In *CAMRa '10*, 34-40.

Zheng, V. W., Zheng, Y., Xie, X., and Yang, Q., 2010. Collaborative location and activity recommendations with GPS history data. In *WWW '10*, 1029-1038.