Photo2Trip: Exploiting Visual Contents in Geo-tagged Photos for Personalized Tour Recommendation

Pengpeng Zhao
School of Computer Science and Technology, Soochow University
Suzhou, China
ppzhao@suda.edu.cn

Xiefeng Xu
School of Computer Science and Technology, Soochow University
Suzhou, China
20154227037@stu.suda.edu.cn

Yanchi Liu
Management Science and Information Systems
Rutgers University
yanchi.liu@rutgers.edu

Victor S. Sheng
Department of Computer Science
University of Central Arkansas
Conway, USA
ssheng@uca.edu

Kai Zheng*
School of Computer Science and Technology, Soochow University
Suzhou, China
zhengkai@suda.edu.cn

Hui Xiong
Management Science and Information Systems
Rutgers University
hxiong@rutgers.edu

ABSTRACT

Recently accumulated massive amounts of geo-tagged photos provide an excellent opportunity to understand human behaviors and can be used for personalized tour recommendation. However, no existing work has considered the visual content information in these photos for tour recommendation. We believe the visual features of photos provide valuable information on measuring user / Point-of-Interest (POI) similarities, which is challenging due to data sparsity. To this end, in this paper, we propose a visual feature enhanced tour recommender system, named 'Photo2Trip', to utilize the visual contents and collaborative filtering models for recommendation. Specifically, we first extract various visual features from photos taken by tourists. Then, we propose a Visual-enhanced Probabilistic Matrix Factorization model (VPMF), which integrates visual features into the collaborative filtering model, to learn user interests by leveraging the historical travel records. Moreover, user interests together with trip constraints are formalized to an optimization problem for trip planning. Finally, the experimental results on real-world data show that our proposed visual-enhanced personalized tour recommendation method outperforms other benchmark methods in terms of recommendation accuracy. The results also show that visual features are effective on alleviating the data sparsity and cold start problems on personalized tour recommendation.

CCS CONCEPTS

•**Information systems** → **Data mining;** *Recommender systems;*

KEYWORDS

Tour Recommendation, Collaborative Filtering, Visual Content

*Contact Author

MM'17, October 23–27, 2017, Mountain View, CA, USA.
© 2017 ACM. ISBN 978-1-4503-4906-2/17/10…$15.00
DOI: http://dx.doi.org/10.1145/3123266.3123336

1 INTRODUCTION

Recent years have witnessed a revolution in location-based social network services. As a consequence, large amounts of geo-tagged photos have been accumulated from users. These footprints (or check-ins) provide an excellent opportunity to understand human behaviors and can be used in many fields, including personalized tour recommendation. Tour recommendation aims to find a trip route visiting several POIs that maximize the utility of users according to their trip constraints and their specific interests on POIs. It can help tourists narrow down candidate POIs to visit, and plan an appropriate visit order and corresponding duration at each POI in an unfamiliar place.

Tour recommendation is complex because tourists have different interests and trip constraints, such as time limitation, the popularity of POIs, and travel time between POIs [11]. Therefore, how to learn user interests plays an important role in personalized tour recommendation. Brilhante et al. [4, 5] used visit frequency in a POI category as user visit preference. Lim et al. [18] used average visit duration of all users in a POI category as user interest and took personal visit duration into consideration in tour recommendation, which got better results than frequency-based approaches. However, if a user has not visited any POIs in a category yet, the above methods are not able to make personalized tour recommendation. A straightforward solution is leveraging collaborative filtering to predict user interest of each unvisited POI.

Nevertheless, the check-in data is extremely sparse since most users are not residents in their tour destinations. And the sparsity issue causes difficulties for collaborative filtering methods to learn effectively. Besides, the cold start problem (no historical check-in records for new users or new POIs) is even more severe in personalized tour recommendation. Therefore, additional information needs to be incorporated to address these issues. We find that the visual features in geo-tagged photos taken by users can provide important context information for predicting user visit interests. From these photos, the POI information can be inferred, also users' behaviors and preferences can be revealed. For example, Figure 1 shows three pairs of POI photos from two different users. A tourist who favors the POIs in the first column might also be interested in the second one since they exhibit similar visual appearances. These observations motivate us to leverage the visual information,

Figure 1: Three pairs of POI photos from six different POIs and visited by two users having similar visual appearances.

which is overlooked by existing methods, in addition to others for personalized tour recommendation.

To this end, in this paper, we leverage the visual contents of geo-tagged photos together with collaborative filtering models for personalized tour recommendation. Specifically, we first extract various visual features from photos taken by tourists, and utilize them to understand the styles of the POIs and the visual preferences of users. Then, we propose a Visual-enhanced Probabilistic Matrix Factorization model (VPMF), which integrates visual features into the collaborative filtering model, to learn user interests by leveraging the historical travel records of peer users. After that, user interests together with trip constraints are formalized to an optimization problem for trip planning. Our experimental results on real-world Yahoo! Flickr datasets show that VPMF significantly improves the performance of visit interest prediction for tour recommendation. On average, it improves by over 5% on trip planning with respect to F_1 and over 25% on visit duration prediction with respect to Root Mean Square Error (*RMSE*), comparing with the state-of-the-art methods. To summarize, our contributions are listed as follows.

- To the best of our knowledge, this is the first work that utilizes visual features of geo-tagged photos to learn user interests for personalized tour recommendation.
- A VPMF model is proposed to integrate visual features into the collaborative filtering model for enhanced performance. The model uses the content of user-generated photos to improve the prediction accuracy. Moreover, it reduces the negative impacts of the data sparsity problem and the cold start problem.
- The proposed method has been evaluated on large-scale real-world data for tour recommendation. The experimental results show that our method outperforms state-of-the-art methods in terms of different metrics, such as precision, recall, F_1, and RMSE.

The rest of this paper is organized as follows. In the next section, we discuss the related work. Section 3 provides the problem definition and preliminaries. We introduce our system framework and propose visual feature enhanced recommendation algorithm VPMF in Section 4. Experimental results and discussions are presented in Section 5, followed by the conclusion in Section 6.

2 RELATED WORK

This paper makes a forward step for tour recommendation, which is rooted in POI recommendation. POI recommendation is to recommend a list of top k most relevant POIs to a user, based on user implicit feedback, such as check-in frequency. Collaborative filtering is widely used in POI recommendation. The state-of-the-art collaborative filtering (CF) is based on matrix factorization and its variants [15, 16, 23, 27]. Salakhutdinov & Mnih [23] proposed a PMF model in a Bayesian probabilistic framework to include Gaussian noise in observations. Under the Gaussian assumption, maximizing the posterior probability over latent features is equivalent to minimizing the square error.

Recently, more advanced models have been proposed to exploit additional information for POI recommendation [1, 13], such as check-in locations, social influence, temporal information and transition between POIs. Ye et al. [28, 29] considered the social influence under the framework of a user-based CF model and modeled the geographical influence by a Bayesian CF model. Moreover, both Yuan et al. [30] and Gao et al. [10] introduced temporal preference to enhance the efficiency and effectiveness of Ye et al.'s solution. Cheng et al. [7] considered more comprehensive information, such as the multi-center of user check-in patterns, and skewed user check-in frequency. Moreover, Liu et al. [19] proposed a bi-weighted low-rank graph construction model, which integrates users' interests and their evolving sequential preferences with temporal interval assessment to provide POI recommendations for a specific time period. However, most of these methods did not explicitly consider the visual content in user-generated photos in POI recommendation. Besides, they evaluated each venue independently without considering other information and ignored the order of visits. Moreover, there are no overall time constraints, and traveling time is not considered. In this paper, we focus on tour recommendation which recommends relevant POIs as well as order the POIs into a trip to satisfy different constraints, e.g., the maximum travel time budget.

Tour recommendation has become very important in recent years. A large number of public available traveler e-footprints (such as geo-tagged photos and blogs) make automatic trip planning possible. Arase et al. [2] proposed a photo trip pattern mining framework to detect users' frequent trip patterns extracted from public geo-tagged photos, i.e., typical sequences of visited cities and visit duration as well as trip themes that characterize the trip patterns. Lu et al. [22] leveraged existing travel clues from geo-tagged photos to suggest customized route plans according to users' preferences. They used geo-tagged photos to discover the tour paths within a destination and travel routes between destinations. Cheng et al. [6] further proposed a probability-based personalize travel recommendation model based on user's profiles (such as gender, age, and race) by leveraging users' attributes in user-generated photos. Although they also utilized visual features in geo-tagged photos, they only used facial visual content to infer user's profiles and did not take advantage of general visual features. A recent work named PersTour [18] is closely related to our work and reflects the levels of user's interest based on visit durations, which are obtained from real-life travel sequences based on geo-tagged photos. PersTour uses average POI visit duration as user interest and has not employed collaborative filtering to predict user interest. However, the major difference between our work and related research described above is that we extract visual features from user-generated photos and consider these visual features with the mobility pattern of tourists in our personalized tour recommendation framework.

3 PRELIMINARIES

In this section, we first introduce some basic concepts of tour recommendation, followed by the correlation analysis between the visual features of POIs and users' ratings. At last, we introduce a basic collaborative filtering model and three visual features will be used in our tour recommendation system.

3.1 Basic Concepts

Popularity. The popularity of a POI p (denoted as $Pop(p)$) is defined as the number of times that the POI p has been visited.

Time-based User Interest. We define the interest of a user u in a POI p (denoted as $Int(p)$) as the ratio between the personal visit duration and the average visit duration of all users.

Personalized POI Visit Duration. With the definition of time-based user interest, we can define the personalized visit duration of POI p (denoted as $T^{Visit}(p)$) as the multiplication of user interest and the average time spend at POI p.

Travel Time. Travel time is the time cost moving from POI p_i to POI p_j (denoted as $T^{Travel}(p_i, p_j)$), which is based on the distance between two POIs p_i and p_j and the given moving speed.

Tour Recommendation. Given m POIs in a city, each POI p_i has a category label $Cat(p_i)$ and a latitude/longitude location, a user u with the starting POI p_1 and the ending POI p_n, and a time budget B, we want to find an optimal trip route $I = (p_1, ..., p_n)$ that maximizes user utility under the following constraints: (1) it starts at location p_1 and ends at location p_n; (2) it completes within the time budget. The utility of visiting POI p_i is represented by the popularity and the user interest of this POI, which are denoted as $Pop(p_i)$ and $Int(p_i)$ respectively. The cost of traveling from p_i to p_j (denoted as $Cost(p_i, p_j)$) is calculated as the summation of the travelling time $T^{Travel}(p_i, p_j)$ and the personalized visit duration of POI $T^{Visit}(p_j)$. That is, $Cost(p_i, p_j) = T^{Travel}(p_i, p_j) + T^{Visit}(p_j)$.

3.2 Correlation Analysis

Before designing a tour recommendation model, it is important to understand tourist visit behaviors. In other words, we try to answer the question: "do tourist visit behaviors correlate with the visual style and appearance of POIs?" To answer this question, we analyzed the correlation between visual contents in photos of POIs and time-based user interest. First, we predict user personalized POI visit duration using the average visit duration of all users of the category of the POI. Then given $top\text{-}k$ most similar POIs on visual appearance as neighborhoods of a POI, we predict user visit duration at the POI using the average visit duration of its neighborhoods. Finally, given $top\text{-}k$ most similar users on the visual content of photos posted by users as neighborhoods of a user, we predict user visit duration at the POI using the average visit duration of the POI taken by his/her neighborhoods. Our analysis results are shown in Figure 2, in which *RMSE* metric is used to measure the prediction error and the smaller value the better. From Figure 2, we can see that the prediction errors have reduced about 6.3% and 2.4% under given visually similar neighbors of a POI and under given neighbors of a user with similar visual taste, respectively. From the results, we can see the answer to the above question is "yes".

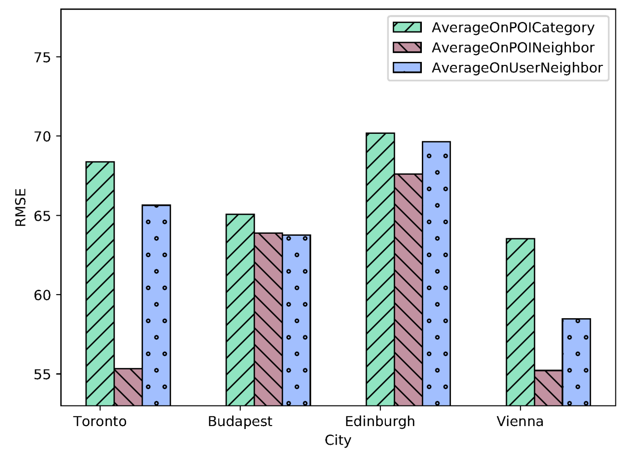

Figure 2: The effects of user interest prediction under given visual neighborhoods of POIs and users respectively.

3.3 PMF Model

Probabilistic Matrix Factorization (PMF) [23] is a simple, accurate, and efficient model among collaborative filtering methods. PMF not only can deal with very large datasets, but also has the ability to make recommendations for users only make few ratings in recommender systems [12, 26]. We will later show how to improve PMF with visual features in Section 4.

Matrix factorization methods construct a latent low-rank dimensional space to represent each user and POI. From the linear combination of the latent features, the missing relationships of users and POIs can be estimated. Using the initial matrix $R \in \mathbb{R}^{N \times M}$ as training data, where R_{ij} is the time-based user interest of user u_i in POI p_j, matrices $U \in \mathbb{R}^{D \times N}$ and $V \in \mathbb{R}^{D \times M}$ can be learned using Matrix Factorization techniques, so that they can approximate matrix R with matrix $\hat{R}$, such that $R \approx \hat{R} = U^T V$. PMF expresses the process of learning in the Bayes probabilistic framework, where the user-POI relation in R is an observation, and U and V describe the system inner characters and need to be evaluated. The observation R is modeled as a draw from a Gaussian distribution, where the mean of R_{ij} is $U_i^T V_j$. And U and V are drawn from the zero-mean normal distribution.

$$p(R|U,V) = \mathcal{N}(\hat{R}, \sigma^2) = \mathcal{N}(U^T V, \sigma^2) \tag{1}$$

$$p(U) = \mathcal{N}(0, \sigma_U^2), p(V) = \mathcal{N}(0, \sigma_V^2) \tag{2}$$

The likelihood of observing a specific user-POI relation in R can be expressed as follows.

$$p(R|U,V,\sigma^2) = \prod_{i=1}^{N}\prod_{j=1}^{M}[\mathcal{N}(R_{ij}|U_i^T V_j,\sigma^2)]^{I_{ij}} \quad (3)$$

where $\mathcal{N}(x|\mu,\sigma^2)$ denotes the normal distribution with mean μ and variance σ. I_{ij} is an indicator function, in which $I_{ij} = 1$ if R_{ij} is known and 0 otherwise. Now, through a Bayesian inference, we can obtain the posterior probability of U and V as follows.

$$p(U,V|R,\sigma^2,\sigma_U^2,\sigma_V^2) \propto p(R|U,V,\sigma^2)p(U|\sigma_U^2)p(V|\sigma_V^2) \quad (4)$$

To calculate U and V, so as to maximize the posterior probability given observation R, we can learn the latent feature U and V of users and POIs purely based on the observation R using Equation 4.

3.4 Visual Features in Geo-tagged Photos

There are lots of different types of visual features in geo-tagged photos. In order to improve recommendation accuracy, we should choose visual features in a proper way. We assume that tourists are attracted by the visual effects of POIs, such as colors, abstract features, and visual contents, as shown in Figure 1. More specifically, given two POIs p_i and p_j, we could calculate the similarity $s(p_i, p_j)$ between the two POIs by measuring their visual correlation through extracted visual features. Next we introduce some widely used visual features.

Color Histogram. In POI photos, color is the first impression to people. For example, POI photos with large color areas, such as blue sky, golden beaches, and blue sea water, have a deep impression on users. Color histogram is a widely used visual feature. We adopt a standard color histogram feature and extract a 512-dimensional color feature vector for each photo. And a joint histogram in RGB color space has 8 bins in each channel.

Scale-Invariant Feature Transform (SIFT). For point description, the SIFT descriptor [21] is known as scale-invariant features and widely used in object recognized and content-based image search for its good classification accuracy [8]. The SIFT finds interest points and captures the local shape around it using edge orientation histograms. SIFT features are also robust to changes in lighting, noise, and minor differences in viewpoint. Because many of the photos are taken from the same scene but different angles, SIFT will be useful in this scenario. We extract a 128-dimensional SIFT feature after resizing each POI photo to 256×256 pixels.

Convolutional Neural Networks. Different from above handcrafted visual features, convolutional neural network (CNN) can automatically discover high-level visual features of photos by learning from training data. It has been shown that CNN performs well in image classification and object detection. The features extracted by CNN can reflect a photo globally, regionally, and locally. Intuitively, these features (or some of them) should be useful for visual recommendation as we will show in our later experiments. In this paper we use the VGG16 model [24], which is the state-of-the-art architecture, to extract features from user-generated geo-tagged photos. Specifically, we resize each photo to 224*224 pixels as the input of VGG16 and obtain a 4096 dimension visual feature vector as the output of the second fully-connected layer.

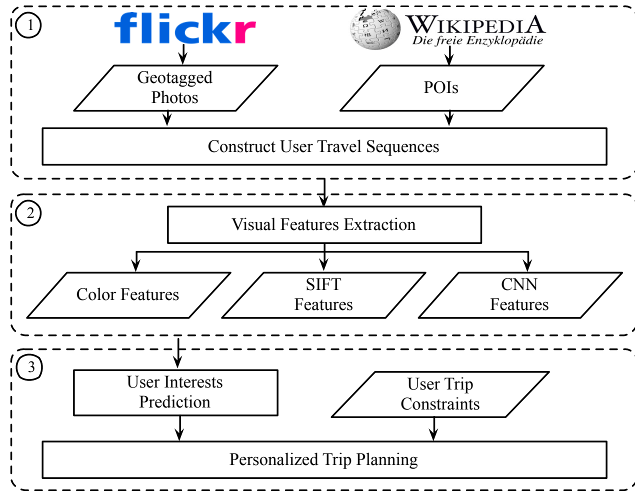

Figure 3: Framework of Photo2Trip Recommender System

4 PHOTO2TRIP TOUR RECOMMENDATION

4.1 System Framework Overview

Figure 3 shows the framework of our Photo2Trip personalized tour recommendation system, which is composed of three main parts. First, we crawl the photos from the public photo-sharing web site (i.e., Flickr). With the same approach described in [17], we obtain a list of POIs from Wikipedia and map these photos to user-POI visits. And we construct user travel sequences based on them. Second, after mining the travel patterns of users' trip sequences, we extract the visual features in the user-generated photos using the visual toolbox, and then propose a visual-enhanced probability matrix factorization (VPMF) model to predict user visit interests. Third, with user's input trip constraints, including travel time limitation, the starting POI, and the ending POI, the trip planning module generates a personalized trip route that maximizes user utility while adhering to the user's trip constraints. Trip planning is further modeled as an orienteering problem and solved using linear programming.

In the following subsections, we will introduce two essential modules (i.e., user interest prediction and personalized trip planning) in our framework in detail.

4.2 VPMF Model

As observed in Section 3.2, user visit behaviors are related to the visual appearance of POIs, and the visual contents in user-posted photos reflect the user visit preferences. According to the idea of neighbor-based collaborative filtering, it is natural to assume that the visit behavior and the visual taste of a user are similar to that of his/her neighbors, and the interests of a POI are similar to those of its similar visual POIs. Based on the above analysis, we propose a visual-enhanced PMF model to improve user interest prediction accuracy. We first select *top−k* nearest neighbors for each POI and for each user respectively based on the visual content similarity of the photos of POIs and the photos taken by users. And then we incorporate the constructed visual neighborhoods into the learning process of PMF.

The similarity of two POIs is measured by the cosine similarity of the visual feature vectors. For the reason of each POI has more than one photos, to get a representative visual features vector of a POI, we merge each dimension visual vectors extracted from POI photos using the maximum pooling method. After that, we linearly combine three similarities of different visual features to get the eventual similarity $s(p_i, p_j)$ of two POIs. The similarity $s(u_i, u_j)$ of two users is also calculated by the cosine similarity of the visual vectors of photos posted by the user in the same way.

Inspired by neighborhood MF [15], in the probability matrix factorization process, the latent features of users u_i and POIs p_j should be close to their neighborhoods N_{u_i} and N_{p_j} respectively. Based on this intuition, we add Gaussian priors to user's and POI's latent feature vectors to ensure that U_i and V_j are centered around the mean of their neighborhood and formulate the following equations.

$$U_i = \sum_{l \in N_{u_i}} s(i,l) \times U_l + \tilde{U}_i, \quad \tilde{U}_i \sim \mathcal{N}(0, \sigma_U^2 \mathbf{I}) \tag{5}$$

$$V_j = \sum_{l \in N_{p_j}} s(j,l) \times V_l + \tilde{V}_j, \quad \tilde{V}_j \sim \mathcal{N}(0, \sigma_V^2 \mathbf{I}) \tag{6}$$

In the above two equations, the latent feature vector of each user and each POI comprise of two terms. The first term characterizes the neighborhood related feature of the user or the POI. For notation convenience, we normalize the similarities to ensure $\sum_{l \in N_{u_i}} s(i,l) = 1$ and $\sum_{l \in N_{p_j}} s(j,l) = 1$. The second term emphasizes the unique feature of each user and each POI, which could diverge from their neighborhood. The variance parameter σ_U^2 and σ_V^2 are used to control the divergence. The lower the variance, the less diverges the feature vector from that of the neighbors. With the visual neighborhood incorporated, the conditional distributions of the observed R, as shown in Equation 3, does not change. Based on the Bayesian formula, the posterior distribution over the latent factors of users and POIs is given as follows (Equation 7).

$$\begin{aligned} p(U, V|R, \sigma^2, \sigma_U^2, \sigma_V^2) &\propto p(R|U, V, \sigma^2) \times p(U|S, \sigma_U^2) \times p(V|S, \sigma_V^2) \\ &= \prod_{i=1}^{N} \prod_{j=1}^{M} [\mathcal{N}(R_{ij}|U_i^T V_j, \sigma^2)]^{I_{ij}} \\ &\times \prod_{i=1}^{N} \mathcal{N}(U_i| \sum_{l \in N_{u_i}} s(i,l) \times U_l, \sigma_U^2 \mathbf{I}) \\ &\times \prod_{j=1}^{M} \mathcal{N}(V_j| \sum_{l \in N_{p_j}} s(j,l) \times V_l, \sigma_V^2 \mathbf{I}) \end{aligned} \tag{7}$$

Given the hyperparameters σ^2,σ_U^2 and σ_V^2, maximizing the log posterior to find U, V in Equation 7 is equivalent to minimizing the following objective function.

$$\begin{aligned} \mathcal{L}(U, V, R) =& \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{N} I_{ij} (R_{ij} - U_i^T V_j)^2 \\ &+ \frac{1}{2} \lambda_U \sum_{i=1}^{M} ||(U_i - \sum_{l \in N_{u_i}} s(i,l) \times U_l)||_{Fro}^2 \\ &+ \frac{1}{2} \lambda_V \sum_{j=1}^{N} ||(V_j - \sum_{l \in N_{p_j}} s(j,l) \times V_l)||_{Fro}^2 \end{aligned} \tag{8}$$

where $\lambda_U = \sigma^2/\sigma_U^2$, $\lambda_V = \sigma^2/\sigma_V^2$ and $||.||_{Fro}^2$ denotes the Frobenius norm. The objective function given by equation 8 is smoothed by the parameter λ_U and λ_V, which control the neighborhood influence of users and POIs based on the error objective function. The lower the values of λ_U and λ_V, the less the visual neighborhood information relies on.

A local minimum of the objective function (Equation 8) can be found by performing stochastic gradient descent(SGD) on U_i and V_j. The update formula is given as follows:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial U_i} =& \sum_{j=1}^{N} (R_{ij} - U_i^T V_j)(-V_j) + \lambda_U (U_i - \sum_{l \in N_{u_i}} s(i,l) \times U_l) \\ &- \lambda_U \sum_{i \in N_{u_l}} s(i,l)(U_l - \sum_{j \in N_{u_l}} s(j,l) * U_j) \end{aligned} \tag{9}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial V_j} =& \sum_{i=1}^{M} (R_{ij} - U_i^T V_j)(-U_i) + \lambda_V (V_j - \sum_{l \in N_{p_j}} s(j,l) \times V_l) \\ &- \lambda_V \sum_{j \in N_{p_l}} s(j,l)(V_l - \sum_{i \in N_{p_l}} s(i,l) * V_i) \end{aligned} \tag{10}$$

4.3 Trip Planning

Trip Planning can be modeled using a bi-criteria generalization of travelling salesman problem (TSP) with two conflicting objectives: maximizing the collected utility and minimizing the travel cost. The orienteering problem (OP) is a variant of TSP that seeks for a trip that maximizes the total collected utility while maintaining the travel cost under a given value. That is, the travel cost objective is turned to a constraint. OP can be formulated as an integer programming problem as follows [11, 20]. Let n be the number of POIs, where the starting POI is denoted as p_1 and the destination POI is denoted as p_n. The utility of visiting POI p_i is represented by the popularity $Pop(p_i)$ and the user interest $Int(p_i)$ of this POI. The cost of traveling from p_i to p_j is calculated as the summation of the travelling time and the personalized visit duration of POI p_j. One main difference between our work and prior works is that we personalize the visit duration at each POI predicted by VPMF, instead of using the average visit duration for all users. With the time budget B, we want to find an itinerary $I = (p_1, ..., p_n)$ that satisfies the following constraints.

$$Max \sum_{i=2}^{N-1} \sum_{j=2}^{N} x_{i,j} (\eta Int(Cat_i) + (1-\eta) Pop(i)) \tag{11}$$

$$\sum_{j=2}^{N} x_{1,j} = \sum_{i=1}^{N-1} x_{i,n} = 1 \tag{12}$$

$$\sum_{i=1}^{N-1} x_{i,k} = \sum_{j=2}^{N} x_{k,j} \leq 1, for\ all\ k = 2, ..., N-1 \tag{13}$$

$$\sum_{i=1}^{N-1} \sum_{j=2}^{N} Cost(i,j) x_{i,j} \leq B \tag{14}$$

$$2 \leq p_i \leq N, for\ all\ i = 2, ..., N \tag{15}$$

$$p_i - p_j + 1 \leq (N-1)(1 - x_{i,j}), for\ all\ i,j = 2, ..., N \tag{16}$$

The objective function (i.e., Equation 11) is to maximize the total popularity and the interest score of visited POIs in the trip, where η is the weight given to balance the popularity and the interest. For a path from p_1 to p_n, if POI p_i is followed by POI p_j, we set the variable $x_{i,j} = 1$. Otherwise, we set $x_{i,j} = 0$. Constraint 12 ensures that the trip starting at POI p_1 and ending at POI p_n. Constraint 13 ensures that the trip is connected and each POI is visited at most once. Constraint 14 ensures that the trip meets the time budget B, based on the function $Cost(p_i, p_j)$ that considers both the traveling time and the personalized POI visit duration. Constraints 15 and 16 ensure that there are no sub-tours in the proposed trip, adapted from the sub-tour elimination used in the travelling salesman problem [9]. The orienteering problem is NP-hard. Hence, exact solutions for the orienteering problem are not feasible for a large number of POIs. The orienteering problem can be formulated as an integer programming problem. For solving this integer programming problem, we use the lpsolve linear programming package [3] to obtain optimal solutions.

5 EXPERIMENTS

5.1 Dataset

We apply the proposed photo2trip method on the Yahoo! Flickr Creative Commons 100M (YFCC100M) dataset [25], the largest public multimedia collection released, which consists of 100 million photos and 0.8 million videos posted on Flickr with relevant meta information, such as the date/time taken, geo-location coordinates and geo-graphic accuracy. The geo-graphic accuracy ranges from the world level to the street level.

From this dataset, we use geo-tagged photos that were taken in four cities, namely Toronto, Budapest, Edinburgh, and Vienna. More details regarding this dataset are shown in Table 2. The dataset was previously used for tour recommendation by Lim et al. [18]. As described in [17], we first obtain a list of POIs from Wikipedia and then map these photos to user-POI visits. After that, we construct user travel sequences and evaluate our proposed approach.

5.2 Comparison Methods

In our experiments, we compare our proposed approaches with three popular baseline approaches and a recently proposed approach PersTour [18]. A brief introduction of each of them is shown as follows.

- **Random Selection (Rand).** Iteratively and randomly choose a POI p_j from unvisited POIs as next POI.
- **Greedy Nearest (GNear).** Iteratively and greedily choose the nearest POI p_j with the least value $T^{Travel}(p_i, p_j)$ from unvisited POIs as next POI .
- **Greedy Most Popular (GPop).** Iteratively and greedily choose the most popular POI p_j with the most value $Pop(p_j)$ from unvisited POIs as next POI.
- **PersTour and η = 0.5 (PT-.5).** PersTour [18] with balanced emphasis on both POI popularity and user interest. That means the objective function is to maximize the total popularity and the interest score of POIs in the trip.
- **PersTour and η = 1 (PT-1).** Perstour [18] with full emphasis user interest. In other words, the objective function is to maximize the total interest score of POIs in the trip.

As described in Section 3, instead of using the average POI visit duration as user interest in PersTour [18], we chose the PMF [23] model to predict user visit interests in terms of different granularity. We first use the PMF model to predict the user visit interests on the category of a POI, and then we predict user visit interests on a specific POI. Our approaches are listed as follows.

- **Photo2Trip using PMF on POI Category level and η = 0.5. (PT-PMFC-.5)** Based on PersTour [18] with balanced emphasis on both POI popularity and user interest, we add the PMF model to predict user interests on the category of a POI. That means the prediction interests in one unvisited category is the same.
- **Photo2Trip using PMF on POI Category level and η = 1. (PT-PMFC-1)** Based on PersTour [18] with full emphasis on user interest, we add the PMF model to predict user interests on the category of a POI.
- **Photo2Trip using PMF on POI level and η = 0.5. (PT-PMF-.5)** Based on PersTour [18] with balanced emphasis on both POI popularity and user interest, we add the PMF model to predict user interest on a specific POI, more detail than the category level.
- **Photo2Trip using PMF on POI level and η = 1. (PT-PMF-1)** Based on PersTour [18] with full emphasis on user interest, we add the PMF model to predict user interest on a specific POI.

Again, the user-generated geo-tagged photos provide important contexts for predicting user visit interest for personalized tour recommendation. To integrate these photos into personalized tour recommendation, we extract three different visual features from these photos and incorporate the visual features into the PMF model (i.e., VPMF). Since we noticed the advantage of predicting user visit interest on a specific POI, we use VPMF to predict user visit interest on a specific POI, instead of predicting user visit interest on the category of POIs. Therefore, we have following two approaches based on VPMF.

- **Photo2Trip using VPMF on POI level and η = 0.5. (PT-VPMF-.5)** This is the model described in Section 4, personalized tour recommendation using the VPMF model to predict user interest on a specific POI. In this case, the objective function is to maximize the total popularity and the interest score of POIs in the trip.
- **Photo2Trip using VPMF on POI level and η = 1. (PT-VPMF-1)** The objective function, in this case, is to maximize the total interest score of POIs in the personalized tour recommendation using the VPMF model to predict user interest on each POI.

5.3 Evaluation Metrics

We evaluate the popular baseline approaches, PersTour [18], and our proposed photo2trip approaches based on PMF and VPMF using leave-one-out cross-validation [14]. When evaluating a specific travel sequence of a user, we use the user's other travel sequences as training data. We evaluate the performance of each algorithm using the following metrics.

- **Tour Precision.** The precision of POIs recommended in the trip is the proportion of POIs recommended in a trip

Table 1: Performance comparison of tour recommendation in terms of Precision, Recall and F1-score

Algo.	Toronto			Budapest			Edinburgh			Vienna		
	Pre.	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*
GNear	.464±.010	.544±.008	.484±.012	.359±.021	.477±.008	.393±.011	.386±.005	.501±.021	.422±.012	.385±.024	.530±.026	.426±.011
GPop	.611±.015	.389±.037	.466±.016	.544±.037	.350±.035	.413±.037	.592±.015	.459±.008	.503±.009	.543±.005	.364±.021	.423±.023
Rand	.451±.002	.274±.028	.336±.019	.401±.035	.237±.038	.289±.024	.450±.031	.271±.014	.325±.016	.487±.006	.285±.018	.351±.008
PT-1	.720±.015	.755±.021	.728±.018	.772±.021	.777±.018	.768±.031	.604±.020	.662±.011	.616±.029	.618±.002	.660±.013	.625±.014
PT-.5	.704±.014	.774±.025	.729±.011	.781±.009	.788±.010	.777±.008	.631±.014	.742±.019	.671±.014	.646±.006	.715±.009	.666±.006
PT-PMFC-1	**.724±.021**	.755±.024	**.731±.019**	**.807±.012**	**.792±.018**	**.792±.020**	**.605±.017**	**.663±.018**	**.618±.021**	**.631±.019**	**.664±.020**	**.635±.021**
PT-PMFC-.5	**.718±.011**	**.779±.015**	**.739±.020**	**.816±.020**	**.801±.035**	**.801±.031**	**.640±.007**	**.750±.009**	**.680±.010**	.645±.012	.715±.020	**.667±.025**
PT-PMF-1	**.746±.011**	**.769±.012**	**.751±.009**	**.813±.021**	**.797±.026**	**.795±.031**	**.620±.025**	**.674±.016**	**.631±.035**	**.654±.008**	**.676±.012**	**.651±.015**
PT-PMF-.5	**.725±.012**	**.791±.015**	**.749±.021**	**.821±.025**	**.806±.021**	**.803±.030**	**.643±.009**	**.756±.012**	**.685±.011**	**.655±.017**	**.725±.015**	**.676±.020**
PT-VPMF-1	**.749±.021**	**.805±.011**	**.765±.019**	**.812±.002**	**.808±.012**	**.809±.007**	**.621±.015**	**.678±.013**	**.634±.021**	**.660±.008**	**.685±.028**	**.676±.019**
PT-VPMF-.5	**.728±.022**	**.828±.023**	**.762±.029**	**.831±.011**	**.809±.012**	**.819±.023**	**.645±.025**	**.768±.016**	**.696±.035**	**.672±.011**	**.751±.009**	**.709±.023**
PT-VPMF-1 over PT-1	4.03%	6.62%	5.08%	7.72%	3.98%	5.34%	2.81%	2.41%	2.92%	6.78%	3.78%	8.16%
PT-VPMF-.5 over PT-.5	3.41%	6.98%	4.53%	6.40%	2.66%	5.41%	2.22%	3.50%	3.73%	4.02%	5.03%	6.45%

Table 2: Dataset description

City	# Photos	# Users	# POI Visits	# Travel Sequences
Toronto	157,505	1,395	39,419	6,057
Budapest	145,364	954	18,513	2,361
Edinburgh	82,060	1,454	33,944	5,028
Vienna	461,905	1,155	34,515	3,193

that was also in a user's real-life travel sequence, defined as $\frac{|P_r \cap P_v|}{|P_r|}$, where P_r and P_v are the set of POIs recommended in the tour and visited by the user in real-life, respectively.

- **Tour Recall.** The recall of POI recommendation in the trip is the proportion of POIs in a user's real-life travel sequence that was recommended, defined as $\frac{|P_r \cap P_v|}{|P_v|}$, where P_r and P_v are the set of POIs recommended in the trip and visited by the user in his/her real-life travel sequence, respectively.
- **Trip F_1-score.** It combines both precision and recall of a recommended trip together with the harmonic mean.
- **Root-Mean-Square Error (RMSE) of POI Visit Duration.** RMSE is a frequently used to measure the difference between a value predicted by a model and the value actually observed. Let p be a POI in recommended itinerary I, which was visited in real-life. Let D_r be the recommended duration and D_v be the duration in real-life respectively. Then, RMSE is defined as follows.

$$RMSE = \sqrt{\frac{\sum_{p \in I}(D_r - D_v)^2}{|I|}}$$

5.4 Results and Discussion

5.4.1 Effectiveness of PMF on the Category Level of POIs. We first evaluate the performance of incorporating PMF into trip planning to predict user visit interests on the category level of POIs. As shown in Table 1, both PT-PMFC-1 and PT-PMFC-.5 in most cases outperform the state-of-the-art PersTour, in terms of precision, recall and F_1. As expected, PT-PMFC consistently outperforms the greedy and random methods. This observation shows the effectiveness of integrating collaborative filtering into predicting visit interests in trip planning, which indicates that using the PMF model to predict user interests on the category level is more accurate than using the average visit time of all user in a category as user interest.

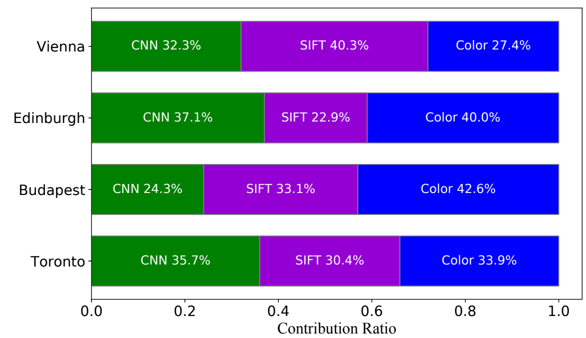

Figure 4: Performance improvement contribution of different visual features.

5.4.2 Effectiveness of PMF on the POI Level. We then evaluate the performance of incorporating PMF into trip planning to predict user visit interests on the POI level, a lower granular level than the category level. As shown in Table 1, both PT-PMF-1 and PT-PMF-.5 consistently outperform both PT-PMFC-1 and PT-PMFC-.5, in terms of precision, recall and F_1. The results indicate that predicting user visit interests on the POI level is more accurate and more effectiveness in trip planning, comparing with predicting on the category level of POIs.

5.4.3 Effectiveness of VPMF on the POI Level. We further evaluate the performance of integrating VPMF into trip planning to predict user visit interests on the POI level by leverage visual content in geo-tagged photos. As shown in Table 1, both PT-VPMF-1 and PT-VPMF-.5 consistently outperform both PT-PMF-1 and PT-PMF-.5, in terms of precision, recall and F_1. The results indicate that predicting user interests by integrating visual content inside the PMF model is more accurate than the approaches based on the PMF model, and show significant effectiveness in trip planning. Overall, PT-VPMF-1 outperforms the existing popular approach PT-1 5.38%, and PT-VPMF-.5 outperforms the existing popular approach PT-.5 5.03% with respect to the average F_1 value on the four cities.

5.4.4 Effectiveness of Different Visual Features. As shown in Table 1, integrating visual features extracted from user-generated photos into the PMF model improves the performance of predicting user interests. The contributions of different visual features in our

Table 3: Performance comparison of visiting duration prediction in terms of RMSE

Algo.	*RMSE*			
	Toronto	**Budapest**	**Edinburgh**	**Vienna**
PT-1	145.20 ± 9.25	65.35 ± 6.31	73.39 ± 9.53	62.99 ± 5.28
PT-.5	143.55 ± 9.88	57.27 ± 5.12	91.48 ± 5.07	68.93 ± 5.69
PT-PMFC-1	127.29 ± 7.14	52.53 ± 5.01	70.17 ± 4.52	59.90 ± 6.04
PT-PMFC-.5	121.87 ± 8.59	50.52 ± 8.25	84.23 ± 9.35	61.26 ± 6.28
PT-PMF-1	110.90 ± 9.99	44.19 ± 9.18	66.68 ± 5.35	52.47 ± 5.87
PT-PMF-.5	104.67 ± 6.78	47.37 ± 9.21	**73.72 ± 8.53**	51.31 ± 6.21
PT-VPMF-1	**109.76 ± 6.51**	**32.71 ± 5.35**	**65.72 ± 8.05**	**48.61 ± 7.25**
PT-VPMF-.5	**101.85 ± 7.68**	**41.87 ± 8.38**	82.12 ± 9.94	**44.88 ± 6.01**
PT-VPMF-1 over PT-1	24.41%	49.95%	10.45%	22.83%
PT-VPMF-.5 over PT-.5	29.04%	26.89%	17.76%	34.89%

Table 4: Performance comparison with cold start scenario in terms of Precision, Recall and F1-score

Algo.	**Toronto**			**Budapest**			**Edinburgh**			**Vienna**		
	Pre.	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*	*Pre.*	*Rec.*	*F1-score*
PT-1	.678±.004	.682±.011	.672±.002	.572±.020	.582±.008	.567±.017	.522±.010	.584±.006	.539±.019	.602±.012	.566±.011	.567±.012
PT-.5	.635±.004	.741±.015	.676±.011	.488±.004	.681±.011	.548±.018	.522±.011	.703±.032	.588±.014	.486±.003	.630±.009	.533±.023
PT-VPMF-1	**.703±.012**	**.703±.008**	**.695±.009**	**.611±.008**	**.607±.011**	**.596±.012**	**.580±.004**	**.607±.018**	**.593±.012**	**.653±.009**	**.584±.014**	**.592±.011**
PT-VPMF-.5	**.691±.023**	**.808±.015**	**.736±.027**	**.528±.014**	.677±.010	**.573±.023**	**.583±.018**	**.722±.002**	**.630±.005**	.484±.017	**.662±.024**	.524±.012

proposed VPMF model are shown in Figure 4. We can observe that color histogram feature is better than other features in two cases, and SIFT and CNN are the best in one case respectively. This indicates that the integrating three types of features together are useful to express user's visit behavior and visual taste. Overall, we obtained the best performance through integrating three types of visual features.

5.4.5 Visiting Duration Prediction Accuracy. With the availability of user interest predictions, we can personalize the POI visit duration more accuracy for each user. Apart from the accuracy of POIs recommended in a trip, recommending the appropriate amount of time to spend at a specific POI is another important consideration in tour recommendation. Visit duration at each POI is important in trip planning. In general, users intend to spend less time on uninteresting POIs to save time budget for interesting POIs. This matches user's behaviors that users usually prefer visiting a few POIs with high interest using all time budget to visiting many POIs with less interest. As shown in Table 3, the recommended personalized POI visit duration of PT-VPMF outperforms state-of-the-art personalized methods PT over 10% in all case and over 25% on average in terms of RMSE. This shows that personalized user visit duration prediction at a POI using VPMF more accurately reflects the real-life POI visit duration of users.

5.4.6 Cold Start Scenario. A cold start user means a user without any travel history data. To investigate the performance of VPMF for cold start users, we adapted the concept of leave-one-out cross-validation [14] in our experiments. That is, we leave one user out for testing. Specifically, we removed all historical travel data of this user and only kept his/her photos with all geo-tags removed. As we lack the check-in history of this user, this user is considered as a cold start user. Only visual content in photos can help reveal user interest. Therefore, the model must have the ability to address the inherent cold start nature and to recommend trip plan accurately to achieve acceptable performance. As shown in Table 4, the performance of all methods decreases comparing with warm start shown in Table 1. PT-PMF has no results in Table 4 since it cannot handle cold start users. On this cold start scenario, PT-VPMF-1 outperforms PT-1 5.74%, and PT-VPMF-.5 outperforms PT-.5 5.15% with respect to the average F_1 value. This demonstrates the significant benefits of incorporating visual features of photos to alleviate the cold start problem.

6 CONCLUSION

In this paper, a tour recommender system leveraging geo-tagged photos, named 'Photo2Trip', was proposed to recommend not only suitable POIs to visit but also visit duration at each POI. Specifically, we proposed a Visual-enhanced Probabilistic Matrix Factorization model (VPMF), which integrated visual features into the collaborative filtering model, to learn user interests by leveraging the historical travel records. Our work improved existing tour recommendation research in two ways: (i) we introduced collaborating filtering into trip planning to predict user visit preferences of non-visited POIs, instead of using the average visit duration of each category of POIs for all users as individual interest; and (ii) we extracted and integrated visual features in user-generated photos of POIs into the collaborative filtering model PMF to further improve user interests prediction.

Using the Yahoo! Flickr dataset across four cities, we evaluated the effectiveness of our proposed approache against various baseline methods. The experimental results showed that: (i) using collaborative filtering to predict user interest resulted in accurate prediction to the real-life travel sequences of users, in terms of both precision and F_1-score; (ii) incorporating visual features into the PMF model could further improve the accuracy of prediction; (iii) our proposed VPMF approaches predicted personalized POI visit duration more accurately, and (iv) incorporating visual features into PMF significantly alleviated the cold start problem.

7 ACKNOWLEDGEMENT

This research was partially supported by the Natural Science Foundation of China under grant No.71329201, 61502324 and 61532018.

REFERENCES

[1] Ryan Prescott Adams, George E. Dahl, and Iain Murray. 2010. Incorporating Side Information in Probabilistic Matrix Factorization with Gaussian Processes. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI'10)*. AUAI Press, 1–9.

[2] Yuki Arase, Xing Xie, Takahiro Hara, and Shojiro Nishio. 2010. Mining people's trips from large scale geo-tagged photos. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 133–142.

[3] Michel Berkelaar, Kjell Eikland, Peter Notebaert, and others. 2004. lpsolve: Open source (mixed-integer) linear programming system. *Eindhoven U. of Technology* (2004).

[4] Igo Brilhante, Jose Antonio Macedo, Franco Maria Nardini, Raffaele Perego, and Chiara Renso. 2013. Where shall we go today?: planning touristic tours with tripbuilder. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 757–762.

[5] Igo Ramalho Brilhante, Jose Antonio Macedo, Franco Maria Nardini, Raffaele Perego, and Chiara Renso. 2015. On planning sightseeing tours with TripBuilder. *Information Processing & Management* 51, 2 (2015), 1–15.

[6] An-Jung Cheng, Yan-Ying Chen, Yen-Ta Huang, Winston H Hsu, and Hong-Yuan Mark Liao. 2011. Personalized travel recommendation by mining people attributes from community-contributed photos. In *Proceedings of the 19th ACM international conference on Multimedia*. ACM, 83–92.

[7] Chen Cheng, Haiqin Yang, Irwin King, and Michael R Lyu. 2012. Fused Matrix Factorization with Geographical and Social Influence in Location-Based Social Networks.. In *Aaai*, Vol. 12. 1.

[8] Gabriella Csurka, Christopher Dance, Lixin Fan, Jutta Willamowski, and Cédric Bray. 2004. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, Vol. 1. Prague, 1–22.

[9] Dominique Feillet, Pierre Dejax, and Michel Gendreau. 2005. Traveling salesman problems with profits. *Transportation science* 39, 2 (2005), 188–205.

[10] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2013. Exploring temporal effects for location recommendation on location-based social networks. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 93–100.

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

[12] Yong Ge, Qi Liu, Hui Xiong, Alexander Tuzhilin, and Jian Chen. 2011. Cost-aware travel tour recommendation. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 983–991.

[13] Quanquan Gu, Jie Zhou, and Chris Ding. 2010. Collaborative filtering: Weighted nonnegative matrix factorization incorporating user and item graphs. In *Proceedings of the 2010 SIAM International Conference on Data Mining*. SIAM, 199–210.

[14] Ron Kohavi and others. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *IJCAI*, Vol. 14. Stanford, CA, 1137–1145.

[15] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 426–434.

[16] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.

[17] Kwan Hui Lim. 2015. Recommending tours and places-of-interest based on user interests from geo-tagged photos. In *Proceedings of the 2015 ACM SIGMOD on PhD Symposium*. ACM, 33–38.

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

[19] Yanchi Liu, Chuanren Liu, Bin Liu, Meng Qu, and Hui Xiong. 2016. Unified Point-of-Interest Recommendation with Temporal Interval Assessment. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1015–1024.

[20] Yanchi Liu, Chuanren Liu, Nicholas Jing Yuan, Lian Duan, Yanjie Fu, Hui Xiong, Songhua Xu, and Junjie Wu. 2014. Exploiting heterogeneous human mobility patterns for intelligent bus routing. In *Data Mining (ICDM), 2014 IEEE International Conference on*. IEEE, 360–369.

[21] David G Lowe. 1999. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, Vol. 2. Ieee, 1150–1157.

[22] Xin Lu, Changhu Wang, Jiang-Ming Yang, Yanwei Pang, and Lei Zhang. 2010. Photo2trip: generating travel routes from geo-tagged photos for trip planning. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 143–152.

[23] Ruslan Salakhutdinov and Andriy Mnih. 2007. Probabilistic Matrix Factorization.. In *Nips*, Vol. 1. 1257–1264.

[24] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[25] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2015. The new data and new challenges in multimedia research. *arXiv preprint arXiv:1503.01817* 1, 8 (2015).

[26] Xinxi Wang and Ye Wang. 2014. Improving content-based and hybrid music recommendation using deep learning. In *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 627–636.

[27] Le Wu, Enhong Chen, Qi Liu, Linli Xu, Tengfei Bao, and Lei Zhang. 2012. Leveraging tagging for neighborhood-aware probabilistic matrix factorization. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM, 1854–1858.

[28] Mao Ye, Peifeng Yin, and Wang-Chien Lee. 2010. Location recommendation for location-based social networks. In *Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems*. ACM, 458–461.

[29] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik-Lun Lee. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, 325–334.

[30] Quan Yuan, Gao Cong, Zongyang Ma, Aixin Sun, and Nadia Magnenat Thalmann. 2013. Time-aware point-of-interest recommendation. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 363–372.