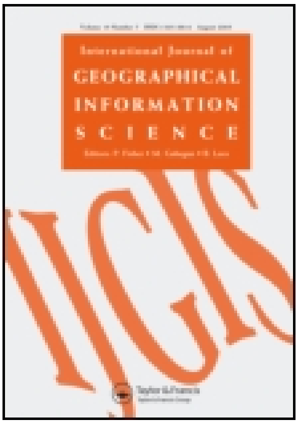

International Journal of Geographical Information Science

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/tgis20

A context-aware personalized travel recommendation system based on geotagged social media data mining

Abdul Majid [a], Ling Chen [a], Gencai Chen [a], Hamid Turab Mirza [a], Ibrar Hussain [a] & John Woodward [b]

[a] College of Computer Science, Zhejiang University, 38 Zheda Road, Hangzhou, 310027, PR China

[b] School of Computer Science, The University of Nottingham Ningbo China, Ningbo, 315100, PR China

Published online: 24 Jul 2012.

International Journal of Geographical Information Science, 2013
Vol. 27, No. 4, 662–684, http://dx.doi.org/10.1080/13658816.2012.696649

Taylor & Francis
Taylor & Francis Group

A context-aware personalized travel recommendation system based on geotagged social media data mining

Abdul Majid[a], Ling Chen[a]*, Gencai Chen[a], Hamid Turab Mirza[a], Ibrar Hussain[a] and John Woodward[b]

[a]*College of Computer Science, Zhejiang University, 38 Zheda Road, Hangzhou 310027, PR China;* [b]*School of Computer Science, The University of Nottingham Ningbo China, Ningbo 315100, PR China*

(*Received 15 January 2012; final version received 20 May 2012*)

The proliferation of digital cameras and the growing practice of online photo sharing using social media sites such as Flickr have resulted in huge volumes of geotagged photos available on the Web. Based on users' traveling preferences elicited from their travel experiences exposed on social media sites by sharing geotagged photos, we propose a new method for recommending tourist locations that are relevant to users (i.e., personalization) in the given context (i.e., context awareness). We obtain user-specific travel preferences from his/her travel history in one city and use these to recommend tourist locations in another city. Our technique is illustrated on a sample of publicly available Flickr dataset containing photos taken in various cities of China. Results show that our context-aware personalized method is able to predict tourists' preferences in a new or unknown city more precisely and generate better recommendations compared to other state-of-the-art landmark recommendation methods.

Keywords: spatiotemporal data mining; geographical gazetteer; trip planning; context-aware query; geo-referenced photographs

1. Introduction

In recent years, popularity of digital cameras and camera phones has contributed to evolving approval of sharing on Internet communities such as Flickr (*flicker.com*) and YouTube (*youtube.com*). Using these online community sites, users tend to expose more and more about their experiences on the Web through rich media data such as photos and videos. The development of location-based social media such as Facebook (*facebook.com*) and Gowalla (*gowalla.com*) is not only transforming the landscape of computing but also stimulating social changes of various kinds, and this phenomenon has moved social media from cyberspace to real place (Sui and Goodchild 2011). The growing size of individual and community footprints on the Web and fast-evolving Internet communities provides evidence about the extent to which the information has pervaded in our lives. In some aspects, humans have transformed from social beings into e-social beings (Luo *et al.* 2011). Photos and videos that constitute a huge proportion of information available on the Web, and are added or exchanged every second, have provided new research opportunities and challenges for multimedia, data mining, and geographic-related research and applications.

*Corresponding author. Email: lingchen@cs.zju.edu.cn

These multimedia data such as photos contain not only contain textual information such as tags, title, notes and description but are also tagged with temporal context (i.e., time at which the photo was taken) and spatial context (i.e., the location in terms of latitude and longitude) where the photo was taken. In the last few years, based on the simple assumption that tourist attractions (landmarks or interesting locations) are those places that are often photographed, it has been a hot research topic to explore common wisdom in photo sharing community to find popular landmarks (Ahern *et al*. 2007, Kennedy *et al.* 2007, Rattenbury *et al*. 2007, Crandall *et al*. 2009, Serdyukov *et al*. 2009, Yang *et al*. 2011) and for travel recommendations (Popescu and Grefenstette 2009, Popescu *et al*. 2009, Lu *et al*. 2010, Yin *et al*. 2011). Collection of tourists' geotagged photos is assumed as a sequence of visited locations to build traveling histories of users, and different methods are proposed to find the popular locations or representative travel sequences to address traveling-related queries. Furthermore, these vast amount of data provide a unique opportunity to explore ways in which users engage and perceive geographical areas. It helps to understand the attitude, attention, and interest of a person or community in these geographical areas (Naaman 2011).

Because trip-planning is a time-consuming task, there is always a need for a system that can recommend tourist locations to match the tourists' interests. According to Resnick and Varian (1997), recommendation systems use opinions of a community of users to help individuals in that community more effectively by identifying contents of interest from a potentially overwhelming set of choices. An enormous amount of users generated content in the form of social media that exhibit their traveling experiences, which provides, a great opportunity to build a recommendation system for travel assistance with the following features.

1.1. Collective wisdom

A social approach (e.g., ask people who know about the area to be explored), adopted by inexperienced travelers in a new area, can provide more up-to-date and accurate information but it takes time for travelers to digest and put together the collected information for use (Yoon *et al*. 2011). Considering user-supplied geotagged photos as source of social traveling experiences, we can explore collective (social) wisdom to (1) compile tourist locations in a city by grouping photos using their associated geotags and (2) determine popularity of locations among tourists in different contexts.

1.2. Personalization

Making a simple assumption that users have specific travel preferences and therefore visit locations that have similar features and taking a photo at a visited location is a sign that the user likes that location, we can get users' specific travel preferences by building users' similarity model from their travel histories (exposed by their contributed photos on sharing sites) and use it to recommend personalized tourist locations to plan trips in different and unknown regions.

1.3. Context awareness

Tourists' preferences in terms of visiting a location or multiple locations in a certain sequence could be affected by their current spatial, temporal, and environmental contexts. For example, a mobile user searching for tourist location may be willing to visit a location near to his/her current location as detected by the mobile device with which he/she is equipped. In terms of weather context, it is intuitively clear that weather conditions may influence where we wish to visit. For example, on a sunny day, we may prefer to visit a park, whereas on a rainy day, we may prefer to visit a museum. It will be useful to restrict recommendation to certain times or dates. For instance, there is no point to recommend a tourist location at the weekend, which is open only from Monday to Friday. From geotagged photos, temporal information (photo taken time) and spatial information (geotags that describe the locations where photos were taken) can be used to estimate the popularity of tourist locations in different temporal contexts. Moreover, various online weather Web services such as Weather Underground (*wundgerground.com*) provide not only the current weather condition of a particular geographical area but also offer its historical weather data. Therefore, current weather conditions provided by these services can be used to augment the query with current weather context, and historical weather data can be used to filter the relevant tourist locations to address the weather context-driven query.

In this article, we focus on context-aware personalized landmark recommendation based on geotagged photos. The method we propose is designed to be deployed in an application scenario that leverages the collective wisdom of people from community-contributed geotagged photo collection to provide a set of tourist locations that match the user's interests and current context given a city that is new to that user. To the best of our knowledge, this is the first work that exploits the users' traveling preferences from their contributed photos and uses the contexts in the photo (i.e., spatial and temporal), in combination with weather context, retrieved from online weather Web services, to support the context-aware personalized tourist recommendation framework. Our contributions in this article are summarized as follows:

(1) A novel system architecture is presented that is capable of addressing dynamic queries for semantically meaningful and personalized tourist location recommendations using geotagged social media. More specifically, the queries may include any or all of the contexts (i.e., temporal and weather).
(2) Illustration is given about (a) how to group photos from collections of user-supplied photos using their associated geotags to find tourist locations and (b) how to aggregate clustered photos' textual information and enrich with supplementary information provided from Web services (i.e., Google Places) to provide more semantic meaning to aggregated locations.
(3) We show how to synergize disjointed contexts and sparse social contents together with online information sources to enrich primitive contexts and contents with higher levels of semantic meanings, that is, profiling locations. We categorize context data associated with aggregated tourist locations to support more complex context-based queries, enabling users to receive more relevant recommendations.
(4) A new method that uses the popularity of tourist locations in different contexts as profile matching criteria to filter the locations according to users' current context and then rank the locations in collaborative filtering manner for personalized recommendations is proposed.
(5) A conceptual foundation is laid down for the analysis of spatiotemporal data of places (tourist locations) obtained from community-contributed geotagged photo collections. We use it to provide location-aware tourist information. It can also be utilized by local authorities, service providers, and tourist agencies for building user-centric applications and to provide location-based services.

The remainder of the article is organized as follows: We begin by discussing the related work in Section 2. Preliminaries and a formal problem definition are given in Section 3.

Section 4 provides the architecture of proposed context-aware tourist location recommendation system and details all its modular tasks. Section 5 reports on the experimental study. Section 6 concludes and discusses future research directions.

2. Related work

This section considers related research about geotags exploration and mapping to landmarks, trips deduction and recommendations from geotagged contents, and personalization in tourist recommendations.

2.1. Geotags exploration and mapping to landmarks

The Flickr database is open to everyone via the FlickrAPI, which allows users' program to search Flickr photo databases for geotagged images. Recently, a number of methods have been proposed to map these geotags to geographical regions and exploit the other information annotated to photos such as title, description, and time, to infer the knowledge to describe and symbolize these geographical regions. Ahern *et al*. (2007) created a World Explorer that used tags on Flickr-geotagged photos to map well-liked tags to geographical locations, resulting in a scale-dependent map overlaid with semantic information on the original data. This work is extended in Kennedy *et al*. (2007) by applying content- and context-based analysis for ranking clusters and finding representative images in a cluster. Rattenbury *et al*. (2007) further investigated the place and event semantics of geotags, in addition to the representativeness. The proposed approach can automatically determine whether a tag corresponds to a 'place' or an 'event'. A 'place' tag is defined as one that exhibits significant spatial patterns, whereas an 'event' tag is defined as one that exhibits significant temporal patterns. Visual representations of landmarks in social media are used in Kennedy and Naaman (2008) to create a visual summary in a response to a geographic query, for augmenting and improving the user interaction. Hays and Efros (2008) used the nearest neighbor method to predict the locations of photos based on their tags and the features. Serdyukov *et al*. (2009) divided the map using a grid so each cell represents a location. They defined a language model to describe the relation of a tag and a place and estimated a cell as the place where a photo was taken by using tags annotated to the photo. Crandall *et al*. (2009) used another approach for location prediction; some primary locations are extracted using the photos annotated with the geotags. Next, they classify the non-geotagged photos from each major location by using support vector machines based on visual, textual, and temporal features of the photos. They suggested not to use a fixed number of clusters and proposed a mean shift algorithm to find the most prominent landmarks and representative photos. Another work based on spectral clustering about identifying location as Point of Interests (POIs) was proposed by Yang *et al*. (2011). They proposed a self-tuning approach based on the cut cost similarity to eliminate the effect of parameters from spectral clustering. Tourist activities inferred from geotagged Flickr photos were derived by Hecht and Gergle (2010) to detect landmarks. Using global positioning system (GPS) traces, Zheng *et al*. (2010) defined a method to extract interesting locations from these data and proposed a matrix factorization method to suggest locations and activities.

To infer the knowledge about location found using spatial proximity of photos, existing works used visual features of photos or photos' metadata such as title, tags, and description. Due to the unrestricted nature of photo sharing applications, one or more of the aforementioned metadata fields might be missing or incorrect. In our work, first we apply a density-based clustering algorithm to photos' geotags to extract tourist locations, and then we aggregate photos' textual tags and enrich with supplementary information provided from a Web service (i.e., Google Places) to give semantic meaning to aggregated locations. Furthermore, to summarize the locations aggregated from photos and to derive the dynamics of users' interests to these locations, temporal tags annotated to photos are exploited to infer users' visits for profiling locations. Profile of each location provides the information about the users who have visited that location and the history of contexts (i.e., weather and temporal) in which the location has been visited. Note that, we identify the temporal context of a visit by exploiting the time-stamps of photos that were taken during that visit and use this visit time to obtain the weather context of the visit from historical weather dataset retrieved from online weather resources.

2.2. Trips deduction and recommendations

To deduce trip-related information, Popescu and Grefenstette (2009) utilized the temporal information associated with photos. For recommendations, they focused the query with temporal constraints in terms of duration of the trip but did not consider the current temporal and weather context of tourist. In Popescu *et al*. (2009), the collective behavior of tourists from aggregated individual trips is represented in a graph and then is used to construct intracity itineraries. An interactive tourist recommendation method is proposed by Lu *et al*. (2010) that took into account a number of factors such as duration of the trip and traveling cost to help the tourist for trip planning. Their approach extracts travel routes from geotagged photos, and then clusters, indexes, and visualizes these routes. Yin *et al*. (2011) explored the photo sharing for recommendation of popular and diversified trajectory pattern. Recently, proliferation of devices equipped with GPS made it possible to collect an individual's movement footprints. For GPS sample point data, Zheng *et al*. (2009) proposed a tree-based hierarchical graph to model location histories of users and used Hypertext Induced Topic Search (HITS)-based inference model to exploit the reinforcement relationship between users and locations for locations and travel sequence recommendations, whereas Wachowicz *et al*. (2011) contributed a method to define and extract moving flock patterns by using the notion of collective coherence.

2.3. Personalized tourist recommendations

Personalization has been identified as an important factor of effectiveness and added value in recommendation systems. Many recommendation algorithms have been proposed based on similarities between objects in a discrete item-space (Wang *et al*. 2006, Sarwar *et al*. 2001), which has proven to be effective in E-commerce applications (Linden *et al*. 2003). Using GPS data, Takeuchi and Sugimoto (2006) used an item-based collaborative filtering method to recommend shops similar to a user's previously visited shops, whereas Horozov *et al*. (2006) used a user-based collaborative filtering method to generate restaurant recommendations. The difference between our work and Horozov *et al*.'s (2006) and Takeuchi and Sugimoto's (2006) works is that we generate recommendations without explicit user ratings. The information we use is a summary of knowledge acquired from community-contributed photos. Using geotags of public photos, Clements *et al*. (2010) presented a method to recommend travel locations based on a user's travel history using collaborative filtering. The proposed approach orders the locations based on their popularity and then linearly combines the popularity score with personalized score weighted by the similarities between the active user and other users. They used the density of geotags in a region to estimate the region's popularity and geotags of photos contributed by users to obtain a rating for that geographical location. As a user can take more than one photo at the same location during the same visit, we use photos' taken time to define visits made by different users to different locations. We build the profile of locations based on users' visits to better understand how locations are engaged and perceived by users in different contexts and to infer users' rating to these locations. Furthermore, for making recommendations, we not only use the interest of users but also consider the users' current context.

Research interests in automatic tourist recommendation guides as discussed above have resulted in numerous methods, techniques, and applications mostly based on models that utilize popularity as desired properties of recommended tourist locations. A little emphasis has been given to personalization for tourist location recommendation. Furthermore, existing tourist location recommendation methods address narrow range of queries, that is, queries with free of context constraints or with a few dimensions of context. We argue that better recommendations may not only depend on the popularity of locations among tourists or tourists' interests obtained from their traveling history but also depend on the user's current contexts (e.g., local weather conditions and current time) and popularity of the tourist locations in different contexts. In this article, we extend the notion of tourist travel recommendation utilizing collaborative filtering techniques while taking into account the contextual information for deriving improved personalized recommendations in geotagged social media.

3. Preliminaries and problem definition

Before we formally define the problem, we give definitions of some basic concepts and terms.

Definition 1: (*Geotagged photo*) A geotagged photo p can be defined as $p = (id, t, g, X, u)$ containing a photo's unique identification, id; its geotags, g; its time-stamp, t; and the identification of the user who contributed the photo, u. Each photo p can be annotated with a set of textual tags, X. Geotags g of photo p are the coordinates of the geographical region where photo p was taken.

Definition 2: (*Photo collection*) Collection of all photos, contributed by all tourists can be represented as $P = \{P_1, P_2, \ldots, P_n\}$, where P_u $(u = 1, \ldots, n)$ is the collection of photos contributed by user u.

Definition 3: (*Location*) A location l can be viewed as geographical region within a city such as a park, a lake, or a museum, which is popular for tourists to visit and take photos.

Definition 4: (*Context-aware query*) A context-aware query Q is defined as $Q = (t, w)$, where t represents temporal context and w denotes weather context.

The problem of personalized location recommendation for trip planning with geotagged social media is formulated as, given a collection of geotagged photos $P = \{P_1, P_2, P_3, \ldots, P_n\}$, where P_u $(u = 1, \ldots, n)$ is the set of photos contributed by user u, how to locate and summarize tourist locations and build a travel history of each user to derive his travel preferences to undertake a context-aware personalized query Q. We aim to address the problem of personalized recommendation by trying to exploit travel history of user to recommend tourist locations that best fit his or her interests.

4. Context-aware personalized travel recommendation system

4.1. System architecture

The architecture behind our approach is configured into various modular tasks to carry out different operations as depicted in Figure 1. We give an overview of these operations here and details are presented in the following sections.

(1) Finding tourist locations from geotag photos (Section 4.2).
(2) Annotation of locations with semantic (Section 4.3).
(3) Profiling locations and building a database of tourist locations (Section 4.4).
(4) Modeling users' traveling preferences and similarities among users based on their traveling preferences (Sections 4.4.1 and 4.4.2).
(5) Making recommendations (Section 4.5).

We find tourist locations using spatial proximity of photos and enrich the aggregated locations with semantic annotations using textual tags annotated to photos in combination with information provided by Web services. Profiles of locations are built to describe the contexts in which they have been visited. To derive temporal context, geotags and temporal tags annotated with photos are exploited, whereas to derive weather context, we query third-party weather Web services to retrieve weather conditions. Relationship between users and locations is drawn to model users' travel preferences. Then, these users' preferences are used to estimate the similarities among users. For making recommendations, first we filter the locations based on contextual constraints, and then rank the locations by personalized score. A measure is defined to identify similar users in previously visited cities and aggregate these users' opinions to obtain personalized score for each location in a target city for the target user.

4.2. Finding tourist locations

Finding tourist locations from a collection of geotagged photos can be viewed as a clustering problem of identifying highly photographed locations. Clustering algorithms such as

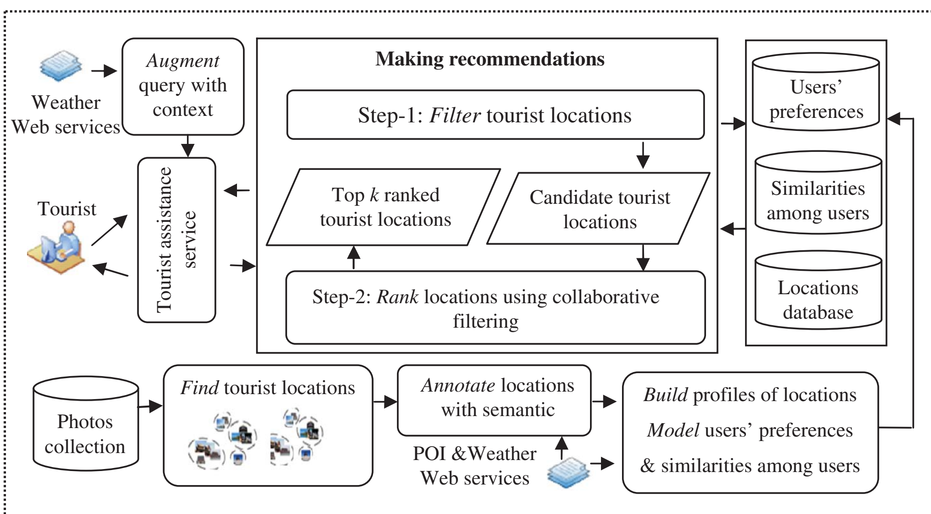

Figure 1. Architecture of proposed system for context-aware personalized travel recommendation.

k-mean and mean-shift have been used to cluster photos using associated geotags for the identification of locations (Kennedy *et al*. 2007, Yin *et al*. 2011). However, density-based clustering algorithms such as DBSCAN (Easter *et al*. 1996) have several advantages over other types of clustering algorithms: they require minimum domain knowledge to determine the input parameters and can find clusters with arbitrary shape. In addition, they can filter outliers and work effectively when applied to large databases. DBSCAN requires only two parameters: ε (*epsilon*) and the minimum number of points required to form a cluster (*minPts*). DBSCAN randomly selects an object and forms a range search with radius ε and iteratively finds subsequent density-reachable objects to make the cluster. DBSCAN clustering works with generic points having a unified density threshold for all clusters; however, the locations extracted by clustering the given collection of photos can have varying sizes and densities. To address this problem, Kisilevich *et al*. (2010) proposed P-DBSCAN, a variant of DBSCAN. They extended the definition of directly density-reachable object by adding an adaptive density technique. In P-DBSCAN, an object O is directly density reachable from another object O' if it is not farther away than a given density radius ε and the ratio of surrounding objects between O and O' must be less than a density ratio ω.

Given a collection of photos P, we use P-DBSCAN, in order to cluster photos to identify tourist locations based on the photos' geotags. The output of a P-DBSCAN is a set of locations (clusters of photos) $L = \{l_1, l_2, \ldots, l_n\}$. Each element $l = \{P_l, g_l\}$, where P_l is a group of geographically clustered photos, and g_l are geographical coordinates to represent the centroid of photos' cluster P_l and are computed from a group of geotags associated with the photos in the cluster P_l.

4.3. Semantic annotation of locations

The tourist locations, identified by using a clustering algorithm on spatial proximity of photos, can be visualized on a map interface where icons or convex hull polygons can be drawn to show the position and boundaries of locations. To give semantic meanings, we provide a method that uses textual tags annotated to photos in combination with the information provided by online Web services, to automatically generate textual descriptions for each tourist location. Our method as described by algorithm-1 contains three steps. In the first step (line 2), we use a method described in Kennedy *et al*. (2007) to derive representative textual tag for each location $l = (P_l, g_l)$. Considering each location $l = (P_l, g_l)$ and set of tags X_l that appear with group of photos P_l, they used a method based on term frequency-inverse document frequency (TF-IDF) to score each tag $x \in X_l$. Note that, TF-IDF is a popular ranking method and is widely used for information retrieval. At the end of step 1, for location $l = (P_l, g_l)$, we have a list of tags X'_l and each tag $x' \in X'_l$ has a score $s(x)$. The higher the score, the more distinctive the tag is within a group X_l. In the second step (line 3), we use Web services, that is, Google Places (*google.com/places*) to extract the information about the POIs in a certain geographical area. These services work in this way: we provide them a geographical coordinate g and a radius r in meters and in response they return the metadata of places that are present within r of g. We use centroid g_l of location l to represent g. The output of step 2 is the set $PLACES = \{place\}$ for location l. Each entry in *PLACES* represented by *place* provides the information to describe a POI. In the last step, we aggregate the results of steps 1 and 2 to get the representative description of tourist location. The aggregation is performed as: we order the set of tags X'_l according to their score that is computed in step 1 (line 4). We iteratively compare each element of X'_l with all elements of *PLACES* (lines 6–9). In the result of comparison,

ALGORITHM-1 Semantic Annotation

Input: $L = \{l\}$ Set of locations
Output: $L' = \{l'\}$ Set of semantic locations
1. **for each** location $l = (P_l, g_l)$ in L **do**
2. **COMPUTE** score s for each tag $x \in X'_l$ belongs to photos' group P_l using TF-IDF
3. **RETRIEVE** *PLACES* from POI Web services.
4. **SORT** X'_l based on score s
5. **CREATE** list *MatchedList*
6. **for each** x in X'_l **do**
7. **for each** *place* in *PLACES* **do**
8. **if** (**MATCH**(x, *place*) = **true**) **then**
9. **ADD** *place* to *MatchedList*
10. **if (LENGTH**(*MatchedList*) > 1) **then**
11. *l.name*←**CLOSEST**(*MatchedList*).*name*
12. **else if** (**LENGTH**(*MatchedList*) = 1) **then**
13. *l.name*←*place.name*
14. **else**
15. *l.name*←**TOP**(x)
16. **ADD** l to L'

(a) If multiple matches are found, then the matched *place* that is closest to geographical coordinate g_l, in terms of spatial distance, we consider its name as a location name (lines 10 and 11).
(b) If a single match is found, then the *place* that is matched, we use its name as location name (lines 12 and 13).
(c) If no match is found, we use the tag with highest score as the location name (lines 14 and 15).

Definition 5: (*Semantic location*) A semantic location l' can be defined as $l' = (l, a)$, where a is the semantic annotation to describe the location l.

4.4. Profiling locations and acquisition of user preferences

Once the photos have been clustered using their spatial proximity to find the tourist locations, and the locations have been annotated with semantic, we are interested in formulating the profiles of locations and build a database of locations. Algorithm-2 illustrates the method for locations' profiling. First step is to identify visits made by different users from photos taken by them on these locations (lines 1–10). For each location $l \in L$, we sort photos of each user u according to photos' taken time. We infer visit v from a photo p taken by a user u at location l at time t. Note that a user u can take more than one photo in same visit at same location. For this, if the difference between the time-stamps of two photos ($p_2.t$–$p_1.t$) taken by same user at same location in less than visit duration threshold $visit_{\text{thr}}$, we consider that both photos belong to same visit. We use the median of time-stamps associated with photos (belong to visit v) as the visit time $v.t$.

Definition 6: (*Visit*) A visit v can be represented as $v = (l, u, t)$, where u is the user who made visit v at location l at time t.

ALGORITHM-2 Profiling Locations

Input: $L = \{l\}$ Set of locations where $l = (P_l, g_l)$
Output: LDB $= \{l'\}$ Database of locations with updated profiles
1. **for each** location l=(P_l, g_l) in L **do**
2.   **CREATE** list V_l
3.   **CREATE** list of users U_{Pl} from P_l and **SORT** photos $P_{ul} \in P_l$ taken by each user $u \in U_{Pl}$ according to photo taken time $p.t$
4.   **for each** user u in U_{Pl} **do**
5.   **CREATE** list T_v
6.   **for each** p in P_{ul} **do**
7.   **if** ($p_i.t$ - $p_{i\text{-}1}.t < visit_{thr}$) **then**
8.   **ADD** $p.t$ to T_v
9.   **else**
10. $v \leftarrow$ **NEW(**visit)
11. $v.t \leftarrow$**MEDIAN**(T_v)
12. $v.w \leftarrow$**RETRIEVE-FROM-WEATHER-DB**($v.t$)
13.   **ABSTRACT**($v.t$, $v.w$)
14.   **ADD** v to V_l
15.   **CLEAR** P_v
16.   **ADD** p to P_v
17. $l.pop(w, t)$ $\leftarrow$**POPULAR**(V_l)
18.   **ADD** l to LDB

The next step is to build the history of contexts in which tourist locations have been visited (lines 11 and 12). To describe this, first we explain some formal notion. Let L be the set of all locations extracted, and let U be the set of all users. For a user $u \in U$, let V_u be the set of visits to different locations made by u so we can derive $V = \bigcup_{u \in U} V_u$. The users who visited location l can be defined as $U_l = \{u \in U\text{: } u \textit{ visited location } l\}$. If we represent the visits made by user u at location l as $V_{ul} = \{v \in V_u\text{: } v.l \in L\}$, then all visits made by all users at location l can be represented as $V_l = \{v \in V\text{: } v.l \in L\}$, where $v = (l, u, t)$. The example shown in Table 1 depicts how the visits are represented in terms of (l, u, t).

To build history of temporal and weather contexts, in which the location l has been visited, the available information for each visit is time. This time-stamp information enables us (1) to induce the temporal context t of each visit and (2) to retrieve weather context (condition) w, when visit $v = (l, u, t)$ was made by user u at location l at time t. Weather services normally publish weather conditions at hourly, daily, or monthly levels that contain different variables such as temperature, precipitation, humidity, etc. Context-related data, such as the time-stamp, and weather variables cannot be directly used as contextual information; thus, we need a context abstraction strategy to obtain abstract context concepts. Various context abstraction methods have been proposed for temporal and weather context abstraction (Shin *et al*. 2009, Lee *et al*. 2010). For example, the raw context (21:30, 25C°)

Table 1. Representation of visits in terms of (u, l, t).

Visits	Locations	Users	Time stamp
v_1	l_1	u_1	18 April 2009 05:20:23
v_2	l_2	u_2	22 January 2011 07:40:05

Table 2. Temporal and weather context concepts.

Temporal context concepts		Weather context concepts	
Day of week: working day, weekend	Time of day: morning, afternoon, night	Temperature: hot, warm, cold	Condition: sunny, cloudy, rainy

Table 3. Representation of visits in terms of (u, l, t, w).

Visits	Locations	Users	Temporal context	Weather context
v_1	l_1	u_1	Weekday, morning	Warm, sunny
v_2	l_2	u_2	Weekday, evening	Cold, raining

can be abstracted to (night, warm). We use context concepts given in Table 2 to represent the temporal and weather contexts of each visit.

For each $v \in V_l$, we transform $v.t$ to temporal context concepts. For weather context, we retrieve weather conditions of location l at $v.t$ from historical weather database and represent it using weather context concepts (line 13). After the context retrieval and abstraction, each visit v belongs to a set of visits $V_l = \{v \in V\text{: } v.l \in L\}$ made to locations l, which can be expressed as $v = (u, l, t, w)$. An example of visits' representation with abstract contextual concepts is depicted in Table 3.

The last step in profile building is to find the popular contexts of each location from the history of contexts derived from visits made to respective location (line 17).

Definition 7: (*Popular temporal and weather context*) Given the set of visits V_l belongs to location l, based on temporal and weather context concepts (i.e., $v.t$ and $v.w$), we consider weather and temporal context concepts with highest frequency as popular context concepts of location l. For example, $pop(l.w) = (warm, sunny)$ depicts that location l has been popularly visited in warm and sunny weather conditions.

After identifying the visits made by different users to different locations and deriving the popularity of each location in terms of different contexts, we build a locations database LDB $= \{l_1, l_2, \ldots, l_n\}$, where each location $l_i = \{V_{li}, pop(w), pop(t)\}$, V_{li} are visits made to location l_i by different users, $pop(w)$ is frequent weather context, and $pop(t)$ is frequent temporal context of location l_i (line 18).

4.4.1. Building user–location matrix

To derive the interest of users U in a set of locations L, we use the links (set of visits V) between users U and locations L to build a weighted undirected graph $G_{UL} = (U; L; E_{UL}; W_{UL})$, where U and L are nodes to represent users and locations, respectively. E_{UL} and W_{UL} are sets of edges and edge weights between U and L to represent users' visits and the number of visits to particular locations. The relationship between users and locations is depicted in Figure 2.

Given m users and n locations, we build an m by n adjacency matrix M_{UL} for graph G_{UL}. Formally, $M_{UL} = [v_{ij}]$, $0 \leq i < m$, $0 \leq j < n$, where v_{ij} represents how many times the ith user has visited the jth location.

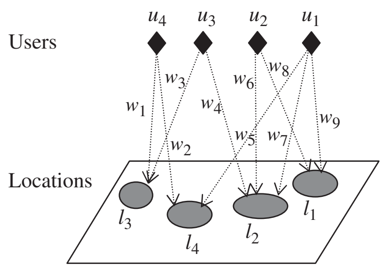

Figure 2. Relationship between users and locations based on visits.

Definition 8: (*User Interest*) From matrix M_{UL}, the user u_p's travel interest can be derived as an array $R_p = <r_{p0}, r_{p1}, \ldots, r_{pn}>$, where r_{pi} is u_p's implicit rating (visits made by u_p) in a location i. $S(R_p)$ is the subset of R_p, $\forall r_{pi} \in S(R_p)$, $r_{pi} \neq 0$, that is, the set of locations that has been preferred (visited) by u_p. The average rating in R_p is denoted as $\bar{R}_p$. For example, $R_1 = <5, 2, 0, 4, 3, 0>$, $R_2 = <1, 0, 3, 2, 0, 1>$ then $S(R_1) = <5, 2, 4, 3>$, and $S(R_2) = <1, 3, 2, 1>$.

4.4.2. Building user–user similarity matrix

We calculate the similarities among users based on their traveling preferences using the Pearson correlation metric as given in Equation (1), and build users' similarity matrix M_{UU}. Note that, we will use this similarity matrix M_{UU} for personalized recommendation based on state-of-the-art user-based collaborative filtering recommendation method. Each entry in M_{UU} represents the similarity between u_p and u_q. A larger value means that both users are more similar in terms of traveling preferences.

$$sim(u_p, u_q) = \frac{\sum_{i\epsilon S(R_p)\cap S(R_q)} (r_{pi} - \bar{R}_p).(r_{qi} - \bar{R}_q)}{\sqrt{\sum_{i\epsilon S(R_p)\cap S(R_q)} (r_{pi} - \bar{R}_p)^2}.\sqrt{\sum_{i\epsilon S(R_p)\cap S(R_q)} (r_{qi} - \bar{R}_q)^2}} \tag{1}$$

5. Recommendations

When a tourist asks for travel assistance using a tourist-assisting service in a new city, a context-aware query is generated on the basis of tourist's current contexts. We assume, when a tourist makes a request, this request can either come from a mobile device or a query is posted against the Web search engine for which local results shall be returned together with regular results. Furthermore, we suppose two scenarios for the probable contexts of query. First, a tourist is in the target city and asking for immediate assistance. Second, he or she is planning to visit target city in the future. For both scenarios, it is required to augment the query with context. For the first case, current system time can be used to represent temporal context, whereas the current weather conditions of geographical area (target city) for which the tourist is making the request can be retrieved from weather Web services. For the latter case, user can provide temporal information explicitly, and on the basis of that temporal information, forecasted weather conditions published by weather Web services can be retrieved. As this augmented query contains contextual information in the form of time-stamp and weather variables, it cannot be addressed directly. Before processing query $Q = (t, w)$, it is required to abstract temporal context t using temporal context concepts and weather context w using weather context concepts as given in Table 2.

For a query Q, processing proceeds as a two-step approach: an initial filtering step retrieves locations of the target city from locations database LDB that meet the contextual constraints given in the query, thus producing a filtered set of tourist locations L'. Temporal and weather contexts are interpreted in terms of temporal and weather context concepts given in Table 2. For example, consider the following instance of query $Q =$ (Weekday-Morning, Warm-Sunny). From LDB, a set of tourist locations L' will be retrieved that have popular associated temporal context concepts 'Weekday-Morning' and 'Warm-Sunny' as popular weather concepts.

In the second step, we use the user–location matrix M_{UL} that represents the users' preference and M_{UU} that represents the similarities among users to personalize the recommendations for active user u_p for the target city. From M_{UU}, we retrieve similarities between u_p and top N most similar users $U' \in U$, who have visited the target city, and use Equation (2) to predict preferences of u_p for each location l_i from L', that is based on collaborative filtering. In collaborative filtering, the user is recommended items that people with similar tastes and preferences liked in the past.

$$\text{Score}\,(l_i) = \bar{R}_p + k \sum_{u_q \in U'} \text{sim}\left(u_p, u_q\right) \cdot \left(r_{qi} - \bar{R}_q\right) \tag{2}$$

$$k = \frac{1}{|U'|} \sum_{u_q \in U'} \text{sim}\left(u_p, u_q\right) \tag{3}$$

$$\bar{R}_p = \frac{1}{\left|S\left(R_p\right)\right|} \sum_{j \in S(R_p)} r_{pj} \tag{4}$$

We use similarity, sim(u_p, u_q), between users u_p and u_q as a weight to calculate the rank score for each location l_i. That is, the more similar u_p and u_q are, the more weight r_{qi} will carry in the prediction of l_i. Instead of using the absolute values of ratings, we use the deviations from the average rating of the corresponding user. One problem with using the weighted sum is that it does not take into account the fact that different users may use the rating scale differently. Therefore, we use an adjusted weighted sum here. In Equation (2), multiplier k serves as a normalizing factor and usually selected as in Equation (3) and average rating of user u_p, $\bar{R}_p$, from locations in his traveling history, is defined according to Equation (4). Note that Equations (2–4) illustrate a well-known method widely used in many recommender systems (Adomavicius and Tuzhilin 2005).

After computing the user's preferences for each location l_i in L', we order the locations based on preference score and return k number of locations as a query result.

6. Experimental evaluation and results

In this section, we cover the details about our experiment set-up and discuss the results.

6.1. Data

6.1.1. Data acquisition

We use the public API of Flickr to collect metadata of 736,383 geotagged photos that were taken in six cities in China between 1 January 2001 and 1 July 2011. Historical weather

Table 4. Sample records from historical weather data.

Time (CST)	Temp (°C)	Wind chill (°C)	Dew point (°C)	Humidity (%)	Pressure (hpa)	Wind speed	Events
12:00 AM	3.0	1.0	−4.0	60	1031	7.2 km/h / 2.0 m/s	Clear
12:30 AM	3.0	0.1	−3.0	65	1030	10.8 km/h / 3.0 m/s	Clear
1:00 AM	3.0	0.1	−4.0	60	1031	10.8 km/h / 3.0 m/s	Clear

data of these cities are collected using the public API of Wunderground. Sample records from historical weather data and photos' metadata are given in Tables 4 and 5.

6.1.2. Data pre-processing

We removed the metadata of (1) photos that were collected in the result of search based on text containing name of a city in their metadata, that is, tags, title, and description but their spatial context (latitude, longitude) did not match the geographical context of that city, and (2) photos with incorrect temporal context. For example, we removed any photo whose upload time was identical to its taken time, because Flickr assigns a default value to photo without the taken time recorded by the camera. Statistics about photos' metadata is given in Table 6, and spatial distribution of photos in different cities is shown in Figure 3.

6.2. Finding tourist locations

To detect locations from photos, we set the value of $minPts$ = 50 photos, ε (*epsilon*) = 100 m, and density ratio $\omega = 0.5$ for P-DBSCAN. Figure 4 shows the boundaries of locations in different cities.

Table 7 summarizes the information regarding the popularity of locations based on unique number of visits and visitors. To detect visits from photo taken activities, we use value of visit duration threshold $visit_{\mathrm{thr}} = 6$ hours.

6.3. Context-aware personalized recommendation

In this section, we describe the effectiveness of our proposed context-aware personalized recommendation method. We explain our evaluation methodology and compare the results of our work with the existing approaches.

6.3.1. Ground truth and methodology for evaluation

6.3.1.1. Ground truth. For evaluation, we select users who have visited at least two distinct cities $\{C_o, C_t\} \in C$, where C_o represents training city and C_t is the test city. To evaluate only those users who have provided a decent amount of preference information, we consider users who have visited at least five locations in training city C_o.

6.3.1.2. Methodology. We predict the locations actually visited by test user $u_p \in U$ in C_t, based on preferences derived from the locations visited by that user in C_o. We use visits made by the test user to tourist locations in C_t to obtain (1) number of relevant locations

Table 5. Sample records from photos metadata.

Photo Id	User Id	Title	Description	Taken time	Upload time	Lat	Long	Tags
303251010	74434506@N00	Shanghai	Shanghai Pearl Tower skyline	5 September 2010 09:22:32	20 September 2010 10:05:12	31.23827	121.487331	Chin, river, Pearl Tower, Shanghai, bund
3057417546	32267947@N06	Bund	Shanghai Bund	23 March 2007 18:25:10	5 April 2007 18:10:32	31.24417	121.486985	river, Asia, Shanghai, Chinese, Pudong, bund, Canada, good

Table 6. Dataset summary.

	Photos			
Cities	Raw	Filtered	Users	Tags
Shanghai	252,768	230,566	80, 530	244,221
Beijing	241,216	220,631	46, 635	232,164
Hangzhou	37,267	28,312	1090	29,715
Chengdu	20,876	18,514	524	19,388
Guangzhou	18,796	17,141	507	18,474
Hong Kong	196,194	185,008	25, 590	192,421

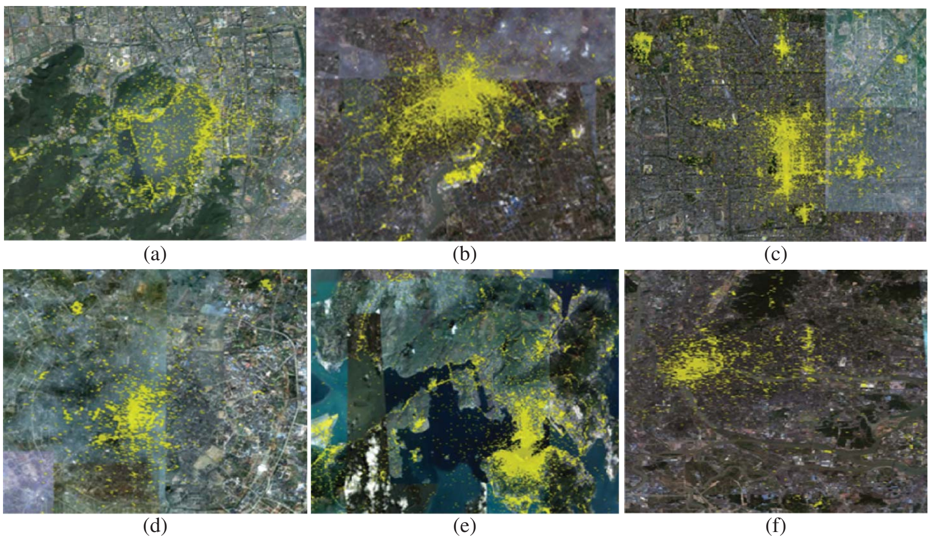

Figure 3. Spatial distribution of photos in cities of China: (a) Hangzhou, (b) Shanghai, (c) Beijing, (d) Chengdu, (e) Hong Kong, and (f) Guangzhou mapped on Google Earth.

denoted as k from total number of visits and (2) temporal and weather contexts associated with visits to build list of contextual constraints. We use these contextual constraints to filter the tourist locations by our context-aware personalized recommendation method. We recommend k number of ranked locations using our and baseline methods. To evaluate the performance of recommendation methods for user u_p, we match the recommended list with the actual list of locations visited by the user in C_t.

6.3.2. *Baseline methods*

We compare the following baseline methods to show the effectiveness of our proposed personalized context-aware recommendation (PCR) method. Two baseline methods, that is, popularity rank (PR) and classic rank (CLR), result in static ranking and generate the same list of tourist locations to all users without considering individual preferences, whereas collaborative filtering rank (CFR) baseline method results in personalized ranking that generates recommendations based on individual's preferences.

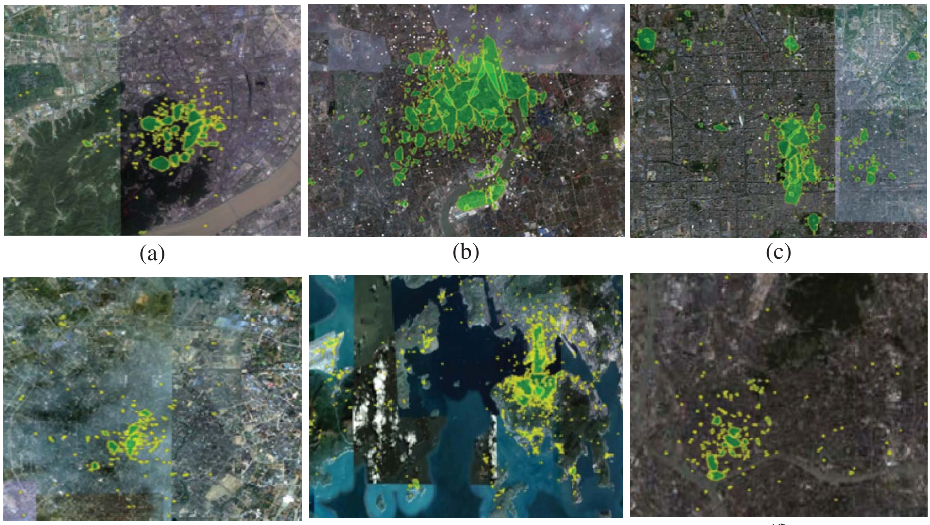

Figure 4. Convex hull polygon are drawn to show the boundaries of locations identified by clustering geotagged photos in different cities: (a) Hangzhou, (b) Shanghai, (c) Beijing, (d) Chengdu, (e) Hong Kong, and (f) Guangzhou.

Table 7. Summary of extracted tourist locations in different cities.

		Locations distribution across visits			Locations distribution across unique visitor		
Cities	Total locations	Visits ≤ 10	$10 <$ Visits < 20	Visits ≥ 20	Visitors ≤ 10	$10 <$ Visitors < 20	Visitors ≥ 20
Shanghai	492	160	138	194	250	110	132
Beijing	411	139	92	180	212	70	129
Hangzhou	128	43	32	53	63	35	30
Chengdu	52	17	13	22	32	9	11
Guangzhou	39	12	14	13	9	14	16
Hong Kong	413	140	122	151	227	60	126

6.3.2.1. Popularity rank. We rank the locations based on the general popularity score that is determined in terms of number of unique visits made to those locations, to use as first baseline method.

6.3.2.2. Classic rank. The idea behind the second approach (Zheng *et al*. 2009), which we cover as a baseline, is to take into account the authority of users instead of treating all users equally. It is thus assumed that popular locations are visited by more authoritative users, and that authoritative users visit more interesting locations. Their approach uses a HITS-based inference method to calculate the locations' interest and users' travel experiences in terms of authority score and hub score by exploiting the reinforcement relationship between users and locations. We use the adjacency matrix M_{UL} that represents the reinforcement relationship between users and locations to apply HITS-based inference model and then rank the locations according to their authority scores.

6.3.2.3. Collaborative Filtering Rank. The third method that we use to compare with our approach is the state-of-the-art user-based collaborative filtering method that exploits evaluations of other tourists with similar interests and potentially provides a ground for the cooperative production of tourist travel recommendations.

6.3.3. Metrics and results

To evaluate the prediction and ranking, we rely on following standard measures from information retrieval.

*6.3.3.1. Precision (*P*).* Precision can be defined as the fraction of correct predictions in total number of predictions made. Figure 5 depicts the performance of our proposed personalized context-aware recommendation and other baseline methods in terms of precision in prediction. Popularity-based ranking and classic rank give better prediction results when compared with collaborative filtering method. The reasons are many users do not have single preferences but visit locations of many types and those locations that are popular and significant when they come to visit a new city. For this reason, evaluation method actually expects us to recommend these. We find that tourists comply more with the general travel preference and are therefore easier to predict by the popularity- or significance-based baseline. This is inherent to the evaluation of recommendations with a train and test set. For users with more number of locations visited in test city, the performance of collaborative filtering improves. These results indicate that collaborative filtering method is effective in recommending places in the case of main tourist tours, whereas it has problems in the case of very short and targeted visit. Moreover, when we exploit context to filter the tourist locations for personalized recommendations, it outperforms all baseline recommendations methods. It shows that we can increase the recommendation accuracy or

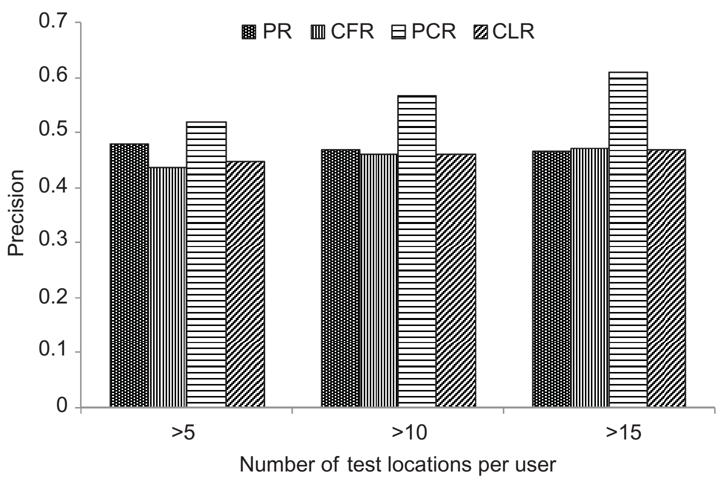

Figure 5. Performance comparison of proposed and baseline methods with different number of test locations per user in terms of precision (P).

relevancy if we consider the context in which the user is requesting and the context in which of the locations have been widely visited while addressing personalized recommendation query.

6.3.3.2. Benefit ratio. Benefit ratio (BR) is the ratio of number of users who get an improved prediction to number of users who get a deteriorated prediction in terms of precision over the baseline. Precision, as discussed before, provides an insight into the recommendation capability of ranking methods in terms of prediction at each visit level. To check the effectiveness of recommendation methods in terms of prediction at user level, we compute BR over all baselines using Equation (5):

$$\mathrm{BR} = \frac{\text{Number of users with improved precision in prediction}}{\text{Number of users with deteriorated precision in presiction}} \tag{5}$$

BR results plotted in Figure 6 show that exploiting context for personalized recommendation can give improved recommendations for most users.

6.3.3.3. Mean average precision. Mean average precision (MAP@*n*) is a widely used evaluation metric to measure the ranking effectiveness that is mean over the precision values after each correct recommendation in the top-*n*. To calculate MAP@50, we recommend 50 locations considering each visit made by each test user in test city as a query and the location visited as one relevant item. We get average precision (AP) for each query $AP = 1/r$, where r is the position of relevant item in ranked list. We obtain the MAP using Equation (6):

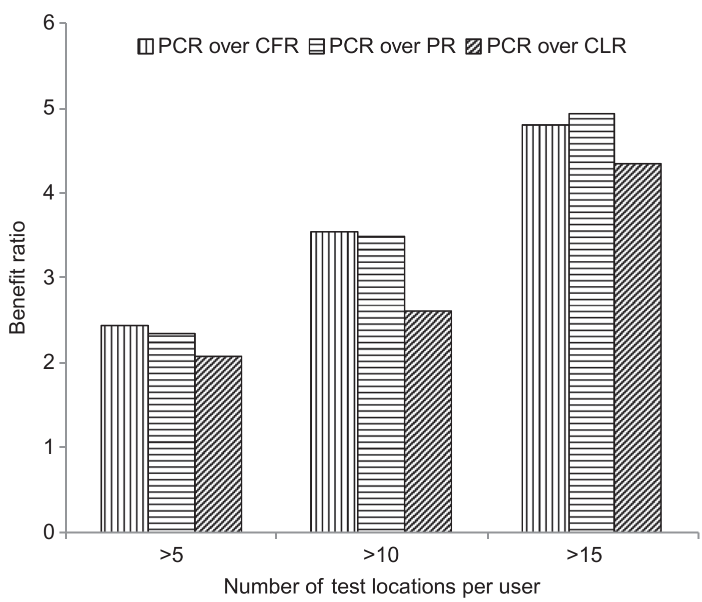

Figure 6. Benefit ratio of proposed personalized context-aware recommendation method over baseline methods.

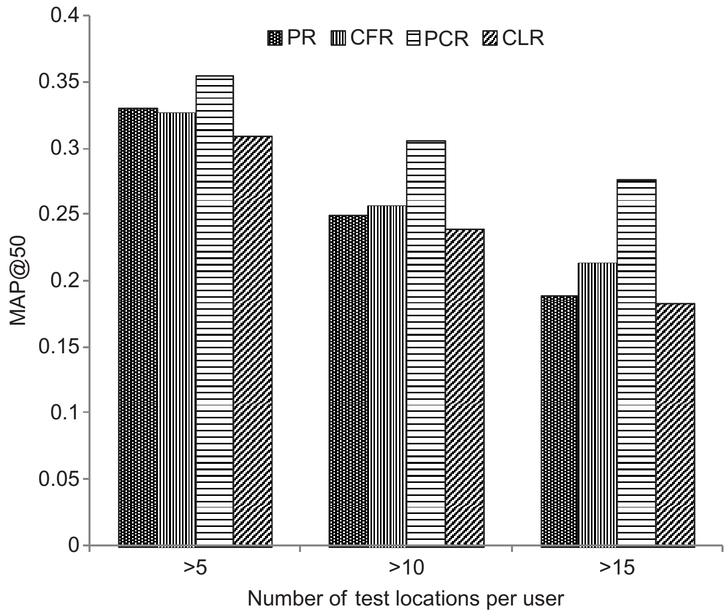

Figure 7. Ranking ability of different methods in terms of MAP@50 across different number of test locations per user.

$$\mathrm{MAP} = \frac{\sum_{i=1}^{Nq} \mathrm{AP}_i}{N_q} \tag{6}$$

where N_q is the total number of queries and AP_i is AP for query i. Figure 7 gives the performance of ranking ability of different ranking methods.

Results obtained from experiments using metric MAP@50 shows that there is a decent improvement in ranking ability of personalized context-aware ranking over other baseline methods. Based on the criteria specified in terms of minimum number of locations visited by each user in the test city, results show that for five locations per user, the effectiveness of personalized ranking over popular ranking is 2.5%, for 10 locations and at 15 locations it outperforms popularity-based ranking by 8.8%.

The results of experiments, in terms of precision BR and mean over precision, show considerable improvement of PCR over all other baseline methods, and this improvement is statistically significant (based on paired t-test with $p < 0.05$).

6.3.4. *Performance evaluation with individual contexts*

Table 8 illustrates the performance of personalized recommendation with individual context factors, that is, temporal context or weather context. We use precision (P) for prediction and MAP@50 to evaluate ranking ability. Results show that exploitation of temporal context concepts produces better recommendation results when compared with weather context concepts.

Table 8. Performance of personalized ranking with different context parameters.

	Precision		MAP@50	
Number of test locations per user	Temporal context (t)	Weather context (w)	Temporal context (t)	Weather context (w)
>5	0.462	0.441	0.319	0.291
>10	0.522	0.480	0.262	0.235
>15	0.533	0.517	0.255	0.214

7. Conclusions and future work

In this article, we put forward an approach to extract semantically meaningful tourist locations from geotagged social media such as photos for tourist travel recommendations. We have contributed a method that applies a collaborative filtering approach to obtain tourist's preferences from his or her publicly contributed photos and takes into account the current context of user for personalized recommendations. We presented the evaluation of our methods on a sample of publicly available photos from the Flickr dataset. It contains metadata of photos taken in various cities in China. Results show that our context-aware personalized method is able to predict tourists' preferences in a new or unknown city more precisely and generate better recommendations compared to other state-of-the-art landmark recommendation methods. We found that people's preferences with short and targeted visits are easier to predict by methods based on popularity. Performance of collaborative filtering methods based on tourist preferences improves in the case of long and real tourist visits. Moreover, considering contexts gives a substantial improvement in the precision of prediction.

This study motivates a number of important directions for further research. We used user-based collaborative filtering for recommendations using users' data with a decent amount of revealed preferences. In real time, when data are incomplete or evolving, a recommender system based on a user-based collaborative filtering model has its own limitations such as new user problem and new location problem. To address these issues, we envisage that other better methods could be investigated to learn user preferences based on user's revealed preferences. For a user with less or unknown preferences, background information such as user features (e.g., age and gender) could be considered or user's stated preferences could be used to initiate the process. Semantic correlations between the locations could be considered to resolve new location problem. Another modification which could be in future versions for better recommendation is to introduce some space–time constraints to the recommendation results, that is, how long it would take to reach the sites, how much time is necessary to visit, and how much time the user has? We also anticipate that the investigation of personalization in combination with context awareness in trips (sequence of locations) for tourist recommendations is valuable. Here, personalization may refer to the types of locations visited, to the trips length, or to the visit rhythm.

Acknowledgments

This research was funded by the Ministry of Industry and Information Technology of China (No. 2010ZX01042-002-003-001), the Natural Science Foundation of China (No. 60703040), the Zhejiang Provincial Natural Science Foundation of China (No. Y107178), and the Department of Science and Technology, Zhejiang Province (Nos. 2007C13019, 2011C13042).

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. *Knowledge and Data Engineering*, *IEEE Transactions on Knowledge and Data Engineering*, 17 (6), 734–749.

Ahern, S., *et al*., 2007. World explorer: visualizing aggregate data from unstructured text in geo-referenced collections. *In*: E.M. Rasmussen, R.R. Larson, E.G Toms and S. Sugimoto, eds., *Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries*, 18–23 June 2007, JCDL '07, Vancouver. New York: ACM, 1–10.

Clements, M., *et al*., 2010. Using flickr geotags to predict user travel behaviour. *In*: F. Crestani, S. Marchand-Maillet, H.H. Chen, E.N. Efthimiadis and J. Savoy, eds., *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, 19–23 July 2010, SIGIR '10, Geneva. New York: ACM, 851–852.

Crandall, D., *et al*., 2009. Mapping the world's photos. *In*: J. Quemada, G. León, Y.S. Maarek and W. Nejdl, eds., *Proceedings of the 18th international conference on World wide web*, 20–24 April 2009, WWW '09, Madrid. New York: ACM, 761–770.

Ester, M., *et al*., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. *In*: E. Simoudis, J. Han, and U.M. Fayyad, eds., *Proceedings of the 2nd international conference on knowledge discovery and data mining*, 2–4 August 1996, KDD '96, Portland, OR: AAAI Press, 226–231.

Hays, J. and Efros, A. A., 2008. IM2GPS: estimating geographic information from a single image. *In*: *Proceedings of IEEE conference on computer vision and pattern recognition*, 24–26 June 2008, *CVPR'08*, Anchorage, AK: IEEE Computer Society, 1–8.

Hecht, B.J. and Gergle, D., 2010. On the "localness" of user-generated content. *In*: K.I. Quinn, C. Gutwin and J.C. Tang, ed., *Proceedings of the 2010 ACM conference on computer supported cooperative work*, 6–10 February 2010, CSCW '10, Savannah, GA New York,: ACM, 229–232.

Horozov, T., Narasimhan, N., and Vasudevan, V., 2006. Using location for personalized POI recommendations in mobile environments. *In*: *Proceedings of the international symposium on applications on Internet applications and the Internet*, 23–27 January 2006, SAINT '06, Phoenix, AZ: IEEE Computer Society, 129.

Kennedy, L., *et al*., 2007. How flickr helps us make sense of the world: context and content in community-contributed media collections. *In*: R. Lienhart, A.R. Prasad, A. Hanjalic, S. Choi, B.P. Bailey and N. Sebe, eds., *Proceedings of the 15th international conference on Multimedia*, 24–27 September 2007, MULTIMEDIA '07, Augsburg. New York: ACM, 631–640.

Kennedy, L. and Naaman, M., 2008. Generating diverse and representative image search results for landmarks. *In*: J. Huai, R. Chen, H.W. Hon, Y. Liu, W.Y. Ma, A. Tomkins and X. Zhang, eds., *Proceedings of the 17th international conference on World Wide Web*, 21–25 April 2008, WWW '08, Beijing. New York: ACM, 297–306.

Kisilevich, S., Mansmann, F., and Keim, D., 2010. P-DBSCAN: a density based clustering algorithm for exploration and analysis of attractive areas using collections of geo-tagged photos. *In*: L. Liao, ed., *Proceedings of the 1st international conference and exhibition on computing for geospatial research & application Article no.* 38, 21–23 June 2010, COM.Geo '10, Washington, DC. New York: ACM,1–38:4.

Lee, D., *et al*., 2010. Exploiting contextual information from event logs for personalized recommendation. *In*: R. Lee, ed. *Computer and information science 2010*, Vol. 317 of s*tudies in computational intelligence*. Berlin: Springer, 121–139.

Linden, G., Smith, B., and York, J., 2003. Amazon.com recommendations: item-to-item collaborative filtering. *Internet Computing*, *IEEE*, 7 (1), 76–80.

Lu, X., *et al*., 2010. Photo2Trip: generating travel routes from geo-tagged photos for trip planning. *In*: A.D. Bimbo, S.F. Chang and A.W.M. Smeulders, eds., *Proceedings of the international conference on Multimedia*, 25–29 October 2010, MM '10, Firenze. New York: ACM, 143–152.

Luo, J., *et al*., 2011. Geotagging in multimedia and computer visiona survey. *Multimedia Tools and Applications*, 51 (1), 187–211.

Naaman, M., 2011. Geographic information from georeferenced social media data. *SIGSPATIAL Special*, 3 (2), 54–61.

Popescu, A. and Grefenstette, G., 2009. Deducing trip related information from Flickr. *In*: J. Quemada, G. León, Y.S. Maarek and W. Nejdl, eds., *Proceedings of the 18th international conference on World Wide Web*, 20–24 April 2009, WWW '09, Madrid.New York: ACM, 1183–1184.

Popescu, A., Grefenstette, G., and Moëllic, P.-A., 2009. Mining tourist information from user-supplied collections. *In*: D.W.L. Cheung, I.Y. Song, W.W. Chu, X. Hu and J.J. Lin, eds., *Proceedings of the 18th ACM conference on information and knowledge management*, 2–6 November 2009, CIKM '09, Hong Kong. New York: ACM, 1713–1716.

Rattenbury, T., Good, N., and Naaman, M., 2007. Towards automatic extraction of event and place semantics from flickr tags. *In*: W. Kraaij, A.P. de Vries, C.L.A. Clarke, N. Fuhr and N. Kando, eds., *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 23–27 July 2007, SIGIR '07, Amsterdam. New York: ACM, 103–110.

Resnick, P. and Varian, H.R., 1997. Recommender systems. *Communications of the ACM*, 40 (3), 56–58.

Sarwar, B., *et al*., 2001. Item-based collaborative filtering recommendation algorithms. *In*: *Proceedings of the 10th international conference on World Wide Web*, 1–5 May 2001, WWW '01, Hong Kong. New York: ACM, 285–295.

Serdyukov, P., Murdock, V., and van Zwol, R., 2009. Placing flickr photos on a map. *In*: J. Allan, J.A. Aslam, M. Sanderson, C. Zhai and J. Zobel, eds., *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, 19–23 July 2009, SIGIR '09, Boston, MA. New York: ACM, 484–491.

Shin, D., *et al*., 2009. Context-aware recommendation by aggregating user context. *In*: B. Hofreiter and H. Werthner, eds., *Proceedings of the 2009 IEEE conference on commerce and enterprise computing*, 20–23 July 2009, CEC '09 Washington, DC: IEEE Computer Society, 423–430.

Sui, D. and Goodchild, M., 2011. The convergence of GIS and social media: challenges for GIScience. *International Journal of Geographical Information Science*, 25 (11), 1737–1748.

Takeuchi, Y. and Sugimoto, M., 2006. CityVoyager: an outdoor recommendation system based on user location history. *In*: J. Ma, H. Jin, L. Yang, and J. Tsai, eds. *Ubiquitous intelligence and computing*, Vol. 4159 of *Lecture Notes in Computer Science*. Berlin: Springer, 625–636.

Wachowicz, M., *et al*., 2011. Finding moving flock patterns among pedestrians through collective coherence. *International Journal of Geographical Information Science*, 25 (11), 1849–1864.

Wang, J., de Vries, A.P., and Reinders, M.J.T., 2006. Unifying user-based and item-based collaborative filtering approaches by similarity fusion. *In*: E.N. Efthimiadis, S.T. Dumais, D. Hawking and K. Jarvelin, eds., *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval*, 6–11 August 2006, SIGIR '06, Seattle, WA. New York: ACM, 501–508.

Yang, Y., Gong, Z., and U, L.H., 2011. Identifying points of interest by self-tuning clustering. *In*: W.Y. Ma, J.Y. Nie, R.A. Baeza-Yates, T.S. Chua and W.B. Croft, eds., *Proceedings of the 34th international ACM SIGIR conference on Research and development in information retrieval*, 25–29 July 2011, SIGIR '11, Beijing. New York: ACM, 883–892.

Yin, Z., *et al*., 2011. Diversified trajectory pattern ranking in geo-tagged social media. *In*: *Proceeding of SIAM international conference on data mining*, 28–30 April 2011, *SDM*, 980–991.

Yoon, H., *et al*., 2011. Social itinerary recommendation from user-generated digital trails. *Personal and Ubiquitous Computing*, 16 (5), 1–16.

Zheng, Y., *et al*., 2009. Mining interesting locations and travel sequences from GPS trajectories. *In*: M. Rappa, P. Jones, J. Freire and S. Chakrabarti, eds., *Proceedings of the 18th international conference on World Wide Web*, 26–30 April 2010, WWW '09, Madrid. New York: ACM, 791–800.

Zheng, V.W., *et al*., 2010. Collaborative location and activity recommendations with GPS history data. *In*: J. Quemada, G. Leon, Y.S. Maarek and W. Nejdl, eds., *Proceedings of the 19th international conference on World Wide Web*, 20–24 April 2009, WWW '10, Raleigh, NC. New York: ACM, 1029–1038.