GoThere: Travel Suggestions using Geotagged Photos

Abdul Majid, Ling Chen, Gencai Chen, Hamid Turab Mirza, Ibrar Hussain
College of Computer Science, Zhejiang University, Hangzhou, 310027, P. R. China
{majid, lingchen, chengc, hamid306, ibrar}@cs.zju.edu.cn

ABSTRACT
We propose a *context* and *preference aware* travel guide that suggests *significant* tourist destinations to users based on their preferences and current surrounding context using contextualized user-generated contents from the social media repository, i.e., Flickr.

Categories and Subject Descriptors
H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms
Algorithms, Experimentation

Keywords
Photo collections, trip planning, context aware query

1. INTRODUCTION
Travel planning is often a difficult and time consuming task for people visiting an unknown destination. The proliferation of digital photo-capture devices equipped with GPS and the growing practice of sharing public photos online using social media sites have resulted in a huge volume of geotagged photos available on the Web. Based on the assumption that tourist attractions are those places that are often photographed, geotags annotated to photos have been exploited for various tasks such as mapping geotags to places[1, 2], and recommendation of travel itineraries [3].

For the selection of locations to visit, tourists may have their own preferences. Furthermore, tourist' interest to visit a location could be affected by his/her current contexts, e.g., time and weather. Using contextualized user-generated contents (photos) from the social media repository such as Flickr, we propose a system that is able to suggest *significant* tourist locations to users based on their *preferences* and current surrounding *context*. Existing methods addressed queries either with free of context constraints or with a few dimensions of context. Our approach has following advantages. 1) We synergize disjointed contexts and sparse social contents together with online information sources to enrich primitive contexts and contents with higher levels of semantic meaning; 2) Context data is categorized to support more complex context based queries; 3) We construct users' interest profiles from their contributed photos to represent their implicit preferences; 4) Locations are ranked for each user considering his/her preferences and location's significance, enabling users to receive recommendations that are more relevant and significant.

2. METHODOLOGY
Figure 1 depicts the different modules comprising our system. We discover locations using spatial proximity of photos and enrich the

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

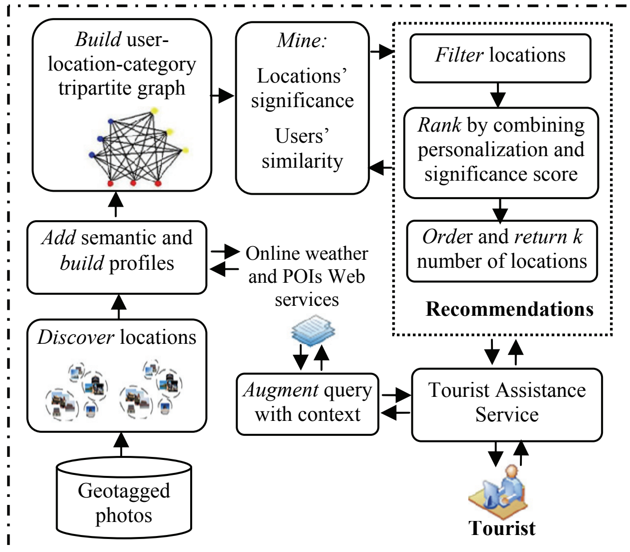

Figure 1. System Architecture

aggregated locations with semantic using textual tags annotated to photos in combination with information provided by online Web services. We build the profiles of locations to describe the contexts in which they have been visited. For temporal context, we exploit the geo and temporal tags annotated to photos. Whereas, to derive weather context we query third party weather Web services to retrieve weather conditions. For making recommendations, first we filter the locations based on contextual constraints, and then rank the locations by combining personalization score and significance score. We define a measure to identify similar users in previously visited cities and aggregate these users' opinions to obtain personalization score for each location in target city for the target user. Significance score for each location is computed using a user-expertise model that exploits the relationships among users, locations, and location categories.

1) Discovering Tourist Locations: To find highly photographed locations as the tourist places in a city, we employ P-DBSCAN[2] to cluster photos using their spatial proximity. Given a collection of photos P, The output of a P-DBSCAN is a set of locations $L=\{l_1, l_2 \ldots l_n\}$. Each element $l=(P_l, g_l)$, where P_l is a cluster of photos and g_l is the geographical coordinates to represent the centroid of location l and is computed from group of geotags annotated to photos in the cluster P_l.

2) Semantic Annotation: We utilize textual tags of photos in cluster P_l, similar as[1] to enrich location l with *"name"*. To infer the *"category"* of location l, we query Web service such as *www.google.com/places* to provide metadata (*name* and *type*) of Point of Interests (POIs) that are within the area with radius $r=200$ *meters* from coordinate g_l. Google Places supports 126 types to describe POIs for search queries. We further generalize these types into 6 categories, i.e., education, shopping, religious, food, transportation, cultural, and entertainment. We select the *"category"* that is highest in frequency in the list of types associated with returned POIs as "*category*" of *l*. After annotation, a tourist location can be represented as $l = (P_l, g_l, name, category)$.

3) Profiling Locations: To build the profile of a location $l = (P_l, g_l, name, category)$, first we identify the visits made to this location by different users. We use photos P_l to infer set of visits V_l for location *l*. As a user can take more than one photo in a same visit, we sort the photos for each user using photos taken time. If the difference between the taken times of two consecutive photos is less than a visit duration threshold $visit_{th}$, we consider both photos belong to the same visit. We use the median of time-stamps of photos that belong to visit *v* as the visit time *v.t*. We use this *v.t* to retrieve weather conditions *w*, when visit *v* was made by user *u* at location *l*. Weather Services normally publish weather data at hourly, daily or monthly level that contain different variable like temperature, precipitation, etc., to describe the weather conditions. We define an abstraction strategy to obtain abstract context concepts from raw contexts, i.e., time stamp and weather variable. E.g., the raw context (21:30, 25C^o) can be abstracted to (night, warm).

Given the set of visits V_l belongs to location *l* with associated context concepts, we consider the context concepts "*popular*" that are higher in frequency. For example, *p(l.w)=(warm-sunny)* depicts that location *l* has been popularly visited in warm and sunny weather conditions. After building the profiles of all locations, we maintain a locations database $LDB=\{l_1, l_2, \ldots, l_n\}$.

4) Building User-Location-Category Graph: We organize three entities (users, locations, and location categories) and relationships (visits) among these entities into a meaningful data structure, i.e., user-location-category tripartite graph $G_{ULC}=(U; L; C; E_{UL}; W_{UL}; E_{UC}; W_{UC}; E_{LC})$, where *U*, *L* and *C* are nodes to represent users, locations and location categories 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. E_{UC} and W_{UC} are sets of edges and edge weights between *U* and *C* to represent users' visits and the number of visits to particular location categories. E_{LC} are edges between *L* and *C* to describe the categories of locations.

5) Mining Significance of Locations: We mine the significance of locations using a user-expertise model, i.e., generate a weighing of location significance through the number of user visits to specific location categories. (1) Given *m* users and *n* categories, we build an $m \times n$ adjacency matrix M_{UC}. Each entry in $M_{UC}(p, q)$, depicts the experience of user u_p in location category c_q. (2) To capture the relationship between users and locations, given *m* users and *k* locations, we build a $m \times k$ adjacency matrix M_{UL}. So in matrix M_{UL}, $I_j=\{i\}$ where $M_{UL}(i, j)\neq 0$, is the set of indices of users who have visited location *j*. We compute significance score of each location $l_j (j=1..k)$, which is of location category $c_q (q=1..n)$ by: $s(l_j) = \sum_{i \in I_j} M_{UC}(i, q)$.

6) Computing Similarity between Users: From matrix M_{UL}, the user u_p's travel preferences can be derived as an array $R_p=< r_{p0}, r_{p1}, \ldots, r_{pn}>$, where r_{pj} is u_p's implicit rating(visits) in a location *j*. $S(R_p)$ is the subset of R_p, $\forall r_{pj} \in S(R_p)$, $r_{pj}\neq 0$, i.e., the set of locations that have been visited by u_p. The average rating in R_p is denoted as $\overline{R}_p$. We compute the similarity between users u_p and u_q using Pearson correlation metric as shown in Equation 1.

$$sim(u_p, u_q) = \frac{\sum_{j \in S(R_p) \cap S(R_q)} (r_{pj} - \overline{R}_p) \cdot (r_{qj} - \overline{R}_q)}{\sqrt{\sum_{j \in S(R_p) \cap S(R_q)} (r_{pj} - \overline{R}_p)^2} \cdot \sqrt{\sum_{j \in S(R_p) \cap S(R_q)} (r_{qj} - \overline{R}_q)^2}} \quad (1)$$

7) Recommendations: Processing of context aware query *Q(t, w)* made by user u_p proceeds as a two-step approach: an initial filtering step retrieves locations of the target city from locations database *LDB* and eliminates tourist locations that do not meet contextual constraints given in the query, thus producing a filtered set of tourist locations *L'*. In the second step: personalization score for each location l_j is linearly combined with its significance score to generate final rank score as: $R(l_j) = (1-\theta)s(l_j) + \theta\,(p(l_j))$. Next the locations in *L'* are ranked using rank score *R* and top k locations are returned as query result. Personalization score is computed using similarities among u_p and other users in previously visited cities by aggregating opinions of similar users as: $p(l_j) = \sum_{u_q \in U} sim(u_p, u_q) \, . \, r_{qj}$.

3. EXPERIMENTS AND RESULTS

Dataset: We use public API of Flickr (*www.flickr.com*) to collect 736383 geotagged photos that were taken in six different cities of China between January 01, 2001 and July 1, 2011. Historical weather data of these cities is collected using public API of an online weather Web service *www.wunderground.com*.

Ground Truth: For evaluation, we select users that have visited at least two distinct cities $\{C_o, C_t\} \in C$. To evaluate only those users who have provided a decent amount of preference information, we consider users who have visited at least 5 locations in "*training*" city C_o.

Evaluation Methodology: We predict the locations actually visited by test user $u_p \in U$ in a "*test*" city C_t, based on preferences derived from the locations visited by that user in "*training*" city C_o. We use visits made by him/her to tourist locations in "*test*" city C_t to obtain; (1) the number of relevant locations denoted as *k*, and (2) the 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 personalized context aware method and other baseline methods. To evaluate recommendation methods for user u_p, we match the recommended list with the actual list of locations visited by him/her in the "*test*" city C_t.

Baselines, Metrics, and Results: As baseline methods, we use (1) Popular-Rank, i.e., ranking locations based on of unique number of visits, and (2) Classic-Rank that exploits the relationship between users and locations to apply HITS based inference model to rank locations according to their authority scores[4]. Table 1 depicts the effectiveness of our and baseline methods in terms of Mean Average Precision (MAP@50) and Precision (P@1) for the users who have visited at least 5 locations in C_t. Results show that, our proposed approach is able to predict user's travel preferences more precisely and generates better recommendations as compared to baselines (paired t-test with $p < 0.05$).

Table 1. Performance Comparison

	Popular-Rank	Classic-Rank	Our-Approach
P@1	0.478	0.447	0.596
MAP@50	0.329	0.308	0.435

4. REFERENCES

[1] L. Kennedy, M. Naaman, S. Ahern, R. Nair, and T. Rattenbury. How flickr helps us make sense of the world: context and content in community-contributed media collections. In *Proc. ACM MM'07*.

[2] S. Kisilevich, F. Mansmann, and D. Keim. P-DBSCAN: a density based clustering algorithm for exploration and analysis of attractive areas using collections of geo-tagged photos. In *Proc. COM.Geo '10*. ACM Press, 2010.

[3] M. D. Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Constructing travel itineraries from tagged geo-temporal breadcrumbs. In *Proc. WWW'10*. ACM Press, 2010.

[4] Y. Zheng, L. Zhang, X. Xie, and W-Y Ma. Mining interesting locations and travel sequences from GPS trajectories. In *Proc. WWW' 09*. ACM Press, 2009.