



Mining Social Media photos to Measure Point-of-Interest Connectivity

M. SHAFI⁺⁺, N. ULLAH, M. I. KHATTAK, I.A. KHAN*, M. SALEEM, N. MINALLAH

Department of CS, University of Engineering and Technology, Peshawar

Received 12th August 2014 and Revised 28th January 2015

Abstract-The capability of social media to allow users to upload their geo-tagged and text tagged photos has led to the creation of a voluminous data, which can be mined to extract useful information. Most of these photos are public and can be accessed and analyzed for extracting useful patterns. This paper presents a prediction of Point of Interest (POI) connectivity through mining these publically available photos on social media. Two POI are considered connected for a particular user, if the user has taken photos of the two points, within certain duration of time. The proposed research can be utilized in (1) building a recommender system that would help the tourists planning their travel sequence and (2) helping the tourism management authorities to predict the tourism statistics of the nearby POI once they have collected statistics of a particular POI. This information can be utilized in efficient resource management for sustainable tourism.

Keywords: Social Media, Mining, Region-of-interest, Geo-tagged, Tourism

1. INTRODUCTION

Most social media applications offer photo-uploading services to users. The users can tag their photos public or private. It has been observed that most of the photos related to points of interest are kept public and hence they could easily be accessed through the Application Programmer Interfaces (APIs). These photos can be analyzed to extract useful information. This paper presents a framework for extracting the POI connectivity through mining these publically available social media photos. Such a framework could be utilized in various ways for tourism improvement such as (1) building an intelligent recommender system for tourists by combining the proposed work with some Artificial Intelligence planning techniques (Sirin, *et al.* 2004; Saleem, *et al.* 2013) and (2) helping the tourism management authorities to predict the tourism statistics of the nearby POI once they have collected statistics of a particular POI which will help in efficient resource management for sustainable tourism.

Tourism is usually performed in new and unfamiliar requirements. In such situations, it is a promising idea to inform the tourists about the interest level they might have in a certain location. A visitor might like or dislike to visit two points together, based on the distance between them, travel time, hurdles and traffic conditions, nature (whether the two points are related or not) and various other parameters. While distance between the two points is easily available, the rest of the parameters are not easily available for most of the places. Hence there is a need for an automated system that could suggest a connectivity-measurement (which could potentially be a combination of all these parameters) of the two points to the tourist. This paper

analyzes the publically available photos on Flickr (<http://www.flickr.com/services/api/>) to find the connectivity between two points. The basic principle for connectivity measurement is straightforward: if a user has taken a photo of point A, and has also taken a photo of point B at the same time or same day; we assume that he has visited both points at the same day.

The rest of the paper is organized as follow: Section 2 states the related work from literature, Section 3 gives an overview of Flickr (an image tagging website) APIs, Section 4 explains the structure of the database for storing different parameters of the photos, Section 5 details the connectivity measurement algorithm, Section 6 presents a case study and Section 7 concludes the paper with future directions.

2. RELATED WORK

Zheng *et al.* have carried out a comprehensive literature survey on the utilization of geo-tagged multimedia (Zheng, *et al.* 2011). They have categorized the geo-tagged multimedia utilization research into four broad categories: organization and summarization of photos (Smith 1996; Camara and Raper 1999; Jung 1999; Lim, *et al.* 2002; Toyama, *et al.* 2003; Naaman, *et al.* 2004; Pigeau and Gelgon 2004; Jaffe, *et al.* 2006; Heuer and Dupke 2007), mining knowledge from geo-tagged multimedia (Jaffe, *et al.* 2006; Kennedy, *et al.* 2007; Hile, *et al.* 2008), learning landmarks (Kennedy and Naaman 2008; Abbasi, Chernov *et al.* 2009), and geographic location estimation from photos (Hays and Efros 2008). The proposed work falls in the category of knowledge mining from geo-tagged data. Following is an overview of some of the closely related existing work.

⁺⁺Corresponding Author email: M. Shafi, shafi@uetpeshawar.edu.pk,

*COMSATS Institute of Information Technology, Abbottabad.

(Zheng, *et al.* 2012) extract the tourist traffic flow information of different POIs from the Flickr database using Markov Chain Model. (Kurata 2012) introduced a method named Potential-Of-Interest. They have developed a mobile-based tool to provide users with potential-of-interest maps, which are a kind of visual aid to the user and help the tourists in decision-making. Five different methods for data collection have been compared in their work, out of which they found photo geo-tagging as the most promising. The geo-tagged photos are uploaded directly from that particular POI. (Kurata *et al.* 2012) suggest that only GPS log that shows the time spent by a visitor by staying near the attraction is not the best estimate of the visitor's interest and enjoyment. Along with the GPS log, videos were recorded at each point to investigate the relationship between tourist activities and their GPS logs. It was concluded that location history of tourists is insufficient knowledge for estimating how much a tourist enjoyed an exhibition. They argued that viewing speed is also necessary for making an estimate about an exhibition's level of attractiveness. (Asakura and Iryo 2007) analyzed behaviors of different tourists by tracking data gathered using mobile instruments (cell phone, GPS, RF ID tags). They suggested a way of finding topological characteristics of tourists' behavior through mining data collected with mobile devices. They used cluster analysis to find out the topological characteristics, measured on Route Topology Index (RTI) for the real data of KOBE (a port city in Japan). The behavioral data connected with mobile instruments are the sequential "dot" in space-time dimensions. "Dot" have various features namely longitude, latitude and time-of-day information. Dots in space-time dimension are combined in a sequence to show the behaviors of a tourist. (Crandall, *et al.* 2009) mined 35 million photos dataset collected from Flickr, and showed that the visual and temporal features improve the accuracy of the estimation of the location of a photo, as compared to the textual information. They demonstrate how to predict locations from visual, textual and temporal features of the photos, to organize a large collection of photos, and to highlight exciting properties about popular locations in the world. (De Choudhury, *et al.* 2010) suggest the construction of travel itineraries from the geo-tagged data using a two-step approach. Photos of a particular user are collected, firstly to find the time a user stayed at a POI, and secondly to find the transit time between the two places. This information for different users is aggregated into a POI graph to suggest travel itineraries to new tourists. They carried out extensive user studies to evaluate the quality of the generated itineraries. They concluded that the itineraries yielded promising results. (Yanai, Kawakubo *et al.* 2009) utilized the geo tagged photos on the image sharing web applications for predicting the relationship between word concepts and

geographical locations. Their first method used geographic location entropy and photo region entropy for the analysis of visual features of an image and location. In the second method, photographs that represent world dimensions were selected. These representative images help in finding cultural differences across the world regarding word concepts with high geo-location entropy. (McKercher and Lew 2004) provide a detailed study of tourist models that have been developed to depict the movement of tourists between destinations. These models indicate that a tourist's movement has a destination and a transit component. They argue that the tourists' movement is affected by time, budget, number of intervening opportunities available, trip variables and socio cultural makeup of the tourists. In (Zheng, Zhang *et al.* 2009), different trajectories by many users generated through GIS, intersecting locations (culturally important places) and classical travel sequences of users are mined. Such information can help in identifying the co-relation between location and tourist (user) and enabling travel recommendations and guidance. A tree-based hierarchical graph (TBHG) is used to model multiple individuals' location histories and Hyper Text Induced Topic Search (HITS) model is used to deduct the user experience and location of interest. Later the travel sequence is also inferred from the user experience and locations of interest. They collected GPS data from 107 users over a period of one year in the real world, and evaluated the system against this large GPS dataset. The results show that the performance of HITS-based inference model is better than the other approaches in the literature.

It's worthwhile to explain two aspects of the proposed research. Firstly, the knowledge mined through social media photos is not completely accurate, however, it gives coarse estimation and this has been the case in almost all the cited research. Secondly, almost all of the closely related work (discussed in the above paragraphs) is user-oriented i.e. a single photographer's trajectories are traced while the proposed work is place-oriented i.e. two places are chosen and their connectivity is calculated based on the proposed scheme.

3. MATERIAL AND METHODS

The Flickr API

Flickr provides a rich set of APIs for uploading and accessing its photo database and the related information of the photos. These APIs are freely available for any non-commercial usage. Flickr supports *REST*, *XML-RPC* and *SOAP* request formats and supports *REST*, *XML-RPC*, *SOAP*, *JSON* and *PHP* response formats. Some of the most common Flickr functionalities that have been utilized in this research are as follows:

Flickr. Photos. Search

This function returns the photos specified by the criteria in the input arguments. Only the photos accessible to that user (including all public photos) are checked for the input criteria and the matching photos are returned.

The input arguments for this function are: *api_key, user_id, tags, tag_mode, text, min_upload_date, max_upload_date, min_taken_date, max_taken_date, license, sort, privacy_filter, bbox, accuracy, safe_search, content_type, machine_tags, machine_tag_mode, group_id, contacts, woe_id, place_id, media, has_geo, geo_context, lat, lon, radius, radius_units, is_commons, in_gallery, is_getty, extras, per_page, page, min_upload_date, max_upload_date, min_taken_date, max_taken_date* are some of the most important arguments for the proposed research. These arguments are used to specify the upload dates and taken dates between which the photos are being searched. (Fig. 1) shows sample JSON response from the search method.

Flickr. Photos. Getinfo

All available details of a particular photo can be obtained using this method. The input argument for this function are: *api_key* (required), *photo_id* (required) and *secret* (optional). The returned parameters are *photo->id, secret, server, isfavourite, license, rotation, original_secret, original_format, owner->nsid, username, realname, location, title, description, visibility, dates-> posted, taken, takengrularity, lastupdated, permissions, editibility, comments, notes, tags* and *urls*.

Further details and examples of these methods can be obtained in the documentation of the Flickr APIs (<http://www.flickr.com/services/api/>).

4. THE PHOTO INFORMATION DATABASE

A database was created for storing the photo information once they have been retrieved from Flickr. This database is used for further analysis and for extraction of connectivity measurement. The database comprises the following two tables:

PhotoInfo: *ID, photoID, ownerID, placeName, thumbnailURL, imageURL, ownerUserName, owner RealName, ownerLocation, title, description, datePosted, dateTaken*

Connectivity: *user_id, place_id1, place_id2, daily Connectivity, weekly Connectivity, monthly Connectivity, all Time Connectivity.*

A row in the Connectivity table with values “1023, place003, place007, true” means that places with id’s *place003* and *place007* are connected for user 1023. The daily connectivity, weekly connectivity, monthly connectivity and all-time connectivity defined as below.

Daily Connectivity: Two places are considered daily connected for a particular user if the user has taken photos of both the places on the same day.

Weekly Connectivity: Two places are considered weekly connected for a particular user if the user has taken photos of both the places in the same week.

Monthly Connectivity: Two places are considered monthly connected for a particular user if the user has taken photos of both the places in the same month.

All-time Connectivity: Two places are considered all-time connected for a particular user if the user has ever taken photos of both the places.

5. CONNECTIVITY MEASUREMENT

The proposed algorithm is straightforward which is based on the principle: if a user has taken photos of two different POI together that suggest that most probably he or she has visited both the places together and we mark these two locations connected for that particular user. The algorithm works as follow: first photos of a particular POI are retrieved from the Flickr then it is searched whether photos of other locations (for which connectivity is measured) have been taken by the same user or not and hence an overall connectivity metric is calculated between different regions of interests. The pseudo code for the proposed algorithm for measuring connectivity is as follow:

Input: Two POI A and B

Output: Percentage connectivity measurement between the two POI

photoIDsofA = search POI A photos on Flickr using Flickr.photos.search

totalPhotos = 0;

connectedPhotos = 0;

for each *id* in *photoIDsofA*

ownerID = get *ownerID* through Flickr.photo.getInfo function and passing the *id* as *photoID*

dateTaken = get *dateTaken* through Flickr.photo.getInfo function and passing the *id* as *photoID*

photosOfB = get photos of B using Flickr.photos.search and passing *ownerID* and *dateTaken* as argument

if (*photosOfB* is not empty)

increment *connected Photos* by 1;

increment *totalPhotos* by 1;

end for

connectivity = (*connectedPhotos*/*totalPhotos*)*100

Here the *totalPhotos* is the total number of unique photos for POI A and *connectedPhotos* represents the subset of *totalPhotos* for which photos of POI B also exist.



Fig.1: Sample JSON Response from FLICKR API



Fig. 2: Geo-mapped locations of selected points

Table 1: Connectivity measures

FROM\TO	London Eye	Madame Tussauds	National Gallery	National History Museum	Science Museum	Tate Modern	The British Museum	Tower of London	V & Albert Museum
London Eye	NA	D:0 W:0 M:2 A:12	D:0 W:0 M:3 A:6	D:0 W:2 M:2 A:5	D:7 W:9 M:11 A:14	D:3 W:4 M:4 A:13	D:1 W:3 M:4 A:8	D:2 W:5 M:8 A:10	D:1 W:3 M:3 A:7
Madame Tussauds	D:3 W:5 M:8 A:12	NA	D:3 W:6 M:7 A:10	D:1 W:3 M:3 A:6	D:2 W:2 M:5 A:7	D:2 W:5 M:5 A:7	D:3 W:6 M:6 A:10	D:5 W:7 M:10 A:14	D:1 W:4 M:5 A:8
National Gallery	D:2 W:3 M:6 A:9	D:4 W:8 M:8 A:12	NA	D:13 W:16 M:19 A:23	D:2 W:3 M:3 A:12	D:2 W:3 M:3 A:6	D:4 W:7 M:10 A:12	D:10 W:13 M:16 A:23	D:2 W:5 M:7 A:8
National History Museum	D:3 W:3 M:5 A:7	D:2 W:3 M:5 A:7	D:12 W:17 M:19 A:22	NA	D:18 W:18 M:20 A:32	D:3 W:3 M:7 A:10	D:14 W:17 M:19 A:20	D:13 W:15 M:16 A:19	D:15 W:16 M:18 A:24
Science Museum	D:12 W:15 M:17 A:23	D:2 W:2 M:6 A:8	D:0 W:2 M:2 A:15	D:16 W:18 M:18 A:28	NA	D:0 W:3 M:3 A:6	D:13 W:17 M:17 A:35	D:7 W:9 M:12 A:13	D:6 W:6 M:6 A:13
Tate Modern	D:7 W:10 M:10 A:34	D:0 W:0 M:0 A:0	D:2 W:2 M:2 A:8	D:4 W:4 M:7 A:11	D:2 W:3 M:5 A:7	NA	D:0 W:0 M:0 A:22	D:7 W:10 M:12 A:16	D:0 W:1 M:3 A:7
The British Museum	D:2 W:5 M:8 A:13	D:3 W:5 M:6 A:9	D:5 W:7 M:9 A:12	D:14 W:18 M:19 A:19	D:15 W:17 M:20 A:37	D:1 W:2 M:3 A:15	NA	D:9 W:10 M:13 A:16	D:12 W:13 M:17 A:23
Tower of London	D:2 W:6 M:8 A:13	D:5 W:8 M:10 A:15	D:11 W:13 M:17 A:23	D:9 W:11 M:14 A:16	D:8 W:9 M:12 A:13	D:5 W:8 M:10 A:13	D:5 W:7 M:9 A:12	NA	D:2 W:4 M:7 A:9
V & Albert Museum	D:2 W:6 M:7 A:9	D:2 W:5 M:5 A:9	D:3 W:5 M:8 A:10	D:13 W:14 M:15 A:20	D:4 W:8 M:9 A:15	D:0 W:1 M:2 A:5	D:13 W:13 M:18 A:24	D:4 W:5 M:8 A:11	NA

6. **LONDON ATTRACTIONS: A CASE STUDY**

Nine regions of interests from London city were selected as case study for the proposed connectivity measurement algorithm. These nine locations (London Eye, Madame Tussauds, National Gallery, National History Museum, Science Museum, Tate Modern, The

British Museum, Tower of London, V& Albert Museum) are among the most visited places by the tourists in London city. Fig. 2 shows the geo-mapped locations of these London attractions. One thousand unique photos (maximum of one photo per single POI per single user) were retrieved for each POI and were

stored in the database described in the earlier section. While calculating connectivity, more photos were then retrieved in real time against these stored photos.

The following table shows the percentage daily connectivity (represented by D), weekly connectivity (represented by W), monthly connectivity (represented by M) and all time connectivity (represented by A) among the selected attraction points in London city.

Following are some interesting facts from the analyses of the above table.

Connectivity A-B is not always equal to connectivity B-A. As shown in the table above the connectivity from a point A to point B is not always equal to connectivity from point B to point A. The reason is that connectivity A-B means that all the people who have visited point A are selected and then only those are selected among them who have also visited point B. Similar definition goes for connectivity B-A

Connectivity to London Eye is greater than that of connectivity from London Eye to all other attraction points. This may be due to the fact that London Eye is an open air place and is visible from quite far, so the visitors who visit other places, most often take photo of London Eye as well because of its visibility from far.

The connectivity of British Museum and Natural History museum is relatively high in spite of reasonable distance between these two places. This could be due to the same nature of these two places.

7. CONCLUSION

This paper presents a technique for the prediction of point of interest connectivity. The prediction is based on mining photos of different locations and regions uploaded by users. We calculate a connectivity measurement to show the level of connectivity between two regions.

The above technique is useful for analysis of tourists' data. For example, a transport company might like to issue a combined ticket package for highly connected regions. Similarly, the connectivity measure will also help tourists to visit related places e.g. if most of the users who visited London Eye also visited Buckingham Palace, then it might be a good choice for the tourists to visit. As stated in the Introduction section, the proposed scheme may be utilized for (1) building an intelligent recommender system and (2) efficient resource management for sustainable tourism.

The presented research has some limitations for example: Flickr uses the text tag information which may

decrease the accuracy of the proposed scheme in case of wrongly tagged photos. This shortcoming may be overcome by removing wrongly tagged photos through intelligent image processing techniques as future work.

REFERENCES:

Abbasi, R., S. Chernov, (2009). Exploiting flickr tags and groups for finding landmark photos. Advances in Information Retrieval, Springer: 654-661.

Asakura, Y. and T. Iryo (2007). "Analysis of tourist behaviour based on the tracking data collected using a mobile communication instrument." Transportation Research Part A: Policy and Practice41(7): 684-690.

Camara, A. S. and J. Raper (1999). Spatial multimedia and virtual reality, CRC Press, Taylor and Francis London.

Crandall, D. J., L. Backstrom, (2009). Mapping the world's photos. Proceedings of the 18th international conference on World wide web, ACM, New York USA, 961-970. .

DeChoudhury, M., M. Feldman, (2010). Automatic construction of travel itineraries using social breadcrumbs. Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, ACM, New York USA, 35-44.

Hays, J. and A. A. Efros (2008). IM2GPS: estimating geographic information from a single image. IEEE Conference on Computer Vision and Pattern Recognition, Anchorage AK, 1-8.

Heuer, J. T. and S. Dupke (2007). "Towards a spatial search engine using geotags." GI-Days, young researchers conference, Insititue of Geoinformatics, 199-204.

Hile, H., R. Vedantham, (2008). Landmark-based pedestrian navigation from collections of geotagged photos. Proceedings of the 7th international conference on mobile and ubiquitous multimedia, ACM, New York USA, 145-152.

Jaffe, A., M. Naaman, (2006). Generating summaries and visualization for large collections of geo-referenced photographs. Proceedings of the 8th ACM international workshop on Multimedia information retrieval, ACM, New York USA, 89-98.

Jung, V. (1999). Metaviz: Visual interaction with geospatial digital libraries. INVISIP-Information Visualisation for Site Planning, International Computer Science Institute Report.

- Kawase, J., Y. Kurata, (2012). When and Where Tourists are Viewing Exhibitions: Toward Sophistication of GPS-Assisted Tourist Activity Surveys. Information and Communication Technologies in Tourism, Springer: 415-425.
- Kennedy, L., M. Naaman, et al. (2007). How flickr helps us make sense of the world: context and content in community-contributed media collections. Proceedings of the 15th international conference on Multimedia, ACM, New York, USA, 631-640.
- Kennedy, L. S. and M. Naaman (2008). Generating diverse and representative image search results for landmarks. Proceedings of the 17th international conference on World Wide Web, ACM, Beijing China, 297-306.
- Kurata, Y. (2012). "Potential-of-interest maps for mobile tourist information services." Proceedings of Information and Communication Technologies in Tourism, Sweden, 248-255.
- Li, X., C. Wu, (2008). Modeling and recognition of landmark image collections using iconic scene graphs. Computer Vision–ECCV Springer: 427-440.
- Lim, E.-P., D. H.-L. Goh., (2002). G-Portal: a map-based digital library for distributed geospatial and georeferenced resources. Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries, ACM.
- McKercher, B. and A. A. Lew (2004). "Tourist flows and the spatial distribution of tourists." A companion to tourism: 36-48.
- Naaman, M., Y. J. Song, (2004). Automatic organization for digital photographs with geographic coordinates. Proceedings of the 2004 Joint ACM/IEEE Conference on Digital Libraries, 53-62.
- Pigeau, A. and M. Gelgon (2004). "Organizing a personal image collection with statistical model-based ICL clustering on spatio-temporal camera phone meta-data." Journal of Visual Communication and Image Representation **15**(3): 425-445.
- Saleem, M., P. W. Chung, (2013). "A cross organisation compatible workflows generation and execution framework." Knowledge-Based Systems, 56, 1-14.
- Sirin, E., B. Parsia, (2004). "HTN planning for web service composition using SHOP2." Web Semantics: Science, Services and Agents on the World Wide Web **1**(4): 377-396.
- Smith, T. R. (1996). "A digital library for geographically referenced materials." Computer **29**(5): 54-60.
- Toyama, K., R. Logan, (2003). Geographic location tags on digital images. Proceedings of the eleventh ACM international conference on Multimedia, ACM. 156-166.
- Zheng, Y. T., Z. J. Zha, and T. S. Chua, (2011). Research and applications on georeferenced multimedia: a survey. Multimedia Tools and Applications, **51**(1), 77-98.
- Zheng, Y. T., Z. J. Zha, and T. S. Chua, (2012). Mining travel patterns from geotagged photos. ACM Transactions on Intelligent Systems and Technology (TIST), **3**(3), 56.