

Constructing Travel Itineraries from Tagged Geo-Temporal Breadcrumbs

Munmun De Choudhury*
Arizona State University
Tempe, AZ, USA
mdechoud@asu.edu

Moran Feldman*
Technion - Israel Inst. of Tech.
Haifa, Israel
moranfe@cs.technion.ac.il

Sihem Amer-Yahia
Yahoo! Research
New York, NY, USA
sihem@yahoo-inc.com

Nadav Golbandi
Yahoo! Research
Haifa, Israel
nadavg@yahoo-inc.com

Ronny Lempel
Yahoo! Research
Haifa, Israel
rlempel@yahoo-inc.com

Cong Yu
Yahoo! Research
New York, NY, USA
congyu@yahoo-inc.com

ABSTRACT

Vacation planning is a frequent laborious task which requires skilled interaction with a multitude of resources. This paper develops an end-to-end approach for constructing intra-city travel itineraries automatically by tapping a latent source reflecting geo-temporal breadcrumbs left by millions of tourists. In particular, the popular rich media sharing site, Flickr, allows photos to be stamped by the date and time of when they were taken, and be mapped to Points Of Interest (POIs) by latitude-longitude information as well as semantic metadata (e.g., tags) that describe them.

Our extensive user study on a “crowd-sourcing” marketplace (Amazon Mechanical Turk), indicates that high quality itineraries can be automatically constructed from Flickr data, when compared against popular professionally generated bus tours.

Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications-Data mining; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Algorithms, Experimentation

Keywords

Flickr, geo-tags, mechanical turk, media applications, orienteering problem, rich media, travel itinerary.

1. INTRODUCTION

Travel itinerary planning is often a difficult and time consuming task for a traveler visiting a destination for the first time. It involves substantial research to identify points of interest (POIs) worth visiting, the time worth spending at each point, and the time it will take to get from one place to another. Without any prior knowledge, one must either

*Part of this research was performed while visiting Yahoo! Research.

rely on (1) travel books, (2) personal travel blogs, or (3) a combination of online resources and services such as travel guides, map services, public transportation sites, and human intelligence to piece together an itinerary.

All these options have shortcomings. Travel books do not cover all cities/locations and, perhaps more importantly, are not free. Personal travel blogs reflect a single person’s view, with no guarantees provided over the writer’s experience or the amount of preparation invested in planning the trip. Finally, compiling an itinerary by selecting individual POIs and researching their to’s and fro’s is a task which is both time consuming and requires significant search expertise.

In this paper, we develop an approach to *automatically* construct travel itineraries at a large scale from photos uploaded by users. More specifically, by analyzing streams of photos taken by users, one can deduce the cities visited by a person, which POIs that person took photos at, how long that person spent at each POI, and what the transit time was between POIs visited in succession. Each such itinerary is comprised of a sequence of POIs, with recommended *visit times* and approximate *transit times* between them. In summary, we make the following contributions:

1. We introduce a novel end-to-end approach that starts with the analysis of latent information reflected in social media sharing sites, and ends with the synthesis of practical information in the form of travel itineraries.
2. As an initial implementation of our approach, we apply a pipeline of multiple heuristics that together extract reliable granular evidence of individual tourists’ trips to a destination from Flickr photos.
3. We aggregate the individual trips to form a graph representing collective touristic behavior, and adapt a solution of the Orienteering problem to efficiently generate intra-city travel itineraries from the graph.

2. RELATED WORK

Our work integrates the two emerging fields of touristic data analysis and touristic information synthesis, and is therefore related to various works in these two fields. For the former, there are a number of studies on analyzing landmark (i.e., POI) visitation patterns from geo-spatial and temporal evidences left by travelers [2, 5, 6]. However, to the exception of [7], those works generally avoid synthesizing or recommending new paths and instead focus solely on the

analysis itself. For the latter, a number of other works construct and recommend tourist itineraries at various granularities [3, 4]. They rely, however, on structured and cleansed data on landmarks, and do not deal with the challenge of analyzing and extracting from noisy data. Our work is also tangentially related to other vast fields such as visualizing geo-spatial data, tracking movements based on sensor networks, and constraint optimization.

3. OUR APPROACH

The first step of our approach is to convert the raw user photos into individual *timed paths* for a given city. Intuitively, these paths, which connect various POIs, are constructed from individual photo streams and describe the movements of individual tourists. The process has three main challenges: (i) pruning irrelevant photos that are not associated with the city of interest or not owned by a *tourist*; (ii) mapping photos to the POIs of the city, and (iii) constructing individual *timed paths*. Each timed path is a sequence of POIs traversed by a user, annotated with the time spent by the user at each POI and the transit times between pairs of successive POIs.

We emphasize here that: 1) while our study focus on leveraging information from a particular rich media sharing site, Flickr, the work is easily extensible to any other social repository, where users can share semantically and geo-temporally tagged rich media; 2) while we process the internal Yahoo! Flickr data repository, the same protocol can essentially be followed by using the open Flickr API.

Given the set of timed paths, our goal is to aggregate the actions of many individual travelers into coherent itineraries while taking into consideration POI popularity. To this effect, we define represented timed paths as a graph and formulate the problem of finding an itinerary between two points given a time constraint. We reduce this problem to the directed Orienteering problem and use a restatement of Chekuri and Pál’s algorithm [1].

4. OUR FINDINGS

We evaluate the quality of travel itineraries constructed by our system in an extensive user study conducted through the Amazon Mechanical Turk (AMT)¹ system. An example one-day itinerary generated by our method is shown in Figure 1. Our experimental study elicited feedback from 250 workers on AMT in order to validate our system’s ability to generate high quality travel itineraries for popular touristic cities, including Barcelona, London, New York City (NYC), Paris, and San Francisco. The questionnaire evaluated diverse aspects of our system generated itineraries such as its overall usefulness as well as its relevance in terms of the transit and visit times to each POI.

We show that users perceive our *automatically* generated itineraries as being as good as (or even slightly better than) itineraries provided by *professional* tour companies. Furthermore, we show that users are satisfied with the recommended transit and visit times for the POIs within the itineraries.

5. CONCLUSION

This paper addressed the question of automatic generation of travel itineraries for popular touristic cities from

Time 09:00	: Start from ground zero
Time 09:00	: Spend 27 minutes at ground zero .
Time 09:27	: Transit to empire state building (estimated travel time: 52 minutes)
Time 10:19	: Spend 1 hour and 13 minutes at empire state building .
Time 11:32	: Transit to new york public library (estimated travel time: 15 minutes)
Time 11:47	: Spend 29 minutes at new york public library .
Time 12:16	: Transit to radio city music hall (estimated travel time: 24 minutes)
Time 12:43	: Spend 51 minutes at radio city music hall .
Time 13:34	: Transit to central park (estimated travel time: 23 minutes)
Time 13:57	: Spend 40 minutes at central park .
Time 14:37	: Transit to rockefeller center (estimated travel time: 33 minutes)
Time 15:10	: Spend 37 minutes at rockefeller center .
Time 15:47	: Transit to grand central terminal (estimated travel time: 22 minutes)
Time 16:09	: Spend 27 minutes at grand central terminal .
Time 16:36	: Transit to chrysler building (estimated travel time: 6 minutes)
Time 16:42	: Spend 31 minutes at chrysler building .
Time 17:13	: Transit to brooklyn bridge (estimated travel time: 32 minutes)
Time 17:45	: Spend 36 minutes at brooklyn bridge .
Time 18:21	: Transit to statue of liberty (estimated travel time: 21 minutes)
Time 18:42	: Spend 42 minutes at statue of liberty .
Time 19:24	: Transit to little korea (estimated travel time: 26 minutes)
Time 19:50	: Spend 31 minutes at little korea .
Time 20:21	: Transit to ground zero (estimated travel time: 38 minutes)

Figure 1: Sample one-day itinerary constructed by our system for the city NYC.

large-scale user contributed rich media repositories. We plan to explore many directions such as applying different filtering and aggregation techniques to accommodate different types of travelers, and constructing “off the beaten track” itineraries that cater to niche audiences rather than mainstream crowds.

6. REFERENCES

- [1] Chandra Chekuri and Martin Pál. A recursive greedy algorithm for walks in directed graphs. In *FOCS*, pages 245–253, 2005.
- [2] David Crandall, Lars Backstrom, Daniel Huttenlocher, and Jon Kleinberg. Mapping the world’s photos. In *Proc. 18th International World Wide Web Conference (WWW’2009)*, pages 761–770, April 2009.
- [3] David Leake and Jay Powell. Mining large-scale knowledge sources for case adaptation knowledge. In *Proc. ICCBR 2007*, pages 209–223, 2007.
- [4] David Leake and Jay Powell. Knowledge planning and learned personalization for web-based case adaptation. In *Proc. ECCBR 2008*, pages 284–298, 2008.
- [5] Adrian Popescu and Gregory Grefenstette. Deducing trip related information from flickr. In *Proc. 18th International World Wide Web Conference (WWW’2009)*, pages 1183–1184, April 2009.
- [6] Tye Rattenbury, Nathaniel Good, and Mor Naaman. Toward automatic extraction of event and place semantics from flickr tags. In *Proc. 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR’07)*, pages 103–110, July 2007.
- [7] Chih Hua Tai, De Nian Yang, Lung Tsai Lin, and Ming Syan Chen. Recommending personalized scenic itinerary with geo-tagged photos. In *Proc. IEEE International Conference on Multimedia and Expo (ICME’2008)*, pages 1209–1212, 2008.

¹<https://www.mturk.com/>