

Recommending Tours and Places-of-Interest based on User Interests from Geo-tagged Photos

Kwan Hui Lim^{*†}

^{*}Department of Computing and Information Systems, The University of Melbourne, Australia

[†]Victoria Research Laboratory, National ICT Australia, Australia

limk2@student.unimelb.edu.au

Supervised by: Shanika Karunasekera^{*}, Christopher Leckie^{*†}, and Jeffrey Chan^{*}

Expected Graduation Date: 2018

ABSTRACT

Photo sharing sites like Flickr and Instagram have grown increasingly popular in recent years, resulting in a large amount of uploaded photos. In addition, these photos contain useful meta-data such as the taken time and geo-location. Using such geo-tagged photos and Wikipedia, we propose an approach for recommending tours based on user interests from his/her visit history. We evaluate our proposed approach on a Flickr dataset comprising three cities and find that our approach is able to recommend tours that are more popular and comprise more places/points-of-interest, compared to various baselines. More importantly, we find that our recommended tours reflect the ground truth of real-life tours taken by users, based on measures of recall, precision and F1-score.

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, Measurement

Keywords

Tour Recommendation, Travel Itinerary, User Interests, Orienteering Problem, Flickr, Wikipedia, Social Networks

1. INTRODUCTION

The prevalence of GPS-enabled camera-phones and photo sharing sites (such as Flickr and Instagram) facilitate users to share geo-tagged photos of interesting places they have visited. The sharing of such photos are increasingly popular in recent years, as illustrated by the 8B existing photos and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMOD'15 PhD Symposium, May 31, 2015, Melbourne, Victoria, Australia.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3529-4/15/05 ...\$15.00.

<http://dx.doi.org/10.1145/2744680.2744693>.

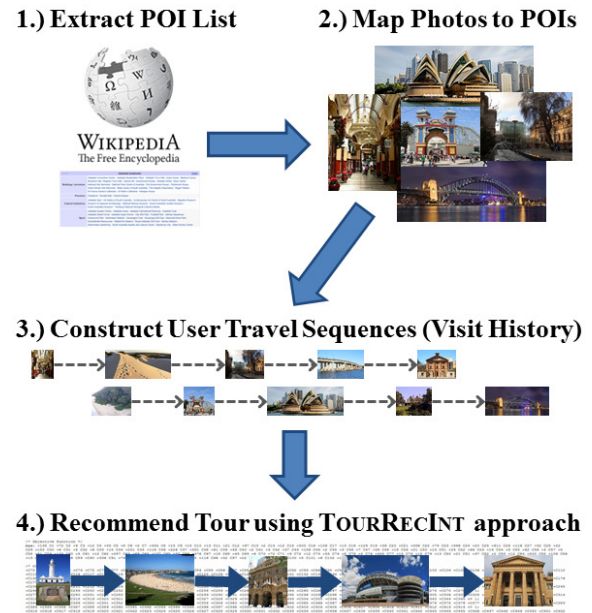


Figure 1: Overall Experimental Framework

3.5M new daily uploads in Flickr [28]. These geo-tagged photos also provide an abundance of location-based information, which can be used to improve the recommendation of tours and Places/Points-of-Interest (POI) to visit. For example, many researchers have utilized these geo-tagged photos to determine the travel history of users and recommend tours based on these travel histories [8, 15, 7, 16, 25].

In this work, we aim to build upon the field of tour recommendation by also considering a user's interest in specific categories of POIs, based on their past travel sequences. Specifically, our main contributions are:

- Outlining the tour recommendation problem in the context of the Orienteering problem [29, 30], and proposing the TOURRECINT approach for tour recommendation based on user interests (Section 3).
- Implementing a framework (Fig. 1) to construct the travel sequences of a user based on his/her geo-tagged photos on Flickr and the list of POIs on Wikipedia (Section 3.2).

- Evaluating our TOURRECINT approach against various baselines using a Flickr dataset that comprises geo-tagged photos taken in three cities (Section 5).

The rest of the paper is structured as follows. Section 2 discusses some related work, while Section 4 describes our experimental methodology. We elaborate on our other related contributions in Section 6, before summarizing the paper in Section 7.

2. RELATED WORK

There is a large body of work that uses geo-tagged photos to recommend tours [8, 15, 7, 16, 25]. In particular, [8] was one of the earlier works to utilize geo-tagged photos to recommend routes in the context of the Orienteering Problem, which involves recommending a budget-constrained tour with a specific starting and destination POI. [15, 16] used a combined topic and Markov model to recommend tours that are based on a user’s interest and recently visited POIs. Using geo-tagged photos, [25] was able to recommend routes that were deemed more beautiful than the baseline shortest route, based on crowd-sourced quantitative and qualitative evaluations. While these earlier works successfully use geo-tagged photos to recommend tours, [8] and [25] do not consider user interests or POI categories, and [15, 16] model user interests using a topic model but do not recommend tours that adhere to a specific starting and destination POI.

More recent work have considered the categories of POIs in recommending tours [6, 23, 4, 11]. Gionis et al. [11] extended the Orienteering Problem for tour recommendation, with the additional constraint that tours have to adhere to a specific POI category visit order (e.g., restaurant → park → shopping → restaurant). On the other hand, the work in [4] is based on the Generalized Maximum Coverage problem [9] but aims to maximize a multi-objective function of POI popularity and user-POI interest. [23] adopts a different approach and recommends tours using random walks on a graph where the nodes are POIs and edges are weighted based on the frequency of common user visits and POI categories. [6] describes a tour planning tool that includes must-visit POI categories but requires users to explicitly state their interests, whereas we implicitly infer interests from their visit history. Although our work is based on the Orienteering Problem, we do not assume any specific POI category visit order or maximize a multi-objective function, but instead we use a must-visit POI category that is based on user interests from his/her visit history.

Location prediction is another research area that is closely related to tour recommendation, in particular location prediction that is based on user interests. Both [12] and [22] determine user interests based on the time and categories of POIs visited, with [12] employing a topic model and [22] using matrix factorization for location prediction. Based on 68 features such as unique POI categories visited and most visited POI categories, [2] performs location prediction using Ranking Support Vector Machines [13] and Gradient Boosted Regression Trees [32].

3. PROPOSED APPROACH

We first frame the tour recommendation problem in the context of the Orienteering problem. Thereafter, we elaborate on our proposed approach to recommend tours based

on user interests, and describe the main steps of our experimental framework.

3.1 Problem Definition

We define our tour recommendation problem based on the Orienteering Problem [29] and use its integer problem formulation from [30, 19], with an additional constraint for a must-visit category based on user interest. Given a set of POIs P , starting POI $p_1 \in P$, and destination POI $p_N \in P$, we want to recommend a tour $T = (p_1, \dots, p_N)$ that adheres to a distance budget B , while maximizing the overall profit of POIs in recommended tour T . Formally, we aim to optimize the following objective function:

$$\text{Max} \sum_{i=2}^{N-1} \sum_{j=2}^N x_{i,j} \text{Pop}(i) \quad (1)$$

such that:

$$\sum_{j=2}^N x_{1,j} = \sum_{i=1}^{N-1} x_{i,N} = 1 \quad (2)$$

$$\sum_{i=1}^{N-1} x_{i,k} = \sum_{j=2}^N x_{k,j} \leq 1, \quad \forall k = 2, \dots, N-1 \quad (3)$$

$$\sum_{i=1}^{N-1} \sum_{j=2}^N \text{Dist}(i, j) x_{i,j} \leq B \quad (4)$$

$$\sum_{i=1}^{N-1} \sum_{j=2}^N x_{i,j} \delta(\text{Cat}_i = c_m) \geq 1, \quad \forall c_m \in C \quad (5)$$

where:

$$x_{i,j} = \begin{cases} 1, & \text{if we visit POI } i, \text{ followed by POI } j \\ 0, & \text{otherwise} \end{cases}$$

$$\delta(\text{Cat}_i = c_m) = \begin{cases} 1, & \text{if } \text{Cat}_i = c_m \text{ (POI } i \text{ is of category } c_m) \\ 0, & \text{otherwise} \end{cases}$$

Eqn. 1 shows the objective function that maximizes the total popularity of all POIs in the recommended tour, where $\text{Pop}(p)$ measures the popularity of POI p based on the total number of visits to POI p . Constraint 2 ensures that the tour starts and ends at POI P_1 and POI P_N , respectively. Constraint 3 ensures tour/path connectivity and that no POIs are re-visited. Constraint 4 ensures that the total distance travelled between consecutive POIs is within the distance budget B (using the function $\text{Dist}(i, j)$ that measures the distance between POI i and POI j).¹

More recently, authors such as [4, 11] further grouped POIs into different categories (e.g., museums, parks, etc). Similarly, we adopt a set of POI interest categories C and represent each POI $p \in P$ by a unique ID, POI name, category, and latitude/longitude coordinates. Given that Cat_p denotes the category of POI p , Constraint 5 ensures that the recommended tour includes at least one visit to a POI belonging to POI category $c_m \in C$. In the next section,

¹While we adopt a simple representation of budget using distance, this can be easily generalized to other representations such as travel time by different transport modes and even include the POI visit duration.

we further elaborate on Constraint 5 and how we determine the POI category $c_m \in C$. In addition, we also included the constraints for sub-tour elimination as described in [24].

3.2 Experimental Approach

Our proposed TOURRECINT approach is based on the Orienteering Problem (outlined in Section 3.1), with additional consideration for user interests based on his/her visit history. In addition to the constraints of a starting POI $p_s \in P$, destination POI $p_d \in P$ and distance budget B , we define an additional constraint of a must-visit POI category $c_m \in C$, as highlighted in Constraint 5. Some examples of must-visit POI categories include Sports, Parks, Entertainment and Shopping. This constraint ensures that the recommended tour T contains at least one POI of the category $c_m \in C$. Specifically, we define the category $c_m \in C$ as the POI category which the user has most frequently visited in his/her other travel sequences. Finally, we solve this tour recommendation problem as an integer program, using the Ipsolve linear programming package [3]. User interest can also be easily generalized to other definitions such as using POI visit times or a relative interest weighting, which we intend to explore as future work.



Figure 2: Places-of-Interests from Wikipedia

As illustrated in Fig. 1, our experimental framework comprises the following steps:

1. **Get POI List.** Extract list of POIs, latitude/longitude coordinates, and interest categories from Wikipedia. Fig. 2 shows an example of a Wikipedia listing of POIs in Melbourne and the corresponding category and coordinates of each POI, which we use for this step.
2. **Get User-POI Visits.** Map Flickr photos to the extracted list of POIs if their coordinates differ by $<100m$ based on the Haversine formula (for earth/spherical distances) [26].
3. **Build User Travel Sequences.** Construct real-life user travel sequences based on the User-POI Visits (from Step 2). Specifically, we group a set of User-POI Visits as an unique travel sequence if the User-POI Visits differ by less than 8 hours.
4. **Recommend Tours.** Use our TOURRECINT approach to recommend tours and evaluate the results based on the real-life travel sequences as ground truth.

More details on our evaluation methodology are given next in Section 4.

4. EXPERIMENTAL METHODOLOGY

In this section, we describe the dataset used in our experiments and elaborate on the various evaluation metrics.

4.1 Dataset

Our experiments were conducted on the Yahoo! Flickr Creative Commons 100M dataset [31]. This dataset comprises 100M photos and videos that were posted on Flickr, accompanied by their relevant meta information such as the date/time taken, latitude/longitude coordinates and geographic accuracy. For our experiments, we only consider photos with the highest geographic accuracy.

Table 1: Description of Dataset

City	No. of POIs	No. of Photos	# POI Visits	# Travel Sequen.
Adelaide	65	26,665	4,022	948
Melbourne	96	61,605	14,607	2,780
Sydney	150	53,832	20,876	3,628

Using this dataset, we then extracted photos that were taken in three major Australian cities, namely: Adelaide, Melbourne and Sydney. Table 1 shows more details regarding this dataset. As described in Section 3.2, we first obtained a list of POIs from Wikipedia (Step 1), then mapped these photos to User-POIs visits (Step 2), next we constructed the user travel sequences (Step 3), before finally evaluating our proposed approach (Step 4).

4.2 Evaluation Metrics

Our evaluation is performed using the following metrics:

1. **Total POIs in Tour:** The total number of POIs recommended in the tour.
2. **Tour Popularity:** The total popularity of all POIs recommended in the tour, where POI popularity is the number of times a POI is visited.
3. **Tour Recall:** The recall of POIs recommended in the tour, defined as: $\frac{|P_r \cap P_v|}{|P_v|}$, where P_r and P_v are the set of POIs recommended in the tour and visited by the user in real-life, respectively.²
4. **Tour Precision:** The precision of POIs recommended in the tour, defined as: $\frac{|P_r \cap P_v|}{|P_r|}$, where P_r and P_v are the set of POIs recommended in the tour and visited by the user in real-life, respectively.
5. **Tour F1-score:** The harmonic mean of both precision and recall, defined as: $\frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$.

Metrics 1 and 2 reflect the objectives of a typical tourist, which are to visit the largest number of POIs and visit the most popular POIs. Metrics 3 to 5 are standard evaluation metrics used in the Information Retrieval field, which we adopt for evaluating our recommended tours to determine how well they perform against the real-life travel sequences that users embark on.

For determining the must-visit interest category $c_m \in C$ in our proposed TOURRECINT (Section 3), we use leave-one-out cross-validation [14] (i.e., when we evaluate a travel

²We approximate the real-life tours (travel sequences) of users based on the photos they have taken.

sequence of a particular user, we define $c_m \in C$ as the POI category which the user has visited the most in his/her other travel sequences). Similarly, we use the starting/destination POIs and distance covered in these real-life travel sequences as input to TOURRECINT.

5. RESULTS AND DISCUSSION

In this section, we discuss some results on the distribution of POI visits. Thereafter, we describe the various baseline algorithms used and discuss the results of our proposed TOURRECINT compared to these baselines.

5.1 Heavy-tailed Distribution of POI Visits




Figure 3: Distribution of POI Visit Count

As shown in Fig. 3, we observed a heavy-tailed distribution for the POI visits in our dataset. In particular, we find that the top 40 POIs (in terms of visit count) represents 99.4%, 87.9% and 88.5% of all POI visits in Adelaide, Melbourne and Sydney respectively. Given the large representation and popularity of these top 40 POIs, we focus on the top 40 POIs of each city for our tour recommendation experiments.

5.2 Baseline Algorithms

In our evaluation, we compare our proposed approach with various baselines, as follows:

- **Greedy Nearest (GREEDNEAR).** From a starting POI $p_s \in P$, iteratively select the *nearest, unvisited POI* as the next POI to visit.
- **Greedy Most Popular (GREEDPOP).** From a starting POI $p_s \in P$, iteratively select the *most popular, unvisited POI* as the next POI to visit.
- **Random Selection (RANDOM).** From a starting POI $p_s \in P$, iteratively select a *random, unvisited POI* as the next POI to visit.

All three baselines will iteratively select a next POI to visit, until either: (i) destination POI $p_d \in P$ is reached; or (ii) distance budget B is exceeded.

Similar to the evaluation of TOURRECINT, we evaluate the three baselines using the real-life travel sequences of users. For each travel sequence with ≥ 3 visited POIs, we use the starting/destination POIs and distance covered in these travel sequence as input to TOURRECINT and the baselines. Thereafter, we measure the performance of each algorithm based on the metrics described in Section 4.2, and repeat the evaluation for all travel sequences with ≥ 3 POIs.

5.3 Tour Recommendation Results

Fig. 4 shows that our proposed TOURRECINT generally out-performs all three baselines (GREEDNEAR, GREEDPOP, and RANDOM) in terms of total POIs, total popularity, precision, recall and F1-score, for all three cities. We now discuss these results in greater detail.

Compared to the baselines, TOURRECINT recommends tours that comprise more POIs and are more popular in most cases³, addressing the typical objectives of tourists to visit as many POIs as possible, with a preference for the more popular ones. GREEDNEAR offers the second best performance as it favours the nearest POI, thus consuming less budget (distance) and is able to cover more POIs. Conversely, GREEDPOP is biased towards the most popular POI regardless of distance. However, reaching this POI typically consumes a large proportion of GREEDPOP’s budget, thus rendering it unable to visit more POIs. Unsurprisingly, RANDOM provides the worst overall performance based on tour popularity.

In terms of recall, TOURRECINT offers the best performance by a large margin, compared to the three baselines. One contributing factor is the consideration of user interest by TOURRECINT, as users are more likely to visit places that they are interested in [1]. On the other hand, the baseline algorithms do not consider user interest, thus resulting in a poorer performance.

TOURRECINT also performs the best in terms of precision, followed by GREEDPOP and GREEDNEAR, with RANDOM performing the worst. Similarly, the results show that TOURRECINT offers the best performance in terms of F1-score, followed by GREEDNEAR, GREEDPOP and RANDOM. The strong performance of TOURRECINT in recall, precision and F1-score shows that TOURRECINT is able to recommend tours that accurately reflect the ground truth of real-life user travel sequences.

6. OTHER RELATED CONTRIBUTIONS

During my PhD, we have also worked on the related areas of predicting next check-in location [18] and detecting location-centric communities [17], which serve as initial steps for our future work in dynamic tour recommendation and group-based tour recommendation, respectively. We discuss about these related contributions next.

6.1 Predicting Next Check-in Location

Location prediction and POI recommendation are two closely related fields and one of our earlier contribution was in improving the prediction accuracy of a user’s next check-in location on location-based social networks (LBSNs) [18]. Our work improved the Social Historical Model (SHM) [10], which uses a language model [27] to predict a user’s next check-in location based on this user’s previous check-ins and that of his/her friends (i.e., social links).

In [18], we showed that LBSN users exhibit a *recency preference* where they are more likely to re-visit recently visited places than those visited in the distant past. Also, we introduced *place-links*, which are links where the two users share a friendship and a common daily check-in. Using two Foursquare LBSN datasets, we showed how the incorpora-

³Except for Sydney where GREEDNEAR recommends more POIs. However, GREEDNEAR under-performs TOURRECINT in terms of precision, recall and F1-score.




Figure 4: Total POI, popularity, recall, precision and F1-score (1st to 5th column) of tours recommended for the Adelaide, Melbourne and Sydney (1st to 3rd row) datasets. For each graph, the x-axis shows the algorithms evaluated, namely: TourRecInt, GreedNear, GreedPop and Random (left to right).

tion of both *recency preference* and *place-links* can improve the performance of next check-in prediction over the original SHM and various baseline location prediction algorithms.

6.2 Detecting Location-centric Communities

One of our future research direction is to recommend different tours for tourists based on whether they are traveling alone or as part of a group. This distinction is important as a tourist might want to visit certain POIs if he/she was travelling alone, but may visit a different set of POIs if he/she is travelling with his/her family or friends (a group). As a progressive step, we first aim to identify groups or communities of users who tend to visit the same set of places [17].

In [17], we proposed an approach to find such location-centric communities by augmenting social ties (i.e., friendships) with temporal and spatial information. We evaluated our proposed approach on two Foursquare datasets and were able to detect communities of users that display strong similarities in terms of the places they visit and reside in. Previously, we have also found that communities of users with the same interests tend to share strong similarities in their residential city [20]. Similarly, other researchers have found that city-level social networks comprise a large number of user triads that visit common places [5].

7. CONCLUSION AND FUTURE WORK

We examined tour recommendation in the context of the Orienteering problem, and developed the TOURRECINT approach for recommending tours based on user interest. Our

proposed framework involves first using geo-tagged photos and a POI list to construct user visit sequences, then using TOURRECINT for tour recommendation based on these user visit sequences. TOURRECINT is based on a variant of the Orienteering problem, with an additional constraint for a must-visit category based on user interest (i.e., the most visited POI category). Using a Flickr dataset across three cities, we evaluated TOURRECINT against various baselines and observe that TOURRECINT recommends tours that are more popular and comprise more POIs. More importantly, we find that TOURRECINT is able to recommend tours that reflect the ground truth of real-life travel sequences, as indicated by high values of recall, precision and F1-score.

In this work, we focus on one aspect of tour recommendation by considering user interest in specific POI categories. For future work, we intend to explore the following:

- We intend to personalize tour recommendations based on both relative user interest and POI visit durations. Like our earlier work on Twitter [21], we first determine user interest in a specific category, relative to his/her interests in other categories. Thereafter, we personalize tours by recommending POIs related to the user’s interests and tailoring the recommended POI visit duration based on the level of user interest.
- Another future direction is to recommend tours based on whether a user is travelling alone or part of a bigger group (e.g., a couple or family). As each member of the group will have their own unique preferences, the

main challenge is in aligning the individual preferences for such group tours.

- As travel plans are subject to changes due to various circumstances (e.g., bad weather, human fatigue, POI/road closure, traffic congestion), another possibility for future work is to develop dynamic tour recommendation algorithms that consider these changing context during the course of a pre-planned tour.

Acknowledgments. National ICT Australia (NICTA) is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. The author thanks Shanika Karunasekera, Christopher Leckie, Jeffrey Chan and the anonymous reviewers for their useful comments.

8. REFERENCES

- [1] J. Bao, Y. Zheng, and M. F. Mokbel. Location-based and preference-aware recommendation using sparse geo-social networking data. In *Proc. of SIGSPATIAL'12*, pages 199–208, 2012.
- [2] R. Baraglia, C. I. Muntean, F. M. Nardini, and F. Silvestri. Learnnext: learning to predict tourists movements. In *Proc. of CIKM'13*, 2013.
- [3] M. Berkelaar, K. Eikland, and P. Notebaert. Ipsolve: Open source (mixed-integer) linear programming system, 2004. <http://lpsolve.sourceforge.net/>.
- [4] I. Brilhante, J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso. Where shall we go today?: planning touristic tours with tripbuilder. In *Proc. of CIKM'13*, pages 757–762, 2013.
- [5] C. Brown, A. Noulas, C. Mascolo, and V. Blondel. A place-focused model for social networks in cities. In *Proc. of SocialCom'13*, pages 75–80, 2013.
- [6] L. Castillo, E. Armengol, E. Onaindía, L. Sebastián, J. González-Boticario, A. Rodríguez, S. Fernández, J. D. Arias, and D. Borrajo. SAMAP: An user-oriented adaptive system for planning tourist visits. *Expert Systems with Applications*, 34(2), 2008.
- [7] A.-J. Cheng, Y.-Y. Chen, Y.-T. Huang, W. H. Hsu, and H.-Y. M. Liao. Personalized travel recommendation by mining people attributes from community-contributed photos. In *Proc. of MM'11*, pages 83–92, 2011.
- [8] M. D. Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Automatic construction of travel itineraries using social breadcrumbs. In *Proc. of HT'10*, pages 35–44, 2010.
- [9] R. Cohen and L. Katzir. The generalized maximum coverage problem. *Info. Process. Lett.*, 108(1), 2008.
- [10] H. Gao, J. Tang, and H. Liu. Exploring social-historical ties on location-based social networks. In *Proc. of ICWSM'12*, pages 114–121, 2012.
- [11] A. Gionis, T. Lappas, K. Pelechrinis, and E. Terzi. Customized tour recommendations in urban areas. In *Proc. of WSDM'14*, pages 313–322, 2014.
- [12] B. Hu, M. Jamali, and M. Ester. Spatio-temporal topic modeling in mobile social media for location recommendation. In *Proc. of ICDM'13*, 2013.
- [13] T. Joachims. Training linear SVMs in linear time. In *Proc. of KDD'06*, pages 217–226, 2006.
- [14] R. Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proc. of IJCAI'95*, pages 1137–1145, 1995.
- [15] T. Kurashima, T. Iwata, G. Irie, and K. Fujimura. Travel route recommendation using geotags in photo sharing sites. In *Proc. of CIKM'10*, 2010.
- [16] T. Kurashima, T. Iwata, G. Irie, and K. Fujimura. Travel route recommendation using geotagged photos. *Knowledge and information systems*, 37(1), 2013.
- [17] K. H. Lim, J. Chan, C. Leckie, and S. Karunasekera. Detecting location-centric communities using social-spatial links with temporal constraints. In *Proc. of ECIR'15*, pages 489–494, 2015.
- [18] K. H. Lim, J. Chan, C. Leckie, and S. Karunasekera. Improving location prediction using a social historical model with strict recency context. In *Proc. of CaRR'15*, 2015.
- [19] K. H. Lim, J. Chan, C. Leckie, and S. Karunasekera. Personalized tour recommendation based on user interests and points of interest visit durations. In *Proc. of IJCAI'15*, pages 1778–1784, 2015.
- [20] K. H. Lim and A. Datta. Tweets beget propinquity: Detecting highly interactive communities on twitter using tweeting links. In *Proc. of WI-IAT'12*, 2012.
- [21] K. H. Lim and A. Datta. Interest classification of twitter users using wikipedia. In *Proc. of WikiSym'13 + OpenSym'13*, pages 22:1–22:2, 2013.
- [22] X. Liu, Y. Liu, K. Aberer, and C. Miao. Personalized point-of-interest recommendation by mining users' preference transition. In *Proc. of CIKM'13*, 2013.
- [23] C. Lucchese, R. Perego, F. Silvestri, H. Vahabi, and R. Venturini. How random walks can help tourism. In *Proc. of ECIR'12*, pages 195–206, 2012.
- [24] C. E. Miller, A. W. Tucker, and R. A. Zemlin. Integer programming formulation of traveling salesman problems. *Journal of the ACM*, 7(4):326–329, 1960.
- [25] D. Quercia, R. Schifanella, and L. M. Aiello. The shortest path to happiness: Recommending beautiful, quiet, and happy routes in the city. In *Proc. of HT'14*, pages 116–125, 2014.
- [26] R. W. Sinnott. Virtues of the Haversine. *Sky and telescope*, 68(158), 1984.
- [27] Y. W. Teh. A hierarchical Bayesian language model based on Pitman-Yor processes. In *Proc. of COLING/ACL'06*, pages 985–992, 2006.
- [28] The Verge. The man behind Flickr on making the service 'awesome again', 2013. <http://www.theverge.com/2013/3/20/4121574/flickr-chief-markus-spiering-talks-photos-and-marissa-mayer>.
- [29] T. Tsiligrirides. Heuristic methods applied to orienteering. *Journal of the Operational Research Society*, 35(9):797–809, 1984.
- [30] P. Vansteenwegen, W. Souffriau, and D. V. Oudheusden. The orienteering problem: A survey. *Euro. Jour. of Operational Rsch.*, 209(1):1–10, 2011.
- [31] Yahoo! Webscope. Yahoo! Flickr Creative Commons 100M dataset (YFCC-100M), 2014. <http://webscope.sandbox.yahoo.com/catalog.php?datatype=i&did=67>.
- [32] Z. Zheng, H. Zha, T. Zhang, O. Chapelle, K. Chen, and G. Sun. A general boosting method and its application to learning ranking functions for web search. In *Proc. of NIPS'07*, pages 1697–1704, 2007.