Time-Dependent Recommendation of Tourist Attractions using Flickr

Steven Van Canneyt [a]   Steven Schockaert [b]   Olivier Van Laere [a]

Bart Dhoedt [a]

[a] *Department of Information Technology, IBBT, Ghent University, Belgium*
[b] *Department of Applied Mathematics and Computer Science, Ghent University, Belgium*

Abstract

We propose a system that recommends tourist attractions based on the moment that the user visits a given city. We start from a large collection of georeferenced photos on Flickr, and use Mean Shift clustering to determine points of interest within a city. We then estimate the probability that a random user would visit a given place within a given temporal context. This system is compared against a baseline system whose only criterion is the overall popularity of the place. Our experimental results show that significant improvements over this baseline can be obtained.

1   Introduction

Geographic location plays an increasingly important role on the Web 2.0. As an important example, currently there are about 150 million[1] photos on Flickr[2] that have been annotated with geographical coordinates. This data can be leveraged in several ways, yielding rich geographical information that is well beyond what we can find in traditional gazetteers. For instance, [12] demonstrates how names and locations of places can be discovered from Flickr, by looking for correlations between the tags that users assign to their photos and geographic location. Another way in which the coordinates of Flickr photos can be used, is to analyze the behavior of Flickr users. We pursue this latter strategy, with the aim of recommending tourist attractions in a city.

Planning a tourist trip to an unknown city may be time-consuming, hence there is a clear interest in systems that can point the user to potential points of interest. Making the simplifying assumption that interesting places are those places that are often photographed, from georeferenced Flickr photos we can readily compile a list of the most interesting places within a given city. In a subsequent step, we may identify those places in the list that are relevant to the given user (i.e., personalization) in the given context (i.e., context-awareness). While the personalization aspect has already been investigated [4], the influence of context factors such as time has not yet received much attention. Nonetheless, it is intuitively clear that such context factors may influence where we want to go: on a sunny summer day we may prefer to go to a park, while a rainy autumn day may rather lead us to a museum. In particular, in this paper we investigate the influence of the month of the year, the day of the week, and the hour of the day on the popularity of tourist attractions.

The paper is structured as follows. In the next section, we provide an overview of related work. In Section 3 we subsequently explain how places-of-interest can be identified from georeferenced Flickr photos, and how their popularity can be assessed. Section 4 then discusses how points of interest can be ranked according to their suitability within a given context. This section also presents our experimental results, which demonstrate that the considered context factors have a statistically significant, but small, impact on the quality of recommendations. This observation is in line with the results of [4] on the personalization of recommendations from Flickr, where it was also found that in this setting people conform more to global

---

[1] See `http://www.flickr.com/map/`, accessed on June 25, 2011.
[2] `http://www.flickr.com/`

popularity rankings than in traditional recommendation settings, as e.g. in the area of e-commerce. These findings suggest that statistical techniques alone might not suffice for this setting. Section 5 elaborates on this, and presents our conclusions.

2   Related work

Different sources of information can be used to acquire insight in the behavior of users in general, and tourists in particular [7]. Traditionally, such information has been obtained from questionnaires [11]. More recently, data collected from GPS devices [19, 10, 23] and mobile phones [1, 16, 13] has often been used as well. In this paper, as another source of information, we use geographic information that has explicitly been disclosed by users on the Web 2.0. Although several websites with a large number of georeferenced resources exist, Flickr has been most often used for this purpose.

The idea of using Flickr to identify points-of-interest (POIs) has already been used by various authors [5, 2, 20, 8, 6, 14]. In [5] the POIs are determined which tourists have found most interesting, while [2] rather deals with the problem of determining POIs whose visual features resemble that of a photo or textual description provided by the user. The task of recommending routes along POIs has been addressed in [20, 8, 6, 14], dealing with issues such as time constraints or explicit user feedback.

The techniques that are used for recommending POIs are often based on collaborative filtering [23, 4], although some context-aware systems have been proposed as well. For instance, [18] takes the location and history of the user into account: the more recently a given place has been visited by the user, the less likely a place from the same semantic category will be recommended. Other systems [3, 9] take into account background information about tourist attractions, such as opening times. This differs from the approach we take in this paper, as we do not assume that semantic background information is available, and we aim at determining the best time to visit an attraction (within its opening times, if appropriate). Along similar lines, [10] analyzes the San Francisco nightlife, based on GPS logs from taxis. The use of Flickr as a source of information, however, has the potential of developing a system at world-scale.

3   Identifying points of interest using Flickr

We collected our dataset by downloading photo metadata (user id, geotag, date and time taken) from the photo-sharing site Flickr. We crawled all georeferenced photos from the top 100 Best Travel Destinations according to traveleye[3] that were taken between January 2000 and September 2010 and which contain a geotag with street level precision. In addition, whenever a user has taken more than one photo on the same place, all but one were removed. This is because we want to determine how many tourists have visited a place and not how much photos were taken at a place. As we are particularly interested in tourist behavior we also remove all photos from users that are not recognized as tourists. Similar as in [4], a user was qualified as tourist if all her photos in one city were taken in at most two periods of 14 days, with at least 1 month between both periods. The dataset thus obtained contains 664 330 photos and was split in three parts: the photos of two thirds of the users were used as training data to determine POIs (called the training set R, 443 942 photos), the photos of one sixth of the users were used to find optimal values of the parameters in our methods (called the development set D, 109 798 photos), and the remaining sixth was used for evaluation (called the test set T, 110 590 photos).

We have then clustered the geotags of the photos in the training set to identify POIs. The Mean Shift algorithm [21] is particularly suitable for this task, as it is scalable, does not require us to specify the number of clusters, and allows us to adapt the scale at which clusters should be identified. Moreover, Mean Shift clustering has already been successfully applied to detect POIs from Flickr photos [5]. We used a Gaussian Kernel for a smooth density estimation. After running the Mean Shift algorithm, we obtained a number of clusters, each of which is assumed to correspond to a POI. However, as some clusters may span a wide area, only the most centrally located photos of the clusters will be associated to this POI. In particular, a POI is characterized as a circle, whose center is determined by the Mean Shift algorithm, and whose radius is initially equal to the 80th percentile of the distances from the geotags in the cluster to its center. The radius of a POI is then further decreased until the circle does not contain any geotags from other clusters anymore. This process is illustrated in Figure 1, which shows the geotags (blue dots) of three different clusters. The figure shows the initial radius of one cluster (the outer circle), whose geotags are located in the shaded area.

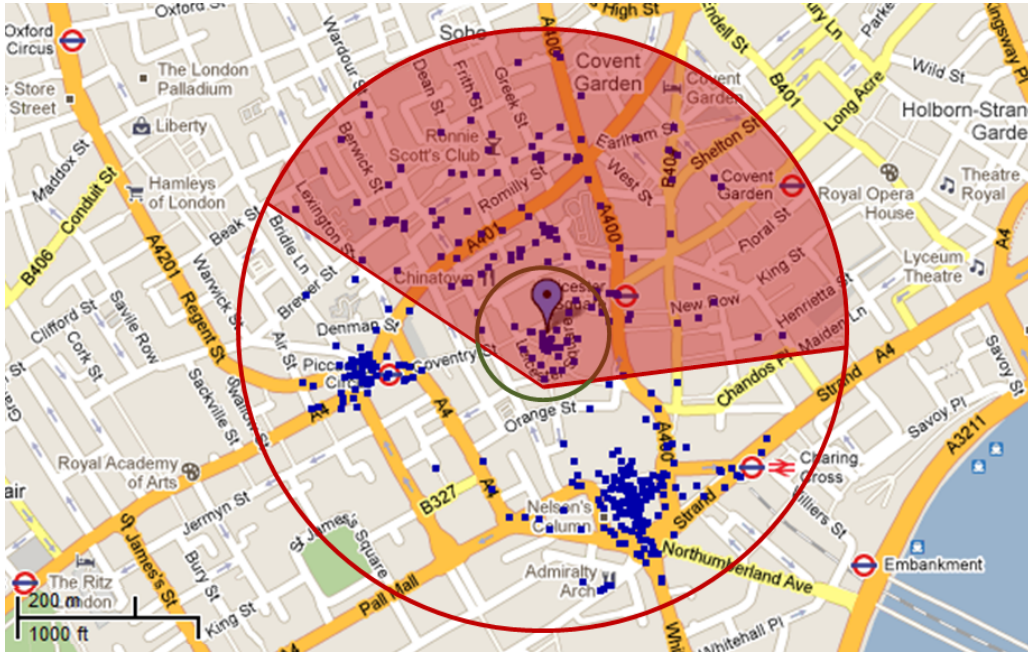[3]http://www.traveleye.com/client/top100.php, accessed in September 2010.

Figure 1: Determining the appropriate radius for the point of interest corresponding to Leicester Square.

Nr.	City	# POIs (N)	# Photos	# Users
1.	London, England, United Kingdom	63	131 904	26 170
2.	New York, NY, United States	62	137 874	26 035
3.	Paris, France	50	109 527	17 923
4.	Berlin, Bundesland Berlin, Germany	30	56 798	9 040
5.	Barcelona, Catalunia, Spain	28	48 841	8 734
6.	Chicago, Illinois, United States	28	45 122	9 864
7.	Washington, District of Columbia, United States	27	44 057	8 442
8.	Rome, Lazio, Italy	23	53 823	8 624
9.	Madrid, Comunidad de Madrid, Spain	20	22 541	5 421
10.	Vancouver, British Columbia, Canada	16	13 843	3 288
	total:	347	664 330	123 541

Table 1: Cities from our dataset where more than 15 POIs were found in which at least 60 users have taken a photo.

The inner circle corresponds to the final radius. It can be verified that this inner circle indeed only contains geotags from the shaded area.

For each city we only retained the POIs where at least 60 users have taken a picture. We evaluated our algorithms on those cities where the number of such POIs was more than 15. This resulted in the 10 cities that are listed in Table 1. Figure 2 displays the 10 most popular POIs that were found in London. These POIs correspond to 1: Big Ben (3050 users); 2: London Eye (2173 users); 3: National Gallery (2018 users); 4: Tower of London (1972 users); 5: Tate Modern (1757 users); 6: Piccadilly Circus (1342 users); 7: Buckingham Palace (1326 users); 8: St Paul's Cathedral (1143 users); 9: British Museum (1085 users); 10: London Bridge (664 users).

4 Recommending points of interest

4.1 Recommendation as a ranking problem

In this section, we consider the task of ranking all POIs within a city, according to their likelihood of being of interest to a user. As a baseline system, we rank the POIs according to their overall popularity. To this
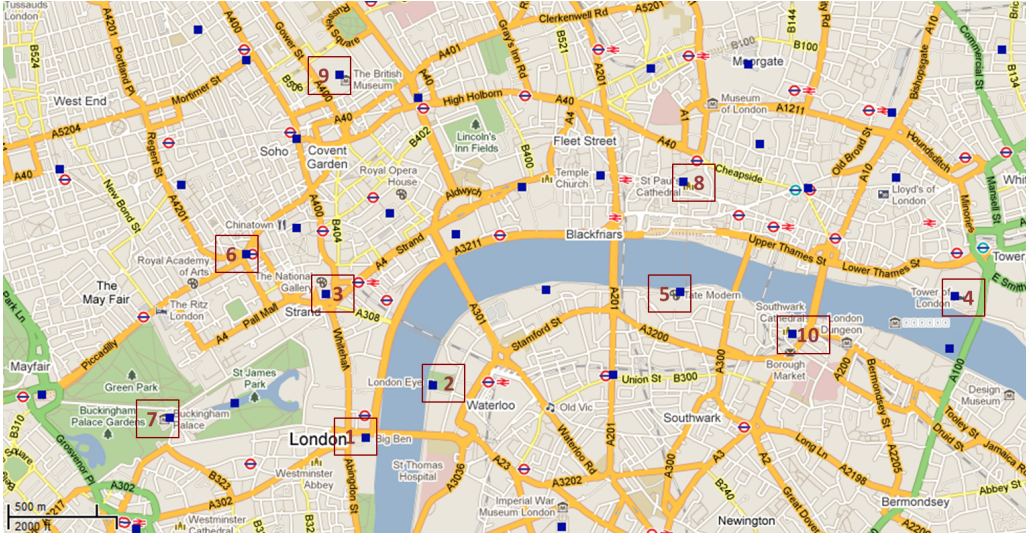
Figure 2: Most popular points of interests that were found for London.

end, we assume a probabilistic approach, where the popularity of a POI poi is identified with the probability that a random user, who is visiting the city, selects poi as her next destination. This probability can be estimated using maximum likelihood:

$$Pr[poi] = \frac{|\text{Users}_{poi}|}{\sum\limits_{p \in Pois} |\text{Users}_p|} \qquad (1)$$

where Users_{poi} is the set of users in R who have taken a photo in poi and $Pois$ is the set of all POIs in the city under consideration.

To improve the baseline, we now consider the context in which the tourist selects a POI, i.e., we are interested in estimating the probability $Pr[poi|co]$ for each POI poi and context co. This probability can be estimated as:

$$Pr[poi|co] = \lambda \cdot \frac{|\text{Users}_{poi,co}|}{\sum\limits_{p \in Pois} |\text{Users}_{p,co}|} + (1 - \lambda) \cdot Pr[poi] \qquad (2)$$

where $\lambda \in [0, 1]$ and $\text{Users}_{poi,co}$ is the set of users in R who have taken a photo in poi within the context co. Note that we used Jelinek-Mercer smoothing to correct the inaccuracy of a maximum likelihood estimation due to data sparseness [22]. When several context parameters are considered (e.g. both day of the week and hour of the day), (2) cannot be reliably used anymore, again due to data sparseness. In such a case, we make an independence assumption between different context factors. In particular, from Bayes' rule, we get

$$Pr[poi|co_1 \wedge co_2] = \frac{Pr[co_1 \wedge co_2|poi] \cdot Pr[poi]}{Pr[co_1 \wedge co_2]}$$

Given that the context parameters remain constant, and making the independence assumption, we get

$$Pr[poi|co_1 \wedge co_2] \propto Pr[co_1|poi] \cdot Pr[co_2|poi] \cdot Pr[poi]$$

and similar for three or more context parameters. The probability $Pr[co_i|poi]$ can be estimated as

$$Pr[co_i|poi] = \lambda \cdot \frac{|\text{Users}_{poi,co_i}|}{|\text{Users}_{poi}|} + (1 - \lambda) \cdot \frac{\sum\limits_{p \in Pois} |\text{Users}_{p,co_i}|}{\sum\limits_{p \in Pois} |\text{Users}_p|}. \qquad (3)$$

where we again use Jelinek-Mercer smoothing. The parameters λ involved in Jelinek-Mercer smoothing were chosen using the development set D. In particular, their values were varied from 0 to 1 with 0.05 step increments to find the value that maximizes Mean Average Precision (see below).

city	MAP					#queries
	baseline	*hour*	*day*	*month*	*combined*	
London	0.2691	0.2718	0.2693	0.2697	**0.2724**	7356
New York	0.2748	**0.2828**	0.2749	0.2774	0.2826	8512
Paris	0.3232	0.3448	0.3402	0.3342	**0.3497**	7283
Berlin	0.3213	**0.3296**	0.3218	0.3200	0.3288	2868
Barcelona	0.2820	**0.2976**	0.2852	0.2825	0.2953	3035
Chicago	0.4431	0.4427	0.4425	**0.4434**	0.4432	2382
Washington	0.2983	0.3026	0.3011	0.3010	**0.3027**	2325
Rome	0.4220	**0.4252**	0.4220	0.4224	0.4248	3270
Madrid	0.3516	**0.3520**	0.3509	0.3508	0.3517	1117
Vancouver	0.3679	0.3656	0.3716	**0.3822**	0.3821	592
total	0.3146	0.3233	0.3183	0.3177	**0.3243**	38740

Table 2: The MAP values of the different ranking algorithms.

To evaluate the rankings, we rely on standard measures from information retrieval. In particular, each user from the test set T who has visited a given city is treated as a query, and the POIs she has visited are considered as the relevant items of the ranking. In this way, we can use Mean Average Precision (MAP) to evaluate how well different rankings correspond to actual user behavior.

Four time-dependent recommendation systems were evaluated. The first system depends on the hour of the day that the tourist visits the city (*hour*). We obtained the best results for three-hour windows, e.g. 'between 8 a.m. and 11 a.m.'. The second system depends on the day of the week (*day*), with each day corresponding to a different context, and the third depends on the month of the year (*month*), with each month corresponding to a different context. Finally, we combined these three systems in a system called *combined*. The MAP scores of these systems are given in Table 2. The λ values involved in the Jelinek-Mercer smoothing were 1 for *hour*, 0.7 for *day*, 0.3 for *month*, and three times 0.6 for *combined*. Based on these results, we can conclude that our time-based systems outperform the baseline, although the differences are rather small. The results also indicate that the hour of the day has a stronger influence on the tourist behaviour than the day of the week or the month of the year. The combination of the three context parameters leads to the best results. All four time-dependent systems perform significantly better than the baseline (Wilcoxon signed-rank test, $p < 0.001$).

To better understand why the improvements are small, we manually inspected the results for London. As could be expected, several attractions remain popular, independent of the hour, day of the week, or month; e.g. the Big Ben, London Eye, Tower of London, etc. are always among the most popular POIs. In fact, there is very little variation in the optimal rankings of the top-10 POIs. On the other hand, for some of the other POIs, the popularity does crucially depend on the context. One example is *Portobello Road*, which is mainly visited on Saturday when there is the famous Portobello Road Market. The *Royal Observatory Greenwich* presents an example of a POI whose popularity is highly dependent on the hour of the day, determined among others by its opening hours. Finally, the popularity of *Wembley stadium* depends on the day of the week, hour of the day, and month of the year, being strongly influenced by the moments when football games and concerts are organized.

4.2 Recommendation as an assignment problem

To gain further insight into the problem of generating time-dependent recommendations for POIs, we also considered a second task, where recommendation is seen as an assignment problem rather than a ranking problem. In particular, we consider a scenario where a user has indicated n POIs she wishes to visit and n time slots during which she will be available. The problem then becomes to map each POI to a time slot, such that the POIs are visited in the best possible contexts. To evaluate different systems, for each user who has visited a city, we try to assign the POIs where she has taken a photo to the context in which she has taken it. Users who have only photographed in one POI have been excluded from the evaluation. Moreover, whenever two or more POIs were photographed within the same context, only one of the POIs was retained. Similarly, when a POI was photographed in different contexts, only one context was retained.

Table 3: Comparing random mappings with a system dependent on the hour of the day ($\lambda = 0.2$).

city	baseline	hour	#queries
London	0.3915	**0.4731**	798
New York	0.4239	**0.4571**	969
Paris	0.4004	**0.4148**	807
Berlin	0.4050	**0.4241**	322
Barcelona	0.4056	**0.4756**	329
Chicago	0.4580	**0.4990**	257
Washington	0.3781	**0.4698**	284
Rome	0.3846	**0.4545**	378
Madrid	0.4716	**0.5101**	134
Vancouver	0.3731	**0.4141**	65
total	0.4071	**0.4552**	4343

A given system can then be evaluated by calculating the percentage of POIs that are assigned to the correct context. Our baseline system simply considers a uniform random assignment. It can be shown that, on average, the percentage of correctly assigned POIs is given by $\frac{1}{n}$. A context-aware system can be obtained by considering that the probability that an assignment that maps poi_i to co_i is given by

$$Pr[(poi_1, co_1) \wedge ... \wedge (poi_n, co_n)] = Pr[(poi_1, co_1)] \cdot ... \cdot Pr[(poi_n, co_n)] \qquad (4)$$

where

$$Pr[(poi_1, co_1)] = Pr[co_i|poi_i] \cdot Pr[poi_i]$$

and $Pr[co_i|poi_i]$ and $Pr[poi_i]$ can be estimated according to (3) and (1) respectively. Moreover, the assignment maximizing the right-hand side of (4) will also be the assignment maximizing

$$\log(Pr[(poi_1, co_1)]) + ... + \log(Pr[(poi_n, co_n)])$$

which means that we have an instance of the linear assignment problem, where the costs are of the form $log(Pr[(poi_i, co_i)])$. An optimal assignment to this problem can be obtained in polynomial time using the Hungarian algorithm.

The results are summarized in Table 3, where a clear improvement over the baseline is witnessed. We only considered the hour of the day as this was found to be the most influential context parameter in Section 4.1. The improvement is found to be statistically significant (Wilcoxon signed-rank test, $p < 0.001$). What is particularly noteworthy is that the improvement varies considerably from city to city. In the cases of London, Barcelona, Washington and Rome, for instance, a substantial improvement is obtained, while for New York, Paris and Berlin, the improvement is rather small.

5 Conclusions and future work

In this paper, we have explored the possibility of using temporal context factors to better predict which POIs might be of interest to a given user. Our experimental results show improvements over baseline techniques that are statistically significant, but small. Similar results were reported in [4], where the idea of personalizing POI recommendations was considered. In general, it seems that tourists tend to visit the most popular POIs at any time, i.e., there is a large number of POIs whose popularity is not time-dependent. Nonetheless, most cities have some POIs whose popularity does strongly depend on the hour, day, or month. In future work, we will investigate the effectiveness of the techniques introduced in this paper w.r.t. other context parameters, such as the weather.

Regarding the practical use of POI recommendations, a stronger interaction between semantic and statistical information is needed. Semantic information, such as opening times of tourist attractions, can clearly help to generate better recommendations. Moreover, semantic information can also be used to generate explanations for these recommendations, e.g., *visit museum X on Monday, because then it is free*. The disadvantage of semantic information, on the other hand, is that it might be difficult to acquire all relevant

information, and this information might, moreover, quickly become outdated. From a statistical analysis of Flickr photos, on the other hand, we can automatically find out when it is appropriate to visit certain attractions. To interface between these two types of approaches, it is of interest to link the POIs that were discovered with Wikipedia pages. When the appropriate Wikipedia page is already georeferenced, this can be done by comparing the coordinates of the page with the center of the POI. However, in general, we may also look at the tags that users have provided with their photos to train language models for POIs [15, 17]. We may then look for the Wikipedia page which is most likely to have been generated by the language model of the POI. Once a Wikipedia page has been assigned to a POI, the corresponding Wikipedia categories provide us with valuable information regarding the semantic type of the POIs (museum, park, historical building, etc.). This may, in turn, be used to derive association rules, which could further help users in deciding where to go.

Acknowledgements

References

[1] Y. Asakura and T. Iryo. Analysis of tourist behaviour based on the tracking data collected using a mobile communication instrument. *Transportation Research Part A: Policy and Practice*, 41(7):684–690, 2007.

[2] L. Cao, J. Luo, A. Gallagher, X. Jin, J. Han, and T. S. Huang. A worldwide tourism recommendation system based on geotagged web photo. *Proceedings of the IEEE International Conference on Acoustics Speech and Signal Processing*, pages 2274–2277, 2010.

[3] K. Cheverst, N. Davies, K. Mitchell, and A. Friday. Experiences of developing and deploying a context-aware tourist guide : The GUIDE project. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*, pages 20–31, 2000.

[4] M. Clements, P. Serdyukov, A. P. D. Vries, and M. J. T. Reinders. Using Flickr geotags to predict user travel behaviour. In *SIGIR10: Proceedings of the 33th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 851–852, 2010.

[5] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg. Mapping the world's photos. In *Proceedings of the 18th International Conference on World Wide Web*, pages 761–770, 2009.

[6] M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Automatic construction of travel itineraries using social breadcrumbs. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia*, pages 35–44, 2010.

[7] F. Girardin, F. D. Fiore, C. Ratti, and J. Blat. Leveraging explicitly disclosed location information to understand tourist dynamics: a case study. *Journal of Location Based Services*, 2:41–56, 2008.

[8] S. Jain, S. Seufert, and S. Bedathur. Antourage: Mining distance-constrained trips from Flickr. In *Proceedings of the 19th International Conference on World Wide Web*, pages 1121–1122, 2010.

[9] A. Kamar. Mobile Tourist Guide (m-ToGuide). Technical report, Deliverable 1.4, Project Final Report. IST-2001-36004, 2003.

[10] M. Loecher and T. Jebara. CitySense: multiscale space time clustering of GPS points and trajectories. In *Proceedings of the Joint Statistical Meeting*, 2009.

[11] A. Pizam and S. Sussmann. Does nationality affect tourist behavior? *Annals of Tourism Research*, 22(4):901 – 917, 1995.

[12] T. Rattenbury and M. Naaman. Methods for extracting place semantics from flickr tags. *ACM Transactions on the Web*, 3(1):1–30, 2009.

[13] C. Ratti, R. M. Pulselli, S. Williams, and D. Frenchman. Mobile landscapes: Using location data from cell phones for urban analysis. *Environment and Planning B: Planning and Design*, 33(5):727–748, 2006.

[14] S. Roy, G. Das, S. Amer-Yahia, and C. Yu. Interactive Itinerary Planning. In *27th International Conference on Data Engineering*, pages 15–26, 2011.

[15] P. Serdyukov, V. Murdock, and R. van Zwol. Placing Flickr photos on a map. In *Proceedings of the 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 484–491, 2009.

[16] T. Sohn, A. Varshavsky, A. Lamarca, M. Y. Chen, T. Choudhury, I. Smith, S. Consolvo, J. Hightower, W. G. Griswold, and E. D. Lara. Mobility detection using everyday GSM traces. In *Proceedings of the 8th International Conference on Ubiquitous Computing*, pages 212 – 224, 2006.

[17] O. Van Laere, S. Schockaert, and B. Dhoedt. Finding locations of Flickr resources using language models and similarity search. In *Proceedings of the ACM International Conference on Multimedia Retrieval*, 2011.

[18] M. Van Setten, S. Pokraev, J. Koolwaaij, T. Instituut, and P. O. Box. Context-aware recommendations in the mobile tourist application COMPASS. In *Proceedings of the Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 235–244, 2004.

[19] J. Wolf, R. Guensler, and W. Bachman. Elimination of the travel diary: An experiment to derive trip purpose from GPS travel data. In *Proceedings of the Transportation Research Board 80th Annual Meeting*, pages 7–11, 2001.

[20] Z. Yin and L. Cao. Diversified trajectory pattern ranking in geo-tagged social media. In *Proceedings of the SIAM Conference on Data Mining*, pages 334 – 342, 2011.

[21] Yizong Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:27–799, 1995.

[22] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems*, 22(2):179–214, 2004.

[23] Y. Zheng, L. Zhang, Z. Ma, X. Xie, and W. Ma. Recommending friends and locations based on individual location history. *ACM Transaction on the Web*, 6(3), 2009.