Visiting the Same Place but Seeing Different Things: Place Models of Touristic Behavior

Christoph Schlieder and Dominik Kremer

Research Group on Computing in the Cultural Sciences,
Laboratory for Semantic Information Technology
University of Bamberg, Germany
{christoph.schlieder, dominik.kremer}@uni-bamberg.de

Abstract GPS tracks and geo-referenced photographs permit to closely monitor the spatial decisions of tourists exploring an urban environment. Different methods for identifying places of high touristic popularity from such data have been proposed, some based on the time tourists spend at a place, others based on their photographing activity. Little is known about how these two ways of defining touristic popularity relate. The article describes a data set and a method for defining the extent of places of touristic interest. Within that methodological framework the two approaches are compared. Finally, the issue of individual differences is addressed showing that we do not all visit the same places to see the same things as a naïve analysis of social place production might suggest.

1 Producing and Consuming Places

The social production of touristic places clearly predates the web with its location-based social networks. It can be traced back at least to the Hellenistic period were the first list with the Seven Wonders of the World was compiled (Lomine, 2005). As transportation speed has increased by orders of magnitude since antiquity, it is not surprising to find that today's lists have grown longer and consist of up to "A 1000 places to see before you die" (Schultz, 2010).

Although touristic travel constitutes an activity in which social and spatial choices closely interact, it has, at least so far, not received the same interest from research on location-based social networks as, for instance, disaster management or social serendipity services (Zhou et al., 2009; Zhou et al., 2010). On the other hand, geographers and social scientists who study the mechanisms of place production have paid little attention to the recent advances in information and communication technology in general, and mobile web applications in particular (Urry, 1995; Cresswell, 2004; Urry and Larsen, 2011). The use of smartphones equipped with a GPS receiver and a camera permits to collect behavioral data at a spatial and temporal resolution that was not accessible to earlier empirical studies of tourism (Girardin et al., 2008, Schlieder and Matyas, 2009).

Does the social production of touristic places imply that we all visit the same places to see the same things? Computational documentation methods certainly help to collect and interpret data that will ultimately allow answering that question. Most computational approaches that have been proposed focus on the problem of characterizing touristic popularity (e.g. Girardin et al. 2009). We argue that although much insight can be gained by looking at popularity, especially visual popularity, an equally important aspect of touristic consumption should not be ignored: time. Little is known about how the time tourists spend at a certain place relates to their photographing activity and what individual differences we should expect. Note that our objective is not the ambitious one of identifying the social mechanisms of place production. We focus on the spatial behavior of individual tourists, that is, on the consumption of places rather than their production. The paper reports about ongoing work and makes the following main contributions:

- We describe a data set of GPS tracks and geo-referenced photographs that documents the spatial behavior of tourists in an urban environment. Based on the data, we identify problems with aggregation methods for defining places of interest and show how to solve them using a network-based place model.
- We generalize the concept of visual popularity of a place introduced by Schlieder and Matyas (2009) to account for a process of touristic place consumption that is characterized by the time spent at the place rather than by the photographic sights taken of that place. The popularity rankings resulting from the original and the modified version are compared.
- Finally, we analyze the data set with respect to individual differences giving a first answer to the question whether everybody visits the same places to see the same things.

2 Related Work on Place Models

Place amounts to more than topography plus land use. Nations as well as neighborhoods are defined in terms of conventional boundaries which result from social negotiation processes. Place production includes such diverse activities as nation building or gentrification and has been studied by human geographers for more than three decades (Relph, 1976; Yuan, 1977; Urry, 1995). This line of research does not aim at developing computational models. Sometimes it positions itself even in contrast to spatial analysis, especially analysis that uses computational methods (Cresswell, 2004). However, we may retain from geographic research that places are permanently defined and redefined. When studying touristic places, we would be ill-advised to use an existing data set with geographic footprints for places of touristic interest since we could not be certain that tourists actually use these spatial conceptualizations. We will instead refer to a list of place names from touristic marketing – a typical and powerful place production activity – and try to infer appropriate place models from analyzing the behavioral data of tourists.

Recently, researchers from geographic information science became interested in place (Winter et al., 2009). Models of vague places are also of concern in the field of geographic information retrieval (Schockaert 2011). A number of formal place models have been proposed, for instance, supervaluation semantics and qualitatively augmented fuzzy footprints. However, the connection between place models and user communities has been studied very little. Schlieder and Henrich (2011) have argued that there are scenarios where the classical membership problem (does the point *X* belong to the region *R*?) transforms into more complex problems: does user *A* believe *X* belongs to *R*? Do users *A* and *B* share similar beliefs about *X* belonging to *R*? The analysis of touristic data is a typical example for such a scenario.

Most relevant for our analysis are the place models that have been recently described in the context of computational analyses of touristic behavior by Griardin et al. (2008; 2009) who analyzed geo-referenced photographs and GPS tracks of tourists. Photo log data consists of a sparse stream of geo-referenced images, while a much denser stream of geographic positions represents the GPS track data. For the purpose of analyzing such data in terms of places, some spatial aggregation of the positions is needed. The simplest aggregation method uses a grid of fixed mesh size and considers all photographs taken within a grid cell to refer to the same place. Girardin et al. (2008) used this approach to generate density maps of photographic activity of tourists visiting Rome ("heatmaps").

Spatial clustering, on the other hand, uses a distance criterion for aggregating geographic positions. An example is the incremental clustering method which Schlieder and Matyas (2009) designed for the same purpose, the computation of density maps of photographic activity. Both approaches, however, lack a meaningful place concept. In a number of cases they aggregate pictures which show different objects while pictures of the same object may end up in different clusters.

3 Data and Methodological Issues

The data set we collected consists of GPS tracks and photographs from tourists who explored the city of Bamberg, Germany, a listed UNESCO world heritage site. A number of features make this city a good choice for studying spatial decision processes in the context of touristic place production. (1) Bamberg is a well-known tourist destination which attracts about 2 million visitors a year. As a consequence, even first-time visitors are likely to have been exposed to some prior information about the city. (2) Although there are museums to visit as well as a couple of breweries, Bamberg is best known for its built heritage which can be observed from the outside making the spatial decisions of the tourists observable in their GPS tracks. (3) The entire old part of the city constitutes the world heritage district, not a single building or group of buildings, giving tourists a broad range of options to choose from.

The data set presented in this paper is based on behavioral data from 17 tourists who volunteered to participate in our study. Contact with the participants was established in front of the tourist information office. Only tourists who said they planned to stay longer than 2 hours in the city were admitted as participants. They were handed a camera equipped with GPS and a magnetic compass, as well as a second GPS receiver with better positional accuracy for recording the track data. The tourist office served as the starting and ending point of the tour. No other constraints were imposed. The participants were told to explore the city in whatever way they liked and for how long as they pleased. The participants were only instructed to try to take at least one photograph every 10 minutes to avoid effects found in an exploratory study in which the photographing task was sometimes forgotten or squeezed into the very last minutes of the tour.

It took the 17 participants between 120 min and 420 min to complete their individually chosen tour. On average they spent 212 min on a tour. The length of the tours ranged between 2990 m and 10000 m with an average of 5440 m. The longer tours are those of tourists who spent more time on their visit. In other words, there were no surprising differences in locomotion speed which averaged at 1.59 km/h. The total number of photographs collected on a tour varied between a minimum of 15 images and a maximum of 234 images with an average of 58 images. An indicator of the photographic activity level is the number of photographs taken per hour which ranged between 3.46 and 33.43. Only the participant with the minimal photographic activity stayed under the 6 images per hour that the instruction requested. All others photographed more with an average of 15.6 images per hour.

Part of the problem with place modeling is due to the fact that, generally, the point from which a photograph is taken does not uniquely identify the object(s) depicted. Of course, photographic vantage points exist. Fig. 1 shows a rare case where such a vantage point has been identified by a physical marker. The marker indicates where photographers should place their feet and recommends a focal length, that is, a specific angle of sight. If the directions are followed, the resulting images of the Nicolai church in Leipzig look virtually identical, showing only variations of daylight and weather.

Fig. 1. Instructions for photographers at the Nikolai church, Leipzig

What looks like a prime example for place production by signage, actually shows something quite different: there are not many vantage points from which a photographer can capture the entire church.[1] In the Bamberg data set, most vantage points are associated with a particular historical building or panoramic sight. The remaining cases of ambiguity can, at least in principle, be resolved by using the information from the camera's compass. However, this does not help with analyzing the GPS tracks in terms of places. In order to determine the time that a tourist spends at a place, we need a spatial characterization of the place.

[1] Note that the marker works poorly as a sign since it cannot be perceived from a distance. It is rather the dramatic placement of a replica of a column from the interior of the church that turns the square into a lieu de mémoire for the peaceful revolution of 1989.

4 Places of Touristic Interest

The tourist office of the city of Bamberg distributes a map that lists 42 places of interests, mostly historical buildings and green spaces, and displays them as point features on the map. For the purpose of our study it is of no importance that the map does not specify a ranking of the places or suggest a specific visiting order. We assume that the tourists possess at least the prior information given by the map. Therefore, we use the place names appearing on the map as a starting point for our analysis. An inspection of the data set reveals a number of characteristic spatial features of the places of interest. Fig. 2 shows some of the points from which the Old Town Hall has been photographed.

Because the tourists move as pedestrians along the street network, the place cannot be described by a simple polygon feature. Inaccessible buildings and water bodies cause holes in the region. We argue that in such cases a network-based model is more appropriate. The Old Town Hall is situated on an island in the river Regnitz and connects to the city via bridges (dotted lines on the map). Note that the map shows only part of the vantage points from which photographs of have been taken. Although all photographs show the Old Town Hall, they do so in rather different ways: image *A* is seen while approaching the building whereas image *B* shows a popular distance view.

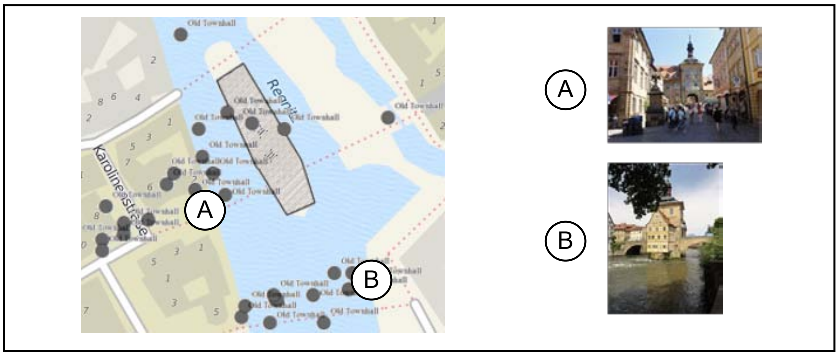

Fig. 2. Photographic vantage points for the Old Town Hall, Bamberg

The network-based place model is determined in four steps which determine the footprint of a touristic place of interest using the network hull of the vantage points from which tourists have taken photographs. Basically, this amounts to define the place associated with a particular historical building by locating the positions of all photographs of that building and by adding adequate network connections between these positions. The network connections need to be included since any data set contains only a finite sample of the infinitely many possible vantage points from which the building can be photographed. Remember that a subnetwork *S* of a planar network *N* is considered *convex* if the shortest path between any pair of points on an edge or vertex of *S* is also a shortest path in *N*. The *network hull* of a set of points on an edge or vertex of *N* is the smallest convex subnetwork containing the points. The place model is computed as follows:

(1) *Preprocessing*: photographs which have been incorrectly geo-referenced due to GPS errors are excluded from the analysis.
(2) *Adjustment*: the vantage points of the photographs are fitted onto the street network.
(3) *Place building*: the place of touristic interest is computed as the network hull of the vantage points.
(4) *Visit timing*: from the GPS track data, we determine the time that a tourist spends on the subnetwork defining the place

The street network is based on OSM data. It includes the outlines of some buildings as paths. This is very useful since tourists often try to walk around a building even if this means leaving the classical street network. Snapping has been computed with a standard GIS tool (ArcGIS 10 with the network analyst extension). For the sake of simplicity, an approximation of the network hull was used that could readily be computed using the tool. The approximation consists of connecting a random permutation of the vantage points within the network as shown in Fig. 3. A more refined analysis will not only compute the exact hull but also integrate a viewshed analysis which checks for visibility of the building.

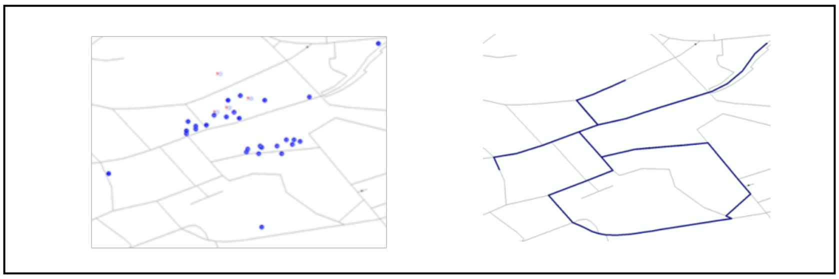

Fig. 3. Determining a place of interest: Old Town Hall

5 Image-based versus time-based popularity

An aspect that is often overlooked in the analysis of touristic popularity is the ambiguity of visual content associated with a geographic position. The number of vantage points from which more than object of interest can be observed is not as rare as one might think – at least in the urban environment we studied. In our data set we found several cases where a tourist has photographed a place *A* followed by a place *B* without changing position, sometimes the images even come in the form of alternating place sequences *ABAB*. As a consequence, place models should admit that regions of touristic interest overlap. Such overlapping places – subnetworks in our model – are in fact present in the data. For the 11 places of touristic interest for which we generated place models according to the method described in the previous section, we found 7 pairs of overlapping subnetworks. Fig. 4 illustrates the most complex situation found, a mutual overlap of networks representing four places.

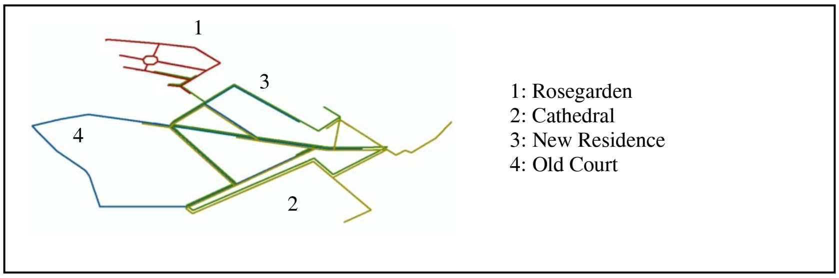

Fig. 4. Overlapping touristic footprints

In order to determine the time tourists spend at a place, we need to deal with overlapping places. Without photographs or additional sensor information it is not possible to tell for a geographic position which belongs to more than one network, which of the places has been visited. The solution adopted consists in simply dividing the time spent in overlapping places among those places (e.g. a third of the time for positions in a triple overlap network).

We are now in a position to determine for each of the 11 places in the data set four basic parameters related to touristic popularity: (1) the total number of images taken at the place, (2) the number of visitors who have taken at least one image at the place, (3) the total time that visitors spent at the place, (4) the number of people who spent at least 1 min at the place. The threshold for determining (4) has been chosen to exclude tourist that just pass by a place without devoting time to a closer visit. Visual popularity is determined according to the measure proposed by Schlieder and Matyas (2009). This image-based popularity combines the number of photographers of a place with the number of photographs they have taken according to a logarithmic law of diminishing returns. The effect of an image on the popularity score becomes smaller, the more images a photographer produces of the place. Tab. 1 shows the image-based popularity scores and the implied ranking of the places.

Place[2]	**OC**	**OT**	**CA**	**GP**	**GC**	**LV**	**MP**	**MM**	**NR**	**UP**	**RO**
images	16	34	40	4	3	30	23	17	19	4	36
visitors	7	12	14	4	3	11	6	7	11	3	10
pop.	16,1	28,8	32,5	8,0	6,0	25,8	14,9	16,3	24,0	6,3	24,3
rank	7	2	1	9	11	3	8	6	5	10	4

Tab. 1: Image-based touristic popularity

At the maximum level of photographing activity which we found in the data set, an image is taken on average every two minutes. In contrast, the GPS track data measures the geographic position every 5 sec. The high resolution of the GPS track data provides a first argument for considering the time a tourist spends at a place. Another maybe more important argument is that not every tourist feels comfortable as a pho-

[2] OC = Old Court, OT = Old Townhall, CA = Cathedral, GP = Geyerswörth Park, GC = Geyerswörth Castle, LV = Little Venice, MP = Michaelsberg Park, MM = Michaelsberg Monastery, NR = New Residence, UP = Upper Parish, RO = Rosegarden

tographer. Some tourists visit a city without photographing at all, while others just photograph people, not places.

The total time that tourists have spent at a place does not constitute a good indicator of popularity because it is heavily affected by individual tourists with a very long visit. We propose to generalize the approach for image-based popularity to a time-based popularity measure which uses two parameters, the number of tourist that have visited the place and the amount of time that they have spent there. This time-based popularity measure controls the effects of individual long visits by applying a law of diminishing returns. The first minute spent at the place increases the popularity by 1, while the second minute increases popularity by log 2, and the *n*-th minute by log $n/(n - 1)$. With T_p denoting the set of the k tourists that have visited place p, we define the time-based popularity score as

$$pop_{time}(p) = \sum_{x \in T_p} \left(1 + \log t(x)\right) = k + \sum_{x \in T_p} \log t(x)$$

Tab. 2 lists the time-based popularity scores together with the implied rankings. Some differences between image-based and time-based ranking are immediately obvious. For instance, image-based rank 3 is downgraded to a time-based rank 6. This can be explained by the fact that the place (Little Venice) allows photographing a scenic river front but is easily explored in a few minutes. On the other hand, image-based rank 7 is upgraded to time-based rank 4. Again, a geographic explanation can be given. The place (Old Court) covers a vast area which cannot be overseen from a single vantage point and therefore invites visitors to engage in exploration. At the same time, this place does not offer spectacular buildings or vistas.

Place	OC	OT	CA	GP	GC	LV	MP	MM	NR	UP	RO
time	98,0	270,2	150,0	52,7	15,7	44,3	90,9	80,6	113,9	11,9	151,4
visitors	15	16	16	6	3	9	6	7	15	4	12
pop.	41,6	51,7	45,7	23,4	11,7	24,5	18,0	20,5	43,0	14,5	36,7
Rank	4	1	2	7	11	6	9	8	3	10	5

Tab. 2: Time-based touristic popularity

The examples for rank downgrade and upgrade reflect environmental differences between places that can be expected to appeal to different touristic interests. This suggests that image-based and time-based popularity scores cover to a certain extent complementary aspects of places. A Spearman rank correlation coefficient of 0.85 shows that the image-based and the time-based ranking nevertheless show a good overall agreement. For the purpose of comparison we consider a third rather naïve, but computationally inexpensive popularity score: the number of images (no matter of what content) that fall within a 50 m buffer around the point features associated with the 11 places on the tourist office map. Tab. 3 shows the results of this popularity ranking method.

Place	OC	OT	CA	GP	GC	LV	MP	MM	NR	UP	RO
Images	16	41	17	9	13	10	4	3	38	2	40
Rank	5	1	4	8	6	7	9	10	3	11	2

Tab. 3: Naïve popularity score based on photo count within a 50m buffer

The Spearman correlation coefficient for the naïve ranking and the image-based ranking is 0.63, that for the naïve ranking and the time-based ranking 0.79. Note that the image-based and the time-based rankings show more agreement among each other than each of them with the naïve ranking. We interpret this as evidence that the image-based and the time-based score provide a better, though complementary, assessment of the touristic popularity of places.

6 Individual Differences and the Social Production of Places

Our model for a place *P* of touristic interest within an urban environment can be written as a triple $P = (n, s, r)$ where *n* denotes the place's name, *s* the subnetwork of streets that acts as the geographic footprint, and *r* the rank of the place according to either the image-based or the time-based popularity. A model *C* for the city consists of the vector of the *k* place models $C = (P_1, \ldots, P_k)$. No matter which popularity ranking is used, the model always reflects the choices of many tourists. This raises the question of individual differences. Instead of the question "Is *P* popular in *C*?" we now ask: "Does tourist *T* consider *P* an interesting place in *C*?" or "Do tourists *T* and *U* share similar attitudes towards the places in *C*?".

We consider time-based popularity. The time that a tourist spends at a place can be considered a vote for that place. A visit that lasts longer than his or her average visit is a positive vote, a shorter visit a negative vote. The similarity of the votes can be compared using an approach first suggested by Resnik et al. (1994) in their recommender system GroupLens which essentially amounts to compute the linear correlation of the votes. Once similarity is determined, a nearest neighbor clustering determines groups of tourists which show similar attitudes towards the 11 places. Fig. 5 shows the dendrogram of the clustering for the 17 tourists from our data set.

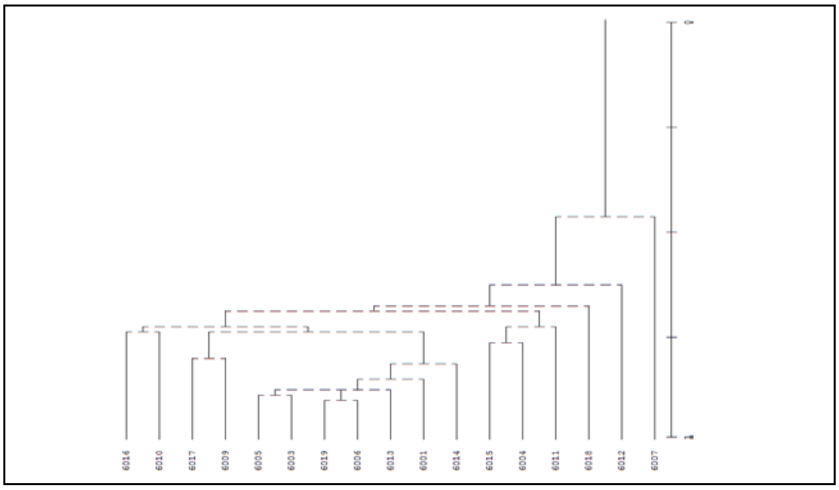

Fig. 5: Dendrogram of user similarity (GroupLens, nearest neighbour) based on consumption time at the different places

A high level of overall similarity is observed which can be explained, at least partly, by the fact that we included only the top 11 attractions from the tourist office's map in the analysis. The data shows that the top 2 attractions, the Old Town Hall and the Cathedral, are visited by all tourists with one notable exception (Tab. 2). Tourist 6007 focused entirely on visiting two more remote but culturally interesting places, the monastery St. Michael and its large park. Among the other tourists, a core cluster can be identified consisting of the tourists 6006 and 6019 who spent most of their time on the top 3 places. Near subclusters are characterized by longer visits to combinations of parks and public green spaces. Cases 6004 and 6015, for example, spent most of their time in the rose garden, visiting the three main attractions on a more modest level.

To sum up, individual differences are found and can be given geographical interpretations. Such differences are of great practical interest to the design of geographic recommender systems (Schlieder, 2007). It is known that differences relevant for recommendations are likely to be found in the long tail of rank-frequency distribution, that is, by discarding the top items. A closer study of the lesser visited items is going to be a focus of our future research.

7 Discussion and Conclusions

We presented a data set on touristic exploration behavior in an urban environment consisting of GPS tracks and geo-referenced photographs. Methodological problems such as spatial data aggregation were identified and solved. We described a network-based model for places of touristic interest and a method for determining image-based and time-based popularity scores from basic parameters of the data set. Our analysis of the two approaches for measuring popularity showed that they are both better at reflecting the properties of the data than a simpler measure of popularity. We also found that there are cases were the image-based and time-based popularity disagree. Since geographical explanations could be given for the cases, we argue that image-based and time-based express complementary aspects of the data. Finally, we looked into individual differences. Although there is overall agreement regarding the top 2 places to visit, we found differences in the lower ranking places.

Studying tourist behavior means to study processes that complement the social production of touristic places, that is, processes of place consumption. A basic insight of our research consists in the fact that the consumption of places is a process that takes different amounts of time. In a city we may find scenic places that are widely advertized in visual media by touristic marketing, but that turn out to take very little time to visit. The converse is also possible. More important, the amount of time varies significantly between individuals. Introducing a time-based popularity score constitutes just the first step in understanding how tourists differ in their spatial and temporal choices. It could help to overcome an analysis that focuses too much on "the tourist gaze". Even without taking any photographs, a tourist can experience the qualities of a place, say a green space that just invites to rest and listen to the city.

References

Cresswell, T. (2004): *Place. A short introduction*. Malden, MA: Blackwell Publishers

Jones, C.; Purves, R.; Clough, P.; Joho, H. (2008): Modelling vague places with knowledge from the Web. *Int. Journal of Geographical Information Science 22* (10), p. 1045 — 1065.

Girardin, F.; Blatt, J.; Calabrese, F.; Dal Fiore, F.; Ratti, C. (2008): Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing 7* (4), p. 36–43.

Girardin, F.; Vaccari, A.; Gerber, A.; Biderman, A.; Ratti, C. (2009): Quantifying urban attractiveness from the distribution and density of digital footprints. In: *International Journal of Spatial Data Infrastructure Research 4*, pS. 175–200.

Lomine, L. (2005): Tourism in August Society. In: J. Walton (ed.): *Histories of tourism. Representation, identity, and conflict*. Bristol, UK: Channel View Publications, p. 69–87.

Relph, E. C. (1976): *Place and placelessness*. London: Pion.

Resnick, P.; Iacovou, N.; Suchak, M.; Bergstrom, P.; Riedl, J. (1994): GroupLens: an open architecture for collaborative filtering of netnews. In: J. Smith, et al. (eds.): *CSCW-94, Proc. Conf. on Computer Supported Cooperative Work*, New York, NY: ACM Press, p. 175–186.

Schlieder, C.; Henrich, A. (2011): Spatial Grounding with Vague Place Models. *SIGSPATIAL Special Issue on Geographic Information Retrieval*.

Schlieder, C. (2007). Modeling collaborative semantics with a geographic recommender. In J. Hainaut et al. (eds.) *Workshop on Semantic and Conceptual Issues in Geographic Information Systems*, Auckland, New Zealand, Nov. 5–9, 2007, LNCS. Springer: Berlin.

Schlieder, C.; Matyas, C. (2009): Photographing a City: An Analysis of Place Concepts Based on Spatial Choices. *Spatial Cognition & Computation 9* (3), S. 212–228.

Schockaert, S. (2011): Vague Regions in Geographic Information Retrieval. In: SIGSPATIAL Special Issue on Geographic Information Retrieval.

Schultz, P. (2010): *1,000 places to see before you die*. Updated. New York: Workman Pub.

Urry, J. (1995): *Consuming Places*. London: Routledge.

Urry, J.; Larsen, J. (2011): *The Tourist Gaze 3.0*. Thousand Oaks, CA: Sage Publications.

Winter, S.; Kuhn, W.; Krüger, A. (2009): Guest Editorial: Does Place Have a Place in Geographic Information Science? *Spatial Cognition and Computation* 9 (3), S. 171–173.

Tuan, Y. (1977): *Space and place. The perspective of experience*. Minneapolis: University of Minnesota Press.

Zhou, X.; Xie, X. (eds.) (2009): *Proc. Int. Workshop on Location Based Social Networks*. November 3, 2009, Seattle, Washington. New York, N.Y: ACM Press.

Zhou, X.; Lee, W.; Peng, W.; Xie, X. (eds.) (2010): *Proc. 2nd Int. Workshop on Location Based Social Networks*. Nov. 2, 2010, San Jose, CA. New York, N.Y: ACM Press.