Destinations Similarity Based on User Generated Pictures' Tags

Alessandro Inversini
Davide Eynard
Elena Marchiori
Leonardo Gentile

webatelier.net
Università della Svizzera italiana, Switzerland
alessandro.inversini@usi.ch
davide.eynard@usi.ch
elena.marchiori@usi.ch
leonard.gentile@gmail.com

Abstract

Pictures about tourism destinations are part of the contents shared online through social media by travellers. Additional pictures information, such as geo-tags and user description of place, can be used to create groups of similar destinations. This paper investigates the possibility of defining destination similarities based on implicit information already shared on the web. Flickr.com was used as a case study as it represents the most popular picture sharing website. Results show the possibility to group similar destinations based on visual components, represented by the contents of the pictures, and the related tag descriptions.

Keywords: destination similarity, folksonomies, geotagging, social media.

1 Introduction

Tourism has always been recognized as an information intensive domain (Gretzel et al., 2000; Buhalis, 2003) where information gathering, processing and distribution is essential for day to day operations (Poon, 1993). Recently researchers (e.g. Gretzel, 2006) demonstrated that Web2.0 and social media are assuming more and more importance within the tourism online promotion. Destination managers are understanding that beside website communication they should be aware that different web sources, mostly informal, are spreading the same messages with different strategies (Inversini and Buhalis, 2009). These sources can be analysed, monitored and exploited, with a well-defined strategy, to take marketing and selling advantages from them (e.g. Inversini et al., 2009). Wise destination managers are already integrating and exploiting social media within their online communications as separated mean or incorporated them within their websites. Within this scenario user generated pictures are gaining more and more importance (Yoo and Gretzel., 2009) as they informally represent and describe a destination. Further, user generated pictures carry a great amount of information because they are often described by sets of small terms called "tags" and sometimes represent places within a map. These map places are termed geo-located tags or geo-tagged.

2 Literature

Tourism Destinations and Technologies

The continuous development of Information Communication Technologies (ICT) during the last few decades has had profound implications for the tourism industry as a whole (Buhalis, 2000). The increasing importance of technology has influenced not only the way transactions and purchase processes (Werthner and Klein, 1999) have evolved, but also the way communication and promotion of tourism goods have develop on the internet (Buhalis, 2003). Recently the advent of, the so called, Social Media (Blackshaw, 2006) enabled tourists to share information on the internet within the “read/write web”, where the end user has become both information consumer, player (Nicholas, et al., 2007) and provider. Marketing managers and researchers are exploiting new ways to adopt social media in the marketing and promotion arenas in order to take advantage of this “electronic word-of-mouth” (Litvin et al., 2008). Recent studies demonstrated that social media contributes to spreading to web information about destinations using different channels and different strategies (Inversini and Buhalis, 2009) as internet users are in need of communicating their touristic experiences (Inversini and Cantoni, 2009). The web, especially social media, offers a variety of different platforms to share experiences, facts and even rumours (Blackshaw and Nazzaro, 2006). This information published on popular social media is contributing significantly to the massive growth of information on the web, be it relevant for the end user or not. Furthermore, one important role within information spreading by social media is played out by user generated pictures. Following Yoo and Gretzel (2009) one in two tourists view destination photos via UGC in different web communities. The relevance of pictures in travel both to understand culture (Pengiran-Kaha et al., 2010) and to recommend a place to visit (Linanza et al., 2011) has also been investigated by recent studies. Contents of pictures shared online can be fragmented into different topics, such as nature, products, and facilities (Govers and Go, 2005). The amount of these experiential-type images can contribute to the online representation of a tourism destination. Pictures shared online can act as a mediated source of information for a prospective consumer, which may influence his/her decision to visit a destination. As in Govers and Go (2005), DMOs can take advantage from online pictures by learning the meanings of pictures shared about a destination and better improve the rich tourism experience that the tourists are looking for. Within social media, user generated travel pictures carry a lot of information. They are often described by sets of small terms called “tags.” Once collected the tags build a *folksonomy*. Tags, also, often represent places within a map. These are called geo-located tags or geo-tags.

The concept of folksonomies

The term folksonomy was introduced by Vander Wal (2004), by mixing the terms “folk” and “taxonomy”. In practice, users assign a set of terms called tags to an individual piece of content in order to group or classify it for retrieval (Sturtz, 2004). The result is an informal social network of terms based on users’ informal classification of content. The collection of all assigned terms for a piece of content of a single user is called personomy, while the collection of personomies is called folksonomy (Hotho et al., 2006). Some examples of successful folksonomies are delicious.com (formerly known as del.icio.us), Steve (steve.museum), and Flickr (flickr.com). In folksonomies users are not forced to use the same tags; however, users with similar interests tend to converge onto a shared vocabulary with their tags. One of the factors of success for folksonomies is the fact that no specific skills are needed to participate (Hotho et al., 2006). It is therefore possible to argue that folksonomies invite deliberate and idiosyncratic tagging, also called meta noise, which decreases system utility (Wu et al., 2006). For the general purpose of this research folksonomies have been classified into broad and narrow folksonomies: (i) *Broad folksonomies* emerge in systems whose users can tag any resource, and where a resource can be annotated with many identical tags, one for each user tagging it; in (ii) *Narrow folksonomies* where tags are singular in nature for each object and users are allowed to tag only a limited set of resources, typically the ones they, or their restricted circle of friends, have provided.

Knowledge from geotags

Geotagging is the process of annotating objects and online resources with geospatial context information, ranging from specific point locations to arbitrarily shaped regions. These annotations can be explicitly provided by users or extracted automatically (i.e. by analyzing where the user is connecting from, getting GPS data from a mobile device, or extracting geographic metadata from a photo. In online photo sharing communities, user text-based annotation (tags) and location metadata (geotags) often co-exist. Geotag information is typically embedded within picture metadata (stored in EXIF format, see http://www.exif.org/specifications.html). In this report the term "geotag" refers to the geographic annotations (e.g.where a photo was taken) while "tag" is always intended as the textual annotation related to a photo (e.g. "cat"). Several studies have been conducted on extracting knowledge from Flickr georeferenced metadata. Clements et al. (2010) introduces a method to predict similar locations, wormholes, based on human travel behaviour. A wormhole is defined as a similar, but not necessarily spatially close, location on the planet. There are two hypotheses for this: (i) users have specific travel preferences and therefore visit locations that are similar to some extent; and (ii) taking a photo in a specific location is an indication that the user likes that location. From a given target location (L) the algorithm aims to find similar locations around the world. For each user (u), a weight (w) is computed based on the distance of the nearest geotagged photo of the user to the target location. Wormholes are then found by aggregating the geotags of all users with W_u as weight per user, and selecting the most relevant positions on Earth according to this metric.

Ahern et al. (2007) show how to analyze tags associated with geo-referenced Flickr images to generate aggregate knowledge in the form of "representative tags" for arbitrary areas of the world. Tags are used to create a visualization tool, "World Explorer" (http://tagmaps.research.yahoo.com/worldexplorer.php), which displays derived tags and original photo items on a world map. Data analysis algorithms are based on multi-level clustering and the scoring of tags is based on TF-IDF (term frequency, inverse document frequency). A user interface shows, for each map region and zoom level, the best-scoring tags for the generated clusters; these tags are shown as text over the map area where each cluster occurs.

3 Research Design

The aim of this research is to describe destination similarity starting from user generated picture tags from the popular photo sharing social media website Flickr.com. In other words this study tries to define similarities amongst a given number of destinations based on unrestricted user descriptions of the destination itself (i.e. the tags associated to the pictures taken at a specific location). Similarities could be used to compile a destinations recommendations list based on pictures shared on social networks by a given users. Furthermore, the research attempts to understand if additional information related to users (i.e. those who uploaded a given photo) and pictures (i.e. pictures sharing the same tags) is driving the process of recognizing similar destinations. Given the goal, and the above mentioned literature (i.e. Clements et al., 2010; Ahern et al., 2007) the two main research questions are:

RQ1: while finding similar destinations thanks to tags, does picture-related information matter to provide better results?

RQ2: while finding similar destinations, thanks to tags, does users-related information matter to provide better results?

In order to pursue this general goal a holistic methodology, based on the related works, has been designed. Research objectives are similar to the ones exposed in Clements et al. (2010), but they are using a TF-IDF approach similar to the one described in Ahern et al. (2007) in order to extract representative tags describing each place. The designed procedure can be summarized as follows: (i) given a dataset of destination-related tags from Flickr a representative description for each city was extracted; then (ii) the similarities between cities were calculated according to the generated representations; and finally (iii) an interface was provided that, given a city as a query term, showed the cities that were more similar according to the previously calculated similarities. The next section describes the research approach of this study.

4 Research Approach

<u>A tag-based city representation</u>

The first step to define similarity between cities was to find an appropriate representation for them. In order to develop the model presented below the “document-term” matrix concept -typically used by document indexing software and search engines- was exploited. According to this representation, documents are considered in terms of the words that appear within them. Similarly, each city was represented as the collection of all the tags assigned to its pictures. For this experiment the tag distribution of Flickr geotagged photos belonging to the union of two sets (i.e. the ones related to the top 150 tourism city destinations for the years

2007 and 2008, according to "Euromonitor International": http://www.euromonitor.com/Top_150_City_Destinations_London_Leads_the_Way) was analysed for a total count of 233 cities.

Flickr.com easily allows users to get the top tags for a given location (specified as a WOEID) through its flickr.places.tagsForPlace API. However, this API only returns the top 100 unique tags, without any information about the photos taken or the users who uploaded them. For this reason, two distinct datasets were built: The former, called *Top100*, contains the tags retrieved using the aforementioned API; the latter, called *Random*, contains a random sampling of photo metadata obtained by querying Flickr APIs with YQL (Yahoo Query Language: an expressive SQL-like language that lets developers query, filter, and join data across Web services.) Selecting, for each city, 10 photos from 300 random days, taken at random hours, avoids bias due to day- or time-related events. An important advantage of this second approach is that user- and photo-related information is available, providing new dimensions across which tag analysis can be performed.

The resulting collections also contained tags which were not useful or not related to the destinations (e.g. very common terms like `day`, `dog`, and `friends`, or photography-related terms such as `canon`, `nikon`, and `black&white`). So, a blacklist was heuristically created based on the analysis of the top 1000 tags in Flickr. Tags occurring in only one city were also pruned, because in the VSM they represent terms that will never match when computing the similarity scores. After cleaning the tag vocabulary the *Top100* and *Random* datasets counted, respectively, 23'300 and 55'000 distinct tags. Figure 1 shows the distribution of the top 1000 keywords. Categories were chosen iteratively during the analysis.

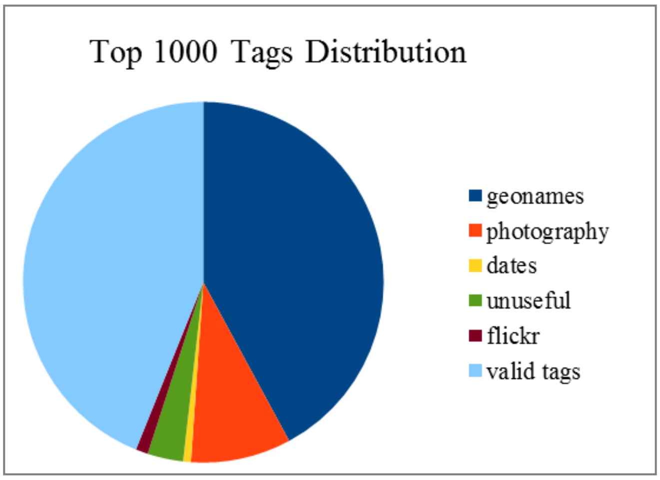

Figure 1- Tag Distribution Top1000

This shows the top 1000 tag distribution to be: valid tags – 43,9%; geotags – 42,1%; photography related tags – 9%; other tags (dates, Flickr related and non-useful tags) –

5%. For the purpose of this study photography related and other tags (dates, Flickr related and non-useful tags) were not considered.

Calculation of similarity between cities

The VSM (Vector Space Model) was used to represent cities in terms of their related tags. In the VSM each city is represented by a vector in an *n*-dimensional space, where *n* is the number of distinct tags, whose components are weighted according to tag frequency. Of course, as tags which are too popular tend to be more widespread and thus less informative, frequency is normalized using the TF-IDF approach, which normalizes a TF (Term Frequency) with an IDF (Inverse Document Frequency) factor. This takes into account the number of documents containing that term (in our case, the number of cities for whose photos a given tag has been used). In the *Top100* dataset the standard IDF has been calculated for normalization:

$$IDF_i = \log_2 \frac{|D|}{|D_i|}$$

where $|D|$ is the number of destinations used in the analysis (233) and $|D_i|$ is the number of destinations containing the tag t_i. For the *Random* dataset two more variants of IDF have been calculated, taking advantage of the additional information we were able to gather. The former variant, called IDFP, exploits the extra information coming from the tagged photos:

$$IDFP_i = \frac{|P|}{|P_i|}$$

where $|P|$ is the total amount of photos in the *Random* dataset and $|P_i|$ is the number of photos tagged with t_i. The latter variant, called IDFU, exploits information about the users who have tagged photos about each specific destination:

$$IDFU_{i,j} = \frac{|P_j|}{|P_{i,j}|} \frac{|U_{i,j}|}{|U_j|}$$

where $|P_j|$ is the number of photos for the destination j, $|P_{i,j}|$ is the number of photos for the destination j tagged with t_i, $|U_{i,j}|$ is the number of users tagging destination j with the tag t_i, and $|U_j|$ is the number of distinct users who tagged at least one photo in j. We calculated tag weights for each city by normalizing tag frequencies in four different ways: (System A) Top100 dataset, standard IDF; (System B) Random dataset, standard IDF; (System C) Random dataset, IDFP, (System D) Random dataset, IDFU. For each different normalization, we then calculated similarities between cities in terms of the cosine distance between their matching vectors:

$$\text{similarity}(a, b) = \cos(\alpha) = \frac{a \cdot b}{|a||b|}$$

Front-end implementation

The interface was based on an existing open source project, Geoplanet Explorer (http://isithackday.com/geoplanet-explorer/). It was implemented as a PHP Web based application accessing a MySQLdatabase.

System Evaluation

To evaluate the system, an online survey was created in order to ask users which of the four similarity measures (i.e. System A - Top100 dataset, standard IDF; System B -Random dataset, standard IDF; System C - Random dataset, IDFP; System D - Random dataset, IDFU) was the best one according to them. The survey was composed by a demographic section and a group of cities to be evaluated. Users were chosen randomly by posting the survey link on popular social networks. After the demographic section the survey proposed to users clear instructions explaining that the similarities between places are not necessarily geographic and are based on place descriptions (i.e. on harvested information – Figure 2).

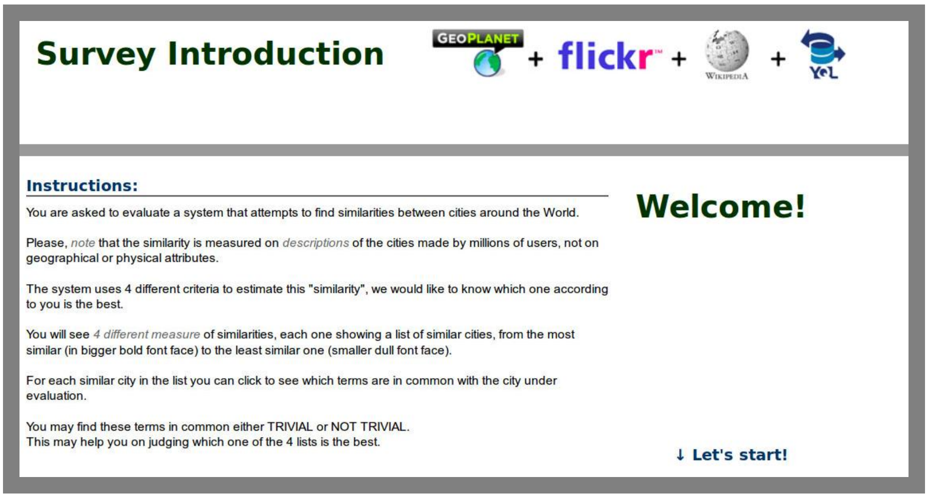

Figure 2 – Welcome Page and Survey Instructions

A small group of cities has been selected for the survey with a high percentage of European tourism destinations. Additionally a small group of cities from Asia, South America and the USA was inserted in the test. Each user judged one (random) city at a time, for a total of five distinct cities. For each of them the user was provided results of the four scoring systems, in the form of four lists containing the top-five related cities.

5 Results

The survey user interface was designed to give basic information about the city (i.e. country, administrative regions, map, etc.), in order to let users easily identify the tested city (in this case Rome, Figure 3). In the front end the sections are clearly defined: (1) information about the place taken from Yahoo! GeoPlanet (http://developer.yahoo.com/geo/geoplanet/), which is used to disambiguate the city from its homonyms (e.g. Rome in Georgia, US); (2) main city pictures from Flickr.com; (3) Map from Yahoo! maps; (4) the four systems to be rated; (5) the most popular 100 tags from Flickr.com.

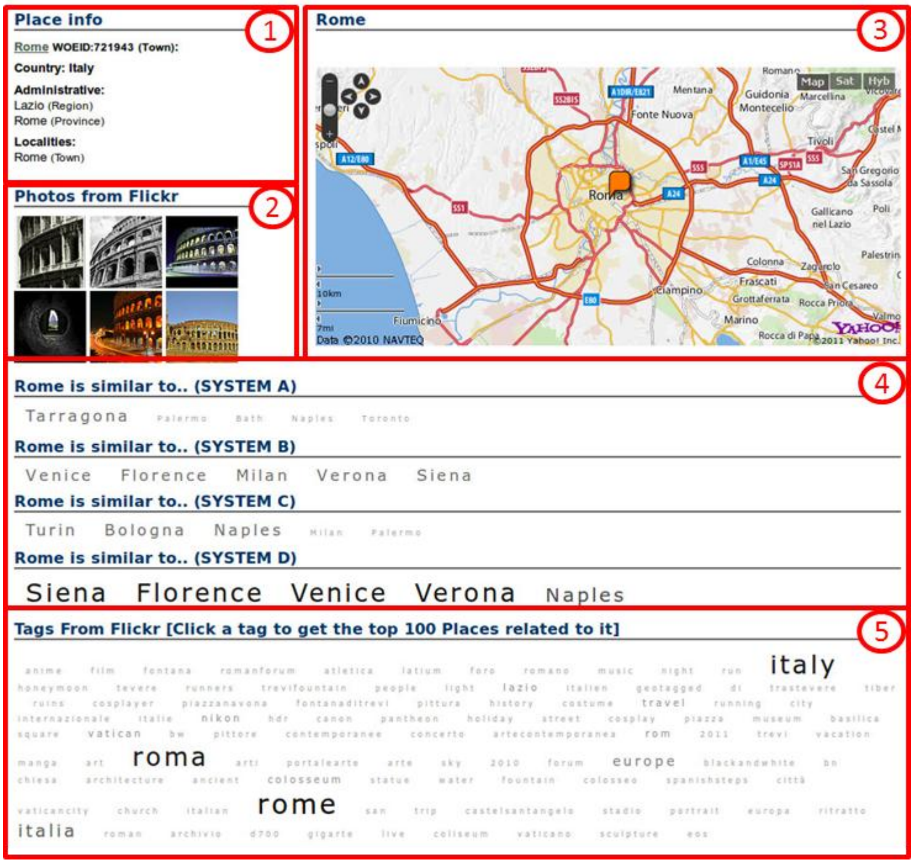

Figure 3- Graphic User Interface

In section 4, users could easily understand not only the similarity ranking, but also (to some extent), the *level* of similarity thanks to the differing font sizes (the bigger the font the higher the similarity explained on the survey starting page). Since the similarity between two places is based on the similarity of their descriptions, users had the possibility to check the tags in common between two cities. This is clearly a simplification, as the adopted similarity metrics were not based on a simple term match. It was considered useful to provide a rough idea of why two cities had been considered similar.

Evaluation

The survey was filled out by 113 users, mostly master students from Università della Svizzera italiana and Politecnico di Milano. Users produced 516 valid answers. The final result (Figure 4) yielded a great deal of useful information.

Figure 4- Survey Results

The best system according to users is System D (n=139 preferences). System D is based on the weighting scheme in which a greater importance was given to the user factor.

System A represents the second most popular answer (n=118 preferences) and it shows that we can extract valuable knowledge gathering information even from the *Top100* dataset. In other words, using only the most 100 popular tags for a city and the classical IDF we can still have good performance, even better than other methods applied to the bigger dataset.

The random dataset based on picture information (System C) ranked third according to users' preferences in the survey. Finally System B, differing from A only for the dataset used, ranked last in users' preferences. "N" represents the "no answer" group (n=84). Returning to the above mentioned research questions, it is thus possible to claim that (RQ2) the additional information coming from user tagging activities (i.e. who tagged a photo with a given tag) is much more relevant, while defining the similarity of two places or destinations, than (RQ1) picture-related information (i.e. which pictures have been tagged with a given tag).

6 Discussion and Conclusion

Surprisingly System A, relying on the much more limited *Top100* dataset, ranked better than other two systems (i.e. Systems B and C) that were using an extended sample based on the random harvesting of picture tags obtained with Yahoo Query Language. This means that the top 100 tags Flickr.com allows to be downloaded with its API are representative enough of the destinations and can be used to compare them with discrete results. Furthermore, the relevance of user-related information seems to be a driver to define similar cities/destinations.

This confirms the validity of other approaches (such as the "Wormholes" one) and provides a useful insight for the development of new tag- and user-based recommender systems. Actually, with the current model and technology this system is already able to suggest and/or recommend to users, that own a collection of pictures of a given place/destination in a popular social network such as Flickr.com, to visit other places/destinations without asking them any additional information about their preferences.

7 Limitation and Future Work

Limitation includes three issues. First, it was assumed that random sample was more precise in shaping destination similarity, but from the results of this research it does not seem to provide more information. Its real advantage, instead, is the additional information about users and photos that can be exploited to provide a better similarity measure. Second, demographic data was asked within the survey but not used in the final evaluation due to a shortfall of the system. Finally, as the current research was a starting point for a more in-depth research a light evaluation methodology with a snowball sampling was used.

Future work will cover these limitations. It will study, in depth, the difference between top 100 samples and random samples using both the case information about users and information about pictures to define if these preliminary findings can be confirmed. The demographic data will be added within the evaluation phase to try and leverage users' travel preferences and history. Finally a better sampling methodology will be used.

References

Ahern,S., Eckles,D., Good, NS., King, S. (2007) Over-exposed?: privacy patterns andconsiderations in online and mobile photo sharing. In Human Factors in Computing Systems, 2007.

Blackshaw, P. (2006). *The consumer-generated surveillance culture*. Retrieved October 13, 2008, from http://www.clickz.com/showPage.html?page=3576076.

Blackshaw, P., Nazzaro, M. (2006). *Consumer-generated media (cgm) 101: Word-of-mouth in the age of the web-fortified consumer*. Retrieved August, 1, 2007.

Buhalis, D. (2000). Marketing the competitive destination of the future, *Tourism Management.* Vol.21(1), 97-116.

Buhalis, D. (2003). *eTourism: Information technology for strategic tourism management*. Prentice Hall, Harlow.

Clements,M., Serdyukov,P., and De Vries, A. (2010). *Finding Wormholes with Flickr Geotags.* Advances in Information, 658–661, 2010.

Govers, R, Go, F. (2005). Projected Destination Image Online: Website Content Analysis of Pictures and Text. *Information Technology and Tourism*, Vol. 7 pp. 73–89

Gretzel, U. (2006). Consumer generated content - trends and implications for branding. e-*Review of Tourism Research*, 4(3), 9-11.

Gretzel, U., Fesenmaier, D. R., & O'Leary, J. T. (2006). The transformation of consumer behaviour. *Tourism Business Frontiers: Consumers, Products and Industry*, 9-18.

Gretzel, U.,Yu-Lan, Y., Fesenmaier, D. (2000). Preparing for the New Economy: Advertising Strategies and Change in Destination Marketing Organizations. *Journal of Travel Research*, Vol. 39, No. 2, 146-156.

Hotho, A., Jaschke,R., Schmitz, C. and Stumme, G. (2006). *Information retrieval in folksonomies: Search and ranking*. Proceedings of the 3rd European Semantic Web Conference, Lecture Notes in Computer Science. Springer.

Inversini, A., Buhalis, D. (2009). *Information Convergence in the Long Tail. The Case of Tourism Destination Information*. In W. Hopken, U. Gretzel & R. Law (Eds.), Information and Communication Technologies in Tourism 2009 - Proceedings of the International Conference in Amsterdam, Netherland (pp. 381-392). Wien: Springer.

Inversini, A., Cantoni, L. & Buhalis, D. (2009). Destinations’ Information Competition and Web Reputation, *Journal of Information Technology & Tourism,* 11 (2009), pp. 221-234.

Inversini, A., Cantoni, L. (2009). Cultural Destination Usability: The Case of Visit Bath. In W. Hopken, U. Gretzel & R. Law (Eds.), Information and Communication Technologies in Tourism 2009 – Proceedings of the International Conference in Amsterdam, Netherland (pp. 319-331). Wien: Springer.

Linaza, A., Agirregoikoa, A. Garcia, Torres, J., and Aranburu, K. (2011*). Image-based Travel Recommender System for small tourist destinations.* In R. Law, M Fuchs, Ricci, F (Eds.), Information and Communication Technologies in Tourism 2011 - Proceedings of the International Conference in Innsbruck, Austria (pp. 1-12). Wien: Springer.

Litvin, S. W., Goldsmith, R. E., & Pan, B. (2008). Electronic word-of-mouth in hospitality and tourism management. *Tourism Management*, 29, 458-468.

Nicholas, D., Huntington, P., Jamali, H.J. Dobrowolski, T. (2007). Characterizing and evaluating information seeking behavior in digital environment: spotlight on the bouncer. *Information processing and Management*, 43(4), 1085-1102.

Pengiran-Kahar, D.I.N.F., Syed-Ahmad, S.F., Ismail S.H.S., and Murphy, J.(2010). *Shared Arabian Muslim Travel Photos.* In Gretzel, U., R. Law, M Fuchs, (Eds.), Information and Communication Technologies in Tourism 2010 - Proceedings of the International Conference in Lugano, Switzerland. Wien: Springer, 591-603.

Poon, A. (1993). *Tourism, Technology and Competitive Strategies,* Wallingford, CT: CAB International, Oxford.

Sturtz, D.N. (2004). *Communal Categorization: the Folksonomy*. INFO622: Content Representation, December.

Vander Wal, T. (2004). *You down with Folksonomy?* - Reteived March 28, 2008 from http://www.vanderwal.net/random/entrysel.php?blog=1529

Werthner, H., Klein, S. (1999). *Information Technology and Tourism – A Challenging Relationship*. Wien -New York, Springer Verlag.

Wu, H. Zubair, M., Kurt, M. (2006). *Harvesting social knowledge from Folksonomies*, HT 2006, Odense, Denmark.

Yoo, K.H. and Gretzel, U. (2009). *Trust in travel-related consumer generated media*. In W. Höpken, U. Gretzel, and R. Law (Eds.), Information and Communication Technologies in Tourism, Wien: Springer, 49-59.