

Enriching User Reviews Through An Opinion Extraction System *

Enriqueciendo revisiones de usuarios mediante un sistema de extracción de opiniones

F. Javier Ortega José A. Troyano Fermín L. Cruz Fernando Enríquez
Universidad de Sevilla Universidad de Sevilla Universidad de Sevilla Universidad de Sevilla
javierortega@us.es troyano@us.es fcruz@us.es fenros@us.es

Resumen: Las webs basadas en el contenido generado por usuarios (UGC) tienen una aplicabilidad potencial en un gran número de campos. En este trabajo realizamos un estudio de la utilidad de estos sistemas para determinar la percepción de los usuarios expresada en sus opiniones sobre productos o servicios. Para ello, hemos compilado y analizado opiniones compartidas por usuarios en TripAdvisor, centrándonos en dos aspectos: el contenido estructurado y el no estructurado. Hemos realizado un análisis cuantitativo y cualitativo de la información extraída por un sistema de minería de opiniones, siendo este último especialmente interesante ya que ofrece información valiosa sobre los puntos fuertes y débiles de los hoteles según la percepción de los usuarios, yendo más allá de la información estructurada. Por último, hemos realizado un estudio de la complementariedad de la información estructurada y la no estructurada, observando un gran incremento de la cantidad de información disponible conjuntando ambas.

Palabras clave: minería de opiniones, contenidos generados por usuarios

Abstract: Web sites based on User-Generated Content (UGC) have a potentially valuable applicability in a number of fields. In this work we carry out a study of the usefulness of these systems from the point of view of detecting the perception expressed by users about services or items. We have compiled and analyzed opinions shared by users on TripAdvisor focusing on two aspects: the structured and the unstructured data. We perform a quantitative and a qualitative analysis of the information extracted by an opinion extraction system from our dataset, being the last one especially interesting since it provides valuable knowledge about the strong and weak points of hotels according to user perceptions, going beyond the structured data. Finally, we provide a study on the complementarity of the knowledge extracted from both, the textual opinions and the structured data, observing a noticeable increment of the amount of information available with the conjunction of both sources.

Keywords: opinion mining, user-generated content

1 Introduction

Review websites have become a useful Web 2.0 tool for on-line customers in their decision making process in order to gather information about a specific service or item before purchasing it. These websites usually integrate a recommender system intended to offer the adequate product to each user, based on

the previous opinions of users about similar products encoded in numerical ratings provided by users, or based on the opinions of other similar users about the same products. With the emergence of Opinion Mining, the analysis of textual opinions can be also a useful tool in this field. In this work we try to answer the research question about the extent in which Opinion Mining analysis can contribute to improve the quality of information that this type of systems can provide to their users.

In this sense, we distinguish in this work between the structured and the unstructured

* This work has been partially funded by the research projects AORESCU (P11-TIC-7684, Consejería de Innovación, Ciencia y Empresas, Junta de Andalucía), DOCUS (TIN2011-14726-E, Ministerio de Ciencia e Innovación) and ACOGEUS (TIN2012-38536-C03-02, Ministerio de Economía y Competitividad).

data provided by users. The first one is usually guided by the interfaces of these reviews sites and gathered through “Likes/Dislikes” schemata or “Stars” or any other form of asking the user for a numerical rating about the item being opinionated. The nature of this type of information makes it easier to process. On the other hand, the unstructured data consists mostly of textual opinions written by users in natural language, usually without any kind of predefined pattern. For this reason, we need to pre-process this information in order to extract useful knowledge from it.

Among the diversity of topics being covered by review websites, maybe one of the most relevant in terms of industry and economy of many countries is tourism. This is the main reason to focus on this domain for our study. Another reason is the huge impact of these kind of on-line systems on tourism industry recently, as stated in many works (Buhalis and Law, 2008; Ye, Law, and Gu, 2009; Nieves and Haller, 2014; Yasvari, Ghassemi, and Rahrov, 2012; Noti, 2013; Aye et al., 2012), to the point that a new line of research, e-Tourism, has come up to address the opportunities and challenges arose from it.

Regarding the research on e-Tourism, there is a wide variety of works, covering the analysis of the trustworthiness that these new channels offer to the users (Cox et al., 2009; Munar and Jacobsen, 2013), the study of the eWord-of-Mouth (eWOM) phenomenon (Yasvari, Ghassemi, and Rahrov, 2012; Barbagallo et al., 2012), the analysis of the influence of on-line reviews on the number of hotel room bookings (Ye et al., 2011), recommendation systems on tourism (Kabassi, 2010; Goossen et al., 2013), or even the development of systems for tourism packaging (Agarwal et al., 2013).

In this work we perform a study of a dataset composed by user opinions written in Spanish about hotels in the Canary Islands (Spain) extracted from one of the most relevant websites on this topic: TripAdvisor¹. The study includes the evaluation of both, structured and unstructured data provided by users through TripAdvisor, analyzing the correlation between both types of information and their complementarity, in such way that we can measure the extent in which the

knowledge provided by an opinion extraction system can enrich the user reviews.

The rest of the paper is structured as follows. The compilation of the dataset object of our study is discussed in section 2. In section 3 we briefly introduce TOES (Cruz et al., 2013), the opinion extraction system used in our work. In section 4 we discuss the relations between the structured and the unstructured data from the dataset and the output provided by the opinion extraction system, from three points of view: a quantitative analysis, a qualitative analysis and a study of the complementarity between both types of information. Finally, we point out the conclusions and future work in section 5.

2 Dataset Compilation

The selection of sources for our dataset has been guided by the relevance of the websites in the area and the amount of information that could be retrieved from them. So, we have chosen TripAdvisor over others because it is one of the most widely used tourism-related website. We decided to work with user-generated reviews about hotels in a specific location, in this case the Canary Islands, due to their particularities as the unique subtropical area in Europe and the importance of the tourism industry in their economy, which assures a huge amount of hotels and user-generated reviews of them, with a high variety of tourists with different needs and perceptions. Other relevant characteristic of TripAdvisor is the fact that any user is allowed to write a review about any item in the system (hotels, restaurants, etc.) with the only requirement of indicating (by clicking on a checkbox) that they have been there. Such a relaxed policy guarantees the provision of a high amount of user-generated content, in spite of the possible detriment in the quality and veracity of the reviews which are out of the scope of this work. We performed a search-driven crawling from TripAdvisor, given that our aim is to gather as much information as possible about hotels in a specific location, as follows:

1. Perform a search against the website with the required location.
2. Retrieve the list of hotels registered in the website for the given location.
3. For each hotel, retrieve all the structured information and the opinions of users.

¹www.tripadvisor.com

We repeat these steps for each island, so our crawler obtained all the hotels located in the Canary Islands together with all the opinions in the system about them. Since we are interested in the characterization of the hotels and not in the creation of complete user profiles, we just retrieve the information about the opinions of the users about those hotels, leaving out the opinions of those users about hotels in other places.

The crawler has been implemented in Java using two well-known libraries in this task: *HtmlUnit*² and *WebHarvest*³.

Metrics	Total	Spanish
Hotels	403	381
Reviews	78,535	12,950
Revs./Hotel	194.87	33.98
Users	68,441	11,039
Revs./User	1.14	1.17
Sentences	308,998	90,234
Sents./Rev.	3.93	6.96
Words	7,122,747	2,406,330
Words/Rev.	90.69	185.81

Table 1: Size of our dataset in terms of number of reviews, hotels and users, in addition to the number of sentences and words in the documents. The third column contains the same metrics applied only to those reviews written in Spanish.

The resulting resource after the execution of the above mentioned crawler is a dataset formed by structured and unstructured data about hotels in the Canary Islands and user reviews written in 2012 about those hotels in TripAdvisor. The structured data about the hotels consists of: name of the hotel, category (in the range of 0-5 stars), location, and the average of the scores provided by the users. About the opinions, we have gathered the user who wrote the opinion, the origin of the user, the profile (whether the user has traveled “solo”, i.e. alone, with friends or with family), the textual opinion, and a set of detailed scores given by the users to six specific features: location, service, comfort, cleanliness, rooms and quality of the hotel, in addition to the overall score of each hotel. Table 1 contains some metrics of the resulting dataset distinguishing the whole collection and the subset formed by reviews written by Spanish users. In the table we show the number

²<http://htmlunit.sourceforge.net>

³<http://web-harvest.sourceforge.net>

of hotels, reviews and users, the amount of textual information retrieved in terms of the number of sentences and words within the reviews and their average per review.

As a simple way of validating the compiled dataset, in Table 2 we show a comparison among the origin of tourists in the Canary Islands according to the gathered reviews and an official study carried out by a government institution, ISTAC⁴, in the same period of the reviews in our dataset (2012). This official study consists of a personal interview to tourists in the main airports of the islands.

Origin	ISTAC	TripAdvisor
Germany	25.98%	2.49%
Belgium	2.74%	1.21%
France	2.79%	2.53%
UK	22.04%	51.06%
Netherlands	3.50%	0.51%
Ireland	1.50%	2.95%
Italy	1.95%	3.24%
Spain	22.31%	16.49%
Others	17.18%	19.53%

Table 2: Percentage of opinions in each resource according to the origin of users. The first column corresponds to the official statistics compiled by ISTAC in 2012, the second one corresponds to data from TripAdvisor.

As shown in Table 2, we can see that most of the users that write opinions in TripAdvisor about hotels in the Canary Islands are from the United Kingdom and Spain. On the other hand, the official statistics from ISTAC show that German, Spanish and British tourists add up to about 70% of the total number of opinions. Although there are some differences, these can be caused by the different nature of the compilation methods used by both sources. Nevertheless, in general we can see that our dataset is qualitatively comparable to the one from ISTAC, meaning that it can be considered a good sample of the tourists in the Canary Island.

3 Domain-adaptable Opinion Extraction System: TOES

The aim of this research work is to study the extent in which an opinion extraction system can be useful in order to enrich the

⁴Instituto Canario de Estadística, the Canary Islands Government http://www.gobiernodecanarias.org/istac/temas_estadisticos/sectorservicios/hosteleriayturismo/demanda/

user opinions within a reviews website. To that end, we have compiled a dataset containing a huge amount of user reviews about hotels. Next we need an opinion extraction system focused on this domain. Thus, we take advantage of TOES (Cruz et al., 2013), a domain-adaptable opinion extraction system that can be easily applied to our study.

TOES is intended to detect and classify the opinions in a text. The underlying idea is to capture knowledge about a particular product class and the way people write their reviews on it. This process consists in two phases: first, it detects the pieces of text expressing individual opinions about specific features of the item being opinionated; in the second step, it computes the polarity of each individual opinion and the intensity of the polarity, and assigns a score in the range $[-1,1]$, representing -1 the most negative polarity and 1 the most positive.

TOES needs a training phase where a set of resources adapted to the domain are built. Some resources are automatically induced from a corpus of annotated reviews, while others are manually generated by an expert with some computational assessment. The training corpus is tagged by an expert aided by TOES. A taxonomy is built from the *feature words*, defining the characteristics that users are expected to write about. Using the taxonomy and the annotations of the expert, TOES builds a set of domain-dependent resources which are used for the detection and classification of opinions.

In our case TOES has been trained using a set of user-generated hotel reviews in Spanish extracted from TripAdvisor. This training set is formed by randomly chosen hotels from touristic Spanish cities like Madrid, Mallorca, Seville, etc., explicitly excluding the Canary Islands, so none of the hotels in the training dataset appear in our original dataset. Some metrics of the training set are shown in Table 3, including the number of annotated opinions that users express in their reviews.

Number of Reviews	1,200
Number of Words	213,843
Words per review	178.20
Annotated opinions	7,720

Table 3: Statistics of the reviews in the training set. The annotated opinions are the features commented by users in their reviews.

Once the domain-dependent resources have been created, TOES can extract user opinions from other texts on the same domain and also classify the polarity for each opinionated feature, determining whether the user expresses a positive or a negative opinion. Specifically, TOES provides, for each textual opinion, the set of features within that text in addition to the opinion words referring to the feature and the polarity of the opinion. In Figure 1 we can see a pair of input text and its corresponding output as an example.

INPUT:

Excelente ubicación para olvidarte del mundo. El personal es encantador. La piscina excelente. El restaurante, a pesar de tener una buena cocina falta variedad, por ejemplo en el desayuno no hay ni croissant, no hay opción de bebidas calientes (café) sino no esta abierto el restaurante y está cerrado por la tarde hasta las 19:00.

TOES OUTPUT:

1, 1, 0.050, 0.950, Excelente ubicación para olvidarte del mundo
 1, 1, 0.032, 0.968, El personal es encantador
 1, 1, 0.050, 0.950, La piscina excelente
 1, 0, 0.991, 0.009, El restaurante, a pesar de tener una buena cocina falta variedad, por ejemplo en el desayuno no hay ni croissant, no hay opción de bebidas calientes (café) sino no esta abierto el restaurante y está cerrado por la tarde hasta las 19:00.

Figure 1: Output provided by TOES given the input text. The columns correspond to the identification of the document, the polarity of the opinion and the negative and positive scores, respectively, computed by TOES.

For more details on the performance of TOES on other domains and a thoroughly explanation of its characteristics, the interested reader can review (Cruz et al., 2013).

4 Enrichment of user reviews

Once our dataset is processed by the opinion extraction system, let us proceed to the study of the application of these results. We evaluate the contribution of the opinion extraction system to the user reviews from two points of view: quantitative and qualitative. Finally we study the contribution of the opinion extraction system in terms of knowledge gain. In order to perform these evaluations properly, we have manually mapped the categories offered by TripAdvisor to the taxonomy used by TOES (see Table 4).

4.1 Quantitative evaluation

From a quantitative stance, we show in this section an evaluation based on the compar-

TripAdvisor	TOES
Quality	Building, Hotel, Price
Comfort	Bed
Rooms	Rooms, Television, Bathroom, Facilities
Cleanliness	Cleanliness
Location	Location, Views
Services	Services, Staff, Internet, Food/Drink

Table 4: Mapping between the feature taxonomies of TOES and TripAdvisor.

ision of the information extracted by TOES and the structured information provided by users in the review websites. The aim of this evaluation is to assess the correlation between both types of information

After applying TOES to the textual opinions from our dataset, we can highlight some conclusions. First, we plot in Figures 2 and 3 a comparison between the distributions of frequencies of the scores in TripAdvisor and the textual opinions extracted by TOES, respectively, showing the number of hotels (x-axis) with respect to the number of opinions (y-axis) about each feature.

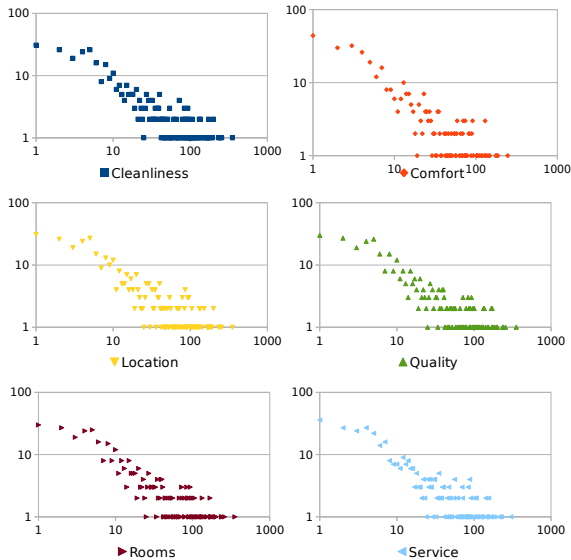


Figure 2: Number of numerical ratings (log-log) provided by users in TripAdvisor for each pre-established feature: Cleanliness, Comfort, Location, Quality, Rooms and Service.

In Table 5 we show a comparison between the scores given by users to each feature in the TripAdvisor taxonomy and the information extracted by TOES from the textual reviews of the users. This table has been computed by aggregating the count of opinions

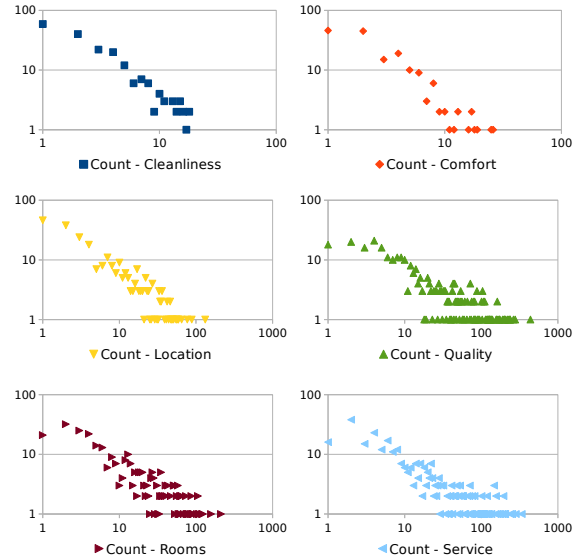


Figure 3: Number of opinions (log-log) provided by users in their textual reviews, according to TOES, for each pre-established feature: Cleanliness, Comfort, Location, Quality, Rooms and Service.

Features	TripAdvisor		TOES	
	Pos.	Neg.	Pos.	Neg.
Cleanliness	79.07%	20.93%	78.98%	21.02%
Comfort	63.95%	36.05%	66.41%	33.59%
Location	75.79%	24.21%	86.51%	13.49%
Quality	79.39%	20.61%	80.19%	19.81%
Rooms	78.72%	21.28%	81.93%	18.07%
Services	82.90%	17.10%	84.83%	15.17%
Average	76.64%	23.36%	79.81%	20.19%

Table 5: Percentage of positive and negative opinions per feature according to the scores in TripAdvisor (columns 2 and 3) and TOES (columns 4 and 5).

per feature for all the hotels in our dataset. Note that TripAdvisor allows its users to provide a score in the range [0-5] (stars) to each feature. We have considered as negative those scores < 3.

The columns corresponding to TripAdvisor scores have been obtained by computing the average of the scores given by users to each feature of the hotels in the dataset, and analogous for the scores in the columns corresponding to TOES. Regarding the data in Table 5, we observe that the overall results obtained by TOES from the textual opinions are fairly close to those expressed by users through the scores. In fact, the average of the differences between the scores from TripAdvisor and the polarity of opinions computed

by TOES is only 3,01%.

4.2 Qualitative evaluation

Through the quantitative evaluation we assess the reliability of TOES by comparing its results to the scores provided by the authors of the reviews to each feature in the TripAdvisor taxonomy. In this section, through the qualitative evaluation we show the capability of TOES of providing a finer-grained information by taking the analysis to the word-level of the specific terms that authors use to express their opinions.

With respect to the vocabulary used by users to express their opinions in the reviews, we can determine those words which are more frequently used in any, positive or negative opinions. One of these sets of words are represented through a word-cloud in Figure 4, corresponding to the most frequent feature words found in negative opinions.



Figure 4: Cloud of the most frequent feature words mentioned in negative user opinions, according to TOES.

From the figure we can infer that “comida” (food), “buffet” and “spa” receive most of the negative comments (obviously the word “hotel” is common in this domain for both, positive and negative opinions). On the other hand, in Table 6 we use another representation in order to highlight the top ten most used words in positive reviews.

These type of analyses can be performed for each hotel, being a useful tool for the providers of items being opinionated, in this case tourist services, in order to detect the pros and cons of the items provided in a finer grain than the one usually offered by the reviews website.

Words	Count
hotel	4,488
personal	2,656
habitacion	3,792
comida	2,006
piscina	1,584
servicio	979
trato	934
buffet	686
zona	624
limpieza	590

Table 6: Top 10 most used words in positive reviews in our dataset according to TOES.

4.3 Complementarity of informations

Given the nature of the structured information and the usability of the methods intended to gather it (usually the user must click on a number of stars or something analogous) in contrast to the more laborious activity of actually writing a text, we expected that most of the reviews contain a numerical value for the features proposed by TripAdvisor, while a smaller amount of them will include a proper written opinion. Table 7 shows the percentage of reviews without scores for each one of the features proposed by TripAdvisor, and also the percentage of reviews without textual opinions extracted by TOES for each feature.

Surprisingly, a higher percentage than expected of users do not provide numerical scores to all the features proposed by TripAdvisor. The case of *Comfort* is shocking: only 19.63% of reviews have a score, and less than 25% of them contain an opinion about it. The question now is: how the unstructured information can help to improve this lack of coverage of the structured one? In Figure 5 we plot a comparison of the percentage of reviews that contain structured information (scores) and written opinions about each feature in TripAdvisor, in addition to the union and intersection of both sets.

The most interesting observation in Figure 5 is provided by the last two columns of each feature: $TOES \cup TripAdvisor$ represents the percentage of user reviews with either, a score or a written opinion about the feature, while the column tagged as $TOES \cap TripAdvisor$ represents the percentage of user reviews that have both types of information. In other words, they correspond

Features	No Scores	No Text. Ops.
Cleanliness	58.80%	75.31%
Comfort	80.37%	76.42%
Location	58.99%	60.59%
Quality	58.53%	24.81%
Rooms	58.49%	47.54%
Service	68.60%	24.49%

Table 7: Percentage of reviews without scores or textual opinions in TripAdvisor, respectively, for the given features over the total of 12,950 reviews in Spanish in our dataset.

to the union and the intersection of those sets, respectively. These metrics highlight the improvement achieved by the inclusion of an automatic opinion mining tool like TOES in the system. Furthermore, the intersection of both sets is smaller than expected with only about 20% of user reviews, which means that, most of the times, users tend to provide only one type of information for each feature. Since there are a higher percentage of written opinions per feature than scores (except for *Cleanliness* and *Location*), we can state that a high percentage of users tend to score those features that they have not commented on.

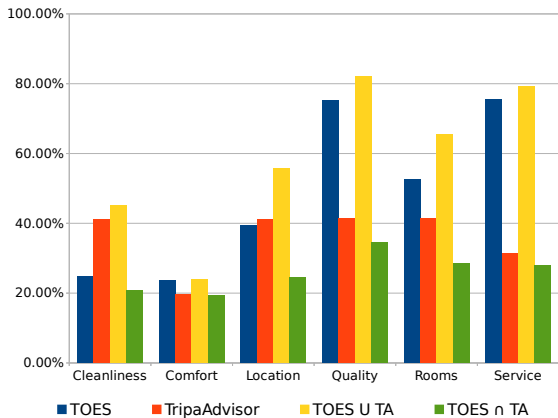


Figure 5: Percentage of reviews with numerical scores (TripAdvisor) and written opinions (according to TOES), together with the conjunction and the intersection of both sets, $TOES \cup TA$ and $TOES \cap TA$, respectively.

Following this idea, we can extract the *complementarity* of both, structured information in the form of scores in TripAdvisor and unstructured information in the form of textual opinions. We define the *complementarity* as the increment of information with respect to the total reviews of each feature. Table 8 contains the results of this metric for each feature in terms of the percentage of re-

Features	Str Inf	Str+Unstr	Compl.
Cleanliness	41.20%	45.21%	5.10%
Comfort	19.63%	23.92%	17.88%
Location	41.01%	55.84%	14.92%
Quality	41.47%	82.20%	26.89%
Rooms	41.51%	65.44%	19.22%
Service	31.40%	79.12%	47.13%

Table 8: Percentage of reviews with scores (Str Inf), reviews with scores or textual opinions (Str+Unst) and the complementarity of both (Compl.): increment of reviews with respect to the total reviews of each feature.

views without scores for the feature considered but with opinions extracted by TOES, with respect to the total number of reviews with relevant information for each feature.

We can see in the table the noticeable complementarity of both sources for all the features proposed by TripAdvisor. According to this, in order to make a reliable recommender system based on user reviews, it should be mandatory to implement an opinion mining tool in order to make the most of the information provided by users, given that this unstructured information supports and even complements the structured data, providing a very useful source of additional knowledge about user opinions.

5 Conclusions

The e-Tourism research has an increasing interest on Opinion Mining and Recommendation Systems, since these fields can provide very valuable advances in the study of customers perceptions about the products or services enjoyed. In this work we have carried out a study intended to highlight in what extent an opinion extraction system can enrich or improve the information provided by users. In this sense, we have performed a quantitative study about the correlation between the numerical ratings and the knowledge extracted from textual opinions of users, in order to check whether these sources express similar perceptions. In our case, the correlation between the ratings in TripAdvisor and the opinions extracted from the textual reviews is very clear, with a 3% of difference in average between both types of information. On the other hand, we have performed a qualitative analysis of the output of an opinion extraction system in terms of the added-value obtained by the analysis of a finer grained and more detailed information

present in the textual opinions and not in the numerical ratings. Finally, we have studied the complementarity of both sources of information, obtaining a metric representing the contribution of the analysis of textual opinions to the structured information, showing that a resource built from both sources contains up to 47% more opinions on some features than using just the numerical ratings.

We plan to further our work by developing a method to automatically integrate structured and unstructured information in an aspect-based recommendation system, in addition to the study of the integration of multilingual opinion extraction systems to take advantage of the huge amount of textual opinions in diverse languages.

References

- Agarwal, J., R. H. Goudar, N. Sharma, P. Kumar, V. Parshav, R. Sharma, and S. Rao. 2013. Cost effective dynamic packaging systems in e-tourism using semantic web. *International Conference on Advances in Computing, Communications and Informatics*, pages 1196–1200.
- Ayeh, J. K., D. Leung, N. Au, and R. Law. 2012. Perceptions and strategies of hospitality and tourism practitioners on social media: An exploratory study. In Matthias Fuchs, Francesco Ricci, and Lorenzo Cantoni, editors, *Information and Communication Technologies in Tourism*, Vienna. Springer Vienna.
- Barbagallo, D., L. Bruni, C. Francalanci, and P. Giacomazzi. 2012. An empirical study on the relationship between twitter sentiment and influence in the tourism domain. In *Information and Communication Technologies in Tourism*. pages 506–516.
- Buhalis, D. and R. Law. 2008. Progress in information technology and tourism management: 20 years on and 10 years after the internet—the state of etourism research. *Tourism Management*, 29(4):609–623.
- Cox, C., S. Burgess, C. Sellitto, and J. Buultjens. 2009. The role of user-generated content in tourists’ travel planning behavior. *Journal of Hospitality Marketing & Management*, 18(8):743–764, oct.
- Cruz, F. L., J. A. Troyano, F. Enríquez, F. J. Ortega, and C. G. Vallejo. 2013. ‘long autonomy or long delay?’ the importance of domain in opinion mining. *Expert Systems with Applications*, 40:3174–3184.
- Goossen, M., H. Meeuwssen, J. Franke, ‘A. Maps Á Tourism, and Land. 2013. Á Destination inspiration using etourism tool. In Zheng Xiang and Iis Tussyadiah, editors, *Information and Communication Technologies in Tourism 2014*. Springer International Publishing, Cham.
- Kabassi, K. 2010. Personalizing recommendations for tourists. *Telematics and Informatics*, 27(1):51–66, February.
- Munar, A. M. and J. Kr. Steen Jacobsen. 2013. Trust and involvement in tourism social media and web-based travel information sources. *Scandinavian Journal of Hospitality and Tourism*, 13(1):1–19, April.
- Nieves, J. and S. Haller. 2014. Building dynamic capabilities through knowledge resources. *Tourism Management*, 40:224–232, February.
- Noti, E. 2013. Web 2.0 and the its influence in the tourism sector. *European Scientific Journal*, 9(20):115–123.
- Yasvari, T. H., R. A. Ghassemi, and E. Rahrovy. 2012. Influential factors on word of mouth in service industries (the case of iran airline company). *International Journal of Learning and Development*, 2(5):227–242, October.
- Ye, Q., R. Law, and B. Gu. 2009. The impact of online user reviews on hotel room sales. *International Journal of Hospitality Management*, 28(1):180–182, March.
- Ye, Q., R. Law, B. Gu, and W. Chen. 2011. The influence of user-generated content on traveler behavior: An empirical investigation on the effects of e-word-of-mouth to hotel online bookings. *Computers in Human Behavior*, 27(2):634–639, March.