Twitter sentiment analysis: Capturing sentiment from integrated resort tweets

Article *in* International Journal of Hospitality Management · May 2016

DOI: 10.1016/j.ijhm.2016.02.001

2 authors:

Kahlil S. Philander
Washington State University

Yunying Susan Zhong
Rosen College of Hospitality Management, University of Central Florida

International Journal of Hospitality Management 55 (2016) 16–24

Contents lists available at ScienceDirect

International Journal of Hospitality Management

journal homepage: www.elsevier.com/locate/ijhosman

Twitter sentiment analysis: Capturing sentiment from integrated resort tweets

Kahlil Philander[a,1], YunYing Zhong[b,*]

[a] William F. Harrah College of Hotel Administration, University of Nevada, Las Vegas, 4505 S. Maryland Parkway, Las Vegas, NV 89154, United States
[b] Hospitality Management, College of Health and Human Services, California State University, Long Beach, 1250 Bellflower Blvd, Long Beach, CA 90840, United States

A R T I C L E I N F O

Article history:
Received 10 October 2014
Received in revised form 26 January 2016
Accepted 2 February 2016
Available online 11 March 2016

1. Introduction

With the rapid development of information and communication technologies, social networking sites and blogs have become an enormous repository for rich user generated content (UGC) (Thelwall et al., 2011). Information system scholars observed that business intelligence has evolved from analysis of traditionally structured business transaction data to studying unstructured and real-time data, driven by the development of online search engines, e-commerce, and social media (Chen et al., 2012). Among the existing collection of social media platforms, the Twitter microblog is one of the most popular mediums that have been adopted by both consumers and companies. Over half a billion Twitter messages (tweets) per day are being recorded (Kirilenko and Stepchenkova, 2014; Krikorian, 2013), and a recent customer engagement technology study reported that Twitter is among the top three most used social media platforms by US hotel and restaurant operators (Kim and Connolly, 2013).

The prevalence of Twitter and its prodigious amount of UGC creates important implications for the hospitality and tourism sector. Research has shown that online UGC has become an important information source that exerts critical influence on customers' brand perception, brand reputation, purchase decision making, and profitability (Browning et al., 2013; Leung et al., 2013; Vermeulen and Seegers, 2009; Ye et al., 2009, 2011; Zhang et al., 2010). Twitter may be particularly susceptible to the effects of electronic words of mouth (EWOM) due to its viral nature. According to Kwak et al. (2010) study, a tweet, if resent by a different user (retweeted), is expected to reach an average of 1000 registered users. In three years following Kwak et al.'s study, the total volume of daily tweets has grown over ten-fold (Krikorian, 2013; Weil, 2010).

While the viral influence of Twitter messages makes it valuable for firms to monitor and identify customer comments, the massive volume and high variance of real-time tweets has made that task extremely time consuming, costly, and often manually impossible (Chiu et al., 2015; Claster et al., 2013). Even though firms have started to embrace Twitter as a marketing platform, because of these obstacles, they often fail to extract much objective quantitative feedback or establish competitive benchmarks (i.e. comparing firm performance in social media against competitor performance) from this available data.

Sentiment analysis, which is still a rather nascent arena, appears to be a promising tool in solving the above mentioned problem. This technique involves using algorithms to automatically extract and classify text data into sentiment categories: positive, negative and/or neutral (Chiu et al., 2015; Pang and Lee, 2008). As compared to more traditional market research methods (e.g. surveys or opinion polls), sentiment analysis not only has the advantage of being more cost and time efficient in many cases, it is a non-intrusive method to extract consumers' opinions and sentiments in "real-time"—avoiding recall biases generally associated with post-consumption self-report measurements (Rylander et al., 1995). Furthermore, sentiment analysis can provide a temporal sentiment profile even on a second by second scale, which is not generally possible for survey based market analysis to achieve.

In this study, we demonstrate the application of sentiment analysis using Twitter data to build low-cost and real-time measures of hospitality customer attitudes/perceptions. Using a popular tourism destination (Las Vegas, NV) as a case study, we create a sentiment index for every Twitter account belonging to an

* Corresponding author. Tel.: +1562 985 4496.
E-mail addresses: kahlil.philander@unlv.edu (K. Philander), yunying.zhong@csulb.edu (Y. Zhong).
[1] Tel.: +1 702 895 1265.

http://dx.doi.org/10.1016/j.ijhm.2016.02.001

integrated-resort property in the Las Vegas metropolitan area. The ensuing sentiment metrics are then used to benchmark these firms against one another, and compare their performance over time. The metrics are also cross checked with external hotel-casino rankings (i.e. TripAdvisor) to establish external validity. Discussion of the usefulness of this business intelligence tool is provided, along with potential limitations.

2. Literature review

2.1. Twitter and its applications

Since its inception in 2006, Twitter has become one of the most influential microblogging platforms on the Internet (Akehurst, 2009; Thelwall et al., 2011). According to a recent article on Huffington Post, Twitter officially reported a total number of 218.3 million "monthly active users," or accounts in its 2013 October filing with the Security and Exchange Commission (Grandoni, 2013). In the United States alone, Nielsen (2012) estimated that in 2012, the microblogging site had over 37 million users via personal computer access, and about another 65 million via mobile devices. In addition to its large user base, Twitter users show a high level of engagement. Nielsen (2012) included Twitter among the top-three sites in terms of users' total time spent (with Facebook and Tumblr). A recent study conducted by the Pew Research Center also found that 46% of Twitter users tweet on a daily basis (Dugann and Smith, 2013).

With its large user base and high engagement, Twitter has increasingly been used by the hospitality and tourism industry for promotion, distribution, marketing management, communication, and market research (Leung et al., 2013). When studying the use of social media by the hotel industry in Hong Kong, Chan and Guillet (2011) found that Twitter and Facebook were the two most widely used social sites. The industry mainly used the sites to promote discounted products and services, answering guests' inquiries, and handling complaints. To a lesser extent, the sites were also employed to engage guests and obtain business intelligence. Twitter is also a popular social medium used by national destination marketing organizations (DMOs) for marketing purpose. Examining the social media activities of eight top international tourism destinations, Hays et al. (2013) uncovered that the DMOs generally posted more frequently on Twitter than on Facebook. The researchers explained that Twitter structure was built for timely updates and massive information broadcast. Due to the constant updates on this platform, tweets by any given users can be rapidly buried. As a result, organizations need to constantly refresh its updates to remain visible.

In the area of market research, Twitter data tends to be utilized more than those from other social media platforms. This is in part because of Twitter's more liberal data availability; Facebook messages are often not publicly accessible due to its stricter privacy settings, and other platforms generally have smaller user pools (Wang et al., 2013). The high utility of Twitter data for research can also be attributed to its "people as sensor" network structure. Twitter has a fundamentally open networking structure, which allows its users to choose who they want to follow without seeking permission (Weng et al., 2010). This open structure of Twitter has made it into a "global pipeline for real time information sharing and broadcasting" (Wang et al., 2011, p. 34). The interactions among Twitter users can be viewed as a large network of sensors that react to external and social events, thus making it particularly suitable for studying the public opinions of these events (Kirilenko and Stepchenkova, 2014).

Aside from its conceptualization as a sensor network, Twitter is conducive for market research due to the internal architecture of the messages. As a microblogging site, Twitter allows its users to post messages of a maximum 140 characters in length. These compact messages, also called "tweets", capture the primary meaning users want to express without much irrelevant content. The unique function of hashtags, "#", enables users to highlight the key words in their messages or find other posts under certain key words. Tweets can be forwarded (or "retweeted") into a social network of friends and contacts by the use of "*RT<@username>*" or via "*<@username>*". The "@" symbol is used to identify tweets which are addressed to a particular registered user (Cha et al., 2010). The condensed structure of tweets and Twitter's unique functions make its data a useful fit for content analysis and sentiment analysis, which involves systematic reduction of content flow for the purpose of quantitative analysis (Kirilenko and Stepchenkova, 2014).

Twitter has the advantage of capturing a rich amount of authentic customers' in-the-moment experience and sentiment (Capriello et al., 2011; Dodds et al., 2011). As compared to traditional research methods, such as surveys and comment cards, Twitter is a non-intrusive research model that suffers less from bias associated with recall and interactions with human subjects (Kirilenko and Stepchenkova, 2014). Asur and Huberman (2010) note that the collective wisdom on Twitter, if analyzed properly in a large enough volume, can be more accurate than other techniques in extracting information, such as surveys and opinion polls. Of course, Twitter data has limitations. For example, Twitter users are not representative of the general population, and Twitter posters are not necessarily representative of the entire Twitter user base–Twitter Inc. noted that 40% of users do not post any content. Boyd and Crawford (2012) caution that researchers should understand, and publicly account for, the limits of Twitter data set analyses.

2.2. Electronic word of mouth and its impacts

The attention social media receive from hospitality and tourism practitioners not only can be attributed to its value for business intelligence, but also to its powerful ability to spread word of mouth (WOM) immediately on a large scale (Leung et al., 2013). WOM is defined as informal communications among consumers concerning a product, service or an organization, and it is distinguished from those communications in which companies pass product knowledge to consumers through massive media (Litvin et al., 2008). Research has long recognized WOM as a trustworthy information source independent from commercial influences perceived by consumers (Leung et al., 2013). Traditional WOM and EWOM share a lot of similarity. However, the latter has the ability to reach beyond one's immediate social circle and transcend the ephemeral nature of offline WOM (Litvin et al., 2008). EWOM, often in the form of online reviews and comments, can be easily archived, linked, and searched. EWOM is increasingly important in shaping consumers' product knowledge and perceptions (Litvin et al., 2008)

A thorough literature review on hospitality and tourism EMOW-online review was conducted by Schuckert et al. (2015). Examining 50 relevant articles between 2004 and 2013, the researchers summarized five thematic clusters within the literature: the effect of online review on consumers' buying, online response management and consumer satisfaction, reviewer's motivation, sentiment analysis, and critical attributes of online review. The researchers noted more attention paid to the first two consumer-concentric topics and more studies (60%) done in the hotel setting. EWOM is also naturally integral with the social media literatures. After a comprehensive review on social media in tourism and hospitality, Leung et al. (2013) pointed out that EWOM serves as a critical source for consumer's pre-trip travel planning, and can be generated throughout their trip and post-trip phase.

As such, prior hospitality and tourism EWOM research can also be classified in two perspectives, listeners' and originators' (Litvin et al., 2008). The studies from the listener viewpoint mainly focus on the effect of EWOM on consumers' attitudes and behavioral intention during the planning phase. As Leung et al. (2013) noted, because tourism and hospitality related products are intangible and often costly, consumers tend to seek others' opinions to reduce risks associated with their decision making. Studies have shown that the valence of the EWOM has significant impacts on consumers' trust (Sparks and Browning, 2011; Ladhari and Michaud, 2015), brand awareness (Vermeulen and Seegers, 2009), and purchase intentions (Mauri and Minazzi, 2013; Vermeulen and Seegers, 2009). The number of reviews and the numerical rating also greatly affect consumers' perception, webpage visits, and booking behaviors (Sparks and Browning, 2011; Ye et al., 2009, 2011; Zhang et al., 2010). Furthermore, other information cues, such as source of reviews, writing style, and reviewer's expertise and personal identifiable information, can play a significant role in affecting consumer's decision making (Sparks et al., 2013; Xie et al., 2011). Based on these observed contributions of EWOM to critical marketing outcomes, Ghose and Ipeirotis (2009) posited that each snippet of UGC has economic value, and organizations with more favorable consumer evaluation can charge a price premium.

From the originator's perspective, many studies focus on consumers' motivations for originating and spreading WOM. Hennig-Thurau et al. (2004) found that the primary reasons for consumers to speak on the Internet are their desire for social interaction, economic incentives, concern for other consumers, and potential to enhance self-worth. Similarly, Litvin et al. (2008) summarized the findings of previous research, and suggested that the motivations of consumers' WOM behaviors include expressing positive feeling, discharging negative emotions, conducting altruistic acts, reciprocating with others, and enjoying shared experiences.

Consumers' active sharing can drastically change the course or meaning of marketing messages that are set forth by the companies. It is important for firms to closely monitor their social communications. This includes both the volume and the content, which enables companies to quickly identify trends or crises (Hennig-Thurau et al., 2010, 2013). Litvin et al. (2008) found that hospitality and tourism companies can take advantage of EWOM sharing to create "buzz" which amplifies their initial marketing efforts. This is achieved by creating or advertising something interesting that can get consumers talking. Research has supported the power of "buzz" with evidence suggesting that exposure to online comments enhance consumers' brand awareness (Vermeulen and Seegers, 2009).

From the listeners' perspective EWOM exerts substantial influences on attitudes, behaviors, and purchase intentions. From the originators' and sharers' perspectives, hospitality and tourism organizations can take advantage of consumer's EWOM to build marketing momentum centered on initial marketing activities. Generally, more positive EWOM has a more positive impact on consumers (Sparks and Browning, 2011; Vermeulen and Seegers, 2009; Ye et al., 2011). Sentiment analysis can be an effective tool for measuring these outcomes.

2.3. Sentiment analysis

To extract consumers' opinions and sentiments from a pool of qualitative data is not a novel task for marketing researchers. Traditionally, such a task can be performed by manual content coding. Researchers examine the data thoroughly and comprehensively, and develop their own sentiment categories independently. Inter-rater reliability is then calculated to determine the degree of homogeneity in their classifications. Divergent findings are discussed as an attempt to resolve disagreement (Gibbs, 2007). While manual analysis can better detect nuances in the text (e.g. sarcasm or implicit sentiments), it is highly time consuming and labor intensive. With the rapidly increasing volume of online data, manual analysis becomes pragmatically challenging, if not impossible. This gives strong impetus to the rise of a new area of research, sentiment analysis.

Sentiment analysis, also referred to as opinion mining, is a technique developed in the fields of artificial intelligence and natural language processing. It is an information retrieval tool that can classify text into subjective categories (negative, positive, or neutral) or measure sentiment strength (Pang and Lee, 2008; Thelwall et al., 2010). There are two major steps in sentiment analysis: opinion extraction and sentiment classification (Pang and Lee, 2008). Opinion extraction is to differentiate subjective texts from factual ones, while sentiment classification focuses on assigning opinion words into different sentiment categories (Chiu et al., 2015). *Opinion words* are words that express desirable (e.g. fantastic, amazing, etc.) and undesirable (terrible, disgusting, etc.) states (Ding et al., 2008).

There are two common methods determining an opinion word's semantic orientation or subjective categories: corpus based and dictionary-based approach (Chiu et al., 2015). Corpus-based approaches involve using the syntactic and co-occurrence pattern of a large corpus (texts that are most representative of a document's content) in identifying sentiment category (Capriello et al., 2011; Chiu et al., 2015; Thelwall et al., 2011). For instance, Turney and Littman (2003) determined a word's semantic orientation by calculating the strength of its association with a set of positive words minus its association with a set of negative words. The associations were estimated by issuing a search engine query, and then noting the query's co-occurrence probability with the positive and negative seed words.

A dictionary-based approach, also called a lexicon-based method, uses a bank of pre-coded words to determine the text's semantic orientation (Chiu et al., 2015; Thelwall et al., 2011). For example, Hu and Liu (2004) generated a set of adjective synonyms and antonyms (opinion words) through bootstrapping process using the WordNet dictionary. They then used this collection of opinion words to predict the sentiment orientation of electronic product reviews at a sentence level. The dictionary-based approach typically counts the numbers of positive and negative opinion words in a sentence. If positive opinion words prevail, the orientation of the sentence is positive and otherwise negative. Sentiment analysis can be conducted not only at a sentence level, but also at others levels: document-, paragraph-, or attribute-level. As the level of granularity increases, so does its complexity. Attribute-level sentiment analysis aims to associate opinions associated with certain features (Chiu et al., 2015).

Sentiment analysis has been successfully applied in various contexts, such as detecting influenza outbreak (Culotta, 2010), determining overall trends in the level of happiness (Dodds et al., 2011), predicting movie box office revenues (Asur and Huberman, 2010), and understanding some consumer opinions on tourism and hospitality related products (Chiu et al., 2015; Claster et al., 2013; Duan et al., 2013). Chiu et al. (2015) developed a Chinese sentiment analysis method to extract opinions toward various hotel attributes in Chinese blogs. Duan et al. (2013) employed the sentiment analysis technique to decompose user reviews into five dimensions of SERVQUAL (i.e. tangible, reliability, assurance, responsiveness, and empathy), and measured consumers' sentiments toward hotel service quality. They also included the five dimensions to determine their effects on consumer's satisfaction and review generating behaviors. Their study results showed that consumers' sentiments on tangibles had the strongest impact on both dependent outcomes. In the tourism context, Claster et al. (2013) utilized the technique to visualize the temporal sentiment fluctuations associated with three tourist destinations: Thailand, Sri Lanka, and Mexico. The researchers advocated sentiment analysis as a viable tool to understand and benchmark tourist's sentiments across destinations.

The rising popularity of sentiment analysis in research can be attributed to its unique advantages. First, as compared to manual coding, computer-aided sentiment analyses are not only more efficient, but also produce comparable results. Capriello et al. (2013) compare the efficiency of manual content coding and two computer-aided sentiment analyses techniques in analyzing 800 travel reviews of former farm stay guests. They found that all three analyses produce similarly reliable results. Second, as compared to traditional methods (e.g. surveys or focus groups), sentiment analysis can effectively reduce cost, time, and manual labor by using automatic algorithms to sort through text (Chiu et al., 2015). Researchers contend that sentiment analysis, when used in analyzing a free flow of user generated content, can discover dynamic and unforeseen insights beyond the numerical rating of pre-determined categories in surveys (Ghose and Ipeirotis, 2009). In addition, self-report measurements, such as in satisfaction surveys, often suffer from recall bias and question framing bias (Peterson and Wilson, 1992). Sentiment analysis of tweets can overcome these shortcomings, as Twitter users tend to post real time expressions that reflect their in-the-moment experience.

Despite its advantages, sentiment analysis also has many drawbacks. As Pang and Lee (2008) pointed out, sentiment analysis is domain and event dependent. Words considered positive in one domain might not be so in another area. Sarcasm is an obvious example of this limitation. Despite its inherent disadvantages, the technique still appears as a promising tool for researchers and industry practitioners. Wang et al. (2013) reviewed previous studies, and reported that the analysis technique yields a rather high accuracy rate (roughly 70% to 80% accuracy rate in training-test data matching tasks). The objective of sentiment analysis is to obtain useful insight from a large quantity of aggregated data, rather than perfect classification of all data points.

3. Methodology

This study uses a dictionary-based method to analyze social media microblogging data from Twitter. Data for this study was acquired from the Twitter application programming interface (API). The Twitter API is a backend server that warehouses all individuals' tweets and enables data collection by the public. The site allows roughly the past week of historical data to be extracted from its database (Twitter, 2015). While there are many methods that can be used to mine this data, we follow an approach adapted from Elder et al. (2012), and use the TwitteR package for the R programming language (Gentry, 2013). We also use the most recent version of Hu and Liu's (2004) sentiment lexicon to determine the sentiment orientation of the online text.

TwitteR provides a well-documented and accessible means to extract data into a commonly used data mining statistical program (KDnuggets, 2012). *TwitteR* is useful for many different data mining procedures, but one form that can be particularly useful to tourism firms is the *searchTwitter* function, which retrieves any tweet containing a specified text string. This can, for example, be used to identify any tweet mentioning a corporate Twitter account. A *searchTwitter* query for "@hotelxyz" will return any tweet, by any Twitter user, who mentioned that hotel's username.

The data mining frame used in this study was to retrieve any tweet that included an active integrated-resort username from the Las Vegas metropolitan area. We define the scope of Las Vegas as the Las Vegas Valley metropolitan area, and include only Twitter verified accounts or non-verified accounts with at least 5000 followers. Las Vegas is a useful market for examining this data mining procedure since it contains a large cluster of different branded hospitality firms that face similar market-wide conditions, and many of these properties have a large number of Twitter followers. In total, 34 properties (and associated Twitter accounts) were used as part of this study. The list of accounts/properties is provided in Appendix A, along with their Twitter account parameters.

To compare the social media satisfaction scores, a sentiment index was built using the data mined tweets. The sentiment score was constructed by scoring the tweet text for positive and negative words using Hu and Liu's (2004) sentiment lexicon. The sentiment score followed method used in Elder et al. (2012), where each positive word added a point to the total score and each negative word removed a point–there are no points added or subtracted for neutral words. The properties were then compared on two criteria: the average (mean) tweet sentiment score for any tweet that mentioned their corporate username (average score) and the ratio of positively scored tweets to negatively scored tweets for any tweet that mentioned their corporate username (ratio score). The average score is defined more formally as,

$$\text{Average Score}_{Rt} = \left(\sum_{i=1}^{n}\text{Pos}_{Rit} - \sum_{i=1}^{n}\text{Neg}_{Rit}\right)/n_{Rt}$$

where the average score at time = t is the difference between the sum of positive (Pos) individual words appearing in tweet i of n total tweets, and negative (Neg) individual words appearing in tweets (also from i of n), divided by the total number of tweets, for any tweet that mentions resort R.

The ratio score is defined more formally as,

$$\text{Ratio Score}_{Rt} = \frac{\sum_{n}^{i=1}\text{PT}_{Rit}}{\sum_{n}^{i=1}\text{NT}_{Rit}}$$

where,

$$\text{PT}_{Rit} = \begin{cases} 1 & \text{if} \quad \text{Pos}_{Rit} - \text{Neg}_{Rit} > 0 \\ 0 & \text{Otherwise} \end{cases}$$

$$\text{NT}_{Rit} = \begin{cases} 1 & \text{if} \quad \text{Pos}_{Rit}\text{-}\text{Neg}_{Rit} < 0 \\ 0 & \text{Otherwise} \end{cases}$$

The ratio score is computed as the ratio of overall positive tweets (PT) i of n at time t, to overall negative tweets (NT) i of n at time t. More simply, in a given time period, the sum of all the positive tweets for a given operator are compared to the sum of all negative tweets, and any neutral tweets are discarded. We illustrate these two scoring methods with a short example of two hypothetical tweets: 1) "Had a great time @AriaLV. Liked the hotel despite the *noise* at night," and 2) "The @AriaLV buffet was *awful*." Both these examples would be linked to the @AriaLV corporate account. The first tweet had two positive words (underlined) and one negative word (italicized), and the second tweet has one negative word. The average score for @AriaLV would therefore be: $((2+0)-(1+1))/2=0$, since there are two positive words in the first tweet and none in the second, along with one negative word in each tweet. The ratio score would be: 1/1 = 1, since there is one overall positive tweet and one overall negative tweet.

The average score provides a useful baseline to understand each firm's relative sentiment, while the ratio score provides a more meaningful overview of the most polarizing opinions, since it ignores tweets with sentiment scores of zero. The internal validity of the scores is compared using Cronbach's alpha, and concurrent validity is tested by comparing the measures to online rankings from travel review site, TripAdvisor, another source of EWOM. While we expect some levels of positive correlation between all of these metrics, we also expect a reasonable amount of variation that reflects different response time, demographics, and indirect measurement involved in a Twitter based interaction.

Table 1
Property twitter sentiment scores.

	Tweet volume	Avg sentiment score	Avg sentiment rank	Pos/Neg ratio score	Pos/Neg ratio rank	TripAdvisor relative rank	TripAdvisor score (%)
Aria Las Vegas	1383	0.69	1	10.72	2	5	84
Tropicana Las Vegas	1033	0.69	2	11.10	1	13	80
Venetian Las Vegas	1295	0.68	3	8.66	6	3	85
Bellagio Las Vegas*	806	0.66	4	9.60	5	2	88
Paris Las Vegas	424	0.64	5	9.90	4	18	73
Caesars Palace	1472	0.63	6	6.08	12	15	72
Red Rock Las Vegas	384	0.61	7	10.22	3	7	85
M Resort Spa Casino	350	0.59	8	5.96	14	N/A**	87
Luxor Hotel & Casino	1701	0.58	9	6.10	11	21	60
Monte Carlo Resort	824	0.57	10	6.81	8	17	70
Wynn Las Vegas	1569	0.56	11	6.79	9	4	84
Treasure Island	454	0.55	12	6.66	10	19	72
Circus Circus	556	0.55	13	5.68	18	20	60
Flamingo Las Vegas	711	0.55	14	5.42	20	29	55
Palazzo Las Vegas	862	0.53	15	7.04	7	1	88
LVH Hotel & Casino	401	0.52	16	5.86	16	32	47
South Point Hotel	780	0.52	17	5.88	15	12	84
Stratosphere Hotel	930	0.52	18	6.05	13	26	66
Harrah's Las Vegas	331	0.51	19	4.03	29	27	63
The Mirage	1418	0.51	20	4.90	25	6	84
The Cosmopolitan	1350	0.48	21	5.20	23	10	77
Rio Las Vegas	467	0.46	22	4.61	27	28	59
Mandalay Bay Resort	1682	0.45	23	5.78	17	11	74
NYNY Vegas	994	0.44	24	5.08	24	9	79
Palms Casino Resort	2059	0.42	25	5.24	21	16	66
Golden Nugget	428	0.42	26	3.13	31	14	79
Riviera Las Vegas	606	0.41	27	5.23	22	30	48
The Quad Las Vegas	196	0.41	28	2.57	34	33	43
Excalibur Las Vegas	633	0.41	29	5.53	19	23	60
Hooters Casino Hotel	187	0.40	30	4.64	26	31	55
Planet Hollywood	730	0.36	31	3.39	30	24	68
MGM Grand Hotel	2832	0.35	32	4.41	28	8	81
Hard Rock Hotel LV	1359	0.34	33	2.74	32	22	64
Bally's Las Vegas	343	0.30	34	2.62	33	25	68

* The Bellagio changed their corporate Twitter account on October 8, from @bellagiolv to @bellagio, which yielded unreliable data for period 2.
** M Resort was not ranked relative to the other listed resorts by TripAdvisor.

A pilot test of the data extraction and cleaning procedure was conducted on August 16, 2013, which led to some analysis modifications. Based on findings from the pilot test, tweets that were sent from certain source platforms were pre-screened, since there was a high percentage of fake accounts that could be identified on certain outgoing platforms (e.g. constant account activity over several 24 h periods). Twitter is similar to other online sites in terms of false reviews, such as Yelp or TripAdvisor, but the richness in data may yield a more effective means to screen some of these false accounts. If some valid accounts were eliminated as a result of this pre-screening, we expect this bias to be consistent across all properties. We expect this outcome to be much less biased than the alternative, as the overly positive "shill" accounts all favored properties belonging to a single corporate entity. We also chose to remove retweeted messages—a single message that was retweeted from many accounts could bias the results toward a single user's opinion. In particular, this bias would lean toward celebrities and corporate accounts with large numbers of followers that can retweet messages. After examining the tweets' scores in the pilot data, several words were then added to the lexicon to reflect commonly used positive/negative words that were absent from the lexicon. These were primarily short-hand abbreviations of full words (e.g. "thx") or emoticons (e.g. ":-)"), which may be a reflection of the 140 character limit of Twitter comments.

Collection dates then continued on August 23, August 30, September 6, and September 13 in 2013 (referred to as period one) as the Twitter Search API allows for roughly the past week of data to be extracted at a given time. After analyzing the period one data, we chose to add an additional set of data collection points on October 25, November 1, November 8, and November 15 in 2013 (period two) for further tests of consistency. In total, 34,315 tweets were mined during this procedure. There were 1083 tweets with errors in decoding that were removed–this was primarily a result of non-standard character use, such as accented letters in foreign languages. A further 2765 duplicate tweets were removed, as the multiple mining points of the Twitter API created the possibility of redundancies. In total, 31,550 tweets were used in the ranking analysis.

4. Results

The property sentiment scores over the eight weeks of data collection are provided in Table 1. In general, positive sentiment outweighs negative sentiment, as all firm had positive average sentiment scores, and positive ratio scores. On average, there were 928 tweets per account, with the MGM Grand account having the most tweets (2832), and Hooter's casino the fewest (1 8 7). While many commonly known luxury brands such as Aria, Bellagio, and Venetian appear in the top five average sentiment score rank, less well known properties, such as Tropicana or the local market serving Red Rock and M Resort, also appear within the top ten. The appearance of these firms suggests that these metrics are capturing more information than simple views on hotel rankings. Indeed, Liu (2012) describes sentiment as opinions, sentiments, evaluations, appraisals, attitudes, and emotions, which is consistent with our observations in analyzing resort tweets.

While the average score and the ratio score ranks had a strong positive correlation with one another, Spearman's $\rho = 0.893$, $n = 34$, $p < 0.001$, a close examination of the scores reveals that several firms had meaningfully different ranks, depending on the metric.

Given that average scores are positive, a higher number of tweets that are neutral (sentiment score of zero) will reduce the average, but have no effect on the ratio score rank. This may explain part of the difference in relative ranking. It is unclear if tweets that score exceptionally high or low (i.e. multiple lexicon words) reflect a stronger emotion, but this certainly could also be an explanation of the difference in rankings. For example, properties that scored relatively higher on the average score than the ratio score may have a subset of customers with disproportionally strong feelings of goodwill.

To examine the effectiveness of the sentiment scoring lexicon in capturing week to week changes, we individually examined tweets that appeared for properties with the highest and lowest average scoring weeks of the study. The tweets in those given weeks were reviewed by the authors for any consistent content themes that may have face validity. To provide better context for the reader, we describe the top three and bottom three average scoring weeks that appear in the study for illustrative purposes–the date provided reflects the data extraction date, and thus includes tweets up to (roughly) seven days prior.

4.1. Positive events

The highest scoring extraction occurred on August 30 for Caesar's Palace (192 tweets). This week captured a promotion by Matt Goss (a Caesar's Palace performer), that allowed the winner to attend his performance and after-party on the property, which produced many positive responses from accounts where the user did not even appear to be in Las Vegas. For example, "@mattgoss @CaesarsPalace Wow what an amazing prize.whoever wins will be so lucky.goodluck everyone!"

The second highest scoring extraction occurred on November 8 for Red Rock Las Vegas (50 tweets). This casino, which generally serves local patrons rather than tourists, consistently scored high during the study, but during this week it had many positive tweets associated with promotion of an onsite event commemorating the 20th anniversary of the Ultimate Fighting Championship, for example, "Celebrate 20 yrs of @UFC at @UltimatePoker Hall of Fame Game w @ChuckLiddell @realroyce & more @redrockcasino 11/1″.

The third highest scoring extraction occurred on September 6 for Harrah's Las Vegas (43 tweets). The scoring this week was buoyed by a series of high scoring tweets regarding, "Toby Keith's I Love this Bar and Grill", which contained positive lexicon words, but may not have been directly used to describe sentiment and instead are simply a reflection of the restaurant name. This is an important example of limitations associated with sentiment analysis, and the need to continue to refine algorithmic approaches in applied settings.

4.2. Negative events

The most negative extraction occurred on November 8 for Planet Hollywood (36 tweets). During this period, there were several tweets remarking a looming feud between Britney Spears (the recently announced Planet Hollywood headline act) and Paris Hilton (whose has the same initials as Planet Hollywood). Unsurprisingly, given the subjects, the substantive meaning behind these passionate observers remains somewhat undecipherable. For example, one tweet said, "so will @ParisHilton so her support for @britneyspears by doing a video or will the be jealousy over @phvegas by stealing intials at PH".

The second most negative extraction occurred on October 25 for Bally's Las Vegas (95 tweets), and provides good evidence of Twitter data analysis' ability to capture real time sentiment, with an opportunity to qualitatively read tweets for more context. The low score this week reflects many negative tweets associated with a shooting that occurred at a nightclub located inside of Bally's, for example, "BREAKING: @LVMPD confirms three men shot inside @BallysVegas. One suspect in custody."

Lastly, the third most negative extraction occurred on November 1 for The Quad (23 tweets). The Quad was among the lowest rated properties overall, with the second lowest tweet volume. As such, a moderate increase in negative reviews associated with poor service and on-site construction appeared to have affected reviews this week. This example demonstrates the importance of having a reasonably sized sample to gauge extreme events.

4.3. Validity tests

To assist in evaluating the content validity of the sentiment metrics, we computed Spearman correlations of the properties' rankings based on the average sentiment rank and the ratio sentiment rank against TripAdvisor's relative rankings of the hotel-casinos and TripAdvisor's average scores of the hotel-casinos from May 2013. The average sentiment ranks have a moderate positive correlation with TripAdvisor ranks, Spearman's $\rho = 0.447$, $n = 33$, $p = 0.009$, and TripAdvisor scores, Spearman's $\rho = 0.432$, $n = 33$, $p = 0.012$. Similarly, the ratio sentiment ranks have moderate positive correlations with TripAdvisor ranks, Spearman's $\rho = 0.513$, $n = 33$, $p = 0.002$, and TripAdvisor scores, Spearman's $\rho = 0.501$, $n = 33$, $p = 0.003$.

The moderate, but statistically significant levels of correlation between the data mined sentiment metrics and the TripAdvisor metrics appears to be in line with the expectation of a reasonable amount of both convergent validity and discriminant validity. The expected convergent and discriminant validities between the Twitter and Tripadvisor metrics lie in the overlapping targeted users and generated content across these two sites. The content on TripAdvisor is primarily hotel experience reviews written by existing customers. However, Twitter circulates not only hotel reviews but also any news regarding current events and marketing campaign. The targeted users are any consumers, existing and prospective, who have an interest in the hotel brands. Therefore, we expect some degree of correlation between these new metrics and online rankings, since both are indicative of the customer experience and sentiment (convergent validity). However, aspects that are unique to the Twitter medium (e.g. the public's reactions to current events and marketing campaigns) lead to the expectation that there will still be some differences in relative rankings (discriminant validity).

Indeed, customer sentiment is a related but somewhat different measure than customer satisfaction. Some tweets may reflect a level of satisfaction from a consumed service, such as this positively scored tweet to Aria's Twitter account, "@AriaLV Absolutely loved the rooms! Luxury!" While other tweets may be more accurately described as an aspirational or perceptional comment, such as the following positively scored tweet by a user seemingly outside of Las Vegas, "@MGMGrand It is 40 degrees in the lovely city of Pittsburgh, so YES, I wish I was in Vegas right about now!! #ilovemgmgrand". As compared to TripAdvisor scores, which appear to mainly reflect customers' hotel experience perceptions and amenities, Twitter tends to capture a broader public opinion (including individuals that have not visited the property) and their attitudes related to wider topics (e.g. events or marketing campaigns indirectly related to the properties).

4.4. Reliability tests

To better understand the reliability of the sentiment metrics, Cronbach's alpha was computed from week to week, and from period one to period two (which each aggregated four weeks of data). Cronbach's alpha for the eight-item weekly average sentiment score ($\alpha = 0.68$), the two-item period sentiment score

(α = 0.59), the eight-item weekly ratio sentiment score (α = 0.53), and the two-item ratio sentiment score (α = 0.50), were all below the commonly used threshold of .70 (Hair et al., 1998)—though the first measure was above the .60 threshold used for exploratory research (Robinson et al., 1991). However, unlike typical scales, we are attempting to develop a metric that can capture dynamic movement of consumer sentiment in real time. We therefore find the Cronbach alpha scores encouraging, in that there is a certain degree of consistency (we are not observing purely statistical noise), but there remains a reasonable amount of periodical inflection. We also note that Cronbach's alpha scores can be inflated by including a higher number of items for comparison, which we have attempted to limit in these analyses.

More stable measures of sentiment (not shown in Table 1) could include metrics such as the net sum of sentiment scores $\left(\sum_{i=1}^{n} \mathrm{Pos}_{\mathrm{Rit}} - \sum_{i=1}^{n} \mathrm{Neg}_{\mathrm{Rit}}\right)$, where α = 0.88 over eight weekly scores and α = 0.79 from period one to period two; the sum of positively scored tweets ($\sum_{i=1}^{n} \mathrm{PT}_{\mathrm{Rit}}$), where α = 0.88 over eight weekly scores and α = 0.80 from period one to period two; and/or the sum of negatively scored tweets ($\sum_{i=1}^{n} \mathrm{NT}_{\mathrm{Rit}}$), where α = 0.82 over eight weekly scores and α = 0.77 from period one to period two. These more stable metrics may be useful by management in evaluating performance longitudinally (e.g. long-term performance of a social media staff or marketing campaigns), rather than as a comparative measure relative to competitors, since the total volume of tweets is an important component of their resulting impact.

5. Discussion

This study described the usefulness of sentiment analysis of social media sites to hospitality operators. In particular, this study demonstrated how Twitter data could be analyzed through a cost-effective application to Las Vegas integrated-resort casino accounts. The study showed that using a relatively straightforward scoring algorithm with a publically available lexicon will yield a set of sentiment metrics that have reasonable reliability and convergent validity with external hotel rankings. As compared to Tripadvisor's ranking, the sentiment scores appear to be reasonably effective at capturing real-time information: not only existing customers' review on their hotel stay, but also the public's buzz toward important short-term events and marketing campaigns.

Our study contributes to hospitality and tourism literature in several ways. First, this study makes a methodological contribution by demonstrating a new and innovative analytical method in understanding the public's sentiments toward hospitality firms, using a large volume of Twitter text data. As social media empowers consumers with more self-expression opportunities, hospitality scholars have shown tremendous interests in consumers' EWOM behaviors. The existing literature mainly focuses on the motivations for and the effects of EWOM. Sparse is the work that attempts to tackle the challenges of distilling business intelligence with large quantities of online data.

While sentiment analysis has been touted as a promising tool to uncover new knowledge from the massive volume of unstructured data in many disciplines, it is rarely explored in the hospitality and tourism context. Most recently, Xiang et al. (2015) applied text mining to analyze a large volume of consumer reviews and identified the key dimensions of hotel guest experience that are associated with their satisfaction ratings. While their study summarized what hotel guests talk about their experience (e.g. room, staff, breakfast etc.), it did not capture how the guests feel about their stay in their reviews. Sentiment analysis goes beyond simply extracting useful information from text to extracting opinions and feelings embedded in the text (Anbananthen and Elyasir, 2013). Using Las Vegas integrated resorts as a case study, our study demonstrated that resort operators can use Twitter sentiment analysis to construct a set of valid and reasonably reliable metrics to capture the public's opinions. Sentiment analysis is not only restricted to Twitter data in the resort context. It is also readily applicable to analyzing text data from other online sources in other industries, such as restaurants or tourism destinations.

From the hospitality operators' perspective, sentiment analysis enables hospitality companies to track public viewpoints on a large scale by generating graphical summarizations of opinions. It is a tool that can be used to visualize a trajectory of the public buzz around a company by comparing changes in scores over time and against other firms. By leveraging the temporal information associated with Twitter data, hospitality marketing managers can effectively identify potential trends and publicity crises and thus take timely actions. Sentiment analysis also helps marketing managers to detect and exclude "flames" (overly heated or antagonistic languages) in social communication, and provides support to enhance antispam systems (Cambria et al., 2013). In this study, the sentiment analysis revealed that some Las Vegas resorts are clearly outperforming others in terms of engagement on social media. Bally's and Tropicana are both Las Vegas Strip properties with similar amenities and TripAdvisor ratings, but Tropicana has clearly outperformed Bally's in terms of its presence on Twitter and users' willingness to engage with the brand through social media.

Sentiment analysis can provide a quantifiable measure of real-time marketing campaign impacts complementary to other traditional research methods. Traditionally, the effectiveness of a marketing campaign is evaluated through direct rating and brand awareness tests using survey method or focus groups. While these methods generate valid results, they can be costly, time consuming, and require customers to rate categories that are pre-determined by researchers or managers. Such a rating system might not necessarily capture the responses that are personally meaningful to customers, and almost certainly will suffer from self-selection and recall bias. While the effectiveness of marketing campaign can be also gauged by sales effect (Kotler et al., 2012), this method is often used in the post-campaign period, and has a time lag between the initial campaign and resultant sales. Unlike these methods, sentiment analysis is able to capture the real-time and naturalistic publicity impacts of current campaigns from a vast consumer base. For example, Caesars' Palace saw a surge of positive sentiment in August 2013 right after a promotion to attend a performers and after-party. Sentiment analysis has the advantage of immediacy and cost efficiency over the traditional methods, and it is a complementary tool to justify the return on investment of marketing campaigns.

There are many limitations to the use of sentiment analysis. We note that the Twitter API provides a limited search capacity regarding returned data volume and time frame. The tweets which are in private setting are not retrievable through the API. Analysis is limited to the extent that social media platform owners are willing to share their data, or that data can be legally mined. Our study focused on a specific category of hospitality firms (integrated casino resorts), within a single market (Las Vegas, NV). Application of this method in other industries or locations may lead to different results, necessitating methodological modifications. For example, small hotel properties may not have sufficient Twitter volume to conduct such analyses, and sentiment analysis may therefore be limited to large corporate accounts with well established brands.

While the lexicon used in this study may have measurement validity across a large sample of tweets, smaller samples or individual comments can be easily miscategorized. Future studies should focus on improving the content validity of these text mining and classification algorithms, and establishing rules of thumb to guide managers toward meaningful sample sizes. Where account mentions are infrequent, a more qualitative/ethnographic approach to examining tweets may yield a richer understanding of consumer attitudes and emotions. We caution that social media, and Twitter in particular, may not be representative of a firm's overall target market. Twitter users skew toward younger and urban groups (Vaynerchuk, 2013), therefore managers should compare demographics captured in Twitter mining to those that are representative of their typical or target consumer. Indeed, Twitter sentiment analysis certainly does not replace traditional market research tools like guest satisfaction surveys, but it can be a useful complementary tool.

Appendix A. Twitter account information

Twitter username	Twitter name	Followers	Following	Tweets	Verified account
@AriaLV	Aria Las Vegas	96,519	2871	5747	Yes
@BallysVegas	Bally's Las Vegas	12,478	3844	2150	No
@BellagioLV*	Bellagio Las Vegas	58,884	126	7071	Yes
@CaesarsPalace	Caesars Palace	82,496	46,389	11,304	Yes
@CircusVegas	Circus Circus	17,151	1559	6505	Yes
@Cosmopolitan_LV	The Cosmopolitan	137,164	74,429	34,039	Yes
@ExcaliburVegas	Excalibur Las Vegas	54,345	14,924	7098	Yes
@FlamingoVegas	Flamingo Las Vegas	26,456	13,921	4691	No
@GoldenNuggetLV	Golden Nugget	11,366	389	1139	No
@hardrockhotellv	Hard Rock Hotel LV	63,108	3142	8342	No
@HarrahsVegas	Harrah's Las Vegas	18,234	6712	2380	No
@HootersCasinoLV	Hooters Casino Hotel	15,079	1606	7860	No
@LuxorLV	Luxor Hotel & Casino	91,911	9910	15,475	Yes
@LVHHotelCasino	LVH Hotel & Casino	29,741	22,503	6970	No
@LVStratosphere	Stratosphere Hotel	9,300	399	6748	No
@MandalayBay	Mandalay Bay Resort	78,788	8690	3788	Yes
@MGMGrand	MGM Grand Hotel	101,395	6765	11,045	Yes
@MonteCarloVegas	Monte Carlo Resort	64,229	143	12,059	Yes
@MResort	M Resort Spa Casino	18,535	3857	12,301	No
@NYNYVegas	NYNY Vegas	43,234	931	7661	Yes
@PalazzoVegas	Palazzo Las Vegas	32,107	1803	4639	No
@Palms	Palms Casino Resort	76,362	6958	9052	Yes
@ParisVegas	Paris Las Vegas	23,944	9671	3860	No
@phvegas	Planet Hollywood	54,983	28,517	8878	Yes
@quadvegas	The Quad Las Vegas	10,064	5005	2393	No
@redrockcasino	Red Rock Las Vegas	17,088	2065	3389	Yes
@RioVegas	Rio Las Vegas	33,388	25,761	4544	No
@RivieraLasVegas	Riviera Las Vegas	12,221	4636	24,229	Yes
@southpointlv	South Point Hotel	15,248	12,674	3563	No
@TheMirageLV	The Mirage	43,226	10,196	10,900	Yes
@TIvegas	Treasure Island	26,682	25,147	3024	Yes
@TropLV	Tropicana Las Vegas	21,907	2612	15,058	No
@VenetianVegas	Venetian Las Vegas	42,111	2295	4818	No
@WynnLasVegas	Wynn Las Vegas	488,061	8583	9100	Yes

Note: Figures were recorded at the beginning of the study.

* The Bellagio changed their corporate Twitter account on October 8, from @bellagiolv to @bellagio, which yielded limited data for period 2.

References

Akehurst, G., 2009. User generated content: the use of blogs for tourism organizations and tourism consumers. Serv. Bus. 3 (1), 51–61.
Anbananthen, K.S.M., Elyasir, A.M.H., 2013. Evolution of opinion mining. Aust. J. Basic Appl. Sci. 7 (6), 359–370.
Asur, S., Huberman, B.A., 2010. Predicting the future with social media. In: Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on, August 2010 vol. 1, pp. 492–499, http://dx.doi.org/10.1109/WI-IAT.2010.63, Retrieved from IEEE Xplore.
Boyd, D., Crawford, K., 2012. Critical questions for big data: provocations for a cultural, technological, and scholarly phenomenon. Inf. Commun. Soc. 15 (5), 662–679.
Browning, V., So, K.K.F., Sparks, B., 2013. The influence of online reviews on consumers' attributions of service quality and control for service standards in hotels. J. Travel Tourism Mark. 30 (1–2), 23–40.
Cambria, E., Schuller, B., Xia, Y., Havasi, C., 2013. New avenues in opinion mining and sentiment analysis. IEEE Intell. Syst. 28 (2), 15–21.
Capriello, A., Mason, P.R., Davis, B., Crotts, J.C., 2011. Farm tourism experiences in travel reviews: a cross-comparison of three alternative methods for data analysis. J. Bus. Res. 66, 778–785.
Cha, M., Haddadi, H., Benevenuto, F., Gummadi, P.K., 2010. Measuring user influence in twitter: the million follower fallacy. In: The Conference Proceeding of the Fourth International AAAI Conference on Weblogs and Social Media, Washington, DC, pp. 10–17.
Chan, N.L., Guillet, B.D., 2011. Investigation of social media marketing: how does the hotel industry in Hong Kong perform in marketing on social media websites? J. Travel Tourism Mark. 28 (4), 345–368.
Chen, H., Chiang, R.H., Storey, V.C., 2012. Business intelligence and analytics: from big data to big impact. MIS Q. 36 (4), 1165–1188.
Chiu, C., Chiu, N.H., Sung, R.J., Hsieh, P.Y., 2015. Opinion mining of hotel customer-generated contents in Chinese weblogs. Curr. Issues Tourism 18 (5), 477–495.
Claster, W., Pardo, P., Cooper, M., Tajeddini, K., 2013. Tourism, travel and tweets: algorithmic text analysis methodologies in tourism. Middle East J. Manage. 1 (1), 81–99.
Culotta, A., 2010. Detecting Influenza Outbreaks by Analyzing Twitter Messages, Retrieved from Cornell University Library ⟨http://arxiv.org/abs/1007.4748⟩.
Ding, X., Liu, B., Yu, P.S., 2008. A holistic lexicon-based approach to opinion mining. In: Proceedings of the 2008 International Conference on Web Search and Data Mining, ACM, pp. 231–240.
Dodds, P.S., Harris, K.D., Kloumann, I.M., Bliss, C.A., Danforth, C.M., 2011. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. PLoS ONE 6 (12), 1–26.
Duan, W., Cao, Q., Yu, Y., Levy, S., 2013. Mining online user-generated content: using sentiment analysis technique to study hotel service quality. In: System Sciences (HICSS), 2013 46th Hawaii International Conference on, IEEE, January 2013, pp. 3119–3128.
Dugann, M., Smith, A., December 2013. Social Media Update 2013. Pew Research Center, Retrieved from ⟨http://pewinternet.org/~/media//Files/Reports/2013/Social%20Networking%202013_PDF.pdf⟩.
Elder, J., Hill, T., Delen, D., Fast, A., 2012. Practical Text Mining and Statistical Analysis for Non-structured Text Data Applications. Academic Press, Waltham, MA, pp. 133–142.
Gentry, J., 2013. Package 'twitteR'. R Core Development Team, Retrieved from ⟨http://cran.r-project.org/web/packages/twitteR/twitteR.pdf.⟩.
Ghose, A., Ipeirotis, P., 2009. The EconoMining project at NYU: studying the economic value of user-generated content on the internet. J. Revenue Pricing Manage. 8 (2), 241–246.
Gibbs, G., 2007. Analyzing Qualitative Data. Sage, Thousand Oaks, London.
Grandoni, D., Dec 2013. One Statistic That Shows How Small Twitter Really is. The Huffington Post, Retrieved from ⟨http://www.huffingtonpost.com/2013/12/19/twitter-statistics_n_4469054.html⟩.
Hair, Joseph F., Anderson, R.E., Tatham, R.L., Black, W.C., 1998. Multivariate Data Analysis. Prentice Hall, Upper Saddle River, NJ, pp. 577–664.
Hays, S., Page, S.J., Buhalis, D., 2013. Social media as a destination marketing tool: its use by national tourism organisations. Curr. Issues Tourism 16 (3), 211–239.
Hennig-Thurau, T., Gwinner, K.P., Walsh, G., Gremler, D.D., 2004. Electronic word-of-mouth via consumer-opinion platforms: what motivates consumers to articulate themselves on the internet? J. Interact. Mark. 18 (1), 38–52.
Hennig-Thurau, T., Malthouse, E.C., Friege, C., Gensler, S., Lobschat, L., Rangaswamy, A., Skiera, B., 2010. The impact of new media on customer relationships. J. Serv. Res. 13 (3), 311–330.
Hennig-Thurau, T., Hofacker, C.F., Bloching, B., 2013. Marketing the pinball way: understanding how social media change the generation of value for consumers and companies. J. Interact. Mark. 27, 237–241.
Hu, M., Liu, B., 2004. Mining and summarizing customer reviews. In: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, August 2004, pp. 168–177, http://dx.doi.org/10.1145/1014052.1014073.
KDnuggets, 2012. Poll Results: Top Analytics, Data Mining, Big Data Software Used, Retrieved from KDnuggets ⟨http://www.kdnuggets.com/2012/05/top-analytics-data-mining-big-data-software.html.⟩.
Kim, J., Connolly, D., 2013. The 2013 customer engagement technology study. Hospitality Technol., 3–22, Retrieved from ⟨http://hospitalitytechnology.edgl.com/reports/143-Findings-on-How-Your-Competitors-are-Engaging-Guests–2013-Customer-Engagement-Tech-Study87653⟩.
Kirilenko, A.P., Stepchenkova, S.O., 2014. Public microblogging on climate change: one year of twitter worldwide. Global Environ. Change 26, 171–182.
Kotler, P., Bowen, J., Makens, J., 2012. Promotion products: communication and promotion policy and advertising. In: Marketing for Hospitality and Tourism, sixth ed. Prentice Hall, Upper Saddle River, NJ, pp. 361–395.
Krikorian, Raffi, 2013. New tweets per second record, and how! Twitter Eng. Blog, Retrieved from ⟨https://blog.twitter.com/2013/new-tweets-per-second-record-and-how⟩.
Kwak, H., Lee, C., Park, H., Moon, S., 2010. What is Twitter, a social network or a news media? In: Proceedings of the 19th International Conference on World Wide Web, ACM, April 2010, pp. 591–600, http://dx.doi.org/10.1145/1772690.1772751.
Ladhari, R., Michaud, M., 2015. eWOM effects on hotel booking intentions, attitudes, trust, and website perceptions. Int. J. Hospitality Manage. 46, 36–45.
Leung, D., Law, R., van Hoof, H., Buhalis, D., 2013. Social media in tourism and hospitality: a literature review. J. Travel Tourism Mark. 30 (1–2), 3–22.
Litvin, S.W., Goldsmith, R.E., Pan, B., 2008. Electronic word-of-mouth in hospitality and tourism management. Tourism Manage. 29 (3), 458–468.
Liu, B., 2012. Sentiment analysis and opinion mining. Synth. Lect. Hum. Lang. Technol. 5 (1), 1–167.
Mauri, A.G., Minazzi, R., 2013. Web reviews influence on expectations and purchasing intentions of hotel potential customers. Int. J. Hospitality Manage. 34, 99–107.
Nielsen, December 2012. State of the Media: The Social Media Report 2012, Retrieved from ⟨http://www.nielsen.com/us/en/reports/2012/state-of-the-media-the-social-media-report-2012.html⟩.
Pang, B., Lee, L., 2008. Opinion mining and sentiment analysis. Found. Trends Inf. Retrieval 2 (1–2), 1–135.
Peterson, R.A., Wilson, W.R., 1992. Measuring customer satisfaction: fact and artifact. J. Acad. Mark. Sci. 20 (1), 61–71.
Robinson, J.P., Shaver, P.R., Wrightsman, L.S., 1991. Criteria for scale selection and evaluation. Meas. Pers. Soc. Psychol. Attitudes 1, 1–16.
Rylander, R.G., Propst, D.B., McMurtry, T.R., 1995. Nonresponse and recall biases in a survey of traveler spending. J. Travel Res. 33 (4), 39–45.
Schuckert, M., Liu, X., Law, R., 2015. Hospitality and tourism online reviews: recent trends and future directions. J. Travel Tourism Mark. 32 (5), 608–621.
Sparks, B.A., Browning, V., 2011. The impact of online reviews on hotel booking intentions and perception of trust. Tourism Manage. 32 (6), 1310–1323.
Sparks, B.A., Perkins, H.E., Buckley, R., 2013. Online travel reviews as persuasive communication: the effects of content type, source, and certification logos on consumer behavior. Tourism Manage. 39, 1–9.
Thelwall, M., Buckley, K., Paltoglou, G., 2011. Sentiment in Twitter events. J. Am. Soc. Inf. Sci. Technol. 62 (2), 406–418.
Turney, P.D., Littman, M.L., 2003. Measuring praise and criticism: inference of semantic orientation from association. ACM Trans. Inf. Syst. (TOIS) 21 (4), 315–346.
Twitter, 2015. FAQ: Twitter developers, Available at: https://dev.twitter.com/faq.
Vaynerchuk, G., 2013. Jab, Jab, Jab, Right Hook: How to Tell Your Story in a Noisy Social World. Harper Collins, New York, NY.
Vermeulen, I.E., Seegers, D., 2009. Tried and tested: the impact of online hotel reviews on consumer consideration. Tourism Manage. 30 (1), 123–127.
Wang, J., Gu, Q., Wang, G., 2013. Potential power and problems in sentiment mining of social media. Int. J. Strateg. Decis. Sci. (IJSDS) 4 (2), 16–26.
Weil, K., 2010. Measuring tweets. In: Twitter Blog, Retrieved from ⟨https://blog.twitter.com/2010/measuring-tweets⟩.
Weng, J., Lim, E.P., Jiang, J., He, Q., 2010. Twitterrank: finding topic-sensitive influential twitterers. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining, ACM, February 2010, pp. 261–270.
Xiang, Z., Schwartz, Z., Gerdes, J.H., Uysal, M., 2015. What can big data and text analytics tell us about hotel guest experience and satisfaction? Int. J. Hospitality Manage. 44, 120–130.
Xie, H., Miao, L., Kuo, P.J., Lee, B.Y., 2011. Consumers' responses to ambivalent online hotel reviews: the role of perceived source credibility and pre-decisional disposition. Int. J. Hospitality Manage. 30 (1), 178–183.
Ye, Q., Law, R., Gu, B., 2009. The impact of online user reviews on hotel room sales. Int. J. Hospitality Manage. 28 (1), 180–182.
Ye, Q., Law, R., Gu, B., Chen, W., 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. Comput. Hum. Behav. 27 (2), 634–639.
Zhang, Z., Ye, Q., Law, R., Li, Y., 2010. The impact of e-word-of-mouth on the online popularity of restaurants: a comparison of consumer reviews and editor reviews. Int. J. Hospitality Manage. 29 (4), 694–700.