

This chapter covers the first part of our business intelligence discussion and gives the reader insight into opinion mining and sentiment analysis. Social media are seen as big data in the sense that they provide a massive amount of online reviews and ratings that can be collected and analyzed in order to consider the impact these data may have on organizations. Particularly, several studies have shown that more positive reviews and higher rates for an organization (and its products or services) may lead to a significantly higher number of desired business actions (e.g., higher sales or more subscriptions to an online newsletter). This chapter explains characteristics such as subjectivity and tone in opinions and shows how a sentiment model can be built. The chapter concludes with challenges faced by this research field today.

This chapter is primarily situated in the IT department of an organization. This means that the technical execution or implementation of business intelligence techniques will be conducted by IT people or engineers rather than business people. Nonetheless input or involvement of business users is still relevant for successful business intelligence applications because fully automated analyses may lead to inappropriate conclusions or business decisions (Fig. 7.1).

7.1 Introduction to Opinion Mining and Sentiment Analysis

One way to illustrate the impact of social media on our daily lives is by looking at online reviews and ratings. For instance, think about the different ways people can look for a restaurant (Minnaert 2012). In the 1990s, i.e., when social media were not omnipresent, people could choose a restaurant after asking advice from other people they know (e.g., relatives, friends, or neighbors). Or for more sophisticated advice, official review institutions published books in which restaurants were scored (e.g., the Michelin Guides: <http://www.viamichelin.com/web/Restaurants>). The first option could result in less experienced advice than the second option, but it was easier and also considered to be trusted. Although both options are still valid

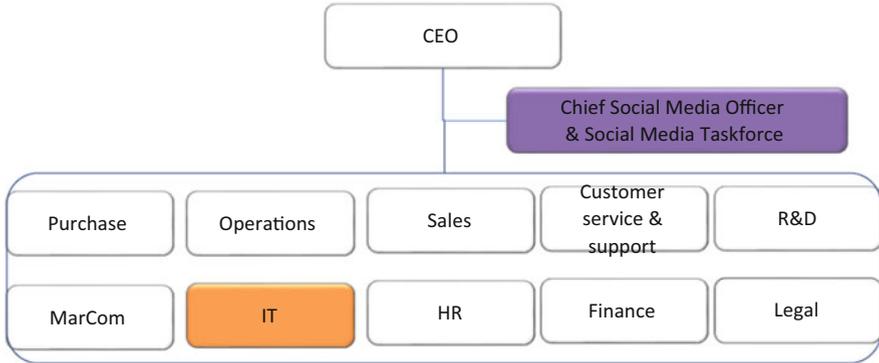


Fig. 7.1 The multidisciplinary approach of opinion mining and sentiment analysis

nowadays, an increasing number of people will rather look for information on the Internet. Additionally, different social media tools now exist in which (unknown) people can share their experience. For instance, regarding restaurants and holiday-related topics, a social media tool such as TripAdvisor™ (<http://www.tripadvisor.com/>) can be consulted to find and/or post reviews and ratings (e.g., as a score out of five on a Likert scale). Similar to the principle of SEO (see Chap. 6), restaurants will then be ranked based on the reviews and ratings posted by Internet users. Even though the users are strangers with subjective opinions, they can influence others, because people are more likely to visit a restaurant with many positive comments and which is ranked in the top 5 or top 10 of restaurants at a certain location.

Take a few moments to reflect on the degree to which you rely on online reviews and/or ratings when buying a certain product or service.

- Would you rely on the reviews of others? If so, rather on the negative reviews or on the positive ones? Would you also take into account the number of reviews? For example, would you buy a product or service with only a few negative reviews (let's say five to ten negative reviews)?
- Would you rely on visual ratings and rankings of others? If so, which ratings do you prefer? For example, regarding a score out of five, would you only buy a product or service if it has 5-star reviews, or do you also allow 4-star or 3-star reviews? Would you also take into account the variation among ratings, or do you only trust products and services with high ratings and without any negative comments?
- Would you rely on other review-related aspects too, e.g., the identity, personal information, or popularity of the reviewers?
- And what about organization-related features, such as the brand name of the product or service?

In general, many studies have shown that organizations with better reviews and ratings have more desired business actions (e.g., higher purchase likelihood, sales,

stock performance, etc.) (Chen et al. 2012; Chevalier and Myzlin 2006; Dellarocas et al. 2007; Duan et al. 2008; Forman et al. 2008; Moe and Trusov 2011; Senecal and Nantel 2004; Tirunillai and Tellis 2012; Zhu and Shang 2010).

Although the studies above agree on the existence of a correlation or relationship between reviews, ratings, and the desired business actions, they still disagree on the direction of this correlation. Particularly, they disagree on whether this relationship is direct or whether some intermediate factors do count (e.g., brand popularity, the number of positive and negative reviews, etc.).

The Edelman trust barometer (2014) shows that people tend to perceive experts and peers as highly trusted or credible sources (rather than a CEO or a government official) when forming an opinion about an organization with regard to employee engagement. Similarly, before deciding whether to buy a certain product or service, the majority of consumers tend to go online to read user reviews and ratings (besides commercial information) and possibly choose for another purchase if negative information is found. The growing power of online reviews and ratings is confirmed by a survey of Cone Communications (2011), which even refers to four out of five US consumers who go online to verify recommendations before buying products or services. The study of Pang and Lee (2008) also illustrates the strong influence of social media by showing that a large majority of Internet users make use of online reviews and ratings, and one-third of them have already written a review or assigned a rating themselves. Moreover, the authors state that Internet users who read and write recommendations are even willing to pay significantly more for a 5-star product or service than for a 4-star (i.e., even up to 99 % more, self-reported). This means that online recommendations count and can possibly result in a higher value per purchase for an organization. On the other hand, negative recommendations can have negative consequences for an organization.

Consequently, for organizations, it is worth monitoring social media tools that collect user reviews and ratings, e.g., TripAdvisor™ (<http://www.tripadvisor.com/>), Yelp™ (<http://www.yelp.com/>), Foursquare™ (<https://foursquare.com/>), or the Internet Movie Database™ (IMDb, <http://www.imdb.com/>). This topic of opinion mining and sentiment analysis is closely linked to social CRM (see Chap. 5), which emphasizes that organizations should constantly keep track of social media posts that deal with keywords related to their brand and products or services. In case of negative reviews or ratings, an organization should react properly and take corrective actions to avoid any reputational damage. While it may take advantage of positive reviews and ratings, an organization can also learn from negative comments to improve its products and services accordingly in order to avoid more negative reviews and to turn them into positive future customer experiences.

7.2 Defining Opinion Mining and Sentiment Analysis

7.2.1 Definitions

Opinion mining and sentiment analysis can be defined as processing “a set of search results for a given item, generating a list of product attributes (quality, features, etc.) and aggregating opinions about each of them (poor, mixed, good)” (Dave et al. 2003: p. 519). Mining is jargon for analyzing and literally refers to digging for information (or distilling knowledge) that is relevant for a business (i.e., by using business intelligence techniques).

Although opinion mining and sentiment analysis are frequently used as synonyms, opinion mining strictly refers to personal opinions (e.g., “I think”), while sentiment analysis rather focuses on feelings or sentiments (e.g., “I like” or “I dislike”). Opinions and sentiments are usually expressed by means of adjectives (e.g., “big” or “boring”). Such opinions or sentiments can be detected in a content analysis of texts, e.g., posts on social media (e.g., Facebook™, Twitter™, blogs, fora) or online newspapers.

The techniques for opinion mining or sentiment analysis can be used for many applications in different domains. One example concerns market share predictions, as an increase in negative buzz about a specific product or service may have a negative impact on the organization’s market shares (see the Antennagate crisis of Apple Inc™, Chap. 5). An example in healthcare is the detection of depressive children based on the number of positive and negative messages on social media tools (e.g., Facebook™). Further on, a study showed how the results of political elections could be predicted by the degree to which political parties and individual politicians are covered in newspapers (Junqué de Fortuny et al. 2012). These authors now investigate tweets for the same purpose.

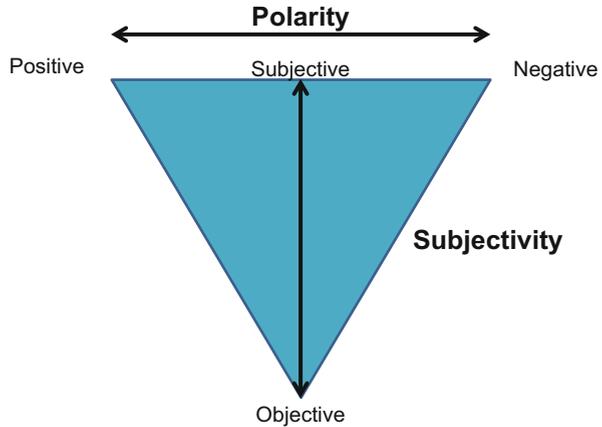
Finally, we must note that opinion mining and sentiment analysis are not limited to social media and online text but that a similar content analysis can also be conducted for offline text. Hence, the techniques discussed in this chapter are situated in the overall domain of text mining.

7.2.2 Characteristics

Based on Liu (2010) and De Smedt and Daelemans (2012), opinions and sentiments can have a different degree of (1) subjectivity and (2) polarity, as shown in Fig. 7.2.

- Subjectivity ranges from objective to subjective. An objective sentence expresses factual information about the world, while a subjective sentence expresses some personal feelings or beliefs.
- The polarity (i.e., tone or orientation) indicates whether an opinion is positive, neutral, or negative.

Fig. 7.2 Subjectivity and polarity in opinions, based on Minnaert (2012)



Subjectivity and polarity are a continuum, possibly expressed in decimals (instead of an “all-or-nothing” approach with binary values). This means that different degrees of subjectivity exist between factual information and personal feelings or beliefs. Similarly, different degrees of polarity exist between positive, neutral, and negative comments. Examples of (subjective) feelings are “I like the color of the device” (positive) or “I really hate the service given by that organization” (negative). On the other hand, examples of (objective) facts are “the newspaper reported on a tremendous increase in sales for the device” (positive), “I am looking for a good device” (neutral), or “the test has shown some technological problems with the device” (negative). Assessing the positive or negative meaning of facts is, however, more difficult than for opinions or feelings of a certain individual, since it depends on the interpretation of the reader. For instance, an increase in sales of a certain product can be positive for an organization, but negative for the competitors.

7.3 Building a Sentiment Model

Subjectivity and polarity depend on the sentiment words or opinion words used within a review. It concerns adjectives (e.g., “good,” “nice,” “fantastic” but also “bad,” “poor,” “awful”) or expressions (e.g., “the more the merrier,” “blood, sweat, and tears”). In general, three approaches can be applied to identify sentiment words (Liu 2012): (1) a manual approach, (2) a dictionary-based approach, or (3) a corpus-based approach.

- **Manual.** The manual approach involves a manual screening of the collected reviews to determine their degree of subjectivity and polarity. As the manual approach requires more effort compared to the other approaches, it is frequently used in combination with the other, more automated approaches.

- **Dictionary based.** The dictionary-based approach starts with a manual screening to find some sentiment words as “seeds,” i.e., words that can be considered as positive or negative. Then, an algorithm searches through online dictionaries to find synonyms and antonyms of these seeds to iteratively complete the list of sentiment words. Afterwards, a manual check is required to clean up the list. This approach is less time consuming compared to the manual approach but still needs a manual check afterwards. It is particularly useful for general topics, rather than for specific (research) domains in which words may have different meanings depending on the context. For instance, in the sentence “the radio plays quietly,” the sentiment word “quietly” can have both a positive and negative meaning depending on where you are (e.g., in a car, at a party, etc.). Examples of general lists with sentiment words are publicly available (e.g., <http://sentiwordnet.isti.cnr.it/>; <http://www.purl.org/net/emolex>).
- **Corpus based.** The corpus-based approach is specific for a certain domain or category. The input are seeds, namely, a small set of category-dependent words (let’s say five seeds), supplemented by a list of references that link to domain-specific sources (called a representative text corpus). The output is a ranked list of words associated with that category. A manual screening is only required for the top-ranked words, because other words are less typical for the specific category under study. The aim is to compose a sentiment lexicon, i.e., the final list of synonyms and antonyms for a given category. Similar to the dictionary-based approach, different iterations are needed. However, too many iterations (let’s say more than eight) will give a higher risk for general words instead of category-dependent words.

Once the sentiment words are identified, the collected reviews can be analyzed. As reviews are subjective, it is important to study reviews of many people in order to generalize the findings. A (statistical) summary can be used to reflect the opinions about a product or service (or the features, which give richer information) by means of a percentage or a score, for example:

- “X% of the reviewers are positive/negative about a product/service or about a feature of that product/service.”
- “A product/service or a feature of that product/service was reviewed with a score of -0.43 ” (i.e., a negative decimal close to one/neutral, thus slightly negative).

7.3.1 Example of a Feature-Based Sentiment Analysis

In order to know how many people are positive or negative about which features of a product or service, each review can be analyzed (or mined) as follows (Liu 2010, 2012).

Example of a user review:

Posted by: Gregory Thurston Posted on: 2014/06/28

“Six months ago, I bought a GStar GPS, particularly the X550 model, whereas my father opted for the TT GO (type 2635). During the first week, I was very pleased with my new gadget. The quality of the screen is very nice with beautiful colors. Now that I am using it more regularly, I experienced that the battery must be recharged quite often and weighs a lot. My father is very enthusiastic about his device, and has no battery problems.”

In order to determine what the review is about, each sentence is separately analyzed. Hence, the user review can be rewritten as follows.

Posted by: Gregory Thurston Posted on: 2014/06/28

(1) “Six months ago, I bought a GStar GPS, particularly the X550 model, whereas my father opted for the TT GO (type 2635).”

(2) “During the first week, I was very pleased with my new gadget.”

(3) “The quality of the screen is very nice with beautiful colors.”

(4) “Now that I am using it more regularly, I experienced that the battery must be recharged quite often and weighs a lot.”

(5) “My father is very enthusiastic about his device, and has no battery problems.”

A feature-based sentiment model can be built by looking at five elements of an opinion: (**e, a, s, h, t**). The letters respectively refer to entity, aspect, sentiment, holder, and time, namely:

- What is the review about (= target entity, e + target aspects or attributes, a)?
- What is the sentiment (s)?
- To whom belongs the sentiment (= source or sentiment holder, h)?
- When was the sentiment formulated (= time stamp, t)?

The concrete (e, a, s, h, t) elements for the example above are as follows.

- **What is the review about (= target entity, e + target aspects or attributes, a)?**

The review deals with two types of navigation systems, namely, GStar X550 and TT GO 2635. More particularly, the text refers to different aspects or attributes of a navigation system. For instance, sentence 3 refers to the screen, whereas sentence 4 refers to the battery duration and the battery weight. Such product features can be used for aspect categorization. Therefore, a distinction can be made between explicit and implicit categorization. Explicit (or direct) categorization refers to aspects that are literally mentioned in the text, e.g., the noun “screen” in sentence 3. On the other hand, implicit (or indirect) categorization

refers to aspects that can be derived, e.g., sentence 4 refers to the battery duration by indirectly mentioning that “the battery must be recharged quite often.” Other examples of implicit categorization are words such as “expensive” (for price), “it fits my pocket” (for size), or “it can be easily installed” (for installation or ease of use). Also synonyms can be used for a navigation system, such as “GPS” (sentence 1), “gadget” (sentence 2), or “device” (sentence 5).

- **What is the sentiment (s)?**

Sentence 1 gives a neutral description of a fact, while sentence 2 expresses a sentiment in the past. The message related to the screen of the GStar is positive (sentence 3), while the message related to the duration and the weight of the battery is rather negative (sentence 4). On the other hand, the message related to the TT GO is generally positive (sentence 5), and also the battery message is positive (sentence 5). The example shows that all sentences must be linked to each other in order to interpret the review correctly. For instance, the father of Gregory Thurston uses another device than Gregory himself.

- **To whom belongs the sentiment (= source or sentiment holder, h)?**

The review under study mentions two persons, namely, Gregory Thurston and his father. To answer this question, it is also important to consider the perspective of the reader. For instance, if the source is an organization, then it rather concerns an advertisement than a user review. Furthermore, a message can be negative for the sentiment holder, but positive for the reader. For instance, a sentence such as “the price of houses is decreased, which has a negative impact on our economy” might be bad news for sellers but good news for buyers.

- **When was the sentiment formulated (= time stamp, t)?**

Finally, the date of the review should be included in the analysis. In this example, the review was written in June 2014. This information can be used, for instance, to monitor the evolution of similar products over time or to identify trends.

Table 7.1 shows how the unstructured text of a review can be summarized as structured data in a database scheme, based on (e, a, s, h, t).

By bundling all information in a table or a database structure (Table 7.1), an overview is created that can be used to verify how many positive items were involved or to detect a trend over time. For instance, if a new product has technical issues, then the number of negative reviews and ratings may significantly increase within a short amount of time. Such information may alarm the organization for a timely intervention.

Alternatively, the data in Table 7.1 can also be written in sentences as follows. Each sentence consists of the same five opinion elements: (e, a, s, h, t).

Table 7.1 An example of a feature-based sentiment model for one review with regular opinions

Target (t): entity (e)	Target (t): aspect (a)	Sentiment/polarity of target (s)	Opinion holder (h)	Date (t)
GStar X550	Screen	Positive	Gregory Thurston	2014/06/28
GStar X550	Battery—duration	Negative	Gregory Thurston	2014/06/28
GStar X550	Battery—weight	Negative	Gregory Thurston	2014/06/28
TT GO 2635	GENERAL	Positive	father of Gregory Thurston	2014/06/28
TT GO 2635	Battery	Positive	father of Gregory Thurston	2014/06/28

(GStar X550, screen, positive, Gregory Thurston, 2014/06/28)
 (GStar X550, battery duration, negative, Gregory Thurston, 2014/06/28)
 (GStar X550, battery weight, negative, Gregory Thurston, 2014/06/28)
 (TT GO 2635, GENERAL, positive, father of Gregory Thurston, 2014/06/28)
 (TT GO 2635, battery, positive, father of Gregory Thurston, 2014/06/28)

The previous example concerned a review with regular opinions. A slightly different approach is used for comparing different entities, i.e., for mining comparative (instead of regular) opinions (Liu 2010, 2012). In this case, opinions express whether one brand, product, or service is better than another, for instance:

- **Non-equal gradable comparison:** “Product A is better than product B.”
- **Equative comparison:** “Product A is about the same as product B.”
- **Superlative gradable comparison:** “Product A is the best.”
- **Non-gradable comparison:** “Product A is different from product B.”

For comparative mining, Table 7.2 shows that the database scheme has one additional column for the second entity and that the sentiment column is replaced by a column for the preferred entity (PE).

7.3.2 Example of a Sentiment Analysis with Reviews and Ratings

While in the previous example the sentences were analyzed in detail, opinions can also be analyzed based on adjectives combined with the corresponding ratings (De Smedt and Daelemans 2012). As an example, suppose that a film (let’s say “In Bruges”) was reviewed and rated by many users with a score out of five (e.g., on a Likert scale). Such individual ratings can then be used to derive a global score for the film.

Table 7.2 An example of a feature-based sentiment model for comparative opinions

Target (t): entity 1 (e1)	Target (t): entity 2 (e2)	Target (t): aspect (a)	Preferred entity (PE)	Opinion holder (h)	Date (t)
Product A	<i>Product B</i>	–	<i>Product A</i>	Gregory Thurston	2014/06/28
Product A	<i>Product B</i>	–	<i>Product A</i> <i>Product B</i>	Gregory Thurston	2014/06/28
Product A	<i>Product B</i>	–	<i>Product A</i>	Gregory Thurston	2014/06/28
Product A	<i>Product B</i>	–	–	Gregory Thurston	2014/06/28

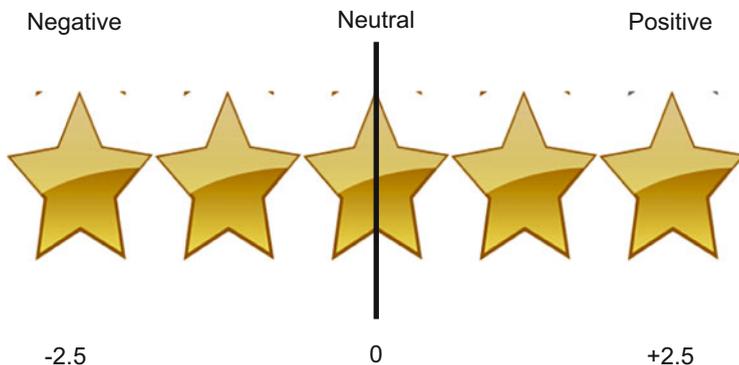


Fig. 7.3 An example of a rating system with five stars

Figure 7.3 shows that the stars of a rating system refer to the notion of polarity and represent a continuum from negative to positive values with a neutral center. For a 5-star rating system, the middle or neutral value is obtained by dividing the total number of stars (i.e., five) by two. So in this example, when a reviewer rates the film with 2.5 stars, he/she actually assigns a polarity score of zero. On the other hand, the most negative rate gets a score of -2.5 , and the most positive one gets a score of $+2.5$. The sign thus refers to a negative or positive opinion and the decimal indicates the degree of polarity.

Suppose that a review looks like this: “In Bruges was good—3/5 stars.” Three stars indicate that the review is slightly positive (Fig. 7.3). The review under study contains only one adjective, namely, “good,” which means that the assigned stars solely refer to this adjective. It can thus be derived that this adjective may express a slightly positive opinion.

The corresponding sentiment model is presented in Fig. 7.4. When translating to a score, three stars would normally refer to a score of $+0.5$ (see Fig. 7.3). However, in this example, Fig. 7.4 illustrates that the adjective “good” is associated with a score of $+1.0$. This can be explained by the fact that many other reviewers have rated films with four stars while describing their opinion with the adjective “good.”

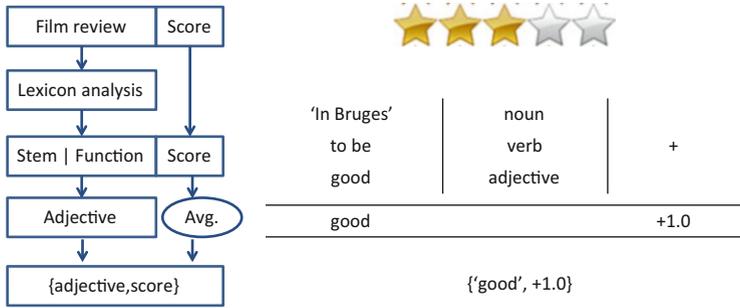


Fig. 7.4 An example of a sentiment model with ratings for a film

Hence, a sentiment model that combines reviews with ratings will first look for a list of adjectives with corresponding scores (based on averages) before assigning global scores, e.g., to particular films. Once such a list of adjectives and scores exists for a certain rating system, it can be reused to conduct a similar sentiment analysis.

One example in another context is to mine newspaper articles in order to predict positive or negative effects on market shares, the results of political elections, etc. (Provost and Fawcett 2013; Junqué de Fortuny et al. 2012). For instance, consider this extract from a (fictive) newspaper article.

“Politician John Spring accuses the newspaper XYZ of inciting hatred after they published an opinion piece on the new housing code together with a picture of a Nigerian mass grave.

The editor explains why they will continue to fight Spring’s love of segregation **hard**. **Outraged**? Yes, we are first and above all **outraged**.”

This extract contains two sentiment words, namely, “hard” and “outraged,” with the latter being more negative than the former. The corresponding sentiment model is shown in Fig. 7.5. It is built by searching for the nouns and verbs in the extract and particularly the adjectives and adverbs.

This newspaper example illustrates a sentiment model with multiple sentiment words. In this case, each sentiment word is assigned an average score (as previously defined), and the overall score is the sum of the individual scores (i.e., -0.43 in this example) (Minnaert 2012).

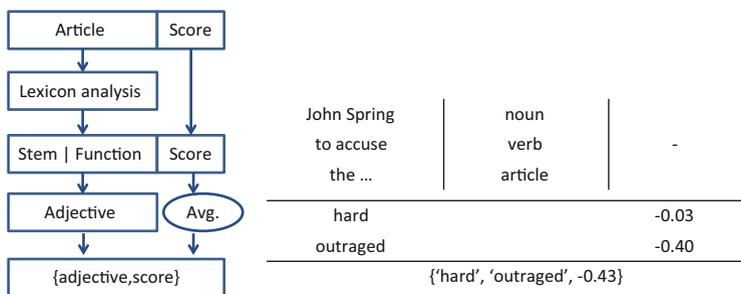


Fig. 7.5 An example of a sentiment model with ratings for a newspaper article

7.4 Big Data Challenges to Opinion Mining and Sentiment Analysis

The introduction section of this chapter already indicated the power of online reviews and ratings. For instance, what would you do as a consumer when reading a review like this?

“We dined in the restaurant and it was horrible! Wouldn’t feed it to my dog! The waiter was rude, and when we told him we were unhappy with our food, we never saw him again.”

On the other hand, the restaurant holder would certainly wish to know when such negative reviews go around (see also Chap. 5 on social CRM).

The question remains however whether all reviews are sincere, and the answer is unsure. Hence, a critical reflection on the use of text mining techniques is needed. This chapter already started by indicating the importance of human input for successful applications, rather than fully automated mining exercises. Furthermore, as writing online reviews and assigning ratings are cheap (e.g., compared to online ads; see Chap. 4) and can highly influence customers and prospects, some organizations might get tempted to conduct unethical behavior (i.e., comparable to the “black hat” techniques of SEO, Chap. 6). For instance, organizations might write reviews themselves or pay someone else to do so in order to create an impression of peer consensus (Luca and Zervas 2014). Or organizations might pay bloggers for removing a blog post containing a negative review. On the other hand, unethical behavior is not necessarily initiated by an organization. For instance, former employees may also launch extremely negative comments or false rumors about a brand after they got fired. Consequently, as a consumer, a critical attitude is required instead of believing everything that is written in user reviews. In order to evaluate a review or rating, it is important to make some of the following considerations:

- Who is the reviewer? Is it written anonymously or is his/her identity known?
- What is the purpose of the purchase? For instance, a purchase that is bad for the reviewer can be good or moderate for another user (e.g., different expectations exist for a family dinner with kids or for a romantic dinner with two).
- If the review is extremely positive, be aware that the organization could have written the review itself.
- Or vice versa, if the review is extremely negative, the review could have been written by the competitors of the organization.

Besides the challenge of unethical behavior, the techniques of opinion mining or sentiment analysis are also difficult to apply since subjective language is complex. For instance, Pang and Lee (2008) give some examples of citations that use negative words to describe something positive (or vice versa) and citations that start with positive wordings and still have a negative ending. The latter is a challenge for the broader domain of opinion mining techniques and thus not limited to a social media context.

In general, opinion mining techniques have to deal with the following challenges (Liu 2010, 2012).

- **Literal versus figurative meaning of words.** Words can have different meanings, which makes automated opinion mining more difficult. For instance, the sentence “The activity really sucks” is negative, whereas the sentence “The vacuum cleaner sucks really good” uses the same verb in a positive way.
- **Neutral meaning of words.** Opinion mining is also challenged by words that are used in a neutral meaning. For instance, in the sentence “If I find a good HD television in the shop, I will buy it,” the adjective “good” is not referring to something positive (although it usually does).
- **Sarcasm.** Opinion mining has to deal with sarcasm, i.e., if the reviewer means the opposite of what he/she is writing. Examples of reviews with sarcasm are “What a great tool! It stopped working after two days.” and “It was great! Yeah right. . . .”
- **Facts without sentiment words.** The fourth challenge indicates that some sentences may look like facts but still have a positive or negative meaning. In the example “The hot tub is heated by much electricity,” the message is negative for consumers who will likely have higher utility bills after using the hot tub. Similarly, the sentence “After a few days sleeping in my new bed, the mattress collapsed” does not contain sentiment words and still refers to a negative opinion.
- **Opinion spamming.** The final challenge covers the previous discussion about unethical behavior and emphasizes the need for reliable sources. Opinion spamming refers to the abuse or sabotage by people working at or being related to a certain organization (e.g., competitors, employees, or former employees). It usually results in extremely negative or extremely positive reviews. Three main types of opinion spamming exist: (1) fake reviews by people who have not used the product or service themselves (and which can be positive or negative); (2) brand reviews instead of reviews about a product or service

(e.g., “I will never buy any product of company XYZ anymore” after having one problem with one product), which is a generalization that can harm an organization even more; and (3) non-reviews, which are ads and other texts that express no opinion.

7.5 Takeaways

This chapter discussed the technical side of social media as big data analytics in the field of opinion mining (or sentiment analysis, which we consider to be synonyms). The related business intelligence techniques can be seen as a content analysis of a massive amount of social media data in order to verify what people are telling about an organization and its products and services. Hence, opinion mining is frequently used within the context of social CRM, namely, for monitoring purposes (see Chap. 5).

An organization can profit from monitoring the opinions of users and professionals. For instance, organizations can monitor messages to find out how their brand is perceived, how their products and services can be improved, or if their marketing strategy is working. Further on, messages can be monitored that deal with the market in which they operate in order to find trends, to monitor any shares the organization might have, or to detect what the competitors are doing (right or wrong). Opinion mining and sentiment analysis can also alert an organization when a crisis hits. Think, for instance, about a crisis of smartphones (see Chap. 5) which may give rise to newspaper articles titled like “Exploding smartphone investigated,” “Exploding smartphone leaves man in shock” (*The Local*, July 2012), or “Exploding smartphone caught on video” (knowyourmobile.com, July 2012). Such messages may frighten customers and prospects of smartphones and are extremely important to organizations that develop or sell smartphones. When analyzing online reviews and ratings, some negative consequences can be tempered by reacting timely and properly.

This chapter particularly focused on online reviews and ratings, which can highly impact on sales and the price that people are willing to spend on a specific product or service. Opinion mining thus requires an investment from an organization but has also the potential to pay off in terms of ROI. Given its relevance, more research is needed to overcome the challenges faced by text mining in general, e.g., to automatically detect the correct meaning of words and opinion spamming.

7.6 Self-Test

- What do opinion mining and sentiment analysis mean, and how do they differ?
- Can you explain how opinion mining can create business value from social media?
- What is a sentiment model and how can it be created?
- Can you explain the challenges faced by opinion mining?

Bibliography

- Chen, Y., Liu, Y., & Zhang, J. (2012). When do third-party product reviews affect firm value and what can firms do? The case of media critics and professional movie reviews. *Journal of Marketing*, 76(2), 116–134.
- Chevalier, J. A., & Myzlin, D. (2006). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*, 43(3), 345–354.
- Cone Communications. (2011). *Cone releases the 2011 online influence trend tracker*. Retrieved July 14, 2014, from: <http://www.conecomm.com/contentmgr/showdetails.php/id/4008>
- Dave, K., Lawrence, S., & Pennock, D.M. (2003). Mining the peanut gallery: Opinion extraction and semantic classification of product reviews (pp. 519–528). *Proceedings of the 12th International Conference on World Wide Web*.
- Dellarocas, C., Zhang, X. M., & Awad, N. F. (2007). Exploring the value of online product reviews in forecasting sales: The case of motion pictures. *Journal of Interactive Marketing*, 21(4), 23–45.
- De Smedt, T., & Daelemans, W. (2012). “Vreselijk mooi!” (Terribly beautiful): a subjectivity lexicon for Dutch adjectives. *Proceedings of the 8th International Conference on Language Resources and Evaluation*.
- Duan, W., Gu, B., & Whinston, A. B. (2008). The dynamics of online word-of-mouth and product sales – An empirical investigation of the movie industry. *Journal of Retailing*, 84(2), 233–242.
- Edelman. (2014). *Trust in employee engagement*. Retrieved November 26, 2014, from: <http://www.edelman.com/insights/intellectual-property/2014-edelman-trust-barometer/trust-in-business/trust-employee-engagement/index.html>
- Forman, C., Ghose, A., & Wiesenfeld, B. (2008). Examining the relationship between reviews and sales: The role of reviewer identity disclosure in electronic markets. *Information Systems Research*, 19(3), 291–313.
- Junqué de Fortuny, E., De Smedt, T., Martens, D., & Daelemans, W. (2012). Media coverage in times of political crisis: A text mining approach. *Expert Systems with Application*, 39(14), 11616–11622.
- Liu, B. (2010). Sentiment analysis and subjectivity. In N. Indurkha & F. J. Damerou (Eds.), *Handbook of natural language processing* (pp. 627–666). Boca Raton: Chapman & Hall, CRC Press.
- Liu, B. (2012). *Sentiment analysis and opinion mining*. San Rafael, CA: Morgan & Claypool Publishers.
- Luca, M. & Zervas, G. (2014). Fake it till you make it: reputation, competition, and Yelp review fraud. *Harvard Business School NOM Unit Working Paper No. 14-006* (<http://dx.doi.org/10.2139/ssrn.2293164>).
- Minnaert, B. (2012). [Guest lecture of Bart Minnaert in the course Creating Value Using Social media at Ghent University, November 2012].
- Moe, W. W., & Trusov, M. (2011). The value of social dynamics in online product ratings forums. *Journal of Marketing Research*, 48(3), 444–456.
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(2), 1–135.
- Provost, F., & Fawcett, T. (2013). *Data science for business. What you need to know about data mining and data-analytic thinking*. California: O’Reilly Media.
- Senecal, S., & Nantel, J. (2004). The influence of online product recommendations on consumers’ online choices. *Journal of Retailing*, 80(2), 159–169.
- Tirunillai, S., & Tellis, G. J. (2012). Does chatter really matter? Dynamics of user-generated content and stock performance. *Marketing Science*, 31(2), 198–215.
- Zhu, F., & Shang, X. M. (2010). Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of Marketing*, 74(2), 133–148.