

1-1-2017

Identifying Major Tasks from On-line Reviews

Feras Al-Obeidat
Zayed University, feras.al-obeidat@zu.ac.ae

Bruce Spencer
University of New Brunswick

Follow this and additional works at: <https://zuscholars.zu.ac.ae/works>



Part of the [Business Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Al-Obeidat, Feras and Spencer, Bruce, "Identifying Major Tasks from On-line Reviews" (2017). *All Works*. 1925.

<https://zuscholars.zu.ac.ae/works/1925>

This Conference Proceeding is brought to you for free and open access by ZU Scholars. It has been accepted for inclusion in All Works by an authorized administrator of ZU Scholars. For more information, please contact scholars@zu.ac.ae.

The 8th International Conference on Emerging Ubiquitous Systems and Pervasive Networks
(EUSPN 2017)

Identifying Major Tasks from On-line Reviews

Feras Al-Obeidat^{a,*}, Bruce Spencer^b

^a*Zayed University, Abu Dhabi, United Arab Emirates*

^b*University of New Brunswick, Fredericton, Canada*

Abstract

Many e-commerce websites allow customers to provide reviews that reflect their experiences and opinions about the business's products or services. Such published reviews potentially benefit the business's reputation, improve both current and future customers' trust in the business, and accordingly improve the business. Negative reviews can inform the merchant of issues that, when addressed, also improve the business. However, when reviews reflect negative experiences and the merchant fails to respond, the business faces potential loss of reputation, trust, and damage. We present the Sentiminder system that identifies reviews with negative sentiment, organizes them, and helps the merchant develop a plan with an end date by which issues will be addressed. In this paper we address the problem of quickly finding subtasks in a large set of reviews, which may help the merchant to identify, from the set of reviews, subtasks that need to be addressed. We do this by identifying nouns that frequently occur only in the reviews with negative sentiment.

© 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the Conference Program Chairs.

Keywords: Social business analysts ; Sentiment Extraction ; Clustering ; Costs Estimation

1. Introduction

On-line reviews can provide merchants valuable information from customers and clients about their products, service, staff, advertising and almost every aspect of their business. The information in on-line reviews potentially includes much marketing intelligence: timely knowledge of product failures and successes, awareness of the client's perception of quality of service, evidence of advertising effectiveness, and competitive pricing. Moreover, this abundant information is available immediately and free of charge. However, evidence suggests that businesses are not achieving the full potential benefit. Part of the reason for this may be the challenging task of reviewing and understanding this content. Reviews widely vary in quality and subject matter, contain biases and contradictions, and cover all aspects of the business – often mixing them together.

In this paper we consider the reviews collected by businesses in a variety of industries, including hotel service, electronic devices and cars, as provided by the Opinosis 1.0 dataset¹. From these publicly available reviews con-

* Corresponding author. Tel.: +971-56-8357521 ; fax: +0-000-000-0000.
E-mail address: Feras.Al-Obeidat@zu.ac.ae

tributed by on-line customers, we organize the information to make it easier for the business owner to exploit the information. First we apply a topic mapping tool to separate the comments into a variety of subject area. Second, we apply sentiment analysis to review the general satisfaction within each subject area. Finally we use information extraction to identify the major tasks that are mentioned in the reviews, so that the business owner can decide what tasks need to be done. We provide an on-line tool called the Sentiminder that automates these steps and creates an initial schedule with an achievable end date, which the business may share with the on-line community. The end goal is to satisfy the expectations that an on-line community of current and potential customers places on the business, helping it remain competitive.

2. Related Work

Related to our work, Thomas and Atish² proposed a text mining model, enhanced using the Regression ReliefF (RReliefF) feature selection method, for predicting the helpfulness of on-line reviews from Amazon.com. Based on their outcomes, they found that RReliefF significantly outperforms two popular dimension reduction methods (BOW, LSA). This study applies text regression for predicting on-line review helpfulness and to show that analysis of the keywords selected by RReliefF reveals meaningful feature groupings.

Htay and Lynn³ proposed a new opinion mining technique in classifying reviews documents by extracting features and opinion words. Their idea is to extract patterns of features and/or opinion phrases. They used linguistic rule to extracting pattern knowledge.

Somprasertsri and Lalitrojwong⁴ proposed an interesting approach on how to identify the semantic relationships between product features and opinions. The technique is used for mining product feature and opinion based on the consideration of syntactic information and semantic information. Their approach uses dependency relations and ontological knowledge with probabilistic based model.

Sánchez-Franco *et al.*⁵ propose a product feature-oriented approach to the analysis of on-line guests reviews, and analyzes the relationship between the most prominent features and guests hotel rating in the on-line travel agencies environment.

3. The Sentiminder Prototype

The *Sentiminder* is a functioning prototype, written in R/Shiny, that helps the owner of any business that needs to respond to on-line reviews.

A wide variety of issues, contributed from many customers, are presented to the Sentiminder. It uses Latent Dirichlet Allocation^{6,7}(LDA) to cluster these comments by topic.

The sentiment of each cluster is estimated using the frequency of words with positive and negative sentiment, using a method common in the literature^{8,9,10}, based on counting the number of positive and negative affect words appearing within. The degree of negative sentiment is considered to be a measure of the priority of addressing these comments. More negative sentiment is also interpreted as a larger potential benefit to the business that will arise from rectifying the issues addressed by the reviews.

These clusters with priorities are presented to the business owner, who can then consider them and interactively assign project management attributes for each task, including estimated cost, duration, earliest start date, and latest end date. The tool optimally balances cost against benefit according to the merchant's relative value of reducing cost versus increasing sentiment. The merchant also provides shared resource constraints in a straightforward way, by specifying any pair of tasks that is constrained from being performed at the same time. This would be the case, for example, if one team of people could perform two different tasks but could not do them simultaneously. The tool selects an optimal set of tasks that meets these scheduling constraints. The merchant can express an overall completion date for all tasks, and the tool then considers only schedules which can be completed within this limit. For further details please see Figure 5.

Sentiminder

Fig. 1. The Start Screen allows the user to select reviews for a business, product or service and to choose a number of topics. Thus user can also select from a precomputed selection of topics from this dataset, which is available within the Opinosis 1.0 dataset.

Topics

| Gibbs | | CTM | | | | |
|------------|-------------|-------------|---------|-----------|-----------|----------|
| Show | 5 | entries | Search: | | | |
| Topic 1 | Topic 2 | Topic 3 | Topic 4 | Topic 5 | Topic 6 | Topic 7 |
| hotel | rooms | location | free | room | service | staff |
| best | clean | wharf | wine | bathroom | parking | friendly |
| although | nice | fisherman's | coffee | tuscan | location, | helpful |
| food | small | perfect | morning | excellent | valet | desk |
| restaurant | comfortable | cable | really | inn | friendly | front |

Fig. 2. Once the number of topics is selected, the user can review the topics that were selected. The user can choose either Gibbs or CTM sampling for partitioning into topics.

4. Identifying Subtasks of Task by Textual Analysis

The Sentiminder has recently been given the ability to suggest subtasks for all of the topics in a given cluster. This provides an aid to the user to help find those subtasks that the user needs to be aware of when estimating the project management attributes, appearing in Figure 5. Each subtask needs to be considered, for example, when assigning a total number of days to address all of the work that makes up the task for a cluster.

Subtask identification is based on textual analysis, and exploits differences between the comments with positive sentiment and those with negative sentiment. The heuristic upon which it relies is quite straightforward. It consists of three steps. (1) Within the comments with most negative sentiment, we identify most frequently occurring nouns. A word's frequency is computed by first creating a term-document matrix and then for each word, represented as a row in this matrix, counting the number of its appearances in the documents, which is represented by adding the occurrence counts in the columns. (2) Likewise from the comments with the most negative sentiment, we find the most frequently

Topic 6 (0.68) 136 p 65 n
 Topic 5 (0.7) 147 p 63 n
 Topic 3 (0.82) 357 p 81 n
 Topic 1 (0.83) 308 p 62 n
 Topic 2 (0.84) 446 p 82 n
 Topic 4 (0.85) 384 p 70 n
 Topic 7 (0.88) 366 p 51 n

Fig. 3. The user can select the topic in a drop-down list. Topic 6 has a sentiment value of 0.52, and contains 136 positive sentiment words and 125 negative sentiment words. The topics with lowest sentiment are at the top of the list

| Job | Earliest start | Duration (days) | Cost | Not same time as |
|-----|---------------------------------|----------------------------------|-----------------------------------|---|
| T1 | <input type="text" value="1"/> | <input type="text" value="50"/> | <input type="text" value="800"/> | <input checked="" type="checkbox"/> T2 <input checked="" type="checkbox"/> T3 <input type="checkbox"/> T4 <input type="checkbox"/> T5 <input type="checkbox"/> T6 <input type="checkbox"/> T7 |
| T2 | <input type="text" value="30"/> | <input type="text" value="70"/> | <input type="text" value="400"/> | <input checked="" type="checkbox"/> T3 <input type="checkbox"/> T4 <input type="checkbox"/> T5 <input type="checkbox"/> T6 <input type="checkbox"/> T7 |
| T3 | <input type="text" value="2"/> | <input type="text" value="35"/> | <input type="text" value="1400"/> | <input type="checkbox"/> T4 <input type="checkbox"/> T5 <input type="checkbox"/> T6 <input type="checkbox"/> T7 |
| T4 | <input type="text" value="20"/> | <input type="text" value="10"/> | <input type="text" value="500"/> | <input checked="" type="checkbox"/> T5 <input type="checkbox"/> T6 <input type="checkbox"/> T7 |
| T5 | <input type="text" value="50"/> | <input type="text" value="30"/> | <input type="text" value="1200"/> | <input type="checkbox"/> T6 <input type="checkbox"/> T7 |
| T6 | <input type="text" value="1"/> | <input type="text" value="100"/> | <input type="text" value="800"/> | <input type="checkbox"/> T7 |
| T7 | <input type="text" value="70"/> | <input type="text" value="50"/> | <input type="text" value="800"/> | |

Fig. 4. After carefully reading all comments, the business owner has decided how much work is needed for each of the seven topics. She has decided which each task may start, and how long each task will take. She has also decided which tasks cannot be performed at the same time as other tasks. In this instance Tasks T1, T2 and T3 are will be done by the same team, so must be done at separate times. Likewise Tasks T4 and T5 cannot be done simultaneously.

occurring nouns. Finally (3) we remove those nouns that occur in the positive comments from the list of nouns in the negative comments. The remaining nouns in the negative list is taken as a set of hints to be used for the subtasks.

The intuition behind the heuristic is as follows. If these two comments appear: “The showers were in poor condition” and “The carpets were in excellent condition.” The nouns “shower” and “condition” would appear as frequent nouns in the negative comments, since “poor” is a word with negative affect. The nouns “carpets” and “condition” would appear as frequent nouns in the positive comments. After removing nouns in the positive comments from those in the negative comments, the remainder of that negative list would be “showers”. In this case, it is the showers that need to be repaired.

5. Initial Results

We show an example from the Opinions comments where the heuristic gives encouraging results. These nouns were identified with subtasks: “navigation”, “push”, “press”. The following comments are assessed as having low sentiment.

- I don't feel the need to push the button ahead of time to prepare for the end of the page at all, which evidently a lot of Kindle 1 owners do since it's a bit more sluggish .
- But holding the darn thing was a chore without pressing that next button page .

- I also found using the navigation button difficult .

In this case, the nouns that have been extracted from these reviews do quite clearly indicate the issues that the customers have reported. Over our dataset, in every topic with an overall low sentiment, this technique generates good candidate subtopics, according to our subjective assessment. However there are examples when subtopics suggestions are less helpful. In these cases, we found that the sentiment analysis has been misled by the appearance of word with negative sentiment in a review that is overall positive.

In Figure 5, we illustrate how the subtask suggestion feature is integrated in the Sentiminder, so that the business owner can select any suggested noun, and review the negative comments that where that noun is mentioned. This represents a significant reduction in the amount of material the business owner needs to consider.

Topics

| Topics | Job Data | Recommend | Dates | |
|--------|---|----------------|-----------------|------|
| Job | Subtasks | Earliest start | Duration (days) | Cost |
| T1 | <div>turn hate push edge miss book</div> | 1 | 50 | 800 |
| T2 | <div>backlight books press interior click run</div> | 30 | 70 | 400 |
| T3 | <div>press hold device sleep position feature</div> | 2 | 35 | 1400 |
| T4 | | 20 | 10 | 500 |
| T5 | | 50 | 30 | 1200 |
| T6 | <div>turn charge thought power pun question</div> | 1 | 100 | 800 |

Fig. 5. Subtasks suggestions for the topics with low sentiment, from reviews of the Amazon Kindle from the Opinosis review collection.

6. Conclusion

The Sentiminder prototype provides a comprehensive set of tools and an integrated work flow that guides a business owner when responding to a set of on-line comments and reviews from their customers. These reviews cover a number of aspects of their business, from many perspectives, so it can be a daunting task to comprehend and condense this information into a set of specific tasks that will address any deficiencies in the product or service. We have argued is it important to do so quickly.

In this paper we add a new feature to the Sentiminder that helps the business owner to rapidly comprehend what needs to be done. The Sentiminder already separates comments into topics using LDA clustering. Given an LDA cluster, we extract the comments with the most negative sentiment and those with positive sentiment. We then identify the nouns that frequently occur in the comments with negative sentiment, but do not frequently occur in those with positive sentiment. When integrated with the Sentiminder, our subtopic extraction provides the business owner with suggestions that are useful, according to our subjective assessment. This allows the business owner to focus quickly

on the negative reviews that pertain to this subtopic, which supports the Sentiminder's goal as a tool for generating rapid responses to on-line reviews.

In future work, we intend to fully integrate the subtasks component into the Sentiminder prototype and perform comprehensive evaluations with authentic users of subtopic extraction, among other features of the Sentiminder.

References

1. K. Ganesan, C. Zhai, J. Han, Opinosis: A graph-based approach to abstractive summarization of highly redundant opinions, in: Proceedings of the 23rd International Conference on Computational Linguistics, COLING '10, Association for Computational Linguistics, Stroudsburg, PA, USA, 2010, pp. 340–348.
URL <http://dl.acm.org/citation.cfm?id=1873781.1873820>
2. T. L. Ngo-Ye, A. P. Sinha, Analyzing online review helpfulness using a regression relief-enhanced text mining method, *ACM Trans. Manage. Inf. Syst.* 3 (2) (2012) 10:1–10:20. doi:10.1145/2229156.2229158.
URL <http://doi.acm.org/10.1145/2229156.2229158>
3. S. S. Htay, K. T. Lynn, Extracting product features and opinion words using pattern knowledge in customer reviews, *The Scientific World Journal* 2013. doi:<http://dx.doi.org/10.1155/2013/394758>.
4. G. Somprasertsri, P. Lalitrojwong, Mining feature-opinion in online customer reviews for opinion summarization, *Journal of Universal Computer Science* 16 (6) (2010) 938–955.
5. M. J. Sánchez-Franco, A. Navarro-García, F. J. Rondán-Cataluña, Online customer service reviews in urban hotels: A data mining approach, *Psychology & Marketing* 33 (12) (2016) 1174–1186. doi:10.1002/mar.20955.
URL <http://dx.doi.org/10.1002/mar.20955>
6. D. M. Blei, A. Y. Ng, M. I. Jordan, Latent dirichlet allocation, *J. Mach. Learn. Res.* 3 (2003) 993–1022.
URL <http://dl.acm.org/citation.cfm?id=944919.944937>
7. I. Feinerer, K. Hornik, D. Meyer, Text Mining Infrastructure in R, *Journal of Statistical Software* 25 (5) (2008) 1–54.
URL <http://www.jstatsoft.org/v25/i05/>
8. P. J. Stone, R. F. Bales, J. Z. Namenwirth, D. M. Ogilvie, The general inquirer: A computer system for content analysis and retrieval based on the sentence as a unit of information, *Behavioral Science* 7 (4) (1962) 484–498. doi:10.1002/bs.3830070412.
URL <http://dx.doi.org/10.1002/bs.3830070412>
9. F. Li, M. Huang, X. Zhu, Sentiment analysis with global topics and local dependency, in: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI'10, AAAI Press, 2010, pp. 1371–1376.
URL <http://dl.acm.org/citation.cfm?id=2898607.2898826>
10. F. Al-Obeidat, E. Kafeza, B. Spencer, Opinions Sandbox: Turning emotions on topics into actionable analytics, in: The 1st EAI International Conference on Emerging Technologies for Developing Countries, Springer, 2017.