Understanding the evolution of fine-grained user opinions in product reviews

1st Peike Xia
College of Computer Science and Electronic Engineering
Hunan University
Changsha, Hunan, China
xiapeike@hnu.edu.cn

2nd Wenjun Jiang†
College of Computer Science and Electronic Engineering
Hunan University
Changsha, Hunan, China
jiangwenjun@hnu.edu.cn

***Abstract*—Online reviews and ratings play an important role in shaping the purchasing decisions of customers in e-commerce. Previous research usually study the opinion extraction and sentiment analysis of online reviews to obtain a personalized score on all aspects of the product. Moreover, most of these studies only statically analyze user's scoring factors, neglecting the fact that user's opinions will change over time. To address the two issues, in this paper, we use review text information to mine user's aspect-based scores and user's opinion changes. We analyze the evolution of the users' opinions over time. First, we use a Latent Aspect Rating Analysis (LARA) method to extract and score aspects. Then, we focus on the evolution patterns, and analyze the changes in the users views over time from several perspectives: the sales volumes of products, the overall scores, the distribution of the overall ratings and aspect ratings, and the proportion of positive ratings.**

***Index Terms*—Review mining, Latent rating analysis, Opinion evolution, Aspect rating**

I. Introduction

In recent years, online shopping has become more and more popular, and many people purchase products and express their opinions online. In China, tens of thousands of consumer reviews are generated everyday, such as online shopping, online ordering. When users are making purchase decisions, they also rely heavily on product reviews made by previous customers. Therefore, these online reviews, as the Internet word of mouth, play an important role in promoting potential consumers to make purchase decisions. Consumers can get detailed reviews of goods and services from other consumers. Businesses and producers can collect information about the pros and cons of their products provided by their customers, and then better understand consumer needs and find directions for improvements. Therefore, it is very valuable to conduct research on online consumer reviews in order to better understand the opinions of reviewers.

When a user posts a comment on a purchased product, he or she will give an overall score of 1 to 5, which largely reflects the user's satisfaction with the product. When the score is 4 or 5, it indicates that the user has a positive attitude towards the product. When the score is 1 or 2, it indicates that the user has a negative attitude to the product. Further, if most users give a high score to a product, we can estimate that the quality of this product is good and it is worth buying, otherwise we will not buy it. Merchants can also know their own product more clearly based on these scores, and make corresponding improvements to meet the needs of users. Some researchers have explored the evolution of user opinions based on user scores [16-17,23]. They use users ratings collected from e-commerce sites and measure their changes to represent changes in opinions and to study changes in behavior related to rating changes, such as changes in the number of comments posted or more extreme opinions.

However, with only such a score, we can only know the user's overall preference for a certain product, and cannot understand the user's specific points of view. For example, if a user gives an overall rating of 5 to a mobile phone, it may not mean that the mobile phone exactly matches the user's preference. It may be just because this user has a special preference for the appearance of this mobile phone, but he may not be satisfied with the energy consumption of the phone. Therefore, more and more researchers are focusing on excavating the fine-grained opinions in the reviews to better understand the needs of users and developers. There are already a lot of work on this, including the topic model-based aspects of extraction [5,6], work based on comments on the emotional polarity of the perspective of mining [7] and so on.

Our work tries to combine review mining with evolution of opinions to explore the changes in user perspectives. For example, when a mobile phone is first introduced, its configuration may be at the top of the market at that time, and users often give relatively high scores on the performance aspect. However, as the time goes on, the software was updated and replaced at a high speed. The demand for mobile phone configuration is getting higher and higher. At that time, the configuration of the mobile phone may not be able to meet the requirements of users. It implies that the user's rating on the performance aspect will slowly decline. We hope that through the technology of review mining, we can analyze the comments, extract the opinions that the users want to express, and then analyze users fine-grained opinions and study their evolution trends.

In this paper, we use the Latent Aspect Rating Analysis (LARA) method to extract the aspects and score reviews, because LARA can well express the opinions mentioned in the user reviews and give each aspect a score. This gives us a good understanding of how satisfied a user is with a product, it also reflects the users' point of view in the comments. Based on this, we consider the factors of time, and analyze the changes in the user's views over time from several perspectives: the sales volumes of products, the distribution of the overall ratings and aspect ratings, and the proportion of positive ratings. Through these fine-grained analysis, it can help us better understand the changes in users' opinions. The main contributions are as follows:

1) Obtain user's evaluation of various aspects of the product by mining product review information.
2) Analyze the differences in user opinions by combining the overall rating and the aspect ratings of the products.
3) Analyze the evolution of the user's perspective over time from a more fine-grained perspective.

The remainder of this paper is organized as follows. Section 2 reviews some of the related work of opinion mining and the evolution of user perspectives. Section 3 details the LARA method for extracting and scoring reviews. Section 4 gives a large number of experimental results to analyze our conjecture. Section 5 provides concluding remarks and future work.

II. Related Work

In this section, we mainly discuss the related work of review mining and evolution of opinions.

A. *Opinion Mining*

Opinion mining and sentiment analysis has received more and more attention in the past few years. Ravi and Ravi [8] present a detailed survey of opinion mining and sentiment analysis, summarizing over one hundred articles on necessary tasks, methods, and applications for opinion mining and sentiment analysis published in the past decade. Muhammad et al. [9] utilize local and global semantic information to improve lexicon-based sentiment analysis approaches. Fujita [10] proposed a hybrid method that uses natural language processing (NLP) basic techniques and enhances the sentiment dictionary with the help of SentiWordNet [11], and fuzzy sets, to the Sentiment Analysis problem at the sentence level. Poria et al. [12] present the first deep learning approach to aspect extraction in opinion mining. proposed the first deep learning method to extract aspect in opinion mining.

With the development of opinion mining techniques based on aspect level, the relationship of ratings and aspect-based sentiment has become increasingly attractive. As an important feature of the text, aspect sometimes referred to as topics has attracted much attention in recent years. The first question is which aspect is the most important, and a large number of literatures have studied the aspects of extracting comments based on the LDA models [13,14]. Parker et al. Yu et al. [15] and Zha et al. [24] used a probabilistic aspect ranking algorithm to determine important aspects after conducting an aspect discovery and sentiment classification for users reviews of a certain product. After aspect weights, aspect rating also becomes focus [18,19-21]. Wang et al [20,21] evaluated a probabilistic rating regression model to analyze potential aspect rating. Moghaddam and Ester [18] and Wang and Ester [19] improved the collaborative filtering model and combined the topic model to identify aspects of the user review text and classified sentiment as positive and negative.

Considering that we need to get specific scores from the comments to do point of view analysis. We used the LARA method [20] to extract and score reviews, since it is very classic and more suitable for our research methods.

B. *Opinion Evolution of Online Consumer Reviews*

According to our survey, in the e-commerce environment, only a few of the research has explored the issue of the evolution of users opinion in product reviews.

They use users ratings collected from e-commerce sites and measure their changes to represent changes in opinions and to study changes in behavior related to rating changes, such as changes in the number of comments posted or more extreme opinions. For example, Moe and Schweidel [1] study how previously posted ratings may affect an individuals posting behavior in terms of whether to post reviews (incidence) and what to post (evaluation), and find that positive ratings environments increase posting incidence, while negative ratings environments discourage posting. Godes and Silva [2] investigate the evolution of online ratings over time and in terms of sequences. They establish that there exist two distinct dynamic processes, one as a function of the amount of time a product has been available for review and the other as a function of the sequences of the reviews themselves. They find that when previous reviews are divided and diversified, subsequent reviews may lead to more purchase errors and lower ratings. Mochon and Schwartz [3] find that a product may be negatively influenced by the quality of the previous product reviewed, but positively influenced by the star rating assigned to that product. Chen et al. find that the relationships between marketing variables and online posting behavior by consumers are different in the early and mature stages. However, the posting behavior they discuss is related to social media rather than e-commerce websites [4].

III. METHOD

We use the LARA method [20] to conduct a aspect-based sentiment analysis of product reviews. There are two main steps. The first is the aspect segmentation algorithm, and the second is the aspect rating. The choice of this solution is mainly due to it can be used to display the users' point of views well with semantics, and it can be scored with clear numerical values in every aspect. Next we give the details of LARA.

LARA can be thought of as a model whose input is a set of comments for a class of commodity entities, each of which has an overall rating. Such a format of reviews is quite common in most of the merchants web site, e.g. Amazon (www.amazon.com) and Epinions (www.epinions.com), and the number of such reviews is growing constantly. Its output is a set of aspect rating related to each comment.

Formally, let $D = \{d_1, d_2, \cdots, d_{|D|}\}$ be a set of review text documents for an interesting entity or topic, and each review document $d \in D$ is associated with an overall rating r_d. We also assume that there are n unique words in the vocabulary $V = \{w_1, w_2, \cdots, w_n\}$.

We further assume that we are given k aspects, which are rating factors that potentially affect the overall rating of the given topic. For example, for phone reviews, possible aspects may include 'price' and 'appear'. An aspect is specified through a few keywords, and provides a basis for latent aspect rating analysis.

Definition (Aspect): An aspect A_i is a (typically very small) set of words that characterize a rating factor in the reviews. For example, words such as price, value, and worth, can characterize the price aspect of a phone. We denote an aspect by $A_i = \{w|w \in V, A_{(w)} = i\}$, where $A_{(.)}$ is a mapping function from a word to an aspect label.

Definition (Aspect Ratings) Aspect rating s_d is a k dimensional vector, where the i-th dimension is a numerical measure, indicating the degree of satisfaction demonstrated in the review d toward the aspect A_i , and $s_{(d_i)} \in [r_{min}, r_{max}]$ A higher rating means a more positive sentiment towards the corresponding aspect.

Specifically, the basic workflow of the proposed aspect segmentation algorithm is as follows: given the initial seed word of each aspect and all the comment texts as input, we associate the defined terms with each of the sentences with the defined aspects, the aspect with the most occurrences is the aspect label of the sentence; based on this initial aspect annotation, we calculate the dependencies between aspects and words by Chi-Square (χ^2) statistic [22], and include the words with high dependencies into the corresponding aspect keyword list. These steps are repeated until the aspect keyword list is unchanged or the number of iterations exceeds the limit. The full description of the algorithm is in Algorithm 1. The Chi-Square (χ^2) statistic to compute the dependencies between a term w and aspect A_i is defined as follows:

$$\chi^2(w, A_i) = \frac{C \times (C1C4 - C2C3)^2}{(C1 + C3) \times (C2 + C4) \times (C1 + C2) \times (C3 + C4)} \tag{1}$$

where C1 is the number of times w occurs in sentences belonging to aspect A_i, C2 is the number of times w occurs in sentences not belonging to A_i, C3 is the number of sentences of aspect A_i that do not contain w, C4 is the number of sentences that neither belong to aspect A_i, nor contain word w, and C is the total number of word occurrences. After aspect segmentation, we get a set of sentences containing aspects. And we would get k partitions of each review d, represent them as a $k \times n$ feature matrix W_d, where $W_{d_{ij}}$ is the frequency of word W_j in the text assigned to aspect A_i of d normalized by the total counts of words in the text of that aspect. Then the emotional polarity of the sentence is calculated for each aspect of the sentence and an emotional score is scored as an aspect score. The scoring is defined as follows

$$s_i = \sum_{\mathrm{j=1}}^{\mathrm{n}} \beta_{\mathrm{ij}} \mathrm{W}_{\mathrm{d}_{ij}} \tag{2}$$

where β_{ij} indicates the word sentiment polarities on aspect A_i. Due to the low aspect coverage in individual reviews (not all reviewers would talk about every aspect of an entity in their reviews), it is infeasible for us to infer the latent aspect ratings of every single review document. As our solution, we gather each every-week's comment and mark the average score of each aspect as the corresponding aspect of the product for this week, which corresponds to various aspects of the user's perspective during this time period.

Algorithm 1 Aspect Segmentation Algorithm

Input: A collection of reviews $\{d_1, d_2, \cdots, d_{|D|}\}$, set of aspect keywords $\{T_1, T_2, \cdots, T_k\}$, vocabulary V,selection threshold p and iteration step limit *I*.

Output: Reviews split into sentences with aspect assignments.

Step 0: Split all reviews into sentences, $X = \{x_1, x_2, \cdots, x_M\}$;

Step 1: Match the aspect keywords in each sentence of X and record the matching hits for each aspect i in Count(i);

Step 2: Assign the sentence an aspect label by $a_i = \arg\max_i Count(i)$. If there is a tie, assign the sentence with multiple aspects;

Step 3: Calculate χ^2 measure of each word (in V);

Step 4: Rank the words under each aspect with respect to their χ^2 value and join the top p words for each aspect into their corresponding aspect keyword list T_i;

Step 5: If the aspect keyword list is unchanged or iteration exceeds *I*, go to **Step 6**, else go to **Step 1**;

Step 6: Output the annotated sentences with aspect assignments.

Although we only analyze aspect opinions at the entity-level, we can still study the detailed aspect-level opinions within each review by applying the learned aspect rating model on the identified aspect segments. Such analysis helps us visualize the detailed review content and extract opinionated sentences for summarizing the items of interest.

IV. Data Analysis

A. Data Set and Preprocessing

Our experimental data comes from the Amazon.com . These data include user's ID, item's ID, the user's comment on products, and a score assigned to the user while writing a review.

We first pre-process these comments. The first step is to convert all words into lowercase. In the second step, we remove the stop words in the comments and remove the words that appear less than 3 in the comments. We believe that the words in the expression often appear multiple times. In the third step, we use the Stemmer tool in NLTK to extract the root of each word.

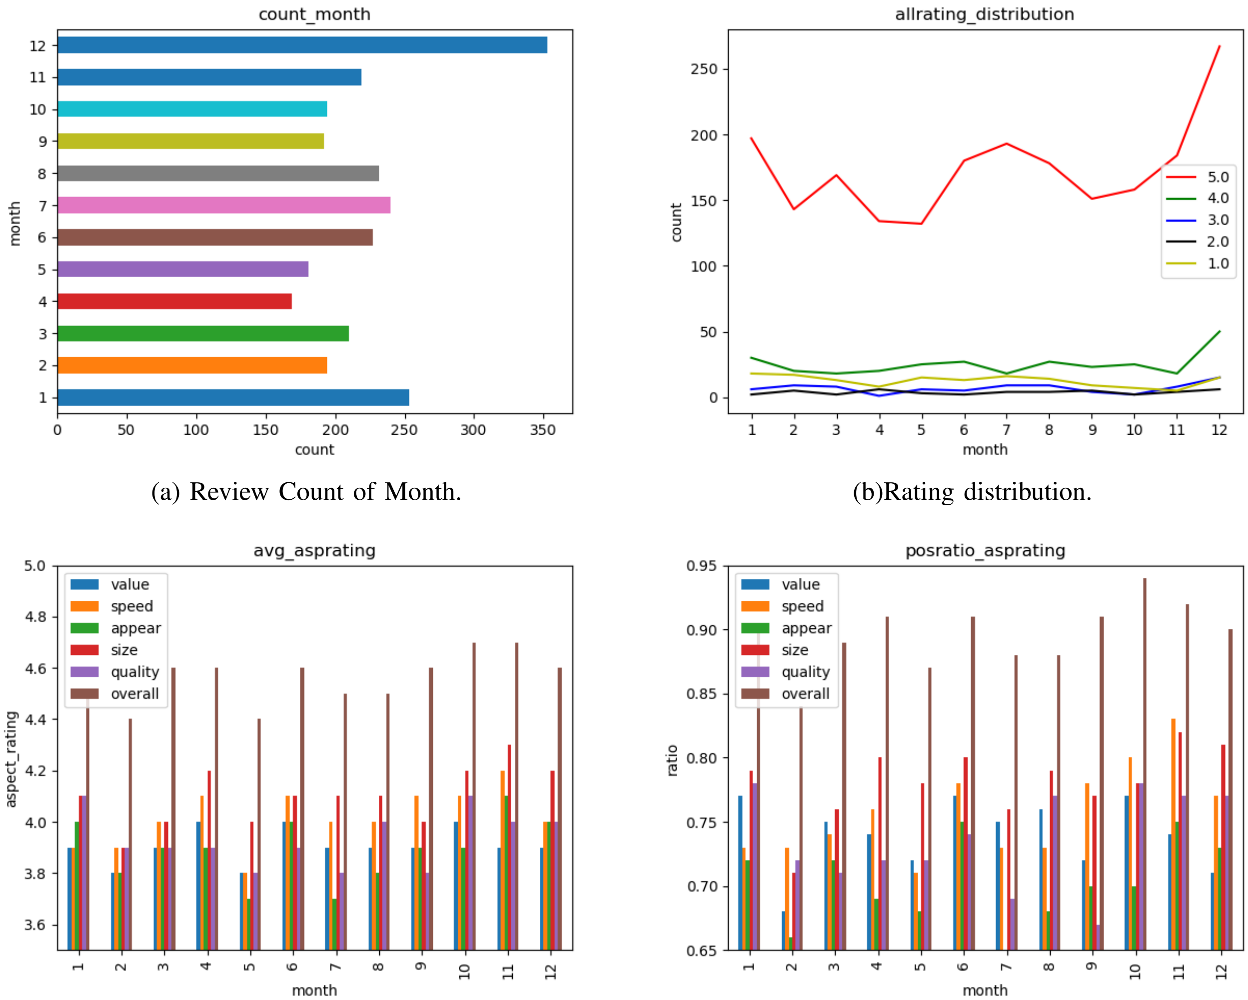

(a) Review Count of Month.

(b)Rating distribution.

(c) Aspect Rating .

(d)Positive Ratio.

Fig. 1. The Change of User's Opinion.

In order to show the trend of opinions evolution over time. We select the 2013 year with the most comments for analysis from the review data we obtained, a total of 592,748 records. Since the characteristics of each product are different, in order to analyze the changing trend of the viewpoints, we select a memory card with a total of 2664 review data as our experimental data.

According to the aspect segmentation algorithm, we initially define five aspects, and artificially select several sets of seed words for each aspect, used as the input to the previously described aspects of the segmentation algorithm, in which we set the parameter p=3 and the number of iterations I=10. Table 1 shows the setting of initial aspect terms.

Table 2 shows an example of aspect rating prediction. We select a comment covering all aspects from all the comments as a presentation. By aspect rating of users' reviews we predicted, it is possible to clearly observe the user's opinions. Unlike overall rating , it can only reflects the overall satisfaction of users. However, in actual commodity reviews, there are very few reviews covering all aspects like this. The average user will only mention a few aspects. Therefore, we calculate the aspect rating of the commodity, which can be based on the average of the users' aspect rating over a period of time.

TABLE I
ASPECT SEED WORDS

Aspects	Seed words
value	value, price, worth, expensive, cheap
speed	speed, fast, transfer, slow
appear	appear, cool, look, small, tiny
size	size, little, storage, space, big
quality	quality, work, heat, run, compatible

B. Opinion Evolution Analysis

In this section, we mainly analyze the changes of users' points of view from multiple perspectives. We believe that changes in the views of users can be reflected in the

TABLE II
ASPECT RATING PREDICTION

OverallRating:	4.0
Review:	**Good size** card and a **great price**. I use this product with the GoPro 3 Black Edition. It **works fantastic**, even at 2.7K resolution, and **imports very fast**, especially over USB 3 card readers. I almost always go with SanDisk these days for flash RAM; they invented much of it and it shows.My only **complaint is its appearance**. It looks really bad, like a hard paper.
Aspect Rating:	value: 5 speed: 5 appear: 2 size: 4 quality: 4

following aspects: the sales volumes of the products, the distribution of the overall ratings, the distribution of the aspect ratings, and the proportion of positive ratings (that is, the number of scores greater than 4 divided by the total number).

As is shown in Figure 1. Figure (a) shows the number of sales per month of merchandise, with the horizontal axis representing the number of sales and the vertical axis representing the month. It can be seen clearly that in the 12th month, the number of sales of goods increase substantially, which is almost 2 times that of other months. In April, sales reach the lowest point of the year. Figure (b) shows the distribution of the overall score. It can be seen that most of the users give a score of 5, a small part of the users give 4 , and only a few users give a score of less than 4. This can only show that the user is generally satisfied with the product, but does not know the user's specific preferences for the product.

Therefore, we combine the rating with the time factors. Since each user's comment may not include all aspects, we aggregate all comments for each month and calculate the average scores for each aspect of these comments as a final aspect rating. Figure (c) and (d) show the average aspect scores for each month and the proportion of positive ratings. Further draw the trend of the overall ratings and the aspect ratings. It can be clearly seen that the users preference for this memory card focuses on the two aspects of 'speed' and 'size', and as the time goes on, the user is getting higher and higher on these two aspects. The highest peak is reached in the 11th month.

For more fine-grained analysis, we further analyze the user's viewpoints at week intervals. In order to show it more clearly, we mainly select two aspects, 'value' and 'quality', and show changes from week 15 to week 21. As is shown in Figure 2, the horizontal axis represents the time, and the vertical axis represents the rating. It can be seen that from the 15th week to the 20th week, the users' scores on the 'value' aspect has always been higher than the 'quality' aspect. But from the 20th week to the 21st week, the overall ratings remains unchanged. Under the circumstances, it can be clearly seen that the score on the 'value' aspect drop sharply. It shows that during this week, users are not satisfied with the 'value' aspect. This is the information we did not find in the overall scores.

Figure 3 shows the count of comments from week 15 to week 20, with the horizontal axis representing the number of comments and the vertical axis representing time. In Week 21, there is a noticeable increase in the number of user reviews, which represent an increase in the number of people who purchase this product during the week. However, as can be seen from Figure 1, from the 19th week to the 21st week, the overall score change is not significant. However, the score on the 'value' aspect decline significantly, and the score on the 'quality' aspect has been improved. Although the rise in commodity prices has affected the user's scoring on the 'value' aspect, the increase in 'quality' has led to a rise in the sales of goods, not to drop. It can also reflect that the user's preference for 'quality' is higher than 'value'. Through the analysis of these data, we find that through the mining of the review, we can better grasp the user's point of view changes.

In this section, we show the results of LARA's aspect extraction and scoring. Taking into account the time factor, the evolution of the users' points of view are analyzed from

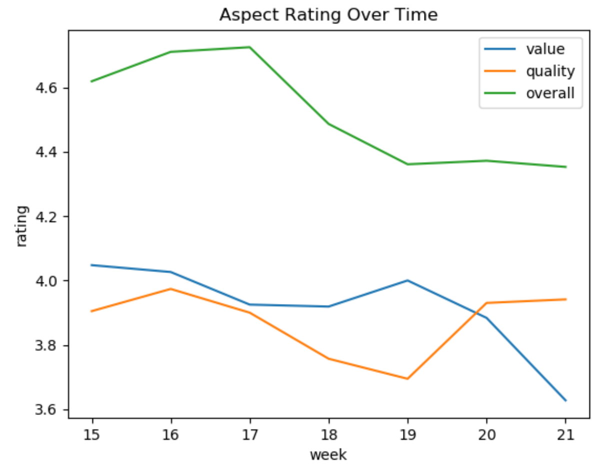

Fig. 2. Aspect Rating Over Time.

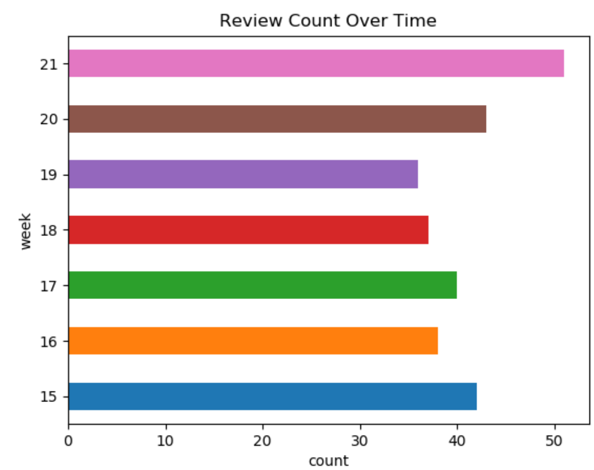

Fig. 3. Review Count Over Time.

the perspective of the number of reviews, and the distribution of the overall ratings and the aspect ratings in terms of weeks and months. Through experimental analysis, our main findings include the following:

- Overall ratings do not respond well to user satisfaction with a product. Fine-grained aspect rating can better reflect the users' opinions of the product.
- The users' points of view tend to change over time. The user originally did not like certain aspects of the product and may be interested in these aspects slowly because of certain factors.
- The change of the users' viewpoints can be directly reflected from the sales of the product, and it can reflect the user's different degree of preference for the product.

V. Conclusion

We try to understand the evolution of fine-grained user opinions in product reviews. To this end, we use the Latent Aspect Rating Analysis (LARA) method to extract and score aspects from products reviews, which takes the comment and the overall rating as the input, and obtains aspect ratings for users according to the initially defined aspect seed words. Then we conduct extensive data analysis on the aspect-based opinions. We believe that user's opinions of a commodity are not immutable, and they will be affected by various factors. Through the analysis of experimental data, we can clearly see that the user's point of view is indeed affected by several factors.

There are many directions for future work. The first is that the amount of data in current work is not sufficient. We only consider a single commodity, and the conclusions we get may be a bit lopsided. Second, we only consider the data for one year. We will study the data for several years for comparison. In our future research, we hope to capture specific factors that affect the evolution of fine-grained user opinions. For example, when a user's score difference between adjacent time intervals is greater than a certain threshold, we can set a model to infer what happened and why.

Acknowledgement

This work is supported by NSFC grants 61632009 and 61502161, NSF grants of Hunan province 2017JJ3035, and the Chinese Fundamental Research Funds for the Central Universities 531107040845. This work is partially supported by the National Science Foundation of China under Grant No. 61472131, 61772191; Science and Technology Key Projects of Hunan Province (2015TP1004, 2015SK2087, 2015JC1001, 2016JC2012)

References

[1] Wendy W Moe and David A Schweidel. Online product opinions: Incidence, evaluation, and evolution. *Marketing Science*, 31(3):372–386, 2009.

[2] David Godes and Jos C Silva. Sequential and temporal dynamics of online opinion. *Marketing Science*, 31(3):448–473, 2012.

[3] Daniel Mochon and Janet Schwartz. The individual dynamics of online reviews. *Advances in Consumer Research*, 2014.

[4] Yubo Chen, Scott Fay, and Qi Wang. The role of marketing in social media: How online consumer reviews evolve . *Journal of Interactive Marketing*, 25(2):85–94, 2011.

[5] Zhiyong Cheng, Ying Ding, Lei Zhu, and Mohan Kankanhalli. Aspect-aware latent factor model: Rating prediction with ratings and reviews. 2018.

[6] Lin Qiu, Sheng Gao, Wenlong Cheng, and Jun Guo. Aspect-based latent factor model by integrating ratings and reviews for recommender system. *Knowledge-Based Systems*, 110:233 – 243, 2016.

[7] Jiangtao Qiu, Chuanhui Liu, Y Inghong Li, and Zhangxi Lin. Leveraging sentiment analysis at the aspects level to predict ratings of reviews. *Information Sciences*, 451, 2018.

[8] Kumar Ravi and Vadlamani Ravi. *A survey on opinion mining and sentiment analysis: Tasks, approaches and applications*. Elsevier Science Publishers B. V., 2015.

[9] Aminu Muhammad, Nirmalie Wiratunga, and Robert Lothian. Contextual sentiment analysis for social media genres. *Knowledge-Based Systems*, 108:92–101, 2016.

[10] Orestes Appel, Francisco Chiclana, Jenny Carter, and Hamido Fujita. A hybrid approach to the sentiment analysis problem at the sentence level. *Knowledge-Based Systems*, 108:110–124, 2016.

[11] Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *International Conference on Language Resources and Evaluation, Lrec 2010, 17-23 May 2010, Valletta, Malta*, pages 83–90, 2010.

[12] Soujanya Poria, Erik Cambria, and Alexander Gelbukh. Aspect extraction for opinion mining with a deep convolutional neural network. *Knowledge-Based Systems*, 108:42–49, 2016.

[13] Yohan Jo and Alice H. Oh. Aspect and sentiment unification model for online review analysis. In *ACM International Conference on Web Search and Data Mining*, pages 815–824, 2011.

[14] Wei Wang. Sentiment analysis of online product reviews with semi-supervised topic sentiment mixture model. In *Seventh International Conference on Fuzzy Systems and Knowledge Discovery*, pages 2385–2389, 2010.

[15] Jianxing Yu, Zheng Jun Zha, Meng Wang, and Tat Seng Chua. Aspect ranking: identifying important product aspects from online consumer reviews. In *Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1496–1505, 2011.

[16] W. Jiang, J. Wu, G. Wang, and H. Zheng. Forming opinions via trusted friends: Time-evolving rating prediction using fluid dynamics. *IEEE Transactions on Computers (TC)*, 65(4):1211–1224, 2016.

[17] W. Jiang, J. Wu, G. Wang, and H. Zheng. Fluidrating: A time-evolving rating scheme in trust-based recommendation systems using fluid dynamics. *Proc. IEEE INFOCOM*, pages 1707–1715, 2014.

[18] Samaneh Moghaddam and Martin Ester. Ilda:interdependent lda model for learning latent aspects and their ratings from online product reviews. pages 665–674, 2011.

[19] Hao Wang and Martin Ester. A sentiment-aligned topic model for product aspect rating prediction. In *Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, 2014.

[20] Hongning Wang, Yue Lu, and Chengxiang Zhai. Latent aspect rating analysis on review text data:a rating regression approach. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 783–792, 2010.

[21] Hongning Wang, Yue Lu, and Cheng Xiang Zhai. Latent aspect rating analysis without aspect keyword supervision. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 618–626, 2011.

[22] M. F. Porter. *An algorithm for suffix stripping*. Morgan Kaufmann Publishers Inc., 1997.

[23] W. Jiang and J. Wu. Active opinion-formation in online social networks. *Proc. IEEE INFOCOM*, pages 1440–1448, 2017.

[24] Zheng Jun Zha, Jianxing Yu, Jinhui Tang, Meng Wang, and Tat Seng Chua. Product aspect ranking and its applications. *IEEE Transactions on Knowledge & Data Engineering*, 26(5):1211–1224, 2014.