

OpinionSeer: Interactive Visualization of Hotel Customer Feedback

Yingcai Wu, Furu Wei, Shixia Liu, Norman Au,
Weiwei Cui, Hong Zhou, and Huamin Qu, *Member, IEEE*

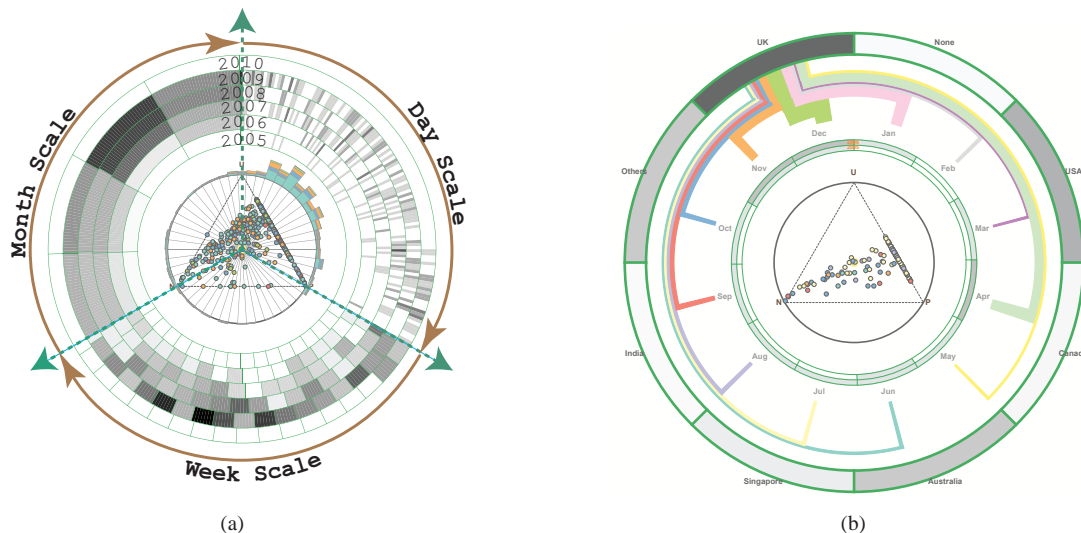


Fig. 1. (a) Temporal rings at different scales (month, week, and day); (b) Temporal and geographic rings where their relationships can be shown on demand by the curved belts.

Abstract— The rapid development of Web technology has resulted in an increasing number of hotel customers sharing their opinions on the hotel services. Effective visual analysis of online customer opinions is needed, as it has a significant impact on building a successful business. In this paper, we present OpinionSeer, an interactive visualization system that could visually analyze a large collection of online hotel customer reviews. The system is built on a new visualization-centric opinion mining technique that considers uncertainty for faithfully modeling and analyzing customer opinions. A new visual representation is developed to convey customer opinions by augmenting well-established scatterplots and radial visualization. To provide multiple-level exploration, we introduce subjective logic to handle and organize subjective opinions with degrees of uncertainty. Several case studies illustrate the effectiveness and usefulness of OpinionSeer on analyzing relationships among multiple data dimensions and comparing opinions of different groups. Aside from data on hotel customer feedback, OpinionSeer could also be applied to visually analyze customer opinions on other products or services.

Index Terms—Opinion visualization, radial visualization, uncertainty visualization.

1 INTRODUCTION

The rapid development of the Internet and e-commerce has brought numerous customer review websites. Prior studies [11, 27] show that positive online reviews have a significant impact on customers' decision-making process. Online customer complaints (e-complaints),

if not handled properly, could easily cause customers to lose loyalty for related products/services, reduce patronage, and create negative word-of-mouth [1]. Thus, online customer feedback of products/service is useful for customer behavior analysis and is important for businesses. For example, when a new service is launched by a hotel chain, the relationship manager would need to know how customers with different backgrounds comment on this new service, and how they compare it with similar services of its competitors. Understanding and tracking this information could help improve customer satisfaction and build customer trust and loyalty over time. As a result, there is a growing need to extract and analyze customer opinions from large collections of online customer reviews.

Recently, much effort has gone into automatic opinion mining [23], making it possible to obtain customer opinions from a large amount of free review text. However, visually examining and analyzing such mining results have not been well addressed in the past. Most existing efforts use basic visualization (e.g., the bar chart used in [20]) to display the final opinion mining results to their audiences. Although existing techniques have achieved certain success, they cannot piece together information from multiple aspects to enable analysts to make a quick decision. In addition, current opinion visualization tools provide scant support for complex opinion analysis, such as identifying

- Yingcai Wu, Weiwei Cui, and Huamin Qu are with the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. Email: {wuyc|weiwei|huamin}@cse.ust.hk.
- Furu Wei and Shixia Liu are with the IBM China Research Lab, Beijing, China. Email: {weifuru@cn.ibm.com, shixia@gmail.com }.
- Norman Au is with the School of Hotel & Tourism Management at the Hong Kong Polytechnic University, Kowloon, Hong Kong. Email: hmnorman@polyu.edu.hk.
- Hong Zhou is with the Shenzhen University, Shenzhen, China. Email: hzhou@szu.edu.cn.

Manuscript received 31 March 2010; accepted 1 August 2010; posted online 24 October 2010; mailed on 16 October 2010.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

underlying factors influencing customer complaint behaviors and analyzing the relationships between demographic characteristics (e.g., age and gender) and complaint behaviors. Moreover, current techniques do not account for uncertainty or inaccuracy, which may lead to wrong conclusions.

In this study, we focus on the visual analysis of online hotel customer feedback. Hotel customers are mostly tourists with diverse cultural backgrounds, coming from different countries. Such diversity may likely cause varied levels of expectations toward the products/service offered, which could be a cause of complaining behavior in the case of product/service failure. For example, Au et al. [1] discovered that mainland Chinese are generally price sensitive, while customers from the US care more about space, cleanliness, and service. Knowing the opinion patterns is important for hotel managers. However, reasoning about customer opinions to detect useful patterns could be time-consuming and difficult for several reasons. First, collected opinion data are high-dimensional and heterogeneous data with structured category dimensions and unstructured review comments, posing a challenge to analysis and visualization. Second, because of the lexical and structural ambiguity of human language, it is difficult for computer systems to determine the exact intended meanings of words. Consequently, effectively modeling the ambiguity and faithfully presenting the information with ambiguity to analysts is also a major obstacle. Finally, no clear boundary exists between positive and negative opinions. Thus, the visualization system should be carefully designed to present all opinions to users with sufficient visual cues, and allow users to determine which subset to further visualize. These features, among others, make opinion data visualization challenging.

We design and develop OpinionSeer to address the need to effectively communicate opinion-mining results and facilitate the analytical reasoning process. In the system, we use a new feature-based opinion mining technique to faithfully model the uncertainty in the review text. In addition, subjective logic [14] is employed to handle and organize multiple opinions with degrees of uncertainty. Moreover, instead of inventing an unfamiliar visual representation, we augment familiar visual metaphors to convey the results from complex opinion analysis. Considering the analytical task and data characteristics of opinion mining, we combine the simplicity and familiarity of radial visualization, scatterplots, and tag clouds while addressing their shortcomings, such as the lack of relationship analysis among multiple facets.

OpinionSeer has two possible uses. Hospitality researchers can use it as a general analysis tool to analyze and detect hidden patterns in raw text data, and provide a user-friendly visual presentation to end users such as hotel managers. For hotel managers, the system allows them to identify useful and meaningful relationships quickly among vast amounts of textual data uploaded by customers on the e-channel, so that an effective decision can be better formulated to give timely and appropriate responses to the customers. Aside from data on hotel customer feedback, OpinionSeer could be applied to the visual analysis of customer opinions on other products or services. The major contributions of our work are as follows:

- We combine an opinion mining technique with subjective logic to model uncertainty in opinions and fuse the opinions.
- We design a new visual representation for customer feedback data to naturally encode the uncertainty information.

2 RELATED WORK

In this section, we discuss related work in two research topics: opinion mining and opinion visualization.

Opinion mining (also known as sentiment analysis) [23] is used to automatically detect relevant opinions within a large volume of review collection. Many approaches have been proposed to mine the overall opinion information at the document level [24] or sentence level [17]. However, a positive review on an object does not always indicate that the opinion holder has positive opinions on all aspects or features of the examined object. To further obtain such detailed aspects, feature-level opinion mining [12] [13] has been proposed and extensively studied on product reviews [25] to find opinions expressed on individual product features. The opinion-mining model in OpinionSeer is built

on the latter method, but is focused on visualizing the opinion mining results, which accounts for uncertainty to effectively model and analyze customer opinions. Moreover, we provide users with visual interaction tools to examine the results from multiple perspectives.

There has been recent growing interest in visualizing opinions extracted from customer reviews posted online. These methods can be classified into two categories: document-level and feature-level opinion visualization. Document-level visualization focuses on visualizing opinion data at the document level. For example, Morinaga et al. [21] suggested a 2D scatterplot called positioning map to show the group of positive or negative sentences. Gamon et al. [8] derived a number of topics and estimated the average sentiment value for each topic. A TreeMap-style user interface called Pulse was designed to visualize the topics and their sentiment values. Chen et al. [3] presented a visual analysis system with multiple coordinated views, such as decision trees and term variation graph, to help users understand the nature and dynamics of conflicting opinions. Gregory et al. [10] suggested an adapted rose plot to display sentiment aspects such as positive, negative, pleasure, pain, and conflict. More recently, Draper and Riesenfeld [6] developed an interactive visualization system to allow users to visually construct queries and view results in real time. Wanner et al. [29] described a concise visual encoding scheme to represent attributes, such as the sentiment, of each RSS news item. The BLEWS system [9] represents the number of documents related to a specific news article as a bar, and then uses an emotionally weighted glow (or halo) around the bars to convey the emotional sentiment.

Although the document-level opinion visualization provides a high-level opinion overview of customer reviews, but not enough details are presented for users to understand customer opinions on certain product/service features (e.g., room, service, and price). With the development of feature-based opinion mining, visualization researchers have developed feature-level opinion visualization. For example, Liu et al. [20] proposed a method to extract feature-level opinions from customer reviews, and augmented traditional bar charts to facilitate visual comparison of extracted feature-level opinions. Oelke et al. [22] introduced several visualization techniques including visual summary reports, cluster analysis, and circular correlation map to facilitate visual analysis of customer feedback data at the feature level. Unlike previous methods, which are either document-level or feature-level opinion visualizations, our method provides a flexible visualization supporting both feature- and document-level opinion visualization using subjective logics. In addition, while existing methods do not consider the uncertainty of opinion extraction, our visualization approach explicitly accounts for uncertainty to reveal faithfully the underlying data. Moreover, we introduce a new visual representation of opinions by augmenting a radial layout. The radial layout enables an integrated visualization of user feedback with multiple dimensions including demographics, temporal, and spatial information, thus allowing analysts to discover opinion patterns more quickly and efficiently.

3 DATA AND TASK ABSTRACTION

In this section, we introduce the selected opinion data, the traditional approach on hotel feedback data analysis, and task abstraction.

3.1 Opinion Data

TripAdvisor¹ is one of the most popular tourism cyber-intermediaries on the Web. Its users are from all over the world, with enormous cultural diversity. Compared with other Websites, the customer profile is relatively more complete. Thus, hotel customer reviews from TripAdvisor are selected as our data samples for our system. The data we obtained from TripAdvisor can be divided into three parts: hotel data, customer data, and review data. Hotel and customer data contain basic information about hotels and customers in the data samples, while the review data include review information such as detailed free-text comments and the review sentiments estimated by our approach.

¹www.tripadvisor.com

3.2 Traditional Analysis Approach

In hospitality research on e-complaints, researchers usually adopt a content analysis procedure or popular qualitative analysis software such as NVivo² to analyze opinion data. Complaints or opinions are first classified into different categories using the grounded theory approach and keyword analysis. Further relationship analysis is conducted using a two-way contingency table analysis. However, dealing with such large-scale, heterogeneous, and high-dimensional data poses a great challenge even for professional hospitality researchers, not to mention hotel managers. Moreover, even if some opinion patterns are found, presenting the findings to a wider audience is another challenge.

3.3 Task Abstraction

To better understand the problem domain and identify the potential uses of the customer feedback data, we compiled a list of detailed questions on customer feedback data that could spike the interest of the end users of our visualization. The end users of the system include hospitality researchers and hotel managers. Through a series of interviews with our target users, we found that hospitality researchers usually study opinion relationships, such as the relationship of opinions and the service category, as well as the hidden patterns related to customers' cultural background. Hotel managers, on the other hand, need to know customer opinions in a short time to take timely actions. The analysis tasks are summarized as follows.

- Q.1 How is the deviation of a group of opinions from the average?
- Q.2 How could several groups of opinions be compared effectively?
- Q.3 How do people's backgrounds affect their opinions on a hotel or a certain group of hotels?
- Q.4 What are the differences in the cultural background of two groups of customers who hold similar or different opinions?
- Q.5 Is there any conflict between free-text comments and the score ratings, e.g., a good review with low ratings?
- Q.6 Are there any localization or geography patterns regarding user opinions on a hotel or a certain group of hotels?
- Q.7 Are there any temporal patterns regarding the users opinions?

4 SYSTEM OVERVIEW

Fig. 2 shows the system overview of OpinionSeer. It contains three major components: *an opinion mining component, a subjective logic component, and an opinion visualization component*. The input of the system is a set of online customer reviews from TripAdvisor.com. The opinion mining component extracts customer opinions from unstructured review comments. It accounts for the ambiguity of human language when analyzing the sentiments of the customer reviews. Thus, in addition to general positive and negative values, the extracted opinions also explicitly contain the uncertainty values to indicate the amount of ambiguous information. Subjective logic is then used to help users organize and handle the extracted subjective opinions with different degrees of uncertainty. OpinionSeer further provides analysts with a tailored opinion visualization built on scatterplot and radial visualization to enable an integrated view of the interactive visual analysis of the complex opinion data. The extracted uncertainty information could be faithfully revealed in the visualization.

5 MINING OPINION FROM ONLINE HOTEL REVIEWS

In this section, we present a feature-based opinion mining approach to extract customer opinions for visual analysis. Subjective logic is further introduced to organize and handle the extracted opinions.

5.1 Feature based Opinion Mining

The collected customer reviews contain customer ratings about the hotels. Although this information is useful for customer opinion analysis, it cannot tell why such ratings are given. The free-text comments of reviews, on the other hand, are more informative (e.g., reasons for the opinions) than the ratings, providing concrete and descriptive information about customer opinions. Nevertheless, analyzing free-text comments manually is time-consuming and tedious. This motivated us

to use an opinion-mining technique to extract customer opinions from the free-text comments automatically.

To analyze customer opinions from different aspects, hotel managers and hospitality researchers usually need to classify customer reviews into different categories (or features) such as *service, space, and cleanliness* [1, 2, 19]. Thus, we use a feature-based opinion mining method [12, 23] to extract opinions from the customer reviews. It works as follows. First, the document to be analyzed is pre-processed and segmented into a collection of sentences from which opinion information is extracted. Second, the opinion information, including the object features and the related opinion scores, is inferred from each sentence. In this step, we define a sentiment keyword dictionary with "positive" and "negative" adjective words commonly found in the hotel customer reviews. We focus on five major hotel features (i.e. *room, location, cleanliness, service and hotel*) in the opinion mining process.

In practice, the customers often use particular words (we call entities) to describe these features. To facilitate feature detection, we define and utilize a feature-entity mapping scheme which maps a set of words (entities) to a given feature. Then, for each sentence, the opinion scores (positive and negative) for the detected feature(s) are measured by counting the number of the sentiment keywords found in the sentence. Please notice that negative expressions in customer reviews are handled specially. For example, a customer may say "*The location was not bad*", from which the customer actually expresses a positive opinion rather than a negative opinion. In this case, we use the opposite sentiment orientation of the sentiment keyword for estimating the opinion score. Finally, the opinion information of the attributes is aggregated to obtain the overall opinion about the hotel.

5.2 Uncertainty Modeling

We introduce a new concept, uncertainty, to augment the results of opinion mining. There is much evidence suggesting the existence of uncertainty in the opinion mining results of hotel reviews. First, it is common that a user may express both positive and negative sentiments on a feature of the hotel. Taking the feature, *room*, as an example, one user may comment: "*The room sure is tiny, yet very clean and comfy*". In previous studies, the positive and negative sentiment information of this example is simply aggregated to obtain the final opinion of the feature, which results in a positive sentiment value. However, this loses the negative opinion information. Positive and negative sentiment indicates the customer's conflict and uncertainty about their opinions. The smaller the difference between the two opinion scores, the more uncertainty the sentence possesses. Second, the detection of the subject of opinion words could not be accurate, which may bring uncertainty into the opinion mining results. Usually, longer sentences likely contain higher degrees of uncertainty.

We model the uncertainty with Gaussian distribution [4]. The overall uncertainty is defined as $u = \alpha \cdot N_{+/-}(\mu_1, \sigma_1) + \beta \cdot (1 - N_{||}(\mu_2, \sigma_2))$ where $N_{+/-}$ indicates uncertainty from the difference between the positive and negative scores, $N_{||}$ denotes the certainty from the sentence length, and $\alpha = \beta = 0.5$. The uncertainty for $N_{+/-}$ and $N_{||}$ is

$$N(x, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (1)$$

where $\mu = \frac{\sum_i x_i}{n}$ and $\sigma = \sqrt{\frac{\sum_i (x_i - \mu)^2}{n-1}}$. Furthermore, in this formula, for $N_{+/-}$, x is defined as $|s_+ - s_-|$, where s_+ and s_- indicate the positive and negative opinion scores respectively; Meanwhile, for $N_{||}$, x is defined as the length of the sentence.

5.3 Opinion Combination Based on Subjective Logic

Every extracted opinion contains positive, negative, and uncertainty scores for each feature. When conveying opinions to a user, we usually need to combine multiple selected opinions for multi-scale visual data exploration. For example, the user is often interested in knowing an overall opinion of selected features. However, because of the uncertainty information, general opinion aggregation approaches do not work. To address this issue, we borrow the concepts and framework

²<http://www.qsrinternational.com/>

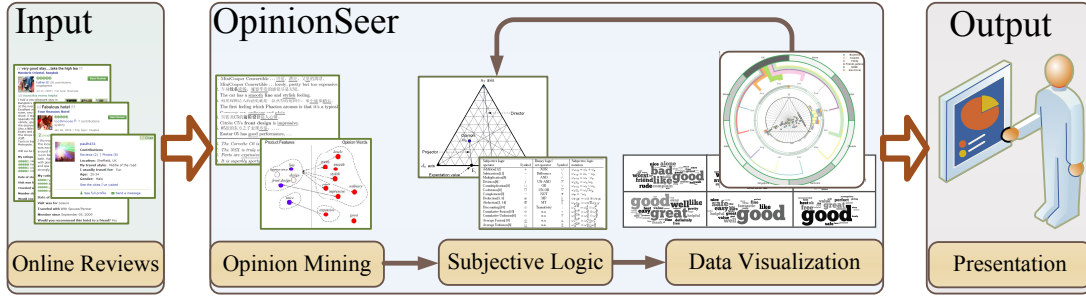


Fig. 2. System overview. The system is built upon opinion mining, subjective logic, and data visualization techniques.

from subjective logic [15] for our multi-scale opinion combination. In subjective logic, the opinion mining results for each feature, such as *room*, are represented by an opinion vector $\langle b, d, u, a \rangle$, where b and d indicate the positive and negative opinion scores, respectively; u denotes the uncertainty; a denotes the base rate which is the a priori probability in the absence of evidence.

A number of operators [14] have been defined in subjective logic. Some operators are generalizations of the binary logic and probability calculus operators, whereas the others are unique to subjective logic because they depend on the belief ownership. In our system, we mainly leverage the AND and FUSION operators. The AND operator corresponds to the binary logic AND, while the FUSION operator combines separate observers's opinions about the same aspect of discernment. The AND operator takes the opinions from distinct aspects of discernment as input and produces an overall opinion as a result. We can view the features (room, location, service, cleanliness, etc.) as the aspects in opinion mining of hotel reviews. Thus, we use the AND operator to combine the opinions of a customer on multiple features (at the feature level). The FUSION operator is used to combine the evidences from different sources, i.e., the opinions from different customers. Hence, we employ the FUSION operator to combine the opinions of multiple customers on the same feature.

Let (b_x, d_x, u_x) and (b_y, d_y, u_y) be two opinion vectors for feature x and y with a_x and a_y as the base rates, respectively, and the combined opinion on x and y can be determined by the AND operator as follows:

$$\omega_{x \wedge y} = \begin{cases} b_{x \wedge y} &= b_x b_y + \frac{(1-a_x)a_y b_x u_y + a_x(1-a_y)u_x b_y}{1-a_x a_y} \\ d_{x \wedge y} &= d_x + d_y - d_x d_y \\ u_{x \wedge y} &= u_x u_y + \frac{(1-a_y)b_x u_y + (1-a_x)u_x b_y}{1-a_x a_y} \\ a_{x \wedge y} &= a_x a_y \end{cases} \quad (2)$$

For the positive score b , only if two related opinions are positive, the resulting score will be positive; while for the negative score d , if any related opinion is negative, the resulting score will be negative. Thus the definitions of b and d are different. With the AND operator, we could combine opinions on different hotel features to estimate the overall sentiment orientation of a free-text customer comment.

Let (b_x^A, d_x^A, u_x^A) and (b_x^B, d_x^B, u_x^B) be two opinion vectors held by two customers, A and B , for the same feature x with a_x^A and a_x^B as the base rates, $\omega_x^{A \diamond B}$ be their cumulative fusion. Additionally, we define $K = u_x^A + u_x^B - u_x^A u_x^B$, when $u_x^A, u_x^B \rightarrow 0$, the relative dogmatism between ω_x^A and ω_x^B is defined by $\gamma = u_x^A / u_x^B$. The FUSION operator is defined as when $K \neq 0$,

$$\omega_x^{A \diamond B} = \begin{cases} b_x^{A \diamond B} &= (b_x^A u_x^B + b_x^B u_x^A) / K \\ d_x^{A \diamond B} &= (d_x^A u_x^B + d_x^B u_x^A) / K \\ u_x^{A \diamond B} &= (u_x^A u_x^B) / K \\ a_x^{A \diamond B} &= \frac{a_x^A u_x^B + a_x^B u_x^A - (a_x^A + a_x^B) u_x^A u_x^B}{u_x^A + u_x^B - 2u_x^A u_x^B} \end{cases} \quad (3)$$

when $K = 0$

$$\omega_x^{A \diamond B} = \begin{cases} b_x^{A \diamond B} &= \frac{(\gamma b_x^A + b_x^B)}{\gamma + 1} \\ d_x^{A \diamond B} &= \frac{(\gamma d_x^A + d_x^B)}{\gamma + 1} \\ u_x^{A \diamond B} &= 0 \\ a_x^{A \diamond B} &= \frac{(\gamma a_x^A + a_x^B)}{\gamma + 1} \end{cases} \quad (4)$$

Given multiple overall opinions from different customers on a hotel, acquired by the AND operator, we can apply the FUSION operator to determine an average opinion of the customers on the hotel. In our system, the base rates a in all the opinion vectors are set to a default value, namely, 0.5, according to [15].

6 OPINION VISUALIZATION

To assist users in visually analyzing the complex opinion data effectively, we developed an opinion visualization system that includes the opinion wheel, the tag cloud spreadsheet, and a set of tailored user interactions. Our design principles include effectiveness, intuitiveness, and attraction. Simplicity or intuitiveness is strongly required because our end users do not have much background on information technology, while the visualization should be aesthetically appealing because the users want to present their findings directly to a wider audience. By working closely with our target users, we developed a visualization system that could convey the results of the opinion mining, from simple to complex, while keeping its intuitiveness.

The system has two major views, an opinion wheel (Fig. 1) and tag clouds (Fig. 5). The opinion wheel seamlessly integrates a scatterplot (opinion triangle) with a radial visualization (opinion ring). The opinion triangle is primarily used for visualizing the extracted opinions, each of which is an opinion vector (b, d, u) with three elements: negative, positive, and uncertainty values. The three vertices of the opinion triangle represent the most negative, positive, and uncertain opinions, respectively. Each customer opinion is plotted in the opinion triangle according to the distance from the three triangle vertices. For example in Fig. 7(a), an opinion shown in the lower left of the triangle means a negative opinion, in the lower right means a positive opinion, and in the top part means an opinion with high uncertainty. The opinion rings surrounding the triangle facilitate the visual exploration of correlations between the customer opinions and other data dimensions. The opinions in the triangle are projected onto the opinion rings to create circular histograms of different data dimensions. Furthermore, to help user examine the real reason of a certain opinion as well as to compare customer reviews, a diagram of tag clouds is synchronized with the opinion wheel. In this section, we will discuss our opinion visualization design and share our experience in collaboration with hospitality researchers for developing the opinion visualization system.

6.1 Opinion Wheel: Integrated Visualization of Customer Opinion Data

The major visualization of OpinionSeer is an opinion wheel, which is a tight integration of a scatterplot and a radial visualization. The opinion posts or features are represented by a scatterplot inside an opinion triangle. In the scatterplot, each point encodes an opinion post or feature.




Fig. 3. (a) The sum of distances from the point P to all three sides is always equal to the height of the equilateral triangle; (b) The opinions are combined by the FUSION operator at the hotel level. Please note that $\omega_1 = a \diamond b$, $\omega_2 = \omega_1 \diamond c$, $\omega_3 = \omega_2 \diamond d$, and $\omega_4 = \omega_3 \diamond e$.

The radial visualization is the bounding wheel of the opinion triangle. We adopt it to illustrate visually the correlations among multiple data dimensions (e.g., age, gender).

6.1.1 Opinion Triangle

Customer opinions are the center of customer feedback data, and play a key role in visual opinion analysis. In hospitality research, the general customer feedback analysis usually starts from customer opinions. Thus, the first step of our design is to determine a reasonable visual representation for the opinions. As described in Section 5.3, each extracted opinion is represented as an opinion vector (b, d, u) , where $b + d + u = 1$. Proper visual encoding of the opinion vector is difficult using traditional information visualization techniques such as parallel coordinates because the important characteristic, $b + d + u = 1$, of the opinion vector cannot be clearly revealed. On the other hand, in an equilateral triangle, the sum of distances from any point in the interior of an equilateral triangle to all three sides is always equal to the height of the triangle. Thus, this triangle property can be used to visually encode the characteristic of the opinion vector (i.e., $b + d + u = 1$).

An opinion vector $\omega_x = (b_x, d_x, u_x)$ could be mapped to a point inside an equilateral triangle $\triangle ABC$ (Fig. 3) whose height is equal to 1. Vertices A , B , and C denote disbelief, uncertainty, and belief, respectively. To achieve this, we draw two lines IJ and DE which are parallel to BC and AC , respectively. Additionally, we make sure that the distance between IJ and BC is equal to d_x . Similarly, the distance between DE and AC is equal to u_x . The intersection point P of IJ and DE is the point that represents the opinion vector, ω_x , inside the triangle. The distances from P to the three sides BC , AB , and AC are d_x , b_x , and u_x , respectively. The sum of the distances is equal to the height of the triangle, that is, $b_x + d_x + u_x = 1$. With the visual encoding method, all opinion vectors could be intuitively shown inside a triangle-style scatterplot, which is also called an opinion triangle in subjective logic [14]. For example, a strong negative opinion could be represented by a point toward the left disbelief vertex of the opinion triangle. Similarly, an opinion with a high degree of uncertainty could be represented by a point toward the top uncertainty vertex of the opinion triangle.

The opinion triangle used together with the subjective logic operators can greatly facilitate visual opinion comparison of different groups of customers. After separately applying the FUSION operator to the opinions of every selected group, we could obtain several fused opinion points inside the triangle; each point represents a fused opinion. By comparing these opinion points inside the opinion triangle, we could readily identify the differences of the customer opinion groups. This capability could then solve Q1 and Q2 described in Section 3.3. Compared with other visual metaphors, the opinion triangle could present the uncertainty information naturally; it is also a scatterplot familiar to and used frequently by our target users. Thus, they can start with a familiar format.

6.1.2 Opinion Rings

Finding opinion patterns regarding categorical information is a fundamental task in hospitality research. In this section, we introduce our adapted visualization approach based on scatterplots, glyphs, and radial visualization layouts to facilitate this task.

Coordinated View versus Integrated View To find opinion patterns and correlations among different dimensions, the extracted opinions need to be analyzed in context, which requires simultaneous visualization of the multidimensional information. One straightforward solution is to provide users with multiple views coordinated with the opinion triangle. Each view focuses on one data dimension. Our initial prototype system includes multiple coordinated views: an opinion triangle view for extracted opinions, five bar chart views of related demographic information and temporal information, a parallel coordinates plot to reveal the relationship between temporal and geographic dimensions, and a map view for geographic information. After presenting and discussing the system to our target users, we did not adopt this approach as the users thought it was difficult for them to relate information scattered in multiple views to find interesting opinion patterns. To address the issue, we attempted to develop a comprehensive visual representation of the data capable of providing an integrated visualization of multidimensional data rather than multiple separate views. Although this would possibly introduce visual clutter when showing too much information simultaneously, we could keep the visual clutter at an acceptable level through proper design and user interactions.

Glyph-based Encoding We started our design from the opinion triangle, which is a triangle-style scatterplot. Each opinion point is associated with one opinion holder (i.e., the customer). Hence, we could simply utilize glyphs, geometric objects with different visual properties, to encode multidimensional categorical information of the opinion holders inside the triangle. Some visual properties of glyphs such as color, shape, and size are available if we require rapid pre-attentive processing [30]. After discussing with our target users, however, we found it was not necessary to show too much information simultaneously in the scatterplot for the following reasons. First, regarding the general analysis tasks (Q3, Q4, and Q5) listed in Section 3.3, users only need to examine the relationship between customer opinions and another categorical dimension one by one, therefore unused dimensions are considered unnecessary. Second, with respect to the tasks related to temporal and geographic dimensions (Q6 and Q7), users may need to analyze multiple dimensions (opinions, demographic, temporal, and geographic information) simultaneously to find temporal and spatial opinion patterns, but the temporal and geographic dimensions cannot be encoded easily by glyphs. While many different locations and time ranges exist, the number of categories that each glyph property could encode is limited [30]. For example, no more than eight colors should be adopted if we want to understand data values quickly. Therefore, inside the opinion triangle, only two pre-attentive visual properties (color and shape) are employed for the glyphs. Color is used to encode the categories of a categorical dimension (e.g., age range), while shape is utilized to represent the groups of the opinions (e.g., room, service, and price).

Categorical Ring Scatterplot with glyphs can show an overall information distribution of a certain dimension such as a distribution of age groups over opinions. However, in our application, a large number of customer opinions could be explored, which may introduce severe visual clutter. Consequently, it is difficult to find opinion relationships with respect to Q3 and Q5, not to mention the visual comparison regarding Q4. To alleviate the problem and improve the scatterplot readability, we incorporated a radial visualization layout into our opinion triangle. Radial visualization is an increasingly prevalent visual metaphor with a compact and aesthetically appealing layout in information visualization and visual analytics [7]. Compared with other existing radial visualization, our approach has two unique features: First, our radial layout supports the subjective logic and accounts for uncertainty. Second, we provide an integrated view of multiple important data dimensions specifically designed for opinion visualization. The basic idea of our approach is to project customer opinions in the

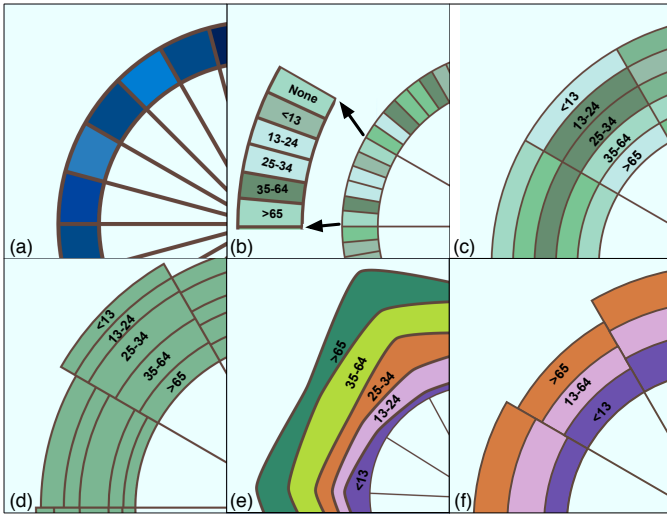


Fig. 4. (a) Color represents the weighted average of the ages of the customers inside the sector; (b)-(c) Color represents the number of customers in each age group inside every sector; (d) Size represents the number of the customers; (e) Stacked graph where the belt width encodes the number, and the color represents different age groups; (f) Our design in which size and color is used to encode the number and the age groups, respectively.

interior of the opinion triangle to its circumscribed ring (called categorical ring), and then visualize the categories of the dimension to be examined on the sectors of the ring.

To ensure effective visualization, we first designed five different layouts using pre-attentive visual properties including color and size to display the category information on the sectors of the categorical ring, as illustrated in Fig. 4(a) - (e). These radial layouts were then presented to our two target users for user evaluation.

Both users rejected the design in Fig. 4(a) because it was difficult for them to associate depth of color with weighted average. They complained that it lost more information than other layouts. The layouts shown in Figs. 4(d) and (e) were received well by a user. He pointed out that size is visually more intuitive to associate with numbers or volumes than color depth, hence the layouts shown in Figs. 4(a), (b), and (c) were not preferred. He also felt that having different colors to represent different categories make it easier to identify than having similar or the same color schemes such as in Figs. 4(b) and (c). Additionally, he suggested that grouping the information neatly into sectors, like in Figs. 4(c) and (d), should be much better than Fig. 4(e). Another user especially like Fig. 4(e), as it is less complicated and the quantity information is width-oriented. In addition, it is easy to identify what information to be communicated in one glimpse. All the others are less preferred by the user because they are all required to read additional chart/table in order to find out what is going on and to understand. To conclude, it is better that the layout use different sizes to indicate the number of customers in a particular category, together with different colors to represent various categories.

Based on the user feedback, we developed a new radial layout shown in Fig. 4(f) in which information of a particular dimension (e.g., age range) is projected to the circumscribed ring. Each sector is divided into multiple parts along the radius direction and each part corresponds to a specific category of customer (an age group in this example). The size of each part is determined by the number of customers that belong to the corresponding category. Different colors are used to differentiate different age groups. This layout could be viewed as circular stacked bar charts. With this design, users can identify how the information dimension examined could affect customer opinions (Q3). If we project customer's score ratings to the ring, we could also examine the relationship between the score ratings and the opinions extracted from the free-text comments (Q5). To enable a side-by-

side visual comparison of the distributions (Q4), we first represent the opinion points using different shapes inside the opinion triangle for different groups of customers. Each sector on the ring is now equally divided into multiple subsectors, and each subsector is associated with one group of the customers. This allows users to visually compare the data distributions readily around the opinion triangle (Fig. 9).

Temporal and Geographic Rings For Q6 and Q7, the temporal information (date of stay) and geographic information (customer location) should be presented to users for analysis. However, this information cannot be conveyed effectively by the categorical ring because they possess special features. The temporal and geographic dimensions usually contain more categories than others. In addition, the temporal information has unique multi-scale periodic patterns, and the geographic information has special directional patterns that cannot be revealed. Nevertheless, radial visualization is still well-suited for revealing both periodic and directional patterns [7]. Thus, we add a temporal ring and a geographic ring to the opinion wheel to visualize effectively the temporal and geographic information, respectively.

The temporal rings can be configured to different styles showing temporal information at different levels of detail based on user requirements, as illustrated in Fig. 1(a). The number of opinions expressed during a specific time range is encoded as the color in the sector associated with the related time range. Fig. 1(b) shows a geographic ring separated into a number of sectors; each sector corresponds to a location, such that the geographic direction of a location could be roughly revealed by the corresponding sector. The number of customer opinions from a location is encoded as a color in the sector associated with the related direction. The luminance (white-black) channel is used to encode the number in the sectors for both temporal and geographic rings because of its capacity to show data detail [26].

Although our design can address Q6, it is still difficult to find the relationships between temporal and geographic information (Q7). Inspired by Parallel Sets [18] which could effectively reveal relationships between category dimensions, we develop a technique to visually relate information between temporal and geographic dimensions. Fig. 1(b) shows the temporal ring and geographic ring simultaneously in the opinion wheel. The relationships could be revealed by connecting related categories using curved belts rather than parallelograms in Parallel Sets. Compared with Parallel Sets which show many-to-many relationships, our technique only shows a one-to-many relationship. Details are shown on demand using connections for only the selected sector on the temporal or geographic rings. This was motivated by explicit feedback from our target users on reducing information overload and visual clutter.

6.1.3 Multi-scale Exploration

The opinion wheel allows users to analyze customer opinions at different levels of detail. For instance, users could analyze customer opinions at the feature level when the opinions on a specific hotel feature or a set of hotel features are analyzed. With this visualization, users could visually compare the opinion distributions of two hotel attributes inside the opinion triangle. The "AND" operator is exploited to combine customer opinions on different hotel attributes to facilitate the exploration at a higher level. If all feature opinions of each customer are combined using the "AND" operator, the overall customer opinions on hotels could be viewed and analyzed by users. Another operator "FUSION" could combine user opinions of different customers. Thus, users can fuse a group of opinions on a particular hotel feature of different customers, or fuse a group of combined opinions (obtained by "AND" at the feature level) of different customers. This allows for visual analysis of customer opinions at multi-scale customer levels.

6.2 Tag Clouds: Detailed Visualization of Customer Opinion Data

To provide rich context that could help the analyst comprehend the major opinion content, tag cloud visualization developed based on Wordle [28] is synchronized with the opinion wheel. Time-varying tag cloud [5] can also be used to track opinion changes over time. For example, when a user selects a subset of opinion points from the scatterplot,

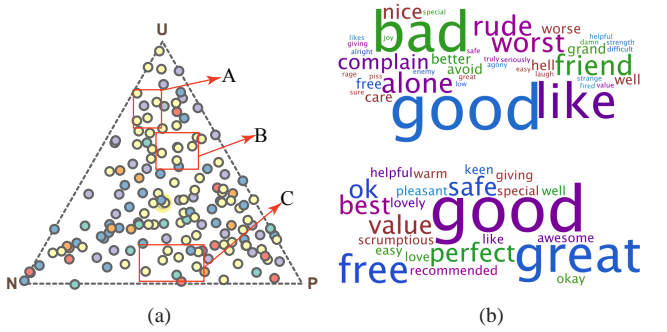


Fig. 7. (a) An opinion triangle where three regions A, B, and C are selected; (b) Top and bottom: two tag clouds of the opinion words associated with Region A and B in (a), respectively.

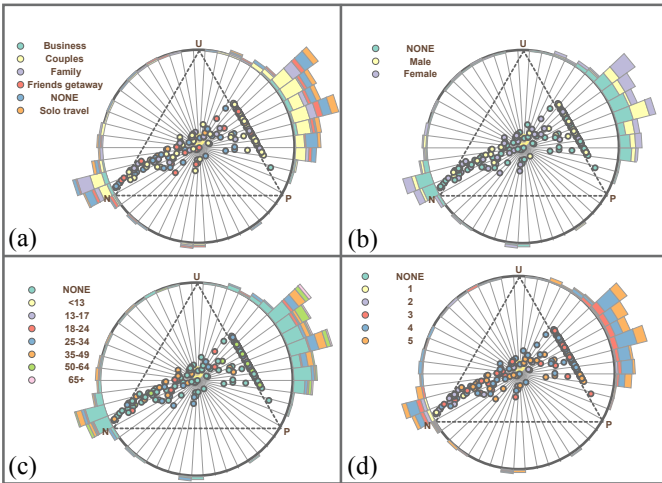


Fig. 8. OpinionSeer results showing how customer opinions are correlated with trip type, gender, age range, and ratings.

inaccuracy may likely be ignored, thus leading to unreliable results. For instance, the opinions in Part A and B would be treated as positive opinions without our method, which may result in a biased conclusion.

Subjective Logic The second experiment was conducted to prove the effectiveness of subjective logic in combining customer opinions with uncertainty. Figure 3(b) presents an example which combines customer opinions at the hotel level with the FUSION operator (\diamond). The opinion a , b , c , and d now represent customer opinions about a specific feature of different hotels. The intermediate results denoted by $\omega_1 = a \diamond b$, $\omega_2 = \omega_1 \wedge c$, $\omega_3 = \omega_2 \diamond d$, and $\omega_4 = \omega_3 \diamond e$ are also shown in the figure. We can observe from the figure that the uncertainty is well considered by the operator. For example, although a and d are quite positive, their high uncertainty values limit their influences on the overall opinion (ω_4), thus resulting in a somewhat negative overall opinion. Similar results could also be obtained by the “AND” operator which also takes uncertainty into account when merging multiple feature-level opinions.

7.2 Case Studies

To show the system usability, we conducted an informal task test. A hospitality researcher was invited to use our system to explore the customer reviews of the top five popular hotels of each hotel class.

For Q1, the participant selected a group of customer opinions inside the triangle by the brushing operation, and then chose the FUSION operator to obtain an overall opinion for this group. The average opinion was estimated by applying the FUSION operator to all customer opinions inside the triangle. By comparing the average opinion with the

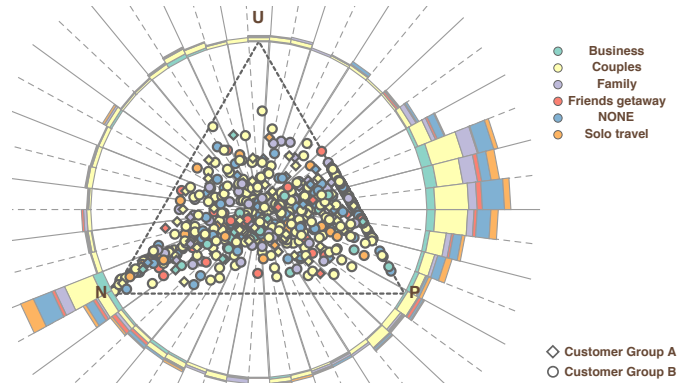


Fig. 9. Visual comparison of the trip types related to two groups of customers. Every sector of the categorical ring is separated into two sub parts (by dashed lines) for showing the distributions of the trip types of the two customer groups.

overall opinion of the selected opinion group, he could easily tell the opinion deviation in Q1. Similarly, the user evaluated the differences among multiple opinion groups (Q2) by visually comparing the overall fused opinions. He felt that the FUSION operator was very helpful for comparing groups of customer opinions, and the fused results were roughly in accordance with his perception. None of other techniques such as standard bar charts can fuse multiple opinions for rapid visual comparison in such an intuitive manner as our method.

Task Q3 is primarily for identifying the relationships between customer opinions and demographic characteristics such as age and gender. The participant used the opinion triangle and the categorical ring for this task. As uncertainty exists and there is no clear boundary between negative and positive opinions, he brushed the opinion triangle interactively to choose appropriate customer opinions for investigation. He felt that the opinion triangle is much more expressive than other conventional visualization approaches such as scatterplots and bar charts. Fig. 8(a)-(c) show his results revealing the relationships between the selected customer opinions and trip type, gender, and age range, respectively. From Figs. 8(b) and (c), we can clearly observe that demographic characteristics, such as age and gender, do influence customer opinions. In general, female customers complained more than male customers in our results (Fig. 8(b)); meanwhile, older customers had fewer complaints than younger customers (Fig. 8(c)). These results are in consistent with previous studies [16]. An unexpected pattern was also discovered by the participant. The trip type also has strong influence on customer opinions. Customers who traveled with family members tended to express negative opinions, while customers of traveling independently had fewer complaints than others, which has never been reported before.

The solution of Q5 is similar to Q3. It was formulated in our previous discussion with the participant. He was quite interested in knowing whether or not the customer ratings are in consistent with the underlying reviews. Figure 8(d) shows the visualization result. Abnormally, a number of reviews with quite negative opinions received high ratings. The participant argued that customers may have different criteria for giving ratings for a hotel. It would also be possible for other visualization means such as bar charts to make similar observations regarding Q3 and Q5. However, as they often could not convey the uncertainty information (from the ambiguity of language or inaccuracy of the sentiment analysis) of the data to users as effectively as our opinion triangle does, the observations would likely be questionable.

In Task Q4, the user was asked to find the differences of the trip types of two customer groups selected from high-class and low-class hotels, respectively. Although common methods such as bar charts could be used for the comparison, the hidden uncertainty information may easily lead to a wrong comparison result, especially when the majority of the extracted opinions are relatively uncertain. Fig. 9 shows

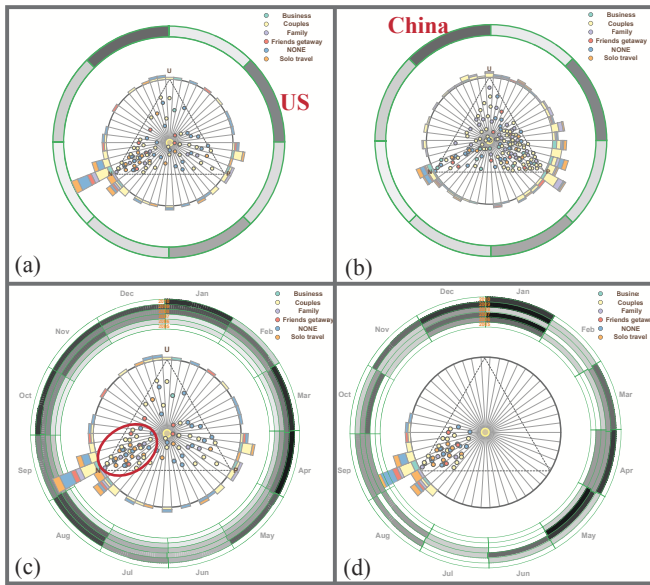


Fig. 10. (a) and (b): OpinionSeer showing the opinions of the customers from US and China; (c) and (d): OpinionSeer showing temporal patterns of customer opinions.

the opinion triangle where the opinions of different customer groups are encoded by different shapes. Every sector of the categorical ring is separated into two sub parts (by dashed lines) for showing the distributions of the trip types of the two customer groups. Additionally, a tag cloud diagram was utilized for providing analysis context to the comparison (Fig. 5). The participant ignored the highly uncertain opinions and compared only the opinions with relatively low uncertainty values to ensure a fair and reliable comparison. He did not find significant difference between the two customer groups in terms of the trip types.

Task Q6 was relatively easy for the participant. When he selected a sector on the geographic ring, the related opinions were highlighted. By examining all sectors on the ring, the participant quickly inspected whether there is any localization pattern with respect to the users opinions. During the test, he immediately found an interesting geographic opinion pattern. Figs. 10(a) and (b) show the results where the sectors of US and China on the geographic ring were selected, respectively. Mainland Chinese generally were found having far less complaints than other customers. This pattern was also reported by Au et al. [2].

Task Q7 was to find any temporal opinion patterns. The participant continued to investigate whether or not the complaints (i.e., negative opinions) from US customers have any temporal patterns. Figure 10(c) shows the opinion data for the US customers selected from Fig. 10(a). He filtered out irrelevant opinions by brushing inside the opinion wheel, and the temporal ring was updated immediately to show how the opinions distributed on the ring (Fig. 10(d)). He identified a possible temporal opinion pattern, namely, there seems to be more complaints in April, May, and December. In these cases, the participant could quickly identify the patterns from our integrated view. He indicated that it would be difficult for him to use other methods with coordinated views to find the patterns.

7.3 User Feedback

OpinionSeer was well received by our end users. One user especially liked the simplicity of OpinionSeer because it was built upon the scatterplot that he knew well. Another user commented, “One of the strengths of OpinionSeer is its ability to analyze and identify the hidden pattern in the raw text data, and provide a user-friendly visual presentation to end users”. He also pointed out that due to insufficient IT training, many hotel managers are reluctant or even resistant to accept new technologies, because they fear that these technologies might affect their ability to provide personalized services to hotel

guests. OpinionSeer can therefore fill up this gap by helping managers to quickly identify useful and meaningful relationships among the vast amount of textual data uploaded by hotel customers on the Internet. This will facilitate the formulation of more effective decisions that can help in providing timely and appropriate response to customers. In addition, the users agreed on the necessity and usefulness of modeling uncertainty for data analysis. One participant pointed out that there is usually no clear boundaries among positive, negative, and uncertain opinions. He appreciated the opinion triangle because it can accurately present underlying information in an intuitive manner. A user stated that, compared with the method of encoding the opinions on a line segment with only positive and negative values, presenting the opinions inside the triangle plane provided more space for opinion selection.

7.4 Discussion

As discussed in the experiments and case studies, the extracted uncertainty information and its visual encoding play important roles in the analysis. The uncertainty information improves the accuracy and correctness of the analysis. Our visual encoding of the uncertainty using the opinion triangle can intuitively convey the uncertainty information to users and enhance the understanding of the extracted customer opinions. The subjective logic operators are also useful and important for opinion analysis. With the operators, OpinionSeer enables users to explore the customer opinions interactively at multiple scales.

Our collaborators provided insightful thoughts on how they view or define uncertainty. One collaborator indicated that different customers with varied cultural backgrounds may have different reaction and judgments on the service or product performance, so uncertainty may exist in this context. For example, “not bad” may imply “quite good” for the Chinese due to their modest characteristics; however, this may not be the case for Westerners. Thus, cultural background may be another moderating variable or uncertainty. The other collaborator suggested that the co-existing positive and negative words often indicate the uncertainty of customer opinion. In our current system, we only considered the second case, as well as the uncertainty introduced by opinion mining. We will improve the uncertainty modeling technique by customers’ different cultural backgrounds in future.

Similar opinions near the opinion triangle center (the default projection center) can be assigned to completely different sector histograms of the categorical ring. This may have negative impact on the analysis, especially when the opinions are dense near the center. Nevertheless, users can avoid the inappropriate aggregation through grouping similar opinions, moving the projection center, and/or applying area-preserving mapping.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we have presented OpinionSeer for interactive visual opinion analysis. We seriously consider the uncertainty information in opinion extraction, combination, and visualization. In opinion extraction, we model the uncertainty from the language ambiguity and opinion mining; in opinion combination, we take the uncertainty into account; in opinion visualization, we create an intuitive visual representation of the uncertainty information. Aside from improving the analysis reliability, this increases the flexibility of the data analysis, since for different applications users can intuitively select customer opinions with different degrees of uncertainty in the opinion triangle for investigation. Our techniques are not limited to the hotel customer feedback data. They can also be useful for visual analysis of customer opinions on other products or services. In the future, we plan to continue our work with domain experts and deploy our system on the Web to make it available to the public. We also want to improve the tag cloud diagram by providing more visual support for comparison, e.g., using the same colors for same terms or roughly the same position.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Rob Law in the School of Hotel & Tourism Management at the Hong Kong Polytechnic University for his help with the system design. This work was supported in part by grant HK RGC GRF 619309 and an IBM Faculty Award.

REFERENCES

- [1] N. Au, D. Buhalis, and R. Law. Complaints on the online environment the case of hong kong hotels. In W. Höpken, U. Gretzel, and R. Law, editors, *Information and Communication Technologies in Tourism 2009*, pages 73–85. Springer-Verlag Wien, 2009.
- [2] N. Au, R. Law, and D. Buhalis. The impact of culture on ecomplaints: Evidence from the chinese consumers in hospitality organization. In U. Gretzel, R. Law, and M. Fuchs, editors, *Information and Communication Technologies in Tourism 2010*, pages 285–296. Springer-Verlag Wien, 2010.
- [3] C. Chen, F. Ibekwe-SanJuan, E. SanJuan, and C. Weaver. Visual analysis of conflicting opinions. In *IEEE Symposium On Visual Analytics Science And Technology*, pages 35 – 42, 2006.
- [4] C. D. Correa, Y.-H. Chan, and K.-L. Ma. A framework for uncertainty aware visual analytics. In *IEEE Symposium on Visual Analytics Science and Technology*, pages 51–58, 2009.
- [5] W. Cui, Y. Wu, S. Liu, F. Wei, M. X. Zhou, and H. Qu. Context preserving dynamic word cloud visualization. In *IEEE Pacific Visualization Symposium*, pages 121–128, 2010.
- [6] G. Draper and R. Riesenfeld. Who votes for what? a visual query language for opinion data. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1197–1204, 2008.
- [7] G. M. Draper, Y. Livnat, and R. F. Riesenfeld. A survey of radial methods for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(5):759–776, 2009.
- [8] M. Gamon, A. Aue, S. Corston-Oliver, and E. Ringger. Pulse: Mining customer opinions from free text. In *International Symposium on Intelligent Data Analysis*, pages 121–132, 2005.
- [9] M. Gamon, S. Basu, D. Belenko, D. Fisher, M. Hurst, and A. C. König. BLEWS: Using blogs to provide context for news articles. In *AAAI Conference on Weblogs and Social Media*, pages 60–67, 2008.
- [10] M. L. Gregory, N. Chinchor, P. Whitney, R. Carter, E. Hetzler, and A. Turner. User-directed sentiment analysis: Visualizing the affective content of documents. In *Workshop on Sentiment and Subjectivity in Text*, pages 23–30, 2006.
- [11] D. Houser and J. Wooders. Reputation in auctions: Theory and evidence from ebay. *Journal of Economics & Management Strategy*, 15(2):353–369, 2006.
- [12] M. Hu and B. Liu. Mining and summarizing customer reviews. In *ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, 2004.
- [13] M. Hu and B. Liu. Mining opinion features in customer reviews. In *AAAI’04: Proceedings of the 19th national conference on Artificial intelligence*, pages 755–760, 2004.
- [14] A. Jøsang. The consensus operator for combining beliefs. *Artificial Intelligence*, 141(1):157–170, 2002.
- [15] A. Jøsang. *Subjective Logic*. draft, available at : http://persons.unik.no/josang/papers/subjective_logic.pdf, 2009.
- [16] K. A. Keng, D. Richmond, and S. Hans. Determinants of consumer complaint behavior: A study of singapore consumers. *Journal of International Consumer Marketing*, 8(2):59–76, 1995.
- [17] S.-M. Kim and E. Hovy. Determining the sentiment of opinions. In *Proceedings of international conference on Computational Linguistics*, pages 1367–1373, 2004.
- [18] R. Kosara, F. Bendix, and H. Hauser. Parallel Sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):558–568, 2006.
- [19] C. C. Lee and C. Hu. Analyzing hotel customers E-complaints from an internet complaint forum. *Journal of Travel & Tourism Marketing*, 17(2 & 3):167–181, 2005.
- [20] B. Liu, M. Hu, and J. Cheng. Opinion observer: analyzing and comparing opinions on the web. In *International Conference on World Wide Web*, pages 342–351, 2005.
- [21] S. Morinaga, K. Yamanishi, K. Tateishi, and T. Fukushima. Mining product reputations on the web. In *ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 341–349, 2002.
- [22] D. Oelke, M. Hao, C. Rohrdantz, D. A. Keim, U. Dayal, L.-E. Haug, and H. Janetzko. Visual opinion analysis of customer feedback data. In *IEEE Symposium On Visual Analytics Science And Technology*, pages 187–194, 2009.
- [23] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, 2008.
- [24] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In *Conference on Empirical methods in natural language processing*, pages 79–86, 2002.
- [25] A.-M. Popescu and O. Etzioni. Extracting product features and opinions from reviews. In *Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 339–346, 2005.
- [26] B. E. Rogowitz, L. A. Treinish, and S. Bryson. How not to lie with visualization. *Computers in Physics*, 10(3):268–273, 1996.
- [27] I. E. Vermeulen and D. Seegers. Tried and tested: The impact of on-line hotel reviews on consumer consideration. *Tourism Management*, 30(1):123–127, 2008.
- [28] F. B. Viégas, M. Wattenberg, and J. Feinberg. Participatory visualization with wordle. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1137–1144, 2009.
- [29] F. Wanner, C. Rohrdantz, F. Mansmann, and D. A. Keim. Visual sentiment analysis of RSS news feeds featuring the us presidential election in 2008. In *Workshop on Visual Interfaces to the Social and the Semantic Web*, 2009.
- [30] K. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, 2nd edition, 2004.
- [31] J. Yang, M. O. Ward, and E. A. Rundensteiner. Interring: An interactive tool for visually navigating and manipulating hierarchical structures. In *IEEE Symposium on Information Visualization*, pages 77–84, 2002.