Product Aspect Ranking and Its Applications

Zheng-Jun Zha, *Member, IEEE*, Jianxing Yu, Jinhui Tang, *Member, IEEE*, Meng Wang, *Member, IEEE*, and Tat-Seng Chua

Abstract—Numerous consumer reviews of products are now available on the Internet. Consumer reviews contain rich and valuable knowledge for both firms and users. However, the reviews are often disorganized, leading to difficulties in information navigation and knowledge acquisition. This article proposes a product aspect ranking framework, which automatically identifies the important aspects of products from online consumer reviews, aiming at improving the usability of the numerous reviews. The important product aspects are identified based on two observations: 1) the important aspects are usually commented on by a large number of consumers and 2) consumer opinions on the important aspects greatly influence their overall opinions on the product. In particular, given the consumer reviews of a product, we first identify product aspects by a shallow dependency parser and determine consumer opinions on these aspects via a sentiment classifier. We then develop a probabilistic aspect ranking algorithm to infer the importance of aspects by simultaneously considering aspect frequency and the influence of consumer opinions given to each aspect over their overall opinions. The experimental results on a review corpus of 21 popular products in eight domains demonstrate the effectiveness of the proposed approach. Moreover, we apply product aspect ranking to two real-world applications, i.e., document-level sentiment classification and extractive review summarization, and achieve significant performance improvements, which demonstrate the capacity of product aspect ranking in facilitating real-world applications.

Index Terms—Product aspects, aspect ranking, aspect identification, sentiment classification, consumer review, extractive review summarization

1 INTRODUCTION

RECENT years have witnessed the rapidly expanding e-commerce. A recent study from ComScore reports that online retail spending reached \$37.5 billion in Q2 2011 U.S. [5]. Millions of products from various merchants have been offered online. For example, Bing Shopping¹ has indexed more than five million products. Amazon.com archives a total of more than 36 million products. Shopper.com records more than five million products from over 3,000 merchants. Most retail Websites encourage consumers to write reviews to express their opinions on various *aspects* of the products. Here, an *aspect*, also called *feature* in literatures, refers to a component or an attribute of a certain product. A sample review "*The battery life of Nokia N95 is amazing.*" reveals positive opinion on the aspect "*battery life*" of product *Nokia N95*. Besides the retail

1. www.bing.com/shopping.

- Z.-J. Zha is with the Institute of Intelligent Machines, Chinese Academy of Sciences, Hefei 230031, China. E-mail: zhazj@iim.ac.cn.
- J. Yu, and T.-S. Chua are with the School of Computing, National University of Singapore, Singapore 117417, Singapore.
- J. Tang is with the School of Computer Science, Nanjing University of Science and Technology, Nanjing 210094, China.
- M. Wang is with the School of Computer and Information, Hefei University of Technology, Hefei 230009, China.

Manuscript received 19 Nov. 2012; revised 28 Apr. 2013; accepted 23 May 2013. Date of publication 4 Aug. 2013; date of current version 7 May 2014.

Recommended for acceptance by J. Pei.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier 10.1109/TKDE.2013.136 Websites, many forum Websites also provide a platform for consumers to post reviews on millions of products. For example, CNet.com involves more than seven million product reviews; whereas Pricegrabber.com contains millions of reviews on more than 32 million products in 20 distinct categories over 11,000 merchants. Such numerous consumer reviews contain rich and valuable knowledge and have become an important resource for both consumers and firms [9]. Consumers commonly seek quality information from online reviews prior to purchasing a product, while many firms use online reviews as important feedbacks in their product development, marketing, and consumer relationship management.

Generally, a product may have hundreds of aspects. For example, iPhone 3GS has more than three hundred aspects (see Fig. 1), such as "usability," "design," "application," "3G network." We argue that some aspects are more important than the others, and have greater impact on the eventual consumers' decision making as well as firms' product development strategies. For example, some aspects of iPhone 3GS, e.g., "usability" and "battery," are concerned by most consumers, and are more important than the others such as "usb" and "button." For a camera product, the aspects such as "lenses" and "picture quality" would greatly influence consumer opinions on the camera, and they are more important than the aspects such as "a/v cable" and "wrist strap." Hence, identifying important product aspects will improve the usability of numerous reviews and is beneficial to both consumers and firms. Consumers can conveniently make wise purchasing decision by paying

^{1041-4347 © 2013} IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.



Fig. 1. Numerous aspects of the product iPhone 3GS.



Fig. 2. Flowchart of the proposed product aspect ranking framework.

more attentions to the important aspects, while firms can focus on improving the quality of these aspects and thus enhance product reputation effectively. However, it is impractical for people to manually identify the important aspects of products from numerous reviews. Therefore, an approach to automatically identify the important aspects is highly demanded.

Motivated by the above observations, we in this paper propose a product aspect ranking framework to automatically identify the important aspects of products from online consumer reviews. Our assumption is that the important aspects of a product possess the following characteristics: (a) they are frequently commented in consumer reviews; and (b) consumers' opinions on these aspects greatly influence their overall opinions on the product. A straightforward frequency-based solution is to regard the aspects that are frequently commented in consumer reviews as important. However, consumers' opinions on the frequent aspects may not influence their overall opinions on the product, and would not influence their purchasing decisions. For example, most consumers frequently criticize the bad "signal connection" of iPhone 4, but they may still give high overall ratings to iPhone 4. On the contrast, some aspects such as "design" and "speed," may not be frequently commented, but usually are more important than "signal connection." Therefore, the frequency-based solution is not able to identify the truly important aspects. On the other hand, a basic method to exploit the influence of consumers' opinions on specific aspects over their overall ratings on the product is to count the cases where their opinions on specific aspects and their overall ratings are consistent, and then ranks the aspects according to the number of the consistent cases. This method simply assumes that an overall rating was derived from the specific opinions on different aspects individually, and cannot precisely characterize the correlation between the specific opinions and the overall rating. Hence, we go beyond these methods and propose an effective aspect ranking approach to infer the importance of product aspects. As shown in Fig. 2, given the consumer reviews of a particular product, we first identify aspects in the reviews by a shallow dependency parser [37] and then analyze consumer opinions on these aspects via a sentiment classifier. We then develop a probabilistic aspect ranking algorithm, which effectively exploits the aspect frequency as well as the influence of consumers' opinions given to each aspect over their overall opinions on the product in a unified probabilistic model. In particular, we assume the overall opinion in a review is generated based on a weighted aggregation of the opinions on specific aspects, where the weights essentially measure the degree of importance of these aspects. A probabilistic regression algorithm is developed to infer the importance weights by incorporating aspect frequency and the associations between the overall opinion and the opinions on specific aspects. In order to evaluate the proposed product aspect ranking framework, we collect a large collection of product reviews consisting of 94,560 consumer reviews on 21 products in eight domains. These reviews are crawled from multiple prevalent forum Websites, such as CNet.com, Viewpoints.com, *Reevoo.com* and *Pricegrabber.com* etc. This corpus is available by request for future research on aspect ranking and related topics. More details of the data are discussed in Section 3. Extensive experimental results on this corpus demonstrate the effectiveness of the proposed product aspect ranking framework.

Product aspect ranking is beneficial to a wide range of real-world applications. In this paper, we investigate its usefulness in two applications, i.e. document-level sentiment classification that aims to determine a review document as expressing a positive or negative overall opinion, and extractive review summarization which aims to summarize consumer reviews by selecting informative review sentences. We perform extensive experiments to evaluate the efficacy of aspect ranking in these two applications and achieve significant performance improvements. Product aspect ranking was first introduced in our previous work [38]. Compared to the preliminary conference version [38], this article has no less than the following improvements: (a) it elaborates more discussions and analysis on product aspect ranking problem; (b) it performs extensive evaluations on more products in more diverse domains; and (c) it demonstrates the potential of aspect ranking in more real-world applications.

In summary, the main contributions of this article include:

- We propose a product aspect ranking framework to automatically identify the important aspects of products from numerous consumer reviews.
- We develop a probabilistic aspect ranking algorithm to infer the importance of various aspects by simultaneously exploiting aspect frequency and

the influence of consumers' opinions given to each aspect over their overall opinions on the product.

 We demonstrate the potential of aspect ranking in real-world applications. Significant performance improvements are obtained on the applications of document-level sentiment classification and extractive review summarization by making use of aspect ranking.

Moreover, the proposed framework and its components are domain-independent and generally applicable in other domains, such as hotel, hawker center, and clothes etc. The rest of this paper is organized as follows. Section 2 elaborates the proposed product aspect ranking framework. Section 3 presents experimental results and analysis, while Section 4 applies the product aspect ranking on real-world applications and demonstrates its usefulness. Section 5 reviews related works and Section 6 concludes this paper.

2 PRODUCT ASPECT RANKING FRAMEWORK

In this section, we present the details of the proposed Product Aspect Ranking framework. We start with an overview of its pipeline (see Fig. 2) consisting of three main components: (a) aspect identification; (b) sentiment classification on aspects; and (c) probabilistic aspect ranking. Given the consumer reviews of a product, we first identify the aspects in the reviews and then analyze consumer opinions on the aspects via a sentiment classifier. Finally, we propose a probabilistic aspect ranking algorithm to infer the importance of the aspects by simultaneously taking into account aspect frequency and the influence of consumers' opinions given to each aspect over their overall opinions.

Let $\mathcal{R} = \{r_1, \ldots, r_{|\mathcal{R}|}\}$ denote a set of consumer reviews of a certain product. In each review $r \in \mathcal{R}$, consumer expresses the opinions on multiple aspects of a product, and finally assigns an overall rating \mathcal{O}_r . \mathcal{O}_r is a numerical score that indicates different levels of overall opinion in the review *r*, i.e. $\mathcal{O}_r \in [\mathcal{O}_{min}, \mathcal{O}_{max}]$, where \mathcal{O}_{min} and \mathcal{O}_{max} are the minimum and maximum ratings respectively. O_r is normalized to [0, 1]. Note that the consumer reviews from different Websites might contain various distributions of ratings. In overall terms, the ratings on some Websites might be a little higher or lower than those on others. Moreover, different Websites might offer different rating range, for example, the rating range is from 1 to 5 on *CNet.com* and from 1 to 10 on *Reevoo.com*, respectively. Hence, we here normalize the ratings from different Websites separately, instead of performing a uniform normalization on them. This strategy is expected to alleviate the influence of the rating variance among different Websites. Suppose there are *m* aspects $\mathcal{A} = \{a_1, \ldots, a_m\}$ in the review corpus \mathcal{R} totally, where a_k is the *k*-th aspect. Consumer opinion on aspect a_k in review *r* is denoted as *o_{rk}*. The opinion on each aspect potentially influences the overall rating. We here assume the overall rating \mathcal{O}_r is generated based on a weighted aggregation of the opinions on specific aspects, as $\sum_{k=1}^{m} \omega_{rk} o_{rk}$ [34], where each weight ω_{rk} essentially measures the importance of aspect a_k in review r. We aim to reveal these important weights, i.e., the emphasis placed on the aspects, and identify the important aspects correspondingly.

In next subsections, we will introduce the aforementioned three components of the proposed product aspect ranking framework. Section 2.1 will introduce the product aspect identification that identifies aspects, i.e., $\{a_k\}_{k=1}^m$ in consumer reviews; Section 2.2 will present the aspect-level sentiment classification which analyzes consumer opinions on aspects i.e., $\{o_{rk}\}_{r=1}^{|\mathcal{R}|}$; and Section 2.3 will elaborate the probabilistic aspect ranking algorithm that estimates the importance weights $\{\omega_{rk}\}_{r=1}^{|\mathcal{R}|}$ and identifies corresponding important aspects.

2.1 Product Aspect Identification

As illustrated in Fig. 3, consumer reviews are composed in different formats on various forum Websites. The Websites such as *CNet.com* require consumers to give an overall rating on the product, describe concise positive and negative opinions (i.e. *Pros and Cons*) on some product aspects, as well as write a paragraph of detailed review in free text. Some Websites, e.g., *Viewpoints.com*, only ask for an overall rating and a paragraph of free-text review. The others such as *Reevoo.com* just require an overall rating and some concise positive and negative opinions on certain aspects. In summary, besides an overall rating, a consumer review consists of *Pros* and *Cons* reviews, free text review, or both.

For the *Pros* and *Cons* reviews, we identify the aspects by extracting the frequent noun terms in the reviews. Previous studies have shown that aspects are usually nouns or noun phrases [19], and we can obtain highly accurate aspects by extracting frequent noun terms from the *Pros* and *Cons* reviews [18]. For identifying aspects in the free text reviews, a straightforward solution is to employ an existing aspect identification approach. One of the most notable existing approach is that proposed by Hu and Liu [12]. It first identifies the nouns and noun phrases in the documents. The occurrence frequencies of the nouns and noun phrases are counted, and only the frequent ones are kept as aspects. Although this simple method is effective in some cases, its well-known limitation is that the identified aspects usually contain noises. Recently, Wu et al. [37] used a phrase dependency parser to extract noun phrases, which form candidate aspects. To filter out the noises, they used a language model by an intuition that the more likely a candidate to be an aspect, the more closely it related to the reviews. The language model was built on product reviews, and used to predict the related scores of the candidate aspects. The candidates with low scores were then filtered out. However, such language model might be biased to the frequent terms in the reviews and cannot precisely sense the related scores of the aspect terms, as a result cannot filter out the noises effectively. In order to obtain more precise identification of aspects, we here propose to exploit the Pros and Cons reviews as auxiliary knowledge to assist identify aspects in the free text reviews. In particular, we first split the free text reviews into sentences, and parse each sentence using Stanford parser². The frequent noun phrases are then extracted from the sentence parsing trees as candidate aspects. Since these candidates may contain noises, we

2. http://nlp.stanford.edu/software/lex-parser.shtml



 A synthesis is ad about the Phone, it is the price. It is a superb phone and even more, but the price or PAYS is a bit much. Is is worth it? Most probably in the long term, but you notice the hole in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after processes. The price of PAYS is a bit much be in your wallet immediately after proceses after proceses. The price of PAYS is a bit much be in your walle

Fig. 3. Sample consumer reviews on various Websites (Best viewed in color).

further leverage the Pros and Cons reviews to assist identify aspects from the candidates. We collect all the frequent noun terms extracted from the Pros and Cons reviews to form a vocabulary. We then represent each aspect in the Pros and Cons reviews into a unigram feature, and utilize all the aspects to learn a one-class Support Vector Machine (SVM) classifier [21]. The resultant classifier is in turn used to identify aspects in the candidates extracted from the free text reviews. As the identified aspects may contain some synonym terms, such as "earphone" and "headphone," we perform synonym clustering to obtain unique aspects. In particular, we collect the synonym terms of the aspects as features. The synonym terms are collected from the synonym dictionary Website³. We represent each aspect into a feature vector and use the Cosine similarity for clustering. The ISODATA (Iterative Self-Organizing Data Analysis Technique) clustering algorithm [14] is employed for synonym clustering. ISODATA does not need to fix the number of clusters and can learn the number automatically from the data distribution. It iteratively refines clustering by splitting and merging of clusters. Clusters are merged if the centers of two clusters are closer than a certain threshold. One cluster is split into two different clusters if the cluster standard deviation exceeds a predefined threshold. The values of these two thresholds were empirically set to 0.2 and 0.4 in our experiments.

2.2 Sentiment Classification on Product Aspects

The task of analyzing the sentiments expressed on aspects is called aspect-level sentiment classification in literature [12]. Exiting techniques include the supervised learning approaches and the lexicon-based approaches, which are typically unsupervised. The lexicon-based methods utilize a sentiment lexicon consisting of a list of sentiment words, phrases and idioms, to determine the sentiment orientation on each aspect [23]. While these method are easily to implement, their performance relies heavily on the quality of the sentiment lexicon. On the other hand, the supervised learning methods train a sentiment classifier based on training corpus. The classifier is then used to predict the sentiment on each aspect. Many learning-based classification models are applicable, for example, Support Vector Machine (SVM), Naive Bayes, and Maximum Entropy (ME) model etc. [25]. Supervised learning is dependent on the training data and cannot perform well without sufficient training samples. However, labeling training data is laborintensive and time-consuming. In this work, the Pros and Cons reviews have explicitly categorized positive and negative opinions on the aspects. These reviews are valuable training samples for learning a sentiment classifier. We thus exploit Pros and Cons reviews to train a sentiment classifier, which is in turn used to determine consumer opinions (positive or negative) on the aspects in free text reviews. Specifically, we first collect the sentiment terms in Pros and Cons reviews based on the sentiment lexicon provided by MPQA project [35]. These terms are used as features, and each review is represented as a feature vector. A sentiment classifier is then learned from the Pros

3. http://thesaurus.com

reviews (i.e., positive samples) and Cons reviews (i.e., negative samples). The classifier can be SVM, Naïve Bayes or Maximum Entropy model [25]. Given a free text review that may cover multiple aspects, we first locate the opinionated expression that modifies the corresponding aspect, e.g. locating the expression "well" in the review "The battery of Nokia N95 works well." for the aspect "battery." Generally, an opinionated expression is associated with the aspect if it contains at least one sentiment term in the sentiment lexicon, and it is the closest one to the aspect in the parsing tree within the context distance of 5. The learned sentiment classifier is then leveraged to determine the opinion of the opinionated expression, i.e. the opinion on the aspect.

2.3 Probabilistic Aspect Ranking Algorithm

In this section, we propose a probabilistic aspect ranking algorithm to identify the important aspects of a product from consumer reviews. Generally, important aspects have the following characteristics: (a) they are frequently commented in consumer reviews; and (b) consumers' opinions on these aspects greatly influence their overall opinions on the product. The overall opinion in a review is an aggregation of the opinions given to specific aspects in the review, and various aspects have different contributions in the aggregation. That is, the opinions on (un)important aspects have strong (weak) impacts on the generation of overall opinion. To model such aggregation, we formulate that the overall rating \mathcal{O}_r in each review r is generated based on the weighted sum of the opinions on specific aspects, as $\sum_{k=1}^{m} \omega_{rk} o_{rk}$ or in matrix form as $\boldsymbol{\omega_r}^T \boldsymbol{o_r}$. o_{rk} is the opinion on aspect a_k and the importance weight ω_{rk} reflects the emphasis placed on a_k . Larger ω_{rk} indicates a_k is more important, and vice versa. ω_r denotes a vector of the weights, and o_r is the opinion vector with each dimension indicating the opinion on a particular aspect. Specifically, the observed overall ratings are assumed to be generated from a Gaussian *Distribution*, with mean $\omega_r^T o_r$ and variance σ^2 as:

$$p(\mathcal{O}_r) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(\mathcal{O}_r - \boldsymbol{\omega}_r^T \boldsymbol{o}_r)^2}{2\sigma^2}\right\}.$$
 (1)

In order to take the uncertainty of ω_r into consideration, we assume ω_r as a sample drawn from a *Multivariate Gaussian* Distribution as:

$$p(\boldsymbol{\omega}_r) = \frac{1}{(2\pi)^{m/2} |\boldsymbol{\Sigma}|^{1/2}} \exp\left\{-\frac{1}{2}(\boldsymbol{\omega}_r - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{\omega}_r - \boldsymbol{\mu})\right\}, (2)$$

where μ and Σ are the mean vector and covariance matrix, respectively. They are both unknown and need to be estimated.

As aforementioned, the aspects that are frequently commented by consumers are likely to be important. Hence, we exploit aspect frequency as the prior knowledge to assist learning ω_r . In particular, we expect the distribution of ω_r , i.e., $\mathcal{N}(\mu, \Sigma)$ is close to the distribution $\mathcal{N}(\mu_0, \mathbf{I})$. Each element in μ_0 is the frequency of a specific aspect: $frequency(a_k) / \sum_{i=1}^{m} frequency(a_i)$. Thus, we formulate the distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ based on its Kullback-Leibler (KL) divergence to $\mathcal{N}(\mu_0, \mathbf{I})$ as

$$p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \exp\left\{-\varphi \cdot \operatorname{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) || \mathcal{N}(\boldsymbol{\mu}_{0}, \boldsymbol{I}))\right\}, \quad (3)$$

where φ is a weighting parameter.

Base on the above formula, the probability of generating overall opinion rating \mathcal{O}_r in review *r* is given as

$$P(\mathcal{O}_r|r) = P\left(\mathcal{O}_r|\boldsymbol{\omega}_r, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma^2\right)$$
$$= \int p\left(\mathcal{O}_r|\boldsymbol{\omega}_r^T \boldsymbol{o}_r, \sigma^2\right) \cdot p(\boldsymbol{\omega}_r|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \cdot p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) d\boldsymbol{\omega}_r, \quad (4)$$

where $\{\omega_r\}_{r=1}^{|\mathcal{R}|}$ are the importance weights and $\{\mu, \Sigma, \sigma^2\}$ are the model parameters. While $\{\mu, \Sigma, \sigma^2\}$ can be estimated from review corpus $\mathcal{R} = \{r_1, \ldots, r_{|\mathcal{R}|}\}$ using the maximumlikelihood (ML) estimation, ω_r in review r can be optimized through the maximum a posteriori (MAP) estimation. Since ω_r and $\{\mu, \Sigma, \sigma^2\}$ are coupled with each other, we here optimize them using a EM-style algorithm. We iteratively optimize $\{\omega_r\}_{r=1}^{|\mathcal{R}|}$ and $\{\mu, \Sigma, \sigma^2\}$ in each E-step and M-step respectively as follows.

Optimizing ω_r given $\{\mu, \Sigma, \sigma^2\}$:

Suppose we are given the parameters $\{\mu, \Sigma, \sigma^2\}$, we use the maximum a posteriori (MAP) estimation to get the optimal value of ω_r . The object function of MAP estimation for review *r* is defined as:

$$\mathcal{L}(\boldsymbol{\omega}_{r}) = \log p\left(\mathcal{O}_{r}|\boldsymbol{\omega}_{r}^{T}\boldsymbol{o}_{r},\sigma^{2}\right)p(\boldsymbol{\omega}_{r}|\boldsymbol{\mu},\boldsymbol{\Sigma})p(\boldsymbol{\mu},\boldsymbol{\Sigma}).$$
(5)

By substituting Eq. (1) - (3), we get

$$\mathcal{L}(\boldsymbol{\omega}_{r}) = -\frac{\left(\mathcal{O}_{r} - \boldsymbol{\omega}_{r}^{T}\boldsymbol{o}_{r}\right)^{2}}{2\sigma^{2}} - \frac{1}{2}(\boldsymbol{\omega}_{r} - \boldsymbol{\mu})^{T}\boldsymbol{\Sigma}^{-1}(\boldsymbol{\omega}_{r} - \boldsymbol{\mu}) - \varphi \cdot \mathrm{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}))|\mathcal{N}(\boldsymbol{\mu}_{0}, \boldsymbol{I})) - \log\left(\sigma|\boldsymbol{\Sigma}|^{1/2}(2\pi)^{\frac{m+1}{2}}\right).$$
(6)

 ω_r can thus be optimized through MAP estimation as follows:

$$\hat{\boldsymbol{\omega}}_{r} = \arg\max_{\boldsymbol{\omega}_{r}} \mathcal{L}(\boldsymbol{\omega}_{r})$$

$$= \arg\max_{\boldsymbol{\omega}_{r}} \left\{ -\frac{\left(\mathcal{O}_{r} - \boldsymbol{\omega}_{r}^{T}\boldsymbol{\sigma}_{r}\right)^{2}}{2\sigma^{2}} -\frac{1}{2}(\boldsymbol{\omega}_{r} - \boldsymbol{\mu})^{T}\boldsymbol{\Sigma}^{-1}(\boldsymbol{\omega}_{r} - \boldsymbol{\mu}) \right\}.$$
(7)

We take the derivative of $\mathcal{L}(\omega_r)$ with respect to ω_r and let it vanish at the minimizer:

$$\frac{\partial \mathcal{L}(\boldsymbol{\omega}_{r})}{\partial \boldsymbol{\omega}_{r}} = -\frac{\left(\boldsymbol{\omega}_{r}^{T}\boldsymbol{o}_{r} - \mathcal{O}_{r}\right) \cdot \boldsymbol{o}_{r}}{\sigma^{2}} - \boldsymbol{\Sigma}^{-1}(\boldsymbol{\omega}_{r} - \boldsymbol{\mu}) = 0, \quad (8)$$

which results in the following solution:

$$\hat{\boldsymbol{\omega}}_{\boldsymbol{r}} = \left(\frac{\boldsymbol{o}_{\boldsymbol{r}}\boldsymbol{o}_{\boldsymbol{r}}^{T}}{\sigma^{2}} + \boldsymbol{\Sigma}^{-1}\right)^{-1} \left(\frac{\mathcal{O}_{\boldsymbol{r}} \cdot \boldsymbol{o}_{\boldsymbol{r}}}{\sigma^{2}} + \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}\right). \tag{9}$$

Optimizing $\{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma^2\}$ given $\boldsymbol{\omega}_r$: Given $\{\boldsymbol{\omega}_r\}_{r=1}^{|\mathcal{R}|}$, we optimize the parameters $\{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma^2\}$ using the maximum-likelihood (ML) estimation over the review corpus R. The parameters are expected to maximize the probability of observing all the overall ratings on the corpus \mathcal{R} . Thus, they are estimated by maximizing the log-likelihood function over the whole review corpus \mathcal{R} as follows. For the sake of simplicity, we denote { μ , Σ , σ^2 } as Ψ .

$$\hat{\Psi} = \arg\max_{\Psi} \mathcal{L}(\mathcal{R}) = \arg\max_{\Psi} \sum_{r \in \mathcal{R}} \log p(\mathcal{O}_r | \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma^2).$$
(10)

By substituting Eq.(1) - (3), we obtain

$$\hat{\Psi} = \arg \max_{\Psi} \sum_{r \in \mathcal{R}} \left\{ -\frac{1}{2} (\boldsymbol{\omega}_{r} - \boldsymbol{\mu})^{T} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\omega}_{r} - \boldsymbol{\mu}) - \frac{\left(\mathcal{O}_{r} - \boldsymbol{\omega}_{r}^{T} \boldsymbol{o}_{r}\right)^{2}}{2\sigma^{2}} - \varphi \cdot \text{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})) ||\mathcal{N}(\boldsymbol{\mu}_{0}, \boldsymbol{I})) - \log\left(\sigma |\boldsymbol{\Sigma}|^{1/2} (2\pi)^{\frac{m+1}{2}}\right) \right\}.$$
(11)

We take the derivative of $\mathcal{L}(\mathcal{R})$ with respect to each parameter in { μ , Σ , σ^2 }, and let it vanish at the minimizer:

$$\frac{\partial \mathcal{L}(\mathcal{R})}{\partial \boldsymbol{\mu}} = \sum_{r \in \mathcal{R}} \left\{ -\boldsymbol{\Sigma}^{-1}(\boldsymbol{\omega}_r - \boldsymbol{\mu}) \right\} - \varphi(\boldsymbol{\mu}_0 - \boldsymbol{\mu}) = 0$$
$$\frac{\partial \mathcal{L}(\mathcal{R})}{\partial \boldsymbol{\Sigma}} = \sum_{r \in \mathcal{R}} \left\{ -(\boldsymbol{\Sigma}^{-1})^T + \left((\boldsymbol{\Sigma}^{-1})^T (\boldsymbol{\omega}_r - \boldsymbol{\mu}) (\boldsymbol{\omega}_r - \boldsymbol{\mu})^T \right)^T \right\}$$
$$(\boldsymbol{\Sigma}^{-1})^T \right\} + \varphi \cdot \left((\boldsymbol{\Sigma}^{-1})^T - \mathbf{I} \right) = 0$$
$$\frac{\partial \mathcal{L}(\mathcal{R})}{\partial \sigma^2} = \sum_{r \in \mathcal{R}} \left(-\frac{1}{\sigma^2} + \frac{(\mathcal{O}_r - \boldsymbol{\omega}_r^T \boldsymbol{o}_r)^2}{\sigma^4} \right) = 0, \quad (12)$$

which lead to the following solutions:

$$\hat{\boldsymbol{\mu}} = \left(|\mathcal{R}| \cdot \boldsymbol{\Sigma}^{-1} + \boldsymbol{\varphi} \cdot \mathbf{I} \right)^{-1} \left(\boldsymbol{\Sigma}^{-1} \sum_{r \in \mathcal{R}} \boldsymbol{\omega}_r + \boldsymbol{\varphi} \cdot \boldsymbol{\mu}_0 \right)$$
$$\hat{\boldsymbol{\Sigma}} = \left(\frac{1}{\varphi} \sum_{r \in \mathcal{R}} \left((\boldsymbol{\omega}_r - \boldsymbol{\mu}) (\boldsymbol{\omega}_r - \boldsymbol{\mu})^T \right) + \left(\frac{|\mathcal{R}| - \varphi}{2\varphi} \right)^2 \cdot \mathbf{I} \right)^{1/2}$$
$$- \frac{(|\mathcal{R}| - \varphi)}{2\varphi} \cdot \mathbf{I}$$
$$\hat{\sigma}^2 = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} (\mathcal{O}_r - \boldsymbol{\omega}_r^T \boldsymbol{o}_r)^2.$$
(13)

We repeat the above two optimization steps until the likelihood value converges. The convergence of this iterative optimization is analyzed as follows. Let $\boldsymbol{\omega}$ denote the parameters $\{\boldsymbol{\omega}_r\}_{r=1}^{|\mathcal{R}|}$. The overall log-likelihood function is denoted as $\mathcal{L}(\boldsymbol{\omega}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma^2)$. At iteration t+1, $\boldsymbol{\omega}^{(t+1)}$ obtained by Eq. (9) is the solution of the optimization in Eq. (7). Thus, we have $\mathcal{L}(\boldsymbol{\omega}^{(t+1)}, \boldsymbol{\mu}^{(t)}, \boldsymbol{\Sigma}^{(t)}, \sigma^{2(t)}) \geq \mathcal{L}(\boldsymbol{\omega}^{(t)}, \boldsymbol{\mu}^{(t)}, \boldsymbol{\Sigma}^{(t)}, \sigma^{2(t)})$. Similarly, $\boldsymbol{\mu}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}$, and $\sigma^{2(t+1)}$ obtained from Eq. (14) are the solutions of the optimization in Eq. (10), leading to $\mathcal{L}(\boldsymbol{\omega}^{(t+1)}, \boldsymbol{\mu}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \sigma^{2(t+1)}) \geq \mathcal{L}(\boldsymbol{\omega}^{(t+1)}, \boldsymbol{\mu}^{(t)}, \boldsymbol{\Sigma}^{(t)}, \sigma^{2(t)})$. These two inequalities indicate that the iterative optimization monotonically increases the log-likelihood function value in each iteration, and finally converges.

After obtaining the importance weights ω_r for each review $r \in \mathcal{R}$, we compute the overall importance score ϖ_k of each aspect a_k by integrating its importance scores over the reviews as $\varpi_k = (\sum_{r \in \mathcal{R}} \omega_{rk})/|\mathcal{R}_k|$, where \mathcal{R}_k is the set of reviews containing a_k . According to ϖ_k , the important product aspects can be identified.

Algorithm 1 Probabilistic Aspect Ranking Algorithm

Input: Consumer review corpus \mathcal{R} , each review $r \in \mathcal{R}$ is associated with an overall rating \mathcal{O}_r , and a vector of opinions o_r on specific aspects.

Output: Importance scores $\varpi_k|_{k=1}^m$ for all the *m* aspects.

while not converged do

Update $\{\omega_r\}_{r=1}^{|\hat{\mathcal{R}}|}$ according to Eq. (9); Update $\{\mu, \Sigma, \sigma^2\}$ according to Eq. (13); end while Compute aspect importance scores $\{\varpi_k\}_{k=1}^m$

3 EVALUATIONS

In this section, we conduct extensive experiments to evaluate the effectiveness of the proposed product aspect ranking framework, including product aspect identification, sentiment classification on aspects, and aspect ranking.

3.1 Experimental Data and Settings

Table I shows the details of our product review corpus, which is publicly available by request. This dataset contains consumer reviews on 21 popular products in eight domains. There are 94,560 reviews in total and around 4,503 reviews for each product on average. These reviews were crawled from multiple prevalent forum Websites, including *cnet.com*, *viewpoints.com*, *reevoo.com*, *gsmarena.com* and *pricegrabber.com*. The reviews were posted between June 2009 and July 2011. Eight annotators were invited to annotate the ground truth on these reviews. They were asked to annotate the product aspects in each review, and also label consumer opinions expressed on the aspects. Each review was labeled by at least two annotators. The average interrater agreement in terms of Kappa statistics is 87% for all the products.

 F_1 -measure was used as the evaluation metric for aspect identification and aspect sentiment classification. It is a combination of *precision* and *recall*, as $F_1 = 2 * precision * recall/(precision + recall)$. To evaluate the performance of

TABLE 1 Statistics of Our Product Review Corpus, # Denotes the Number of Reviews/Sentences

Product Name	Domain	#Review	#Sentence
Canon EOS 450D (Canon EOS)	camera	440	628
Fujifilm Finepix AX245W (Fujifilm)	camera	541	839
Panasonic Lumix DMC (Panasonic)	camera	650	1,546
Apple MacBook Pro (MacBook)	laptop	552	4,221
Samsung NC10 (Samsung)	laptop	2,712	4,946
Apple iPod Touch 2nd (iPod Touch)	MP3	4,567	10,846
Sony NWZ-S639 16GB (Sony NWZ)	MP3	341	773
BlackBerry Bold 9700 (BlackBerry)	phone	4,070	11,008
iPhone 3GS 16GB (iPhone 3GS)	phone	12,418	43,527
Nokia 5800 XpressMusic (Nokia 5800)	phone	28,129	75,001
Nokia N95	phone	15,939	44,379
Sony Handycam HDR (Handycam)	Camcorder	5,692	8,827
GoPro Motorsports Hero SD (GoPro)	Camcorder	5,689	8,828
Sharp AQUOS LC-70 (AQUOS)	TV	1,776	4,605
Samsung LN46D630 (Samsung TV)	TV	1,757	4,542
Vizio E421VA (Vizio)	TV	1,850	5,290
Garmin Nuvi (Garmin)	GPS	1,253	3,474
Tomtom XXL (Tomtom)	GPS	1,995	5,112
Epson Artisan 835 (Epson)	Printer	2,539	4,965
Brother HL-2280DW (Brother)	Printer	755	1,743
HP Officejet 4500 (Officejet)	Printer	895	2,246



Fig. 4. Performance of product aspect identification. The results passed statistical significance test, i.e., T-Test, with p-values $\,<\,$ 0.05.

aspect ranking, we adopted the widely used *Normalized Discounted Cumulative Gain* at top k (NDCG@k) [13] as the evaluation metric. Given a ranking list of aspects, NDCG@k is calculated as

$$NDCG@k = \frac{1}{Z} \sum_{i=1}^{k} \frac{2^{t(i)} - 1}{\log(1+i)},$$
(14)

where t(i) is the importance degree of the aspect at position i, and Z is a normalization term derived from the top-kaspects of a perfect ranking. For each aspect, its importance degree was judged by three annotators as three importance levels, i.e. "Un-important" (score 1), "Ordinary" (score 2), and "Important" (score 3). Ideally, we should invite annotators to read all the reviews and then give their judgements. However, such labeling process is very time-consuming and labor-intensive. Since NDCG@k is calculated with the importance degrees of the top-*k* aspects, we speed up the labeling process as follows. We first collected the top-kaspects from the ranking results of all the evaluated methods in Section 3.4. We then randomly sampled 100 reviews on these aspects, and provided them to the annotators for labeling the importance levels of the aspects. In particular, the annotators were invited to read the reviews and identify the coherence or conflict between the opinion on each aspect and the overall rating in each review. Generally, an aspect with more coherence cases tends to be more important, while an aspect with more conflict cases is likely to be less important. Besides, the frequencies of the aspects in all the reviews were presented to the annotators as another reference for the labeling. The importance ratings from the annotators for each aspect were then averaged to form the final rating.

3.2 Evaluations of Product Aspect Identification on Free Text Reviews

We compared our aspect identification approach with the following two methods: (a) the method proposed by Hu and Liu in [12], which extracts nouns and noun phrases as aspect candidates, and identifies aspects by rules learned from association rule mining; and (b) the method proposed by Wu *et al.* in [37], that extracts noun phrases from a dependency parsing tree as aspect candidates, and identifies aspects by a language model built on the reviews.



Fig. 5. Performance of sentiment classification on product aspects. T-Test, p-values < 0.05.

Fig. 4 shows the performance comparison on all the 21 products in terms of F_1 -measure. From these results, we can see that the proposed approach get the best performance on all the 21 products. It significantly outperforms Hu's and Wu's methods by over 9.0% and 5.3% respectively in terms of average F_1 -measure. This indicates the effectiveness of *Pros* and *Cons* reviews in assisting aspect identification on free text reviews. Hence, by exploiting the *Pros* and *Cons* reviews, our approach can boost the performance of aspect identification.

3.3 Evaluations of Sentiment Classification on Product Aspects

In this experiment, we compared the following methods of sentiment classification: (a) one unsupervised method. The opinion on each aspect is determined by referring to the sentiment lexicon SentiWordNet [23]. This lexicon contains a list of positive/negative sentiment words. The opinionated expression modifying an aspect is classified as positive (or negative) if it contains a majority of words in the positive (or negative) list; and (b) three supervised methods. We employed three supervised methods proposed in Pang et al. [25], including Naïve Bayes (NB), Maximum Entropy (ME), and Support Vector Machine (SVM). The sentiment classifiers were trained on the Pros and Cons reviews as described in Section 2.2. In particular, SVM was implemented by using libSVM [2] with linear kernel, NB was implemented with Laplace smoothing, and ME was implemented with L-BFGS parameter estimation.

Fig. 5 shows the experimental results. We can see that the three supervised methods perform much better than the unsupervised approach. They achieve performance improvements on all the 21 products. In particular, SVM performs the best on 18 products, NB obtains the best performance on the remaining three products. In terms of average performance, SVM achieves slight improvements compared to NB and ME. These results are consistent with the previous research [25].

3.4 Evaluations of Aspect Ranking

In order to evaluate the effectiveness on aspect ranking, we compared the proposed aspect ranking algorithm with



Fig. 6. Performance of aspect ranking in terms of NDCG @5. T-Test, p-values < 0.05.



Fig. 7. Performance of aspect ranking in terms of NDCG @10. T-Test, p-values < 0.05.

the following three methods: (a) **Frequency-based method**, which ranks the aspects according to aspect frequency; (b) **Correlation-based method**, which measures the correlation between the opinions on specific aspects and the overall ratings. It ranks the aspects based on the number of cases when such two kinds of opinions are consistent; and (c) **Hybrid method**, that captures both aspect frequency and the correlation by a linear combination, as $\lambda \cdot$ *Frequency-based Ranking* + $(1 - \lambda) \cdot$ *Correlation-based Ranking*, where λ is set to 0.5 in the experiments.

Figs. 6-8 show the comparison results in terms of NDCG@5, NDCG@10, and NDCG@15, respectively. On average, the proposed aspect ranking approach significantly outperforms frequency-based, correlation-based, and hybrid methods in terms of NDCG@5 by over 9.0%, 7.4% and 8.1%, respectively. It improves the performance over these three methods in terms of NDCG@10 by over 4.6%, 3.6% and 4.0%, respectively, while in terms of NDCG@15 by over 4.6%, 3.3% and 4.0%, respectively. Hence, we can speculate that the proposed approach can effectively identify the important aspects from consumer reviews by simultaneously exploiting aspect frequency and the influence of consumers' opinions given to each aspect over their overall opinions. The frequency-based method only captures the aspect frequency information, and neglects to consider the impact of opinions on the specific aspects on the overall ratings. It may recognize some general aspects as important ones. Although the general aspects frequently appear in consumer reviews, they do not greatly influence consumers' overall satisfaction. Correlation-based method ranks the aspects by simply counting the consistent cases between opinions on specific aspects and the overall ratings. It ignores to model the uncertainty in the generation of overall ratings, and thus cannot achieve satisfactory performance. The hybrid method simply aggregates the results from the frequency-based and correlation-based methods, and cannot boost the performance effectively. Table II shows sample results by these four methods. Top 10 aspects of the product *iPhone 3GS* are listed. From these four ranking lists, we can see that the proposed aspect ranking method generates more reasonable ranking than the other methods. For example, the aspect "phone" is ranked at the top by the other methods. However, "phone" is

a general but not important aspect. To better investigate the reasonability of the ranking results of the proposed approach, we refer to one public user-feedback report, i.e., the "China Unicom 100 customers iPhone user feedback report" [4]. This report shows that the top four aspects of iPhone product, which users most concern about, are "3G Network" (30%), "usability" (30%), "out-looking design" (26%), "application" (15%). We can see that these four aspects are also ranked at top by our proposed aspect ranking approach.

Moreover, Fig. 9 shows the correlations among the importance weights of some aspects obtained from the proposed approach. Due to the page limitation, the correlations among the top 10 aspects of three products are illustrated here. We can see that some aspects are correlated to each other reasonably, for example, the aspects "apps" and "storage" of the product *iPhone 3GS*, "design" and "touch-pad" of *Macbook*, and "focusing" and "speed" of *Cannon Eos* etc.

4 APPLICATIONS

Aspect ranking is beneficial to a wide range of realworld applications. We here investigate its capacity in two



Fig. 8. Performance of aspect ranking in terms of NDCG @15. T-Test, p-values < 0.05.



Fig. 9. Sample aspect correlation of the products iPhone 3GS, Macbook, and Cannon Eos: (a) iPhone 3GS. (b) Macbook. (c) Cannon EOS.

applications, i.e. document-level sentiment classification on review documents, and extractive review summarization.

4.1 Document-level Sentiment Classification

The goal of document-level sentiment classification is to determine the overall opinion of a given review document. A review document often expresses various opinions on multiple aspects of a certain product. The opinions on different aspects might be in contrast to each other, and have different degree of impacts on the overall opinion of the review document. For example, a sample review document of *iPhone* 4 is shown in Fig. 10. It expresses positive opinions on some aspects such as "reliability," "easy to use," and simultaneously criticizes some other aspects such as "touch screen," "quirk," "music play." Finally, it assigns an high overall rating (i.e., positive opinion) on iPhone 4 due to that the important aspects are with positive opinions. Hence, identifying important aspects can naturally facilitate the estimation of the overall opinions on review documents. This observation motivates us to utilize the aspect ranking results to assist document-level sentiment classification.

We conducted evaluations of document-level sentiment classification over the product reviews described in Section 3.1. Specifically, we randomly sampled 100 reviews of each product as testing samples and used the remaining reviews for training. Each review contains an overall rating, which is normalized to [0,1]. We treated the reviews with high overall rating (>0.5) as positive samples, and those with low rating (<0.5) as negative samples. The reviews with ratings of 0.5 were considered as neutral and not used in our experiments. We collected noun terms, aspects, and sentiment terms from the training reviews as features. Note that sentiment terms are defined as those appear in the sentiment lexicon provided by MPQA project [35].

TABLE 2 Top 10 Aspects Ranked by Four Methods for *iPhone 3GS*

	Frequency	Correlated	Hybrid	Our Method
1	Phone	Phone	Phone	Usability
2	Usability	Usability	Usability	Apps
3	3G	Apps	Apps	3G
4	Apps	3G	3G	Battery
5	Camera	Camera	Camera	Looking
6	Feature	Looking	Looking	Storage
7	Looking	Feature	Feature	Price
8	Battery	Screen	Battery	Software
9	Screen	Battery	Screen	Camera
10	Flash	Bluetooth	Flash	Call quality

All the training and testing reviews were then represented into feature vectors. In the representation, we gave more emphasis on the important aspects, and the sentiment terms modifying them. Technically, the feature dimensions corresponding to aspect a_k and its corresponding sentiment terms were weighted by $1 + \varphi \cdot \varpi_k$, where ϖ_k is the importance score of a_k , and φ is a tradeoff parameter and was empirically set to 100 in the experiments. Based on the weighted features, a SVM classifier was learned from the training reviews and used to determine the overall opinions on the testing reviews.

We compared our approach with two existing methods, i.e., Boolean weighting and term frequency (TF) weighting. Boolean weighting represents each review into a feature vector of Boolean values, each of which indicates the presence or absence of the corresponding feature in the review. Term frequency (TF) weighting [24] weights the Boolean feature by the frequency of each feature on the corpus. Table III shows the classification performance on the reviews of all the 21 products as well as the average performance over them. Here, our approach is termed as AR since it incoporates Aspect Ranking results into the feature representation. From Table III, we can see that our AR weighting approach achieves better performance than the Boolean and TF weighting methods. In particular, it performs the best on all the 21 products, and significantly outperforms the Boolean and TF weighting methods by over 4.4% and 5.9% respectively, in terms of average F_1 -measure. It is worthy to



"Excellent phone!



Apple iPhone 4 (16 GB) — This is an excellent phone, regardless of service provider. It is very reliable has the most apps out of any smart phones, and is easy to use and love It features a 3.5 inch touch screen, which is pretty hard to scratch.

They don't handle long falls or water well, so if you are clumsy or forgetful, investing in an otter box case would be wise.

Included when you purchase the phone is a compass, gps, voice control, youtube, safari browser, email, stocks, calculator, weather, voice memos, itunes, notes, alarm clock, camera, and more!

If you take care of this phone, it will take care of you! I have never had any serious issues with my phone. Sometimes there are some minor quirks, but nothing that has affected the performance of the phone overall.

The speakers that it plays music over are a little weak though, but that's to be expected really. The home screen key can be a little difficult but again nothing serious

Fig. 10. Sample review document on product iPhone 4.

Product	Boolean	TF	AR
AQUOS	0.67	0.66	0.68
BlackBerry	0.74	0.73	0.77
Brother	0.66	0.64	0.69
Canon EOS	0.68	0.67	0.71
Epson	0.64	0.63	0.67
Fujifilm	0.69	0.68	0.73
Garmin	0.65	0.64	0.67
GoPro	0.61	0.60	0.64
Handycam	0.65	0.65	0.66
iPhone 3GS	0.78	0.77	0.80
iPod Touch	0.71	0.69	0.74
MacBook	0.72	0.72	0.75
Nokia 5800	0.79	0.78	0.81
Nokia N95	0.71	0.70	0.75
Officejet	0.65	0.63	0.68
Panasonic	0.69	0.67	0.71
Samsung	0.72	0.72	0.75
Samsung TV	0.67	0.66	0.69
Sony NWZ	0.68	0.65	0.71
Tomtom	0.64	0.64	0.65
Vizio	0.63	0.63	0.65
Average	0.68	0.67	0.71

T-Test, p-values < 0.05.

note that Boolean weighting is a special case of AR weighting. When we set all the aspects to be equally important, AR weighting degrades to Boolean weighting. From these results, we can conclude that aspect ranking is capable for boosting the performance of document-level sentiment classification effectively. In addition, the results also show that Boolean weighting achieves slight performance improvement over TF weighting by about 1.5% in terms of average F_1 -measure. This is consistent with previous research [25].

4.2 Extractive Review Summarization

As aforementioned, for a particular product, there is an abundance of consumer reviews available on the internet. However, the reviews are disorganized. It is impractical for user to grasp the overview of consumer reviews and opinions on various aspects of a product from such enormous reviews. On the other hand, the Internet provides more information than is needed. Hence, there is a compelling need for automatic review summarization, which aims to condense the source reviews into a shorter version preserving its information content and overall meaning. Existing review summarization methods can be classified into abstractive and extractive summarization. An abstractive summarization attempts to develop an understanding of the main topics in the source reviews and then express those topics in clear natural language. It uses linguistic techniques to examine and interpret the text and then to find the new concepts and expressions to best describe it by generating a new shorter text that conveys the most important information from the original text document. An extractive method summarization method consists of selecting important sentences and paragraphs etc. from the original reviews and concatenating them into shorter from.

In this paper, we focus on extractive review summarization. We investigate the capacity of aspect ranking in improving the summarization performance. As introduced above, extractive summarization is formulated by extracting the most informative segments (e.g. sentences or passages) from the source reviews. The most informative content is generally treated as the "most frequent" or the "most favorably positioned" content in existing works. In particular, a scoring function is defined for computing the informativeness of each sentence *s* as follows [3]:

$$I(s) = \lambda_1 \cdot I_a(s) + \lambda_2 \cdot I_o(s); \quad \lambda_1 + \lambda_2 = 1,$$
(15)

where $I_a(s)$ quantifies the informativeness of sentence *s* in terms of the importance of aspects in *s*, and $I_o(s)$ measures the informativeness in terms of the representativeness of opinions expressed in *s*. λ_1 and λ_2 are tradeoff parameters. Generally, $I_a(s)$ and $I_o(s)$ are defined as follows:

 $I_a(s)$: Most existing methods regard the sentences containing frequent aspects as important. They define $I_a(s)$ simply based on aspect frequency as

$$I_a(s) = \sum_{aspect \text{ in } s} \text{frequency}(aspect).$$
(16)

 $I_o(s)$: The resultant summary is expected to include the opinionated sentences in source reviews, so as to offer a summarization of consumer opinions. Moreover, the summary is desired to include the sentences whose opinions are consistent with consumer's overall opinion. Correspondingly, $I_o(s)$ is defined as:

$$I_o(s) = \alpha \cdot \text{Subjective}(s) + \beta \cdot \text{Consistency}(s).$$
(17)

Subjective(s) is used to distinguish the opinionated sentences from factual ones, and Consistency(s) measures the consistency between the opinion in s and the overall opinion as follows:

Subjective(s) =
$$\sum_{term \text{ in } s} |\text{Polarity}(term)|,$$

Consistent(s) = $-(overall \ rating - \text{Polarity}(s))^2,$ (18)

where Polarity(s) is computed as

$$Polarity(s) = \sum_{term in s} \frac{Polarity(term)}{Subjective(s) + \varepsilon},$$
(19)

where Polarity(*term*) is the polarity of a particular term and ε is a constant to prevent zero for the denominator.

With the informativeness of review sentences computed by the above scoring function, the informative sentences can then be selected by the following two approaches: (a) sentence ranking (SR) method [29] ranks the sentences according to their informativeness and select the top ranked sentences to form a summarization; and (b) graph-based (GB) method [10] represents the sentences in a graph, where each node corresponds to a particular sentence and each edge characterizes the relation between two sentences. A random walk is then performed over the graph to discover the most informative sentences. The initial score of each node is defined as its informativeness from the scoring function in Eq. (15) and the edge weight is computed as the Cosine similarity between the sentences with *unigram* feature.

As aforementioned, the frequent aspects might not be the important ones and aspect frequency is not capable for characterizing the importance of aspects. This motivates us to improve the above scoring function by exploiting the aspect ranking results, which indicate the importance of aspects. We define the informativeness of sentence s in terms of the importance of aspects within it as:

$$I_{ar}(s) = \sum_{aspect \text{ in } s} \text{ importance}(aspect),$$
(20)

where the importance(*aspect*) is the importance score obtained by our proposed aspect ranking Algorithm 1. The overall informativeness of s is then computed as:

$$I(s) = \lambda_1 \cdot I_{ar}(s) + \lambda_2 \cdot I_o(s); \quad \lambda_1 + \lambda_2 = 1.$$
(21)

We conducted evaluation on the product review corpus introduced in Section 3.1 to investigate the effectiveness of the above approach. We randomly sampled 100 reviews of each product as testing samples. The remaining reviews were used to learn the aspect ranking results. In order to avoid selecting redundant sentences commenting on the same aspect, we adopted the strategy proposed in [39]. Specifically, after selecting each new sentence, we updated the informativeness of the remaining sentences as follows: the informativeness of a remaining sentence s_i commenting on the same aspect with a selected sentence s_i was reduced by $\exp\{-\eta \cdot similarity(s_i, s_i)\}$ where $similarity(\cdot)$ is the Cosine similarity between two sentences with *unigram* feature. η is a tradeoff parameter and was empirically set to 10 in the experiments. We invited three annotators to generate the reference summaries for each product. Each annotator was invited to read the consumer reviews of a product and write a summary of up to 100 words individually by selecting the informative sentences based on her own judgements. We adopted ROUGE (i.e., Recall-Oriented Understudy for Gisting Evaluation) [17] as the performance metric to evaluate the quality of the summary generated by the above methods. ROUGE is a widely used evaluation metric of summaries [17]. It measures the quality of a summary by counting the overlapping N-grams between it and a set of reference summaries generated by human.

$$\text{ROUGE-N} = \frac{\sum_{S \in \{\text{Reference Summaries}\}} \sum_{\text{gram}_n \in S} Count_{match}(gram_n)}{\sum_{S \in \{\text{Reference Summaries}\}} \sum_{\text{gram}_n \in S} Count(gram_n)}$$
(22)

where *n* stands for the length of the n-gram, i.e., $gram_n$. *Count_{match}(gram_n*) is the maximum number of n-grams cooccurring in the candidate summary and the reference summaries. We compared the summarization methods using aspect ranking results as in Eq. (21) with the methods using the traditional scoring function in Eq. (15). In particular, four methods were evaluated: SR and SR_AR, i.e., Sentence Ranking [29] with the traditional scoring function and the proposed function based on Aspect Ranking, respectively; GB and GB_AR, i.e., Graph-based method with the traditional and proposed scoring functions, respectively. The tradeoff parameters λ_1 , λ_2 , α , and β were empirically set to 0.5, 0.5, 0.6, and 0.4, respectively. Here, we reported summarization performance in terms of ROUGE-1 and ROUGE-2 corresponding to unigrams and bigrams, respectively.

Table IV shows the ROUGE-1 performance on each product as well as the average ROUGE-1 over all the 21

TABLE 4 Performance of Extractive Review Summarization in Terms of ROUGE-1

Product	SR	SR_AR	GB	GB_AR
AQUOS	0.21	0.22	0.20	0.21
BlackBerry	0.22	0.24	0.23	0.26
Brother	0.18	0.20	0.19	0.22
Canon EOS	0.21	0.22	0.21	0.24
Epson	0.21	0.22	0.21	0.22
Fujifilm	0.23	0.24	0.23	0.25
Garmin	0.21	0.21	0.21	0.22
GoPro	0.23	0.24	0.23	0.25
Handycam	0.20	0.22	0.22	0.25
iPhone 3GS	0.19	0.21	0.18	0.22
iPod Touch	0.21	0.21	0.21	0.23
MacBook	0.16	0.18	0.18	0.23
Nokia 5800	0.15	0.17	0.17	0.18
Nokia N95	0.24	0.24	0.24	0.27
Officejet	0.21	0.22	0.21	0.22
Panasonic	0.21	0.22	0.22	0.24
Samsung	0.21	0.23	0.20	0.24
Samsung TV	0.27	0.27	0.26	0.28
Sony NWZ	0.23	0.25	0.23	0.24
Tomtom	0.24	0.25	0.23	0.24
Vizio	0.25	0.27	0.26	0.27
Average	0.21	0.23	0.22	0.24

T-Test, p-values < 0.05.

products, while Table V provides the corresponding performance in terms of ROUGE-2. From these results, we can obtain the following observations:

- By exploiting aspect ranking, the proposed SR_AR and GB_AR approaches outperforms the traditional SR and GB methods, respectively. In particular, SR_AR obtains performance improvements over SR by around 9.5% and 16% in terms of average ROUGE-1 and ROUGE-2, respectively. GB_AR achieves around 9.1% and 22.6% improvements over GB in terms of average ROUGE-1 and ROUGE-2, respectively;
- Consider the ROUGE-1 results in Table IV, SR_AR performs better than SR on 17 out of the 21 products and performs the same on the remaining four products, while GB_AR outperforms GB on all the 21

TABLE 5 Performance of Extractive Review Summarization in Terms of ROUGE-2

Product	SR	SR_AR	GB	GB_AR
AQUOS	0.031	0.033	0.029	0.030
BlackBerry	0.061	0.065	0.068	0.071
Brother	0.037	0.045	0.038	0.048
Canon EOS	0.055	0.061	0.065	0.075
Epson	0.041	0.043	0.042	0.047
Fujifilm	0.057	0.058	0.056	0.064
Garmin	0.032	0.038	0.032	0.041
GoPro	0.075	0.079	0.071	0.081
Handycam	0.060	0.070	0.061	0.104
iPhone 3GS	0.057	0.087	0.048	0.092
iPod Touch	0.051	0.058	0.057	0.058
MacBook	0.031	0.034	0.034	0.042
Nokia 5800	0.031	0.034	0.032	0.035
Nokia N95	0.071	0.081	0.091	0.103
Officejet	0.050	0.064	0.056	0.061
Panasonic	0.040	0.053	0.043	0.061
Samsung	0.051	0.061	0.061	0.071
Samsung TV	0.054	0.072	0.055	0.075
Sony NWZ	0.041	0.045	0.050	0.063
Tomtom	0.072	0.077	0.074	0.075
Vizio	0.050	0.068	0.055	0.070
Average	0.050	0.058	0.053	0.065

T-Test, p-values < 0.05.

products. For the ROUGE-2 results in Table V, SR_AR and GB_AR achieve better performance on all the 21 products compared to SR and GB, respectively;

• The graph-based methods, i.e., GB_AR and GB, obtain slight performance improvements compared to the corresponding sentence ranking methods, i.e., SR_AR and SR.

In summary, the above results demonstrate the capacity of aspect ranking in improving extractive review summarization. With the help of aspect ranking, the summarization methods can generate more informative summaries consisting of consumer reviews on the most important aspects. Table VI illustrates sample summaries of the product *Sony Handycam Camcorder*. We can see that the summaries from the methods using aspect ranking, i.e. SR_AR and GB_AR, contain consumer comments on the important aspects, such as *"easy to use*", and are more informative than those from the traditional methods.

5 RELATED WORKS

In this section, we review existing works related to the proposed product aspect ranking framework, and the two evaluated real-world applications. We start with the works on aspect identification. Existing techniques for aspect identification include supervised and unsupervised methods. Supervised method learns an extraction model from a collection of labeled reviews. The extraction model, or called extractor, is used to identify aspects in new reviews. Most existing supervised methods are based on the sequential learning (or sequential labeling) technique [18]. For example, Wong and Lam [36] learned aspect extractors using Hidden Marköv Models and Conditional Random Fields, respectively. Jin and Ho [11] learned a lexicalized HMM model to extract aspects and opinion expressions, while Li et al. [16] integrated two CRF variations, i.e., Skip-CRF and Tree-CRF. All these methods require sufficient labeled samples for training. However, it is time-consuming and labor-intensive to label samples. On the other hand, unsupervised methods have emerged recently. The most notable unsupervised approach was proposed by Hu and Liu [12]. They assumed that product aspects are nouns and noun phrases. The approach first extracts nouns and noun phrases as candidate aspects. The occurrence frequencies of the nouns and noun phrases are counted, and only the frequent ones are kept as aspects. Subsequently, Popescu and Etzioni [28] developed the OPINE system, which extracts aspects based on the KnowItAll Web information extraction system [8]. Mei et al. [22] utilized a probabilistic topic model to capture the mixture of aspects and sentiments simultaneously. Su et al. [32] designed a mutual reinforcement strategy to simultaneously cluster product aspects and opinion words by iteratively fusing both content and sentiment link information. Recently, Wu et al. [37] utilized a phrase dependency parser to extract noun phrases from reviews as aspect candidates. They then employed a language model to filter out those unlikely aspects. After identifying aspects in reviews, the next task is aspect sentiment classification, which determines the orientation of sentiment expressed on each aspect. Two major approaches for aspect sentiment classification include lexicon-based and supervised learning

approaches. The lexicon-based methods are typically unsupervised. They rely on a sentiment lexicon containing a list of positive and negative sentiment words. To generate a high-quality lexicon, the bootstrapping strategy is usually employed. For example, Hu and Liu [12] started with a set of adjective seed words for each opinion class (i.e., positive or negative). They utilized synonym/antonym relations defined in WordNet to bootstrap the seed word set, and finally obtained a sentiment lexicon. Ding et al. [6] presented a holistic lexicon-based method to improve Hu's method [12] by addressing two issues: the opinions of sentiment words would be content-sensitive and conflict in the review. They derived a lexicon by exploiting some constraints. On the other hand, the supervised learning methods classify the opinions on aspects by a sentiment classifier learned from training corpus [25]. Many learning based models are applicable, such as Support Vector Machine (SVM), Naive Bayes and Maximum Entropy (ME) model etc. More comprehensive literature review of aspect identification and sentiment classification can be found in [20]. As aforementioned, a product may have hundreds of aspects and it is necessary to identify the important ones. To our best knowledge, there is no previous work studying the topic of product aspect ranking. Wang et al. [34] developed a latent aspect rating analysis model, which aims to infer reviewer's latent opinions on each aspect and the relative emphasis on different aspects. This work concentrates on aspect-level opinion estimation and reviewer rating behavior analysis, rather than on aspect ranking. Snyder and Barzilay [31] formulated a multiple aspect ranking problem. However, the ranking is actually to predict the ratings on individual aspects.

Document-level sentiment classification aims to classify an opinion document as expressing a positive or negative opinion. Existing works use unsupervised, supervised or semi-supervised learning techniques to build documentlevel sentiment classifiers. Unsupervised method usually relies on a sentiment lexicon containing a collection of positive and negative sentiment words. It determines the overall opinion of a review document based on the number of positive and negative terms in the review. Supervised method applies existing supervised learning models, such as SVM and Maximum entropy (ME) etc. [25], while semisupervised approach exploits abundant unlabeled reviews together with labeled reviews to improve classification performance [30], [33]. The other related topic is extractive review summarization, which aims to condense the source reviews into a shorter version preserving its information content and overall meaning. Extractive summarization method forms the summary using the most informative sentences and paragraphs etc. selected from the original reviews. The most informative content generally refers to the "most frequent" or the "most favorably positioned" content in exiting works. The two widely used methods are the sentence ranking and graph-based methods [10]. In these works, a scoring function was first defined to compute the informativeness of each sentence. Sentence ranking method [29] ranked the sentences according to their informativeness scores and then selected the top ranked sentences to form a summary. Graph-based method [7] represented the sentences in a graph, where each node

TABLE 6

Sample Extractive Summaries of the Product Sony Handycam Camcorder

Method	Summary
	Even though this product was \$20 more expensive than a competitor's price, I chose B&H due to the number of positive feedback ratings for the
	product purchased - and the free shipping option. I would recommend this merchant to anyone looking for a reliable and quality place to buy audio
SR	/ video equipment. You guys have a great reputation and continually back it up with great prices, outstanding inventory, on time delivery of
	merchandise, and great customer service.
	Great service, good prices, true winner in my estimation. Then Sony came out with this cheaper high spec and light weight (mass wise) little cam
	corder, What joy!, small, good in low light, great sound and so easy to use you do not even have to read the manual to understand all the features
SR_AR	(love this!). A reliable and cost competitive supplier of photographic equipment. The Web site is easy to use, easy to navigate, delivery of
	merchandise on time and ordering online is convenient. My husband and friends like your store and products and always recommend to others.
	Always, professional, always courteous, always prompt, and always the best service and products around! Keep up the great customer service, and
	I order from this merchant when they have the items I need. It is the best site to find the best equipment at the best price. Easy access to site and
GB	information, prompt and free delivery of merchandise. Overall, good merchant, no complaints. excellent service on line Excellent on followup for
	customer satisfaction; would recommend others using their service. I have recommended you a few times and will still do it !! thanks
	the website is good amd the most reliable place to purchase camera products the selections are great, the prices are very good. The descriptions
	were very detailed and pictures accurate and helpful. I have always been able to rely on them for good business practices and timely service. The
GB_AR	service provided is great! Very reliable, easy to use, easy to navigate. I would recommend this merchant to anyone looking for a reliable and
	quality place to buy audio / video equipment.

corresponds to a sentence and each edge characterizes the relation between two sentences. A random walk was then performed over the graph to discover the most informative sentences, which were in turn used to compose a summary.

6 CONCLUSION

In this article, we have proposed a product aspect ranking framework to identify the important aspects of products from numerous consumer reviews. The framework contains three main components, i.e., product aspect identification, aspect sentiment classification, and aspect ranking. First, we exploited the Pros and Cons reviews to improve aspect identification and sentiment classification on free-text reviews. We then developed a probabilistic aspect ranking algorithm to infer the importance of various aspects of a product from numerous reviews. The algorithm simultaneously explores aspect frequency and the influence of consumer opinions given to each aspect over the overall opinions. The product aspects are finally ranked according to their importance scores. We have conducted extensive experiments to systematically evaluate the proposed framework. The experimental corpus contains 94,560 consumer reviews of 21 popular products in eight domains. This corpus is publicly available by request. Experimental results have demonstrated the effectiveness of the proposed approaches. Moreover, we applied product aspect ranking to facilitate two real-world applications, i.e., document-level sentiment classification and extractive review summarization. Significant performance improvements have been obtained with the help of product aspect ranking.

ACKNOWLEDGMENTS

This work was supported by the "NExT Research Center" funded by MDA, Singapore, under Grant; WBS:R-252-300-001-490.

REFERENCES

- J. C. Bezdek and R. J. Hathaway, "Convergence of alternating optimization," J. Neural Parallel Scientific Comput., vol. 11, no. 4, pp. 351–368, 2003.
- [2] C. C. Chang and C. J. Lin. (2004). Libsvm: A library for support vector machines [Online]. Available: http://www.csie.ntu.edu.tw/~cjlin/libsvm/

- [3] G. Carenini, R. T. Ng, and E. Zwart, "Multi-document summarization of evaluative text," in *Proc. ACL*, Sydney, NSW, Australia, 2006, pp. 3–7.
- [4] China Unicom 100 Customers iPhone User Feedback Report, 2009.
- [5] ComScore Reports [Online]. Available:
- http://www.comscore.com/Press_events/Press_releases, 2011.
- [6] X. Ding, B. Liu, and P. S. Yu, "A holistic lexicon-based approach to opinion mining," in *Proc. WSDM*, New York, NY, USA, 2008, pp. 231–240.
- [7] G. Erkan and D. R. Radev,"LexRank: Graph-based lexical centrality as salience in text summarization," J. Artif. Intell. Res., vol. 22, no. 1, pp. 457–479, Jul. 2004.
- [8] O. Etzioni et al., "Unsupervised named-entity extraction from the web: An experimental study," J. Artif. Intell., vol. 165, no. 1, pp. 91–134. Jun. 2005.
- [9] A. Ghose and P. G. Ipeirotis, "Estimating the helpfulness and economic impact of product reviews: Mining text and reviewer characteristics," *IEEE Trans. Knowl. Data Eng.*, vol. 23, no. 10, pp. 1498–1512. Sept. 2010.
- [10] V. Gupta and G. S. Lehal, "A survey of text summarization extractive techniques," J. Emerg. Technol. Web Intell., vol. 2, no. 3, pp. 258–268, 2010.
- [11] W. Jin and H. H. Ho, "A novel lexicalized HMM-based learning framework for web opinion mining," in *Proc. 26th Annu. ICML*, Montreal, QC, Canada, 2009, pp. 465–472.
- [12] M. Hu and B. Liu, "Mining and summarizing customer reviews," in *Proc. SIGKDD*, Seattle, WA, USA, 2004, pp. 168–177.
- [13] K. Jarvelin and J. Kekalainen, "Cumulated gain-based evaluation of IR techniques," ACM Trans. Inform. Syst., vol. 20, no. 4, pp. 422–446, Oct. 2002.
- [14] J. R. Jensen, "Thematic information extraction: Image classification," in *Introductory Digit. Image Process.*, pp. 236–238.
- [15] K. Lerman, S. Blair-Goldensohn, and R. McDonald, "Sentiment summarization: Evaluating and learning user preferences," in *Proc. 12th Conf. EACL*, Athens, Greece, 2009, pp. 514–522.
- [16] F. Li et al., "Structure-aware review mining and summarization," in Proc. 23rd Int. Conf. COLING, Beijing, China, 2010, pp. 653–661.
- [17] C. Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in *Proc. Workshop Text Summarization Branches Out*, Barcelona, Spain, 2004, pp. 74–81.
- [18] B. Liu, M. Hu, and J. Cheng, "Opinion observer: Analyzing and comparing opinions on the web," in *Proc. 14th Int. Conf. WWW*, Chiba, Japan, 2005, pp. 342–351.
- [19] B. Liu, "Sentiment analysis and subjectivity," in Handbook of Natural Language Processing, New York, NY, USA: Marcel Dekker, Inc., 2009.
- [20] B. Liu, Sentiment Analysis and Opinion Mining. Mogarn & Claypool Publishers, San Rafael, CA, USA, 2012.
- [21] L. M. Manevitz and M. Yousef, "One-class SVMs for document classification," J. Mach. Learn., vol. 2, pp. 139–154, Dec. 2011.
- [22] Q. Mei, X. Ling, M. Wondra, H. Su, and C. X. Zhai, "Topic sentiment mixture: Modeling facets and opinions in weblogs," in *Proc. 16th Int. Conf. WWW*, Banff, AB, Canada, 2007, pp. 171–180.
- [23] B. Ohana and B. Tierney, "Sentiment classification of reviews using SentiWordNet," in *Proc. IT&T Conf.*, Dublin, Ireland, 2009.

- [24] G. Paltoglou and M. Thelwall, "A study of information retrieval weighting schemes for sentiment analysis," in Proc. 48th Annu.
- Meeting ACL, Uppsala, Sweden, 2010, pp. 1386–1395. [25] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment classification using machine learning techniques," in Proc. EMNLP, Philadelphia, PA, USA, 2002, pp. 79-86.
- [26] B. Pang, L. Lee, and S. Vaithyanathan, "A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts techniques," in Proc. ACL, Barcelona, Spain, 2004, pp. 271-278.
- [27] B. Pang and L. Lee, "Opinion mining and sentiment analysis," in Found. Trends Inform. Retrieval, vol. 2, no. 1-2, pp. 1-135, 2008.
- [28] A. M. Popescu and O. Etzioni, "Extracting product features and opinions from reviews," in Proc. HLT/EMNLP, Vancouver, BC, Canada, 2005, pp. 339–346.
- [29] D. Radev, S. Teufel, H. Saggion, and W. Lam, "Evaluation challenges in large-scale multi-document summarization," in Proc. ACL, Sapporo, Japan, 2003, pp. 375–382.[30] V. Sindhwani and P. Melville, "Document-word co-regularization
- for semi-supervised sentiment analysis," in Proc. 8th IEEE ICDM, Pisa, Italy, 2008, pp. 1025–1030.
- [31] B. Snyder and R. Barzilay, "Multiple aspect ranking using the good grief algorithm," in Proc. HLT-NAACL, New York, NY, USA, 2007, pp. 300–307.
- [32] Q. Su et al., "Hidden sentiment association in chinese web opinion mining," in Proc. 17th Int. Conf. WWW, Beijing, China, 2008, pp. 959-968.
- [33] L. Tao, Z. Yi, and V. Sindhwani, "A non-negative matrix tri-factorization approach to sentiment classification with lexical prior knowledge," in Proc. ACL/AFNLP, Singapore, 2009, pp. 244–252.
- [34] H. Wang, Y. Lu, and C. X. Zhai, "Latent aspect rating analysis on review text data: A rating regression approach," in Proc. 16th ACM SIGKDD, San Diego, CA, USA, 2010, pp. 168–176.
- [35] T. Wilson, J. Wiebe, and P. Hoffmann, "Recognizing contextual polarity in phrase-level sentiment analysis," in Proc. HLT/EMNLP, Vancouver, BC, Canada, 2005, pp. 347–354.
- [36] T. L. Wong and W. Lam, "Hot item mining and summarization from multiple auction web sites," in Proc. 5th IEEE ICDM, Washington, DC, USA, 2005, pp. 797-800.
- [37] Y. Wu, Q. Zhang, X. Huang, and L. Wu, "Phrase dependency parsing for opinion mining," in *Proc. ACL*, Singapore, 2009, pp. 1533–1541. [38] J. Yu, Z.-J. Zha, M. Wang, and T. S. Chua, "Aspect ranking:
- Identifying important product aspects from online consumer reviews," in Proc. ACL, Portland, OR, USA, 2011, pp. 1496–1505.
- [39] B. Zhang *et al.*, "Improving web search results using affinity graph," in *Proc. 28th SIGIR*, Salvador, Brazil, 2005, pp. 504–511.
 [40] Z. Zhang and B. Varadarajan, "Utility scoring of product
- reviews," in Proc. 15th CIKM, Arlington, TX, USA, 2006, pp. 51-57.
- [41] L. Zhao, L. Wu, and X. Huang, "Using query expansion in graphbased approach for query focused multi-document summarization," J. Inform. Process. Manage., vol. 45, no. 1, pp. 35-41, Jan. 2009.



Zheng-Jun Zha is a Professor with the Hefei Institute of Intelligent Machines, Chinese Academy of Sciences, Hefei, China. He previously worked as a Senior Research Fellow in the School of Computing, National University of Singapore (NUS). He received the B.E. degree in automation and the Ph.D. degree in pattern recognition and intelligent system from the University of Science and Technology of China, Hefei, China. His current research interests include user-generated content anal-

ysis, multimedia information retrieval, and social media analysis. He has published a series of book chapters, journal articles, and conference papers in these areas, including TOIS, TMM, TIP, ACM Multimedia, SIGIR, ACL, etc. He received the Best Paper Award in the ACM International Conference on Multimedia (ACM Multimedia) 2009, the Best Demo Runner-Up award in ACM Multimedia 2012, the Best Student Paper Award in ACM Multimedia 2013, and the Best Paper Award in the International Conference on Internet Multimedia Computing and Service 2013. He is a member of the IEEE



Jianxing Yu is a Ph.D. student in the School of Computing, National University of Singapore, Singapore. He received the B.E. degree from the Renmin University of China, Beijing, China, in 2008. His current research interests include sentiment analysis, text mining, and natural language processing. He is the recipient of numerous scholarships and awards.



Jinhui Tang is currently a Professor with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China.. He received the B.E. and Ph.D. degrees in July 2003 and 2008 respectively, both from the University of Science and Technology of China. His current research interests include large-scale multimedia search, social media mining, and computer vision. He has authored over 80 journal and conference papers in these areas. He serves as a Editorial Board Member of pat-

tern analysis and applications, multimedia tools and applications, information sciences, neurocomputing, a Technical Committee Member for about 30 international conferences, and a Reviewer for about 30 presticious international journals. Prof. Tang is a co-recipient of the Best Paper Award in ACM Multimedia 2007, PCM 2011, and ICIMCS 2011. He is a member of the IEEE.



Meng Wang is a Professor with the Hefei University of Technology, Hefei, China. He previously worked as a Senior Research Fellow in the School of Computing, National University of Singapore (NUS). He received the B.E. degree and the Ph.D. degree from the University of Science and Technology of China, Hefei, China. His current research interests include multimedia content analysis, search, mining, recommendation, and large-scale computing. He has authored more than 100 book chapters, jour-

nal, and conference papers in these areas. He received the Best Paper Awards in the 17th and 18th ACM International Conference on Multimedia and the Best Paper Award in the 16th International Multimedia Modeling Conference. He is a member of the IEEE.



Tat-Seng Chua is the KITHCT Chair Professor at the School of Computing, National University of Singapore, Singapore. He was the Acting and Founding Dean of the School during 1998-2000. His current research interests include multimedia information retrieval and social media analysis. In particular, his research focuses on the extraction, retrieval and question-answering (QA) of text, video, and live media arising from the Web and social networks. He is the Director of a multi-million-dollar joint Center (named NExT)

between NUS and Tsinghua University in China to develop technologies for live media search. He received the Ph.D. degree from the University of Leeds, U.K.

> For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.