

ReviewMiner: An Aspect-based Review Analytics System

Derek Wu, Hongning Wang
Department of Computer Science
University of Virginia
Charlottesville, VA 22904, USA
{dxw7za,hw5x}@virginia.edu

ABSTRACT

We develop an aspect-based sentiment analysis system named ReviewMiner. It analyzes opinions expressed about an entity in an online review at the level of topical aspects to discover each individual reviewer's latent opinion on each aspect as well as his/her relative emphasis on different aspects when forming the overall judgment of the entity. The system personalizes the retrieved results according to users' input preferences over the identified aspects, recommends similar items based on the detailed aspect-level opinions, and summarizes aspect-level opinions in textual, temporal and spatial dimensions. The unique multi-modal opinion summarization and visualization mechanisms provide users with rich perspectives to digest information from user-generated opinionated content for making informed decisions.

CCS CONCEPTS

•Information systems →Personalization; *Online analytical processing; Content ranking;*

KEYWORDS

Aspect-based sentiment analysis, review mining, personalization

1 INTRODUCTION

With the emergence and advancement of social media, more and more people freely express their opinions on all kinds of entities, such as products and services on the Internet. Such user-generated opinionated content is useful for other users to make informed decisions and for merchants to improve their services. However, despite abundant studies in opinion mining research [4, 7], there are few practical systems providing ordinary users with easy access to opinions at a fine-grained level of topical aspects. For example, most existing tools or systems have focused on overall sentiment classification in user reviews [6, 9, 16], but with solely a predicted overall rating it is still hard for a user to figure out whether the entity is of high quality in a specific aspect of his/her interest, or why it is better than other comparable entities.

To achieve a deeper and more detailed understanding of user-generated opinionated data, some recent works studied a new text mining problem called Latent Aspect Rating Analysis (LARA) [11, 14, 15, 17]. Given a set of reviews with only overall ratings, LARA aims to analyze opinions at the level of topical aspects to

discover each reviewer's latent rating on each aspect as well as the relative importance he/she has placed on different aspects when forming the overall judgment. Revealing the latent aspect ratings and weights in each individual review enables a wide range of important applications. For example, the identified latent ratings on different aspects immediately support aspect-based opinion summarization; aspect weights are directly useful for analyzing users' rating behaviors; and the combination of latent aspect ratings and weights support personalized ranking of entities by using only reviews from the reviewers who share similar aspect weights to those preferred by an individual user.

In this work, we develop a prototype system called ReviewMiner¹ based on the research in latent aspect rating analysis to grant such analytic power to the end users of this system. ReviewMiner not only provides basic search functions for the users to explore the analyzed entities and reviews in the system, but also personalizes the retrieved results according to the users' input preferences over the identified aspects, recommends similar entities based on the detailed aspect-level opinions, and summarizes aspect-level opinions in textual, temporal and spatial dimensions. The function of personalization and recommendation assists users in identifying results of interest and exploring alternative choices more efficiently. The unique multi-modal opinion summarization and visualization mechanisms provide the system users various perspectives to digest information from user-generated opinionated content for making more informed decisions. In addition, ReviewMiner also actively records users' interactive search behaviors in the system, including their input queries, hovered and clicked results, updates of aspect preferences, highlighted text, and votes of helpfulness on the retrieved review documents. The logged information is then utilized to analyze users' search intent and build accurate user models for assisting him/her in the future.

2 RELATED WORK

Due to the blooming of research in opinion mining and sentiment analysis [4, 7, 11, 14, 15, 17], various practical systems have been developed to analyze user-generated opinionated content. Liu et al. [10] built a prototype system called Opinion Observer to analyze and compare user opinions of competing products, where opinions are summarized by frequent text patterns extracted from pros/cons sections of user reviews. Jin et al. [6] developed the OpinionMiner system to identify opinion expressions in user reviews and classify them into positive and negative classes. Ku et al. [9] developed an opinion analysis system named CopeOpi, which extracts opinions about specific entities from the Web, summarizes the polarity and strength of these opinions, and tracks opinion variations over time. In addition to analyzing overall sentiment in user-generated content, systems that focus on aspect-level sentiment have also been built. The OpinionFinder system identifies subjective sentences

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '17, August 07-11, 2017, Shinjuku, Tokyo, Japan
© 2017 ACM. 978-1-4503-5022-8/17/08...\$15.00
DOI: 10.1145/3077136.3084148

¹<http://hcdm.cs.virginia.edu:8080/ReviewMiner>

and marks various aspects of the subjectivity in these sentences [16]. The OPINE system [12] identifies important product features from user reviews, their evaluation by reviewers, and their relative quality across products.

However, all these existing systems focus on aggregated opinions across reviewers at an overall or aspect level. None of them is able to identify an individual reviewer’s emphasis on different aspects, i.e., aspect weight. Our work aims to analyze both the aspect ratings and weights at the level of individual reviews [14, 15]. It enables system users to utilize such detailed opinionated data to perform complex analytic tasks, including opinion-based entity ranking, rated aspect summarization and comparison, and personalized entity recommendation.

3 SYSTEM DESIGN

There are four major components in our ReviewMiner system: 1) review document crawler, LARA analyzer and analyzed review repository; 2) query parser; 3) multi-modal user interface; and 4) interactive behavior logging system and user modeling. Figure 1 highlights the overview of the system. In the following, we will discuss the implementation details of each component accordingly.

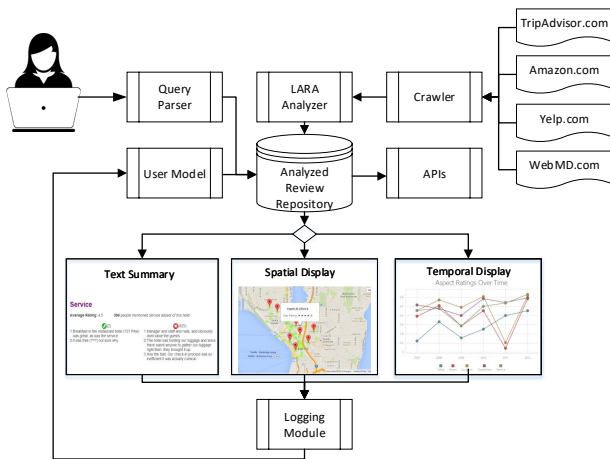


Figure 1: Overview of ReviewMiner system.

3.1 Crawler, Analyzer, and Review Repository

This component forms a pipeline to collect, analyze, and store the online opinionated review documents into a structured database.

• **Crawler:** ReviewMiner keeps crawling four types of opinionated review documents: 1) hotel reviews from TripAdvisor (www.tripadvisor.com); 2) product reviews from Amazon (www.amazon.com); 3) restaurant reviews from Yelp (www.yelp.com); and 4) medication reviews from WebMD (www.webmd.com). In product reviews, we focus on six subcategories including digital cameras, TVs, video surveillance systems, mobile phones, tablets, and laptops. Basic information about an item, e.g., name, image, overall ratings, and short descriptions (i.e., address for hotels, feature specifications for Amazon products, and usage for medications), is collected during crawling. For each review, information about its author, review date, title, and content is collected. Simple filtering is performed at the crawling stage: 1) reviews with fewer than three sentences are discarded; 2) items with fewer than five reviews are discarded.

The crawler is invoked periodically to capture the updates of user reviews on these four different review sites.

• **LARA Analyzer:** We implement the two-step algorithm proposed in [14] to perform aspect-based sentiment analysis on the crawled content. The choice of this solution is mainly due to its computational efficiency for real-time processing and the fixed types of entities analyzed in the system.

Specifically, keyword-based bootstrapping is performed for aspect segmentation, and latent rating regression is used for latent aspect rating and weight prediction. When performing aspect segmentation, we manually chose the number of aspects in each category of entities, and selected the most representative words as seed words for the bootstrapping-based aspect segmentation method. Due to the low aspect coverage in individual reviews (not all reviewers would talk about every aspect of an entity in their reviews), it is infeasible for us to infer the latent aspect ratings and weights of every single review document. As our solution, we first aggregate reviews under each item and estimate the item-level aspect ratings and weights in the system.

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

• **Analyzed Review Repository:** In order to ensure runtime efficiency of front-end execution, the aspect segments, ratings and weights for each item and review are precomputed and stored in a back-end relational database. In addition, to provide flexible search functions over all the analyzed items, keyword-based Lucene indices [1] are built over the item name and description fields for every category of entities.

3.2 Query Parser

Standard keyword search is supported by the inverted indices built on the fields of entity names and descriptions in each category. For example, users can type a specific (or partial) hotel names or locations as their query, such as “hotels in Chicago” to search for hotels in ReviewMiner. BM25 ranking algorithm is applied to both fields, and we give higher weight to query terms matching in the name field than that in the description field to emphasize the importance of name matching in search result ranking.

However, such a simple keyword-based query scheme cannot support users in explicitly expressing their complex information need over the search results. For example, if a user wants to find a hotel in downtown Seattle with good service and a price lower than \$200/night, he/she has to first find all hotels in the Seattle area (e.g., by querying “hotels in Seattle”), then manually filter out the irrelevant results by reading review content. To support these complex search intents in our query parser module, keyword-based aspect identification is also performed to achieve semantic interpretation of the queries.

The basic workflow of our query parser is as follows:

- (1) Segment input free text query into phrases using the Stanford NLP parser [2].
- (2) Classify the phrases into aspects by the learned aspect seed words: e.g., the phrase “around downtown area” will be assigned to the location aspect in hotel search.

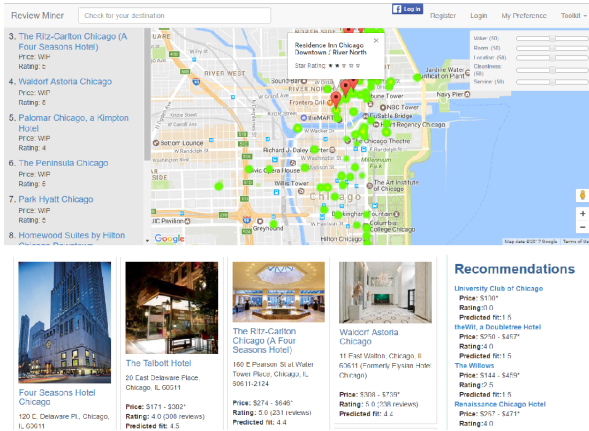


Figure 2: Entity search result page for the hotel domain. Users can directly issue natural language based queries in the search bar, navigate the map to select hotels in the region, or specify preferences over the identified topical aspects to rank the entities accordingly.

- Predict sentiment polarity of each phrase with respect to the identified aspect by the learned rating regression model, and map them into three categories of “low,” “medium” and “high.” For example, the query phrase “with excellent cleanliness condition” would be interpreted as expressing “high” requirement over “cleanliness” aspect.

After these query parsing steps, a user’s input query is compiled into a semantically structured format as $\{aspect \rightarrow specification\}$, which facilitates ReviewMiner’s ranking of the retrieved results. The identified entity name and description from the query are used to retrieve the candidate items from the inverted indices initially; and then each candidate is evaluated against the recognized aspect specification to estimate its relevance quality to the query. For example, if one query specifies “imperative” requirement on the value aspect but “optional” requirement on the service aspect, then the items with good value aspect ratings will be promoted over the items with low value aspect ratings but high service aspect ratings.

Besides the natural language text query input, users can also explicitly specify their preferences over the identified aspects via a dropdown menu on the system interface (see Figure 2). Higher weight will be given to users’ explicit input preferences in determining the final ranking of the retrieved results (in both entity search and review search).

3.3 Multi-modal User Interface

As the output of ReviewMiner, multi-modal result display is enabled by the detailed aspect-level analysis of review documents. In particular, ReviewMiner supports three types of user interaction interfaces: text-based opinion summary and comparison, spatial display of the retrieved results, and temporal display and comparison of the user’s specified items. Users can easily access any of these three interfaces when interacting with the system.

- Text-based Opinion Display:** The review text content from each retrieved item is segmented into aspects and highlighted with different colors for users to quickly digest the opinions (see Figure 3). In the aspect-based opinion summary, opinionated sentences are selected from all the associated review content and ordered according to their sentiment polarities. As a result, comparative

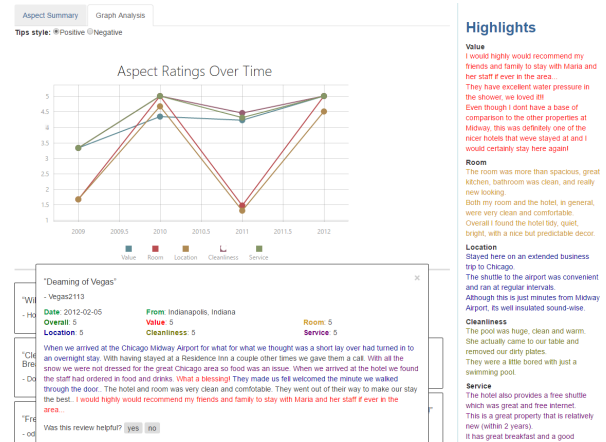


Figure 3: Review result page. In ReviewMiner, we segment reviews by sentences, assign them into corresponding aspects, and highlight them with different colors. We also visualize the results in temporal dimension and construct aspect-based summarization and comparison across entities.

summarization is enabled by listing the top ranked opinionated sentences by aspects across different items. With such functionality, users can easily navigate through the selected item candidates and make informed comparisons.

- Spatial Opinion Display:** For the hotel and restaurant reviews, ReviewMiner visualizes the opinions in the spatial dimension, which helps users quickly find out where “good” choices are located and explore comparative alternatives nearby (see Figure 2). We want to emphasize that although spatial visualization of the retrieved items has been adopted in many practical systems, e.g., TripAdvisor, all of those systems simply list the location of items on a map. From a user’s perspective, in order to assess the quality of candidate items, he/she still has to go to the detailed review page of every item. This leaves the spatial comparison of hotels difficult or impossible.

To solve this deficiency, an extra opinion-based heatmap layer [3] is added in ReviewMiner, representing the overall rating distribution of the retrieved entities at the target location. On the heatmap, areas of red color indicate regions with entities of higher overall ratings, compared to regions of light green color. The markers indicate the top-ranked entities in the visible area, with respect to the users’ aspect preferences. With the support of spatial opinion display, users no longer need to dig deep into detailed reviews for comparison; instead, the heatmap enables them to visually browse the area and thus greatly simplifies their decision making process.

- Temporal Opinion Display:** ReviewMiner also provides visualization of opinions in the temporal dimension by displaying the inferred aspect ratings and weights over time for each selected item (see Figure 3). Specifically, the reviews associated with each selected item are first grouped according to the year when they were published. The inferred aspect ratings were aggregated from the reviews accordingly. We chose year as the time unit to balance the number of reviews in each point and the total number of points for visualization. Based on such visualization of aspect ratings and mentions overtime, users can easily track the dynamics of reviewers’ opinions on this particular item and the development of sentiment towards it over time. In addition, the temporal opinion visualization also enables side-by-side comparison across items in temporal

dimension. This helps users to acquire more comprehensive and detailed assessment to compare the selected items.

3.4 Interactive Behavior Logging and User Modeling

ReviewMiner supports user registration and login in order to accurately keep track of individual users' information-seeking behaviors in the system. All of a user's actions in the system, including typing a query, clicking on a returned result, browsing the analyzed review page, updating his/her aspect preferences, highlighting review text, and voting on the helpfulness of a review, will be logged and analyzed to infer the user's underlying information need. All actions will be logged under the user's unique and anonymized system ID, if they are logged in to ReviewMiner; otherwise, actions will be logged under the user's IP address. Login is provided as a ReviewMiner account registered with an email address and password, or via single sign-on using Facebook Login. Facebook Login grants access to a user's friend list, which makes it possible for ReviewMiner to factor in the user's friends' search and browsing behaviors in this system, i.e., collaborative ranking and recommendation.

The system maintains a unique profile of aspect preferences for each registered user under each category, and updates the profile when the user explicitly inputs his/her aspect preferences or clicks on a result in the search result page. The employed profile updating strategy follows the ranking model adaptation method developed in [13]. In particular, ReviewMiner keeps track of the dwell time of users' browsing behaviors, and treats a clicked result page with a dwell time longer than 30 seconds as positive feedback [5], and those skipped [8] or quickly left (dwell time less than 5 seconds) as negative feedback. The click feedback is represented by a vector of the inferred latent aspect ratings for each corresponding item. Such input is fed into the personalized ranking model adaptation module to estimate the user's aspect preference. For finer-grained behavior analysis, ReviewMiner implements cursor tracking and logs text if the user hovers their cursor over it for 2 seconds or longer. These logs are analyzed as indicators of user interest: for instance, if a user frequently pauses over text pertaining to a certain aspect, ReviewMiner can assume the user weights it heavily.

As a result, in ReviewMiner, we have multiple criteria to rank the retrieved items: 1) a user's explicitly input aspect preference (w_u); 2) inferred aspect specification from the input query (w_p); 3) estimated personalized aspect specification from the user's interaction history (w_q). Those different aspect preferences are linearly combined to get the final aspect weights for candidate ranking,

$$s_i \leftarrow (\lambda_u w_u + \lambda_p w_p + \lambda_q w_q)^T r_i \quad (1)$$

where s_i is the final ranking score for item i , $\{\lambda_u, \lambda_p, \lambda_q\}$ represent the relative importance of those three types of aspect preferences, and r_i is inferred aspect rating vector for item i . In the detailed system implementation, we intentionally bias towards the user's explicitly input aspect preference w_u by setting $\lambda_u = 0.7$, and gives less importance to inferred aspect preference from the user's input query and interaction history, i.e., setting $\lambda_p = 0.2$ and $\lambda_q = 0.1$.

4 CONCLUSION

In this work, we developed a review analytic system named ReviewMiner for multi-modal opinion analysis and decision support. ReviewMiner performs aspect-based opinion analysis in textual,

spatial and temporal dimensions to enable users to digest the opinions conveyed in the review text content from different perspectives. ReviewMiner also automatically adapts to different users' aspect preferences based on their interaction history in the system to perform personalized result ranking and recommendations.

In addition, we want to emphasize that ReviewMiner not only provides easy access to massive opinion data to ordinary users, but also supports business analytics researchers to keep track of customer feedback and understand customer opinions of products and services. For example, ReviewMiner can recognize the inquired item's most commented aspects in its customer reviews, identify the corresponding relative emphasis the reviewers have expressed over those aspects, and track the temporal dynamics of user opinions and emphasis over those aspects. Such analysis can hardly be achieved in any other existing opinion mining or business analytics systems. As our future work, we plan to provide aspect-based opinion analysis APIs for third-party developers.

5 ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments. This paper is based upon work supported by the National Science Foundation under grant IIS-1553568.

REFERENCES

- [1] Andrzej Bia lecki, Robert Muir, Grant Ingersoll, and Lucid Imagination. 2012. Apache lucene 4. In *SIGIR 2012 workshop on open source information retrieval*. 17.
- [2] Marie-Catherine De Marneffe, Bill MacCartney, Christopher D Manning, and others. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC*, Vol. 6. 449–454.
- [3] Google Developers. 2017. Google Maps API. <https://developers.google.com/maps/>. (2017).
- [4] Ann Devitt and Khurshid Ahmad. 2007. Sentiment polarity identification in financial news: A cohesion-based approach. In *ACL*, Vol. 7. 1–8.
- [5] Steve Fox, Kuldeep Karnawat, Mark Mydland, Susan Dumais, and Thomas White. 2005. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)* 23, 2 (2005), 147–168.
- [6] Wei Jin, Hung Hay Ho, and Rohini K Srihari. 2009. OpinionMiner: a novel machine learning system for web opinion mining and extraction. In *Proceedings of the 15th ACM SIGKDD*. ACM, 1195–1204.
- [7] Nitin Jindal and Bing Liu. 2006. Identifying comparative sentences in text documents. In *Proceedings of the 29th ACM SIGIR*. ACM, 244–251.
- [8] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th ACM SIGIR*. ACM, 154–161.
- [9] Lun-Wei Ku, Hsiu-Wei Ho, and Hsin-Hsi Chen. 2009. Opinion mining and relationship discovery using CopeOpi opinion analysis system. *Journal of the American Society for Information Science and Technology* 60, 7 (2009), 1486–1503.
- [10] Bing Liu, Mingqing Hu, and Junsheng Cheng. 2005. Opinion observer: analyzing and comparing opinions on the web. In *Proceedings of the 14th WWW*. ACM, 342–351.
- [11] Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 165–172.
- [12] Ana-Maria Popescu and Orena Etzioni. 2007. Extracting product features and opinions from reviews. In *Natural language processing and text mining*. Springer, 9–28.
- [13] Hongning Wang, Xiaodong He, Ming-Wei Chang, Yang Song, Ryen White, and Wei Chu. 2013. Personalized Ranking Model Adaptation for Web Search. In *Proceedings of the 36th Annual ACM SIGIR Conference*. ACM.
- [14] Hongning Wang, Yue Lu, and Chengxiang Zhai. 2010. Latent aspect rating analysis on review text data: a rating regression approach. In *Proceedings of the 16th ACM SIGKDD*. ACM, 783–792.
- [15] Hongning Wang, Yue Lu, and Chengxiang Zhai. 2011. Latent aspect rating analysis without aspect keyword supervision. In *Proceedings of the 17th ACM SIGKDD*. ACM, 618–626.
- [16] Theresa Wilson, Paul Hoffmann, Swapna Somasundaran, Jason Kessler, Janyce Wiebe, Yejin Choi, Claire Cardie, Ellen Riloff, and Siddharth Patwardhan. 2005. OpinionFinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP*. Association for Computational Linguistics, 34–35.
- [17] Yao Wu and Martin Ester. 2015. Flame: A probabilistic model combining aspect based opinion mining and collaborative filtering. In *Proceedings of the 8th ACM WSDM*. ACM, 199–208.