TipMe: Personalized advertising and aspect-based opinion mining for users and businesses

Dimitris Proios
Department of Informatics and Telematics
Harokopio University of Athens
Athens, Greece
it21034@hua.gr

Magdalini Eirinaki
Computer Engineering Department
San Jose State University
San Jose, CA, USA
magdalini.eirinaki@sjsu.edu

Iraklis Varlamis
Department of Informatics and Telematics
Harokopio University of Athens
Athens, Greece
varlamis@hua.gr

***Abstract*—Online advertisements are a major source of profit and customer attraction for web-based businesses. In a successful advertisement campaign, both users and businesses can benefit, as users are expected to respond positively to special offers and recommendations of their liking and businesses are able to reach the most promising potential customers. The extraction of user preferences from content provided in social media and especially in review sites can be a valuable tool both for users and businesses.**

In this paper, we propose a model for the analysis of content from product review sites, which considers in tandem the aspects discussed by users and the opinions associated with each aspect. The model provides two different visualizations: one for businesses that uncovers their weak and strong points against their competitors and one for end-users who receive suggestions about products of potential interest. The former is an aggregation of aspect-based opinions provided by all users and the latter is a collaborative filtering approach, which calculates user similarity over a projection of the original bipartite graph (user-item rating graph) over a content-based clustering of users and items. The model takes advantage of the feedback users give to businesses in review sites, and employ opinion mining techniques to identify the opinions of users for specific aspects of a business. Such aspects and their polarity can be used to create user and business profiles, which can subsequently be fed in a clustering and recommendation process.

We envision this model as a powerful tool for planning and executing a successful marketing campaign via online media. Finally, we demonstrate how our prototype can be used in different scenarios to assist users or business owners, using the Yelp challenge dataset.

Keywords*—*advertising, personalization, targeted ads, recommender, Yelp

I. INTRODUCTION

The advance of social networking services and the popularity of product or business reviewing sites, has led web users to an era of information abundance. While users may easily sort related searches by average rating or price, this is not sufficient in order to make an informed decision. Users need to explore the product or business further, by reading related reviews, product/business information etc. In "targeted advertising" product promotions reach selected consumers based on their demographics, psychographics and behavior (e.g. purchase history). However, most of the times, the promotions are purely content- or context-based and do not take into consideration the detailed preferences of each individual user. We use the term "***personalized advertising***" in order to describe our approach for locating the most promising potential buyers for a product or service based on review and rating information they provided for similar products or services and on the reviews and ratings provided by other buyers.

The same content of user reviews for products and services may allow businesses to understand the market interests, act proactively and position their products where there is still place in the market. This can be done by understanding: i) how potential clients are grouped, ii) what are the specific interests of each group, iii) how competitive products and services compare to each other, iv) what are their pros and cons and what is their impact to the end-users etc. This allows companies to focus on user groups that will most probably respond to their campaigns, to highlight the aspects of products that will attract users" attention and to improve those product aspects that turn away their customers.

The model that we propose in this paper aims on the in-depth analysis of product reviews and the combination of aspect based opinion mining and personalized recommendations. Our objective is to help both parties involved in a search, namely the business and the end user. Our hypothesis is that the advertisement placement is more successful when it's targeted to each individual user, given their implicitly disclosed preferences, as well as some more traditional signals such as demographics, location, etc. This draws from the fact that different people value different aspects of the same product/business/service. For example, when searching for a digital camera, one might be interested in the price and size, whereas another user may value the ease of use. Similarly, when searching for a good Italian restaurant, one user might value the ambience and wine list of a place, while another might prefer restaurants that are family-friendly. Ideally, businesses would like to target customers who value the qualities the businesses excel in and to highlight the aspects of value to the user in a personalized advertisement context, and at the same time the customers will benefit in receiving such promotions by being recommended relevant alternatives that perfectly match their criteria.

Product aspect evaluation information can be derived by the comments and reviews users leave online. Using text and opinion mining, the aspects of each review as well as the related polarity can be identified and used to form user and item (i.e. business) profiles. Using these profiles, we subsequently match users to businesses, using the content of their reviews (content-based recommendations), the similarity among them (user- and item-based collaborative filtering), as well as clustering techniques to group similar users or businesses together.

We envision this system as a powerful tool to plan and execute a successful marketing campaign via online media. We demonstrate our prototype using the ***Yelp*** challenge dataset. More specifically, we take advantage of the "tips" that Yelp users provide about businesses. Tips are short text comments that mention the good or bad aspects of a business with the use of adjectives that expose good or bad opinion. For example, when a tip for a restaurant is "Great Italian food, mediocre wine... :(" we can easily draw that "Italian food" is the good aspect of the restaurant, whereas "wine" is the bad aspect. In this case, we will be able to suggest this restaurant to another user, who is primarily interested in Italian food but disappointed from the restaurants he/she has already visited. We will also be able to address this tip to the company in order to improve the wine selection and to its competitors in order for them to improve their Italian cuisine recipes.

In this paper, we present a model that supports aspect-based opinion mining and personalized advertisements, along with our ongoing work on developing a user-friendly prototype system that can serve as a useful tool to both users and businesses. The rest of the paper is organized as follows: in section 2 we briefly describe related works on targeted advertising based on user reviews. In section 3, we describe the proposed model and process for collecting, analyzing and exploiting user reviews in order to improve advertisement targeting. In section 4, we provide the details of our prototype application on the Yelp Dataset Challenge and in section 5 we demonstrate some usage scenarios of the prototype application. Finally, section 6 summarizes the next steps of this work.

II. RELATED WORK

Many interesting works exist that focus on extracting the opinions from the customer reviews [8]. One line of work focuses on identifying aspects and their polarity and using them as an additional tool in representing the semantic orientation of a review. This can be done at the word [2], [3] or the category level [10]. This information can then be used to provide intuitive visualizations of the data. Several applications in the analysis of hotel reviews [1], [19] recognize the usability of features/aspects in supporting user decision and suggest useful visualizations of customers' reviews. however, they do not take advantage of this information in order to personalize recommendations. Only recently, researchers in social media analytics and recommendation engines have started employing the power of aspects and the associated opinion scores in the recommendation process. McAuley and Leskovec [11] as well as Ling et al. [7] both used LDA-based models to incorporate latent topics found in the review text in the recommendation process, using Non-negative Matrix Factorization and a Gaussian mixture priors model for the ratings. Diao et al. [20] use partial scores under different aspects in reviews, in combination with ratings and other metadata in a hybrid (content and collaborative filtering) approach, to generate movie recommendations. Having a different objective, our previous work [15], [18] uses the sentiment scores of aspects in reviews to generate personalized restaurant review recommendations.

The work presented here builds on the same foundation of forming user and business profiles using the polarity of their opinions expressed for products" aspects in user-provided content. Instead of the text reviews, we focus on the tips provided by Yelp users. Sklar and Concepcion [17] also focused on tips instead of reviews. However, their objective was to recommend the right tips to the right people via the Foursquare platform, by taking into consideration the timeliness of user-provided tips and the users' tastes and social connections. In a more similar case, the detection of sentiment polarity to the various aspects present in a text, is found to improve content oriented advertisement, since users responded more to ads that matched the positive aspects within a page, instead of those matching any randomly selected aspect of the whole text [5].

So far, and given the information that is disclosed by the industries that use ad targeting, such as Google[1] or Facebook[2], we can infer that there are several common approaches in ad targeting, that take into consideration one or more of the following: location, demographics, interests (as implied by groups joined/apps used in a social network application, keywords used in searches, etc.), (online) activity, cookies (of previously visited websites) and context (device used, web page visited etc.).

In this paper, we bridge together aspect based opinion mining and recommendation systems areas and employ intuitive visualizations of the knowledge extracted. The proposed tool can be used both as a recommender system for the end-user/consumer, which recommends products and in the same time highlights the aspects that are highly valued by the user and as a market analysis tool for the businesses being reviewed, which can get an overview of their pros and cons, compare with their competitors and further drill-down to the opinions of their targeted user groups.

III. PROPOSED MODEL

The main idea behind this work is to take advantage of the comments and reviews provided by the users of a reviewing site in order to identify important factors in their decision making process. The analysis of this content will reveal the aspects of interest to each individual user and the aspects that influenced, positively or negatively, her decision about a product or business. We envision and propose a system that will exploit such information and can benefit both end users, and businesses as outlined here.

Assuming that *the target of the advertisement is a user* that expressed an opinion about a business (e.g. said that a restaurant had a "tasteless white wine"), the system will provide her with:

- more aspects of the same business using all other users' reviews, e.g. this restaurant's white wine is definitely a bad choice, since it was rated negatively by more users, but the restaurant has great "beer on tap", or has a delicious "cheese selection",
- more aspects of the same business using the reviews of users with similar interests, e.g. the users that disliked the restaurant's white wine have suggested the "red wine" or liked the "meat dishes",

[1] https://support.google.com/adsense/answer/9713?hl=en

[2] https://www.facebook.com/business/a/online-sales/ad-targeting-details

- competitive businesses and their (positive or negative) aspects as derived from the reviews of all other users, e.g. suggest another restaurant in the area that has been ranked high for its "white wine".
- competitive businesses and their (positive or negative) aspects as derived from the reviews of users with similar interests, e.g. if the user usually reviews italian restaurants, then "italian food" is on the user's interest so we will advertise only reviews for "Italian restaurants" with "good white wine".

The first two suggestions aim in persuading the user to give a second chance to the business, by ordering the positively rated dishes or finding aspects that she probably missed in a first visit. The last two business recommendations will provide alternative options for the user, which are expected to improve her customer experience.

Similarly, if we assume that *the target of the service is a business*, then the system will:

- provide business owners with the aspects that users commented about their business. This information can be valuable for the business owners since they will be aware of their strengths and weaknesses and target their advertisement campaigns to the appropriate users (i.e. users who value the qualities they excel in). A further analysis of users' demographics will allow business owners to run personalized campaigns or offer coupons/promotions to specific user groups.
- provide business owners with the list of aspects users mentioned in their reviews of the business and its competitors. Using such a list, the business owners will be able to identify opportunities on aspects that are of high interest among their target user group but are rated negatively/not rated in their business reviews and consequently improve their products and services.

A. Recommender workflow

In our problem setting, we assume a set of users U that provide reviews for a set of businesses B. In this work, we focus on the content and ignore information concerning the freshness of the review, other users' rating of the review (available, for example, in reviewing sites such as *Epinions*[3] or reviews at *c|net*[4]), etc. More formally, each review r_k in the set of reviews R ($r_k \in R$) is a text provided by a user for a business and is represented by a triplet as follows:

$$r_k = \prec u_i, b_j, text_k \succ : u_i \in U, b_i \in B$$

The term "aspect" in the opinion mining literature is usually connected with the attributes of the item being reviewed and corresponds to a noun, a set of nouns or a set of verbs and nouns. For example, the aspects extracted from a restaurant review can be the "service", the "food quality", the "food taste" etc. Usually aspects are accompanied by adjectives or adverbs which are located in the same clause and provide a positive, neutral or negative orientation to the aspect, e.g., "reluctant service", "the worst food" etc. In bibliography [9], this is mapped to a pair of <head term, modifier> where "head term" refers to an aspect and "modifier" is the related opinion, e.g. <service, reluctant>, <food, worst>. With the use of an opinion lexicon that maps modifiers to positive or negative values, it is possible to find the polarity and intensity of a review for an aspect.

The textual analysis of the review leads to the detection of aspects with associated opinion values. The information flow for this aspect-based opinion extraction process is depicted in Figure 1. The result of this process is the representation of aspects that appear in the text of the review $r_k \in R$ in the form of triplet as follows:

$$r_k = \prec u_i, b_j, \prec a_l, s_l, v_l \succ \succ$$

where a_l is the aspect term (e.g. a noun or verb extracted from the review and is associated with a positive or negative opinion), s_l is the semantic polarity (i.e. $\pm$) of the opinion and v_l is the value associated with this polarity. The value of the polarity increases when highly polarized opinion terms appear in the same context with the aspect.

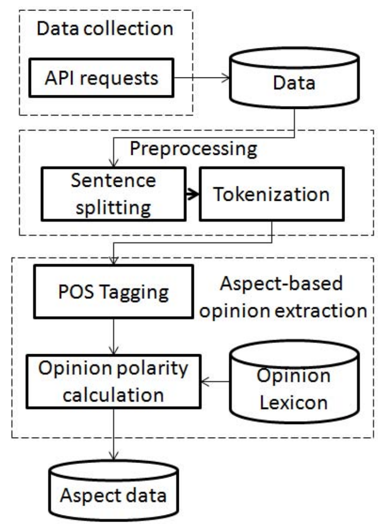

Fig. 1. The aspect-based opinion extraction process.

The exploitation of the collected aspect-based opinion data can be achieved by aggregating aspects and opinions by user or by reviewed item and consequently incorporate the information into the recommender system [14]. If we examine recommendations under the prism of collaborative filtering [4], our aim is to associate user-provided aspects and values with the interests of each user and consequently use these implicit interests to calculate user similarity or item similarity and thus improve recommendations. Similarly, under the content-based filtering prism, the aspects provided by a user and the aspects assigned to an item can be used to find the most appropriate matches. This can be easily implemented by querying the aspects database for each individual user. Finally, if we decide to follow a hybrid recommendation strategy, we will be able to group users or items based on their matching aspects, or any other content-based information (e.g. geo-location information, demographics, explicitly denoted interests etc.) and consequently make suggestions using collaborative filtering or content matching within groups. Figure 2 summarizes these approaches, which we further analyze in what follows.

[3] http://www.epinions.com
[4] http://www.cnet.com/reviews/

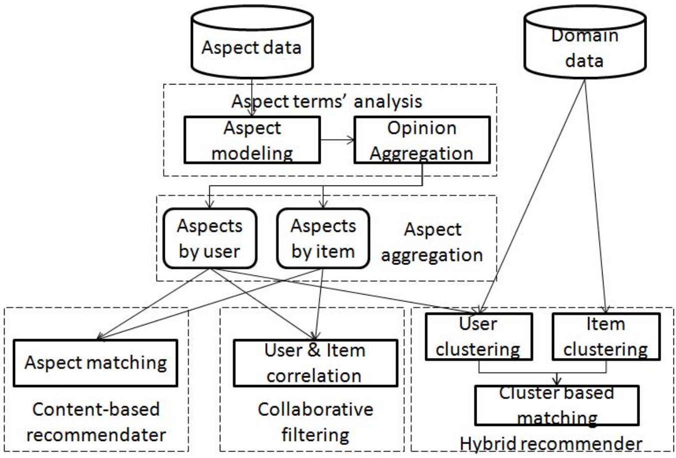

Fig. 2. The analysis of the aggregated aspect-based opinions and the recommendation strategies.

B. Aspect and opinion extraction

The first step in the aspect-based analysis of reviews is to identify the most interesting aspects and cluster them in more coarse-grained aspects, which we call "dominant aspects" in the following. This is usually performed using Topic-Modeling Techniques (e.g. Probabilistic Latent Semantic Analysis (PLSA) [9], Factorised Latent Dirichlet Allocation [13]) or hierarchical clustering [16], [6]. The result of these approaches is a reduced set of dominant aspects and a mapping of all aspect terms found in the reviews to these dominant aspects. As a result, from the analysis of the review text of the review r_k we get a mapping to the set of dominant aspects DA,

$$r_k = \prec u_i, b_j, \prec da_m, s_m, v_m \succ_{m=1..M} \succ$$

where da_m is one of the M dominant aspects found in the previous step.

C. Aspect aggregation

The next step is to aggregate the opinionated aspects for each user, and for each item, in order to get a more concrete image of the user's interests or the item's strong (or weak) aspects. The result for a user u_i will be a set of triplets:

$$interests_{u_i} = \prec da_m, s_{im}, v_{im} \succ_{m=1..M}$$

where v_{im} is the sum of all the values provided by user u_i for aspect da_m in her reviews and s_{im} is the sign of this sum (either + or -).

Similarly, we define the "impression" about item b_j as:

$$impression_{b_j} = \prec da_m, s_{jm}, v_{jm} \succ_{m=1..M}$$

where v_{jm} is the sum of all the values provided by all users for the aspect da_m of item b_j and s_{im} is the sign of this sum (either + or -).

D. Recommendation engine

1) Aspect matching: Based on the aggregated knowledge, we can simply search for items that match a user's interests. For example, if a user u_i is unhappy with an aspect da_m of a the items he reviewed as denoted by the respective <sign, value> combination (i.e. s_{im} is - and v_{im} is big), then we can recommend to the user anitem b_k, which has been evaluated (by other users) very positively for the same aspect s_{km} is + and (v_{km} is high). This simple analysis can be performed directly by querying the aspect database.

2) Aspect-based collaborative filtering: Matrix factorization is a technique which can be used for providing item recommendations using collaborative filtering. Several matrix factorization algorithms, such as SVD, FunkSVD or SVD++, have become very popular among recommender engines for social networks. All the proposed variations decompose the original bipartite matrix that contains user-provided ratings (or implicit user opinions) of items (Singular Value Decomposition - SVD) and then corresponds each item b_j with two latent vectors q_b, y_b and each user u_i with a latent vector p_u. The recommendations for user u_i are based on the predictions for user u_i of all the corresponding values.

Recently, authors in [12] used item and user categorization, achieved through clustering in order to improve SVD++ performance. This latter approach falls in the category of hybrid solutions. Aspect clustering, can be achieved using hierarchical clustering algorithms, such as those applied for taxonomy generation [16].

The general task of collaborative filtering is formally defined as follows: " Given the set of n users U, the set of m businesses B and a matrix $m \times n$ that contains what users like which businesses (or how users rate the businesses), then predict the rating of a user u_i for a business b_j."

When aspect-based analysis information is used instead of explicit ratings then we will have the set of n users U, the set of m businesses B and the set of all aspects A. Consequently the matrix dimensions will be $m \times (n \cdot l)$ where l is the number of aspects and the problem will be to predict the rating of a user u_i for the aspect a_k of a business b_j. In order to recommend businesses to the user, we aggregate the predicted ratings for all aspects of each business and rank businesses accordingly.

IV. Implementation

For testing the applicability of our ideas in a real case scenario, we created a prototype application using data from the *Yelp* platform. The dataset that we employed contains more than 40 thousands businesses and 250 thousands users. The users have provided check-in information (more than 30,000 entries), 1,125,458 reviews and 403,210 short text tips for different businesses.

For each business Yelp provides several features, such as their name, a set of business categories it belongs to, the longitude and latitude, address, opening times, special attributes, star rating etc. Users in Yelp form a social network since apart from their profile information they also have friends within the Yelp ecosystem. Users' opinions about businesses are expressed with tips and reviews. Tips are short, sentence-sized texts provided by a user for a business. They have a date and a number which represents the number of users that liked the tip. Reviews are pieces of text too but usually longer. They also have a date, a star rating (the user gave to the business), a number of votes from other users (depending on the fun, coolness and usefulness of the review). Apart from tips and reviews, the dataset contains check-in information that associates users with the businesses they have visited.

We focused our analysis only to information extracted from the tips, although our approach is easily extended to include the more detailed reviews. The motivation behind this decision is that tips: a) typically provide positive or negative opinions, b) judge only one or two aspects of the targeted business, focusing on the most important ones and c) are faster to process than full reviews. We also kept information about businesses and their categories and we identified users with their user_id in Yelp. The first step of our processing pipeline was to parse every tip with the Stanford Core NLP parser[5] and consequently generate a small part-of-speech (POS) annotated text. For example, the resulting annotation for the tip "Great food, huge portions and a gift shop and showers." is depicted in Figure 3. In a nested, tree-like format this can also be depicted as:

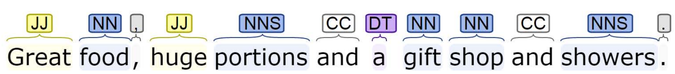

Fig. 3. An example of a POS-tagged tip.

```
(ROOT
 (NP (JJ Great) (NN food)) (, ,)
 (NP (JJ huge) (NNS portions)) (CC and)
 (NP (DT a) (NN gift) (NN shop)) (CC and)
 (NP (NNS showers)) (. .)
)
```

The next step was to process the POS annotated tips and find the polarity of the opinion they convey. In this step we checked if the nouns, verbs or adjectives of a tip were opinionated words contained in a combined lexicon. The combined lexicon was based on "AFFIN-111"[6] and "Harvard General Inquirer" [7] opinion lexicons, which evaluate words with a grade varying from -1 to 1. We checked only adjectives and adverbs for polarity. If one or more such words were spotted inside a sentence, then the respective polarity score is added to the tip.

In the final step, the opinion polarity score of an adjective or adverb inside the tip was associated with a noun or verb within the tip, and more specifically within the same clause (i.e. noun or verb phrase). An illustration of the assigned polarity values to the different aspects of the example tip is provided in Figure 4. Positive aspects are marked with green color, negative with red and neutral with grey. The resulting information, consisting of the aspects and opinion polarity scores for each tip, was stored along with the initial dataset of users, businesses (and their categories) and tips.

For NLP tasks, we employed the NLTK 3.0 package[8] which is also implemented in Python. Our prototype application, which builds on top of the aspects and scores extracted from the user-provided tips, and displays aspects and scores per business, was implemented using the D3 visualization

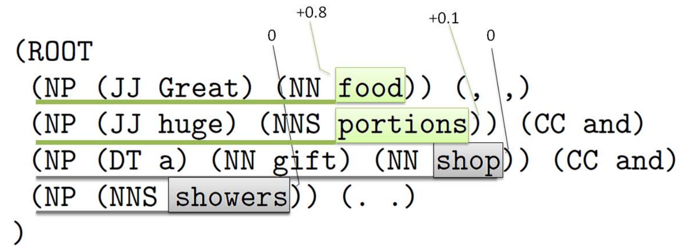

Fig. 4. An example of aspect and opinion polarity detection within a tip.

javascript library[9], designed for the creation of interactive data visualizations. The library allows for easy linking of data attributes to visual encodings within a relatively flexible framework.[10] The preparation of the input JSON formatted file for the graph, was implemented with a combination of Python scripts and MySQL queries.

V. Prototype Application

The prototype application builds on a graph-based visualization which allows us to easily adapt it to the different end-user needs. As we mentioned in Section III, the proposed model can provide recommendations both to clients and business owners. Thus we implement different scenarios that demonstrate the usability of the tool for business owners that wish to improve their services and consequently their reputation and advertisers that wish to increase clients' loyalty and their responsiveness to advertisement campaigns.

In the first scenario, an end-user has provided a negative tip about a business and negatively rated a specific aspect. The application finds the businesses that match this aspect but have a positive score in this aspect. The businesses are ranked in decreasing score order and the top-10 businesses are presented to the user. By clicking on a business, the user can see all the aspects of the recommended business.

In the second scenario, an end-user seeks for the most highly ranked businesses on a certain aspect. By clicking on any of the top-10 businesses, the user can see all its aspects and choose accordingly.

In the third scenario, a business owner or marketeer wants to see the competitors of a certain business. The set of 10 businesses presented in the first visualization can be the result of a query with multiple criteria, whereas the second visualization may uncover the strengths and weaknesses of the competitors.

In all the scenarios that we have implemented, the user of the prototype application can interact with a visualization that depicts one or more recommended businesses (see Figure 5). By clicking on any of the business nodes, the user is presented with all the aspects that have been used in users' tips. The visualization of the cumulative polarity scores, with positive and negative being colored green and red respectively, and the size of the bubble being analogous to the score (see

[5] http://nlp.stanford.edu/software/corenlp.shtml
[6] http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=6010
[7] http://www.wjh.harvard.edu/~inquirer/
[8] www.nltk.org/

[9] http://d3js.org/
[10] Much of our code was drawn from examples by Mike Bostock, a creator of D3.

Figures 5 and 6), gives a very intuitive, user-friendly aspect-based recommendation tool to the business analyst, marketer or end user.

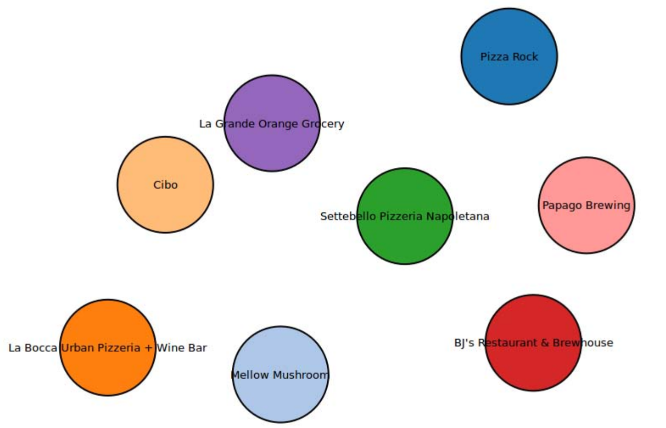

Fig. 5. A sample screen of the top-k recommended businesses for a user.

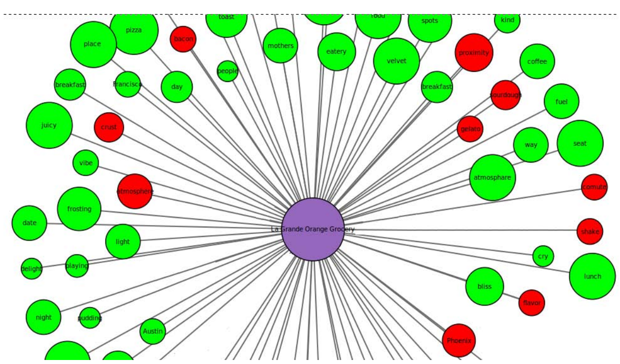

Fig. 6. A sample screen of the positive and negative aspects of a recomended business.

VI. Future work

As part of our future work we plan to focus on the improvement of aspect extraction and aspect-based opinion mining performance. In this direction, we plan to employ a list of positive and negative emoticons in order to increase sentiment detection coverage. We also intend to consolidate synonym words by incorporating a thesaurus and explore the effectiveness of aspect aggregations into more high-level categories. The general purpose of the visualization was to facilitate users in browsing and querying the complicated information that results from the aspect-based analysis of tips and reviews. The prototype application is quite flexible and allows the implementation of different analysis scenarios that can be of use to professional users or end-users and clients. It is on our plans to extend the prototype application with usage scenarios that will support businesses to perform social media analytics, and consequently target the appropriate audience. Overall, the service can be extended to a useful crowd-sourced solution for supporting the SWOT (Strengths, Weaknesses, Opportunities Threats) analysis for the business.

References

[1] E. Bjørkelund, T. H. Burnett, and K. Nørvåg. A study of opinion mining and visualization of hotel reviews. In *Proceedings of the 14th International Conference on Information Integration and Web-based Applications & Services*, pages 229–238. ACM, 2012.

[2] X. Ding, B. Liu, and P. S. Yu. A holistic lexicon-based approach to opinion mining. In *WSDM '08*, 2008.

[3] M. Eirinaki, S. Pisal, and J. Singh. Feature-based opinion mining and ranking. *Journal of Computer and System Sciences (JCSS)*, 78(4):1175–1184, 2012.

[4] M. D. Ekstrand, J. T. Riedl, and J. A. Konstan. Collaborative filtering recommender systems. *Foundations and Trends in Human-Computer Interaction*, 4(2):81–173, 2011.

[5] T.-K. Fan and C.-H. Chang. *Sentiment-oriented contextual advertising*. 2009.

[6] X. Li, H. Wang, G. Yin, T. Wang, C. Yang, Y. Yu, and D. Tang. Inducing taxonomy from tags: An agglomerative hierarchical clustering framework. In *Advanced Data Mining and Applications*, pages 64–77. Springer, 2012.

[7] G. Ling, M. R. Lyu, and I. King. Ratings meet reviews, a combined approach to recommend. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, 2014.

[8] B. Liu. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers, 2012.

[9] Y. Lu, C. Zhai, and N. Sundaresan. Rated aspect summarization of short comments. In *Proceedings of the 18th international conference on World wide web*, pages 131–140. ACM, 2009.

[10] J. P. H. P. I. A. S. M. Maria Pontiki, Dimitrios Galanis. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, RecSys '14, 2014.

[11] J. McAuley and J. Leskovec. Hidden factors and hidden topics: Understanding rating dimensions with review text. In *Proceedings of the 7th ACM Conference on Recommender Systems*, RecSys '13, 2013.

[12] N. Mirbakhsh and C. X. Ling. Clustering-based factorized collaborative filtering. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 315–318. ACM, 2013.

[13] S. Moghaddam and M. Ester. The flda model for aspect-based opinion mining: Addressing the cold start problem. In *Proceedings of the 22nd international conference on World Wide Web*, pages 909–918. International World Wide Web Conferences Steering Committee, 2013.

[14] F. Ricci, L. Rokach, and B. Shapira. *Introduction to recommender systems handbook*. Springer, 2011.

[15] S. Roohi, V. Suresh, and M. Eirinaki. Aspect based opinion mining and recommendation system for restaurant reviews. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, October 2014.

[16] A. Shepitsen, J. Gemmell, B. Mobasher, and R. Burke. Personalized recommendation in social tagging systems using hierarchical clustering. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 259–266. ACM, 2008.

[17] M. Sklar and K. Concepcion. Timely tip selection for foursquare recommendations. In *RecSys 2014 Poster Proceedings*, RecSys '14, 2014.

[18] V. Suresh, S. Roohi, M. Eirinaki, and I. Varlamis. Using social data for personalizing review rankings. In *RSWeb'14*, October 2014.

[19] Y. Wu, F. Wei, S. Liu, N. Au, W. Cui, H. Zhou, and H. Qu. Opinionseer: interactive visualization of hotel customer feedback. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1109–1118, 2010.

[20] Q. Xiao, M. Qiu, C.-Y. Wu, A. Smola, J. Jiang, and C. Wang. Jointly modeling aspects, ratings and sentiments jointly modeling aspects, ratings and sentiments for movie recommendation (jmars). In *ACM KDD'14*, August 2014.