Review-based Entity-ranking Refinement

Panagiotis Gourgaris, Andreas Kanavos, Christos Makris and Georgios Perrakis
Dept. of Computer Engineering and Informatics, University of Patras, Patras, Greece

Keywords: Inference Network, Knowledge Extraction, Opinion Mining, Re-ranking Model, Searching and Browsing, Text Mining, Web Information Filtering and Retrieval.

Abstract: In this paper, we address the problem of entity ranking using opinions expressed in users' reviews. There is an abundance of opinions on the web, which includes reviews of products and services. Specifically, we examine techniques which utilize clustering information, for coping with the obstacle of the entity ranking problem. Building on this framework, we propose a probabilistic network scheme that employs a topic identification method so as to modify ranking of results based on user personalization. The contribution lies in the construction of a probabilistic network which takes as input the belief of the user for each query (initially, all entities are equivalent) and produces a new ranking for the entities as output. We evaluated our implemented methodology with experiments with the OpinRank Dataset where we observed an improved retrieval performance to current re-ranking methods.

1 INTRODUCTION

The rapid development of web technologies and social networks has created a huge volume of reviews on products and services as well as opinions on events and individuals. Opinions are considered as an important part in human activity because of their affect on decision-making. Specifically, consumers are used to being informed by other users' reviews in order to carry out a purchase of a product, service, etc. One other major benefit is that businesses are really interested in the awareness of the opinions and reviews concerning all of their products or services and thus appropriately modify their promotion as well as their further development.

Concerning consumers, one has to refer to many reviews so as to create an overall evaluation assessment for a set of objects of a specific entity. From these reviews, there must be an adequate extraction of several opinions for utilizing an observable conclusion for each one of the objects. The purpose of this extraction is the classification of specific objects and the latter discernment of those that are notable. At this point, it is clear to mention that this multitude of opinions creates a challenge for both the consumer and the entity ranking systems. As an example, suppose that we are interested in purchasing a smartphone device; this drives us in searching for device reviews that are written by users based on their experiences of corresponding products. Potential consumers usually look for specific features (or else named aspects) in a product, e.g. ios or android operating system, battery life, number of camera megapixels, etc. As a result, consumers usually refer to reviews of other users searching for devices with positive opinions regarding these specific features and characteristics.

The above procedure is really exhausting and time-consuming. But as a matter of fact, the development of computational techniques for assisting users to utilize all opinions, is a very important and interesting research challenge. In (Ganesan and Zhai, 2012), authors depict the setup for an opinion-based entity ranking system. The intuition behind their work is that each entity can be represented by all the review texts and that the users of such a system can determine their preferences on several attributes during the evaluation process. Thus, we can expect that a user's query would consist of preferences on multiple attributes. For the previous example regarding smartphone devices, one potential user's query could be "ios system, large battery life, 32 megapixels", expressing user's preferences in three different aspects of the mentioned entity. A solution to the problem of assessing entities can be its transformation into a matching preferences problem, where we can employ any standard information retrieval model. That is, given a query from the user, which consists of keywords and expresses the desired features that an en-

Gourgaris P., Kanavos A., Makris C. and Perrakis G..
Review-based Entity-ranking Refinement.
DOI: 10.5220/0005428604020410
In *Proceedings of the 11th International Conference on Web Information Systems and Technologies* (WEBIST-2015), pages 402-410
ISBN: 978-989-758-106-9

tity should have, we can evaluate all candidate entities based on how well their opinions match user's preferences.

The setup presented in (Ganesan and Zhai, 2012) is an information retrieval approach which uses the importance of aspect keywords on review texts. We investigate the behavior of entity ranking following the information retrieval approach. We strive to use the ratings of aspects in order to identify entities containing similar aspect reviews among the users and use this information to make a better entity ranking. Finally, we consider that not all aspects are equally important to be used in the assessment of the entities.

Furthermore, in our work, we enhance the work presented in (Makris et al., 2013), as we propose a semantically driven Bayesian Inference Network, incorporating semantic concepts (as extracted in (Makris and Panagopoulos, 2014)) so as to improve the ranking quality of documents. Concerning Bayesian Networks, they are progressively being used in a variety of areas like Web Searching (Acid et al., 2003), (Teevan, 2001), Bioinformatics (Niedermayer, 2008) and other. A major subclass of Bayesian Networks is the Bayesian Inference Network (BIN) (Turtle, 1991) that has been employed in various applications (Abdo et al., 2014), (Ma et al., 2006), (Teevan, 2001).

Building on this idea, we utilize schemes that take into account clustering about the opinions emerging in reviews. We also propose a probabilistic network scheme (based on Inference Network modeling), that employs a topic identification method. The rest of the paper is organized as follows. In Section 2, related work as well as contribution is presented. In Section 3, we present the extensions regarding ranking techniques. Subsequently, in Section 4, we describe our re-ranking proposed system. In following, Section 5 presents a reference to our experimental results; we therefore give a presentation of our results. Finally, Section 6 concludes the paper and provides future steps and open problems.

2 RELATED WORK

As we have already stated, in our manuscript, we try to address the problem of creating a ranked list of entities using users reviews and at a latter stage, to present a re-ranked list according to their selections. As a result, the aspect-oriented or feature-based opinion mining as defined in (Ganesan and Zhai, 2012) is employed. Along this line of consideration, each entity is represented as its total review texts and users express their queries as preferences in multiple aspects. Moreover, in (Ganesan and Zhai, 2012), authors presented a setup for entity ranking, where entities are evaluated depending on how well the opinions expressed in the reviews are matched against user's preferences. They studied the use of various state-of-the-art retrieval models for this task, such as the BM25 retrieval function (Robertson and Zaragoza, 2009), the Dirichlet prior retrieval function (Zhai and Lafferty, 2001), as well as the PL2 function (Amati and van Rijsbergen, 2002). Also, they proposed some new extensions over these models, including query aspect modeling (QAM) and opinion expansion; the latter expansion model introduced common praise words with positive meaning for favoring texts and correspondingly entities with positive opinions on aspects.

In (Makris and Panagopoulos, 2014), they further improved the setup by developing schemes, which take into account sentiment and clustering information about the opinions expressed in reviews; also authors propose the naive consumer model as an unsupervised schema that utilizes information from the web so as to yield a weight of importance to each of the features used for evaluating the entities.

Regarding reviews, a great deal of research has been utilized in the classification of reviews to positive and negative ones, based on the overall sentiment information contained. There have been proposed several supervised (Dave et al., 2003), (Pang and Lee, 2004), unsupervised (Nasukawa and Yi, 2003), (Turney and Littman, 2003), as well as hybrid (Pang and Lee, 2005), (Prabowo and Thelwall, 2009) techniques. In addition, there has been much research in the direction of employing users reviews for provisioning ratings in according aspects (Lu et al., 2009), (Wang et al., 2010). These methods are relevant to the one proposed here as with the use of aspect based analysis, the ratings of the different aspects from the reviews can be consequently extracted. However, our approach differs in the applied methodology as we do not explicitly utilize any of the modeling capabilities that these theories provide.

A very related research area is opinion retrieval (Liu, 2012), which aims to identify documents that contain opinions. An opinion retrieval system is usually created on top of the classical recovery models; relevant documents are initially retrieved and concurrently some opinion analysis techniques are being used so as to export documents with emerging opinions. The field of expert finding can be considered as another related research area. Particularly, a ranked list of persons that can be regarded as experts on a certain topic (Fang and Zhai, 2007), (Wang et al., 2010) can be recovered. In particular, we are trying to export a ranked list of entities, but instead of evaluating the entities based on how well they match a topic, we can use the opinions for the entities and as a result to observe how well they match the user's preferences.

Concerning the ranking quality of documents, authors in (Lee et al., 2011) enrich the semantics of user-specific information and documents targeting at efficient implementation of personalized searching strategies. They adopt a Bayesian Belief Network (BBN) as a strategy for personalized search since they provide a clear formalism for embedding semantic concepts. Their approach is different from ours, as they use belief instead of inference networks and then they employ the Open Directory Project Web directory. In (Abdo et al., 2014), the authors enhance the BINs using relevance personalization information and multiple reference structures applying their technique to similarity-based virtual screening, employing two distinct methods for carrying out BIN searching: reweighting the fragments in the reference structures and a group fusion algorithm. Our approach aims at a different application and employs semantic information, as a distinct layer in the applied inference network.

Alongside this line of research, there are approaches that exploit information from past user queries and preferences. Relevant techniques range from simple systems implementing strategies that match users' queries to collection results (Meng et al., 2002), to the employment of the machine learning methods exploiting the outcomes of stored queries, so as to permit more accurate rankings (Liu, 2011). There is also a related, but different to our focus, work (Brandt et al., 2011) combining diversified and interactive retrieval under the label of dynamic ranked retrieval. In contrast, in (Makris et al., 2013), they initially proposed transparent embedding of semantic knowledge bases to improve search engine results re-ranking; for this purpose they created a new probabilistic model which takes as input different semantic knowledge bases. Here in this paper, we insert document clustering based on entity identification in the belief network making it possible to identify the interest of users to thematic groups of results. This is a novel approach on exploiting semantic knowledge on belief networks applied for information retrieval problems.

The main contribution of the proposed method is the incorporation and further examination of clustering techniques regarding the opinions emerging in reviews; also the proposal of a probabilistic network scheme based on Inference Network modeling in order to modify the ranking of results as the users select entities. The method exploits the user's belief from the selected entity through the constructed network, to the other entities that contain the senses of the selected one. The re-ranking of the results is based on a vector which contains a weight for each entity representing the probability of the entity to be relevant for the user. Our method constructs a probabilistic network from the entities, the clusters of the entities as well as the aspects; so when the users select an entity, the weights of the entities, which are in the same cluster, take larger values and thus are ranked higher. Detailed experiments with the proposed method show increased performance in comparison to the one-phase rankings (without re-ranking module) of the result set.

3 EXTENSIONS

We present a methodology for extracting weights and therefore effecting the Ranking List taken from the system. There are two perspectives we try to address. The first one concerns information extracted from the reviews that already exist in our database. The second deals with personalization from the viewpoint of query introduced real time by the user.

3.1 Opinion-aspect Query Expansion

We can utilize opinion expansion in query q with use of the *WordNet*. We consider query q_i as that part of the overall query concerning a single aspect.

In addition, we form Q_{exp} which is the set of $\{q_1, q_2, \ldots, q_n\}$, where each query has value $q_i = \{t_i, st_{i1}, st_{i2}, \ldots, st_{ik}\}$ and t_i is treated holistically in the form of the term q_i. Also, st_{ij} are the synonyms of the opinion term t_i and k is their number.

The aspects are handled in a similar way. As a matter of fact, each time there is a reference in an aspect, the reference would concern its synonyms too.

Finally, we introduce collections Dq, where Dq_i is the collection of reviews that contain terms from q_i.

3.2 Query Personalization

The intuition behind this method is that the aspects do not have the same meaning when applied to the query. More precisely, the same takes place for q_i. The assumption we make is intuitive that the first aspect applied by the user, is the most important. So the aspects introduced ultimately have less importance.

3.3 RRR Weighting

Two weights are used as re-ranking factors, termed as *RRR1* and *RRR2* throughout our study. Initially, for the *RRR1* weight, we formulate the importance of a specific aspect a_i in the original collection C. We use the information $|C_{fi}|$ which is the number of documents containing a_i.

$$0 \leq \log_2(1 + \frac{|C_{fi}|}{|C|}) \leq 1 \quad (1)$$

For the second re-ranking factor, *RRR2*, we utilize the Poisson discrete distribution. It expresses the probability of a given number of events to occur in a fixed interval of time, if these events occur with a known average rate λ and independently of the time since the last event.

$$P_\lambda(X = k) = \frac{\lambda^k}{k!} \times e^{-\lambda} \quad (2)$$

where λ here is constant with value *1.8*. The idea here is that aspects with given position in the query k are introduced to the system. The event of its occurrence in place k is what needs to be evaluated. By selecting a fixed λ, we favour aspects that are between places 1 and 2.

4 RE-RANKING PROPOSED SYSTEM

In order for a user to utilize the services provided by our inference network model, they simply need to give their feedback concerning a specific query, consisting of a combination of different aspects. Every time the user clicks on a result (entity in our example), the inference network is utilized producing the new improved ranking list according to their choice. In particular, the system works in a real-time and in parallel basis, namely re-arranging the initial order and the algorithm introduced is presented in detail in the following sections. Both the re-ranking method and the network were inspired by (Makris et al., 2013).

4.1 The Re-Ranking Method

Our refinement method improves the initial ranking based on the user's selections; our personalization re-organizing step is utilized accompanied by the input produced by each and every user. This step runs iteratively every time the user makes a selection. In particular the proposed network is utilized, either in combination with the initial ranking returned by the search engine or with the previous ranking of the results in the re-ranking process (if we have a series of re-rankings). The above options are expressed by the following equation. These were also used, though not entirely in the same way, in (Makris et al., 2013) (see also (Antoniou et al., 2012)):

$$\begin{aligned} &NewRankingScore_i = \\ &\quad (n - PreviousRank_i + 1) \times (1 + \beta \times R_i) \end{aligned} \quad (3)$$

where R_i denotes the re-ranking weight provided by the network for e_i (its computation is described in the next section), $PreviousRank_i$ stands for the previous rank position for entity e_i utilized from the BM25 procedure algorithm (IR model) incorporating our *RRR_weights*, n is the number of results retrieved and β is a user defined weight factor. Intuitively, when the factor β is changed, the re-ranking process results in major rank changes.

Equation 3 is used for the composite case where the new ranking system is composed with the previous ranking of search results. In the new ranking produced, the results are ranked according to the above calculated score. When the user selects further results, the same procedure is followed with the difference that the ranking produced by the previous phase is used as input for the next reordering.

Even though we assume that most results selected by a user are relevant, our scheme incorporates smoothly the previous ranking, hence it is robust to user misselections. A misselection of a result leads to the inclusion of its relevant information to the ranking process, but still can be made to not affect significantly the produced ranking.

4.2 Re-Ranking Weight Calculation

Our inference network, as depicted in following scheme, consists of four component levels (two of them are the same): the entities level (implementing user's personalization assuming that each entity corresponds to a unique document), the clusters level, the aspects level (query level) and a fourth level that represents the entities as well as the weights they are assigned by the re-ranking procedure. The final level can be considered to play the role of the query layer in the traditional inference network model and its presence signifies that we are interested to model specific re-rankings based on users' queries.

The proposed inference network is implemented once for the specific dataset and its structure does not change during re-ranking. The entities level contains a node e_i for each entity of the query's results. For each entity the user selects, we assign weights to corresponding clusters Cl_j as we incorporate a fuzzy clustering algorithm where each entity can be placed to more than one cluster (with weight corresponding to its contribution to this specific cluster). The cluster

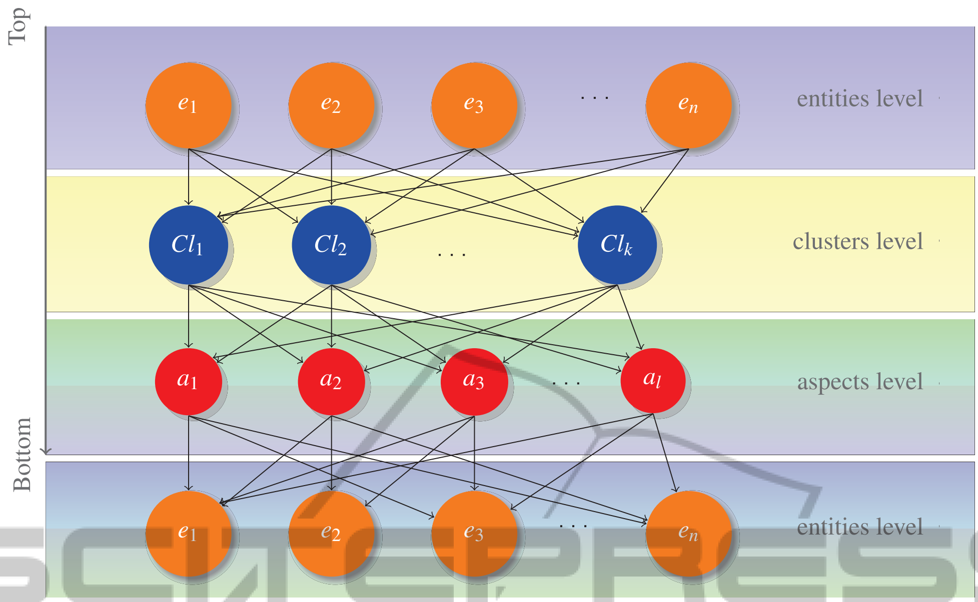

Figure 1: Inference Network.

nodes are implicitly connected with the aspects a_m; more precisely, in this level we take into consideration the probabilities that we will describe in the next section. Finally, from the aspect nodes, we expose information regarding the entity nodes belonging to the last level of our proposed network. The formed network is a four level, unidirectional graph in which the information flows from the initial entity node of the first level to the cluster nodes and then through the aspect nodes to the entity nodes of the last level (from top to bottom).

The contribution in our network is the addition of the aspects level (or else the query expansion) based on users entity's selection. As it is presented in the Figure 1, the entity nodes are connected to their different aspects through directed arcs. The existence of a directed path between a cluster node and an aspect node denotes that this aspect is appearing in the respective cluster and more specifically that the specific aspect has been observed in the specific entity reviews collection. The last level's entity nodes are same with the nodes of the first layer, but with a slight difference; they represent the same entities but regarding different aspects (we present this fact by connecting the aspect nodes with the entity nodes of the last layer with different lines, just to depict the fact that we are interested in different aspects). The aspects level models the needs of a user; that is the aspects the user looks for. The aspect nodes are connected to every node of the last level representing an entity where this aspect appears. In addition, these entity nodes have an accumulated belief probability that is used for the proposed re-ranking procedure. The value of this belief is estimated based on the different emerging aspects of the entity and denotes the conceptual similarity between the entity and the information need of the user.

4.3 Estimation of Probabilities and Rearrangement of Results

The proposed inference network is based on that proposed in (Turtle, 1991) but with a slight difference; there, an information retrieval model is proposed while in our manuscript a re-ranking model is utilized. More specifically, we have only employed it as a weight propagation mechanism using the machinery for computing the beliefs at the last level of the network, thus providing a set of prior weights for the entities of the first level. That is, as a user selects an entity, the corresponding entity is assigned a probability and as a result we compute the alternation of the beliefs at the last level. This belief is transferred through the network to the aspect nodes and then to the final layer representing the entity nodes and alternates the results through the re-ranking process.

For the estimation of the nodes probabilities for the constructed inference network, we begin by assigning weights to the root nodes (entities). We enable the entity node that matches the user's selection; the specific entity will trigger the cluster nodes the entity belongs to and as a result, the arcs starting from this entity node take a prior probability. More specifically, each entity node has an initial probability that denotes the chance of the selection of that corresponding entity from the user. For each entity e_i this initial probability will be:

$$p(e_i) = \frac{1}{n},\ i = \{1,\ldots,n\} \quad (4)$$

where n is the number of the different entities. This probability will change into 1 for the selection denoting user belief. In addition, for our specific dataset and for each entity e_i that participates to all clusters, the probability of a cluster Cl_j will be:

$$p(Cl_j|e_1,e_2,\ldots,e_n) = \frac{w_{all_j}}{\sum_{j=1}^{k} w_{all_j}},\ j = \{1,\ldots,k\} \quad (5)$$

$$w_{all_j} = \sum_{i=1}^{n} p(e_i) \times w_{ij} \quad (6)$$

where k is the number of the different clusters and w_{ij} is the contribution of the entity e_i in the cluster Cl_j. The aspects probabilities at third level are calculated as follows:

$$p(a_m|Cl_j) = \frac{wa_{mj}}{\sum_{j=1}^{k} wa_{mj}},\ m = \{1,\ldots,l\} \quad (7)$$

$$wa_{mj} = \sum_{i=1}^{n} \frac{tf(a_m,e_i)}{\sum_{v=1}^{l} tf(a_v,e_i)} \times p(e_i) \times w_{ij} \quad (8)$$

where l is the number of the different aspects, wa_{mj} is the number of aspects and $tf(a_m,e_i)$ is the term frequency of aspects a_m related to the entity e_i. The entities probabilities at fourth level are calculated as follows:

$$p(e_i|a_1,a_2,\ldots,a_l) = \frac{\sum_{v=1}^{l} tf(a_v,e_i)}{\sum_{i=1}^{n}\sum_{v=1}^{l} tf(a_v,e_i)} \quad (9)$$

Finally, the whole probability of each entity is calculated by the following equation, that is the transfer of user's belief through the network:

$$p(e_i,a_1,\ldots,a_l,Cl_1,\ldots,Cl_k,e_1,\ldots,e_n) =$$
$$p(e_i|a_1,a_2,\ldots,a_l) \times \prod_{m=1}^{l}\prod_{j=1}^{k} p(a_m|Cl_j)$$
$$\times \prod_{j=1}^{k} p(Cl_j|e_1,e_2,\ldots,e_n) \times \prod_{i=1}^{n} p(e_i)$$

This is the belief which is then put in equation 3, providing the user with a re-ordered list.

5 EXPERIMENTS

5.1 Experimental Setting

In order to assess our methods performance, we compared the initial ranking taken from the ideal ranking based on the ratings of the aspects that users had given with the ranking of the results after the application of the proposed methods. The experiments were carried out using the OpinRank Dataset, which was presented in (Ganesan and Zhai, 2012) and consists of entities, which are accompanied by reviews of users from two different domains (cars and hotels); the reviews come from the sites Edmunds.com and Tripadvisor.com respectively. Particularly for our evaluation, we use the reviews from the domain of the cars which includes car models as well as the corresponding reviews for the years 2007 (227 models), 2008 (228 models) and 2009 (143 models). In our set of experiments, we perform 100 queries and we evaluate the performance of our schemas so as to produce the correct entity ranking, calculating the *nDCG* at the first 10, 20 results as well as 50 results.

We used the Normalized Discounted Cumulative Gain (*nDCG*) measure (Kalervo Jarvelin, 2000), which quantifies the gain of a document based on its position in the result list. The *nDCG* measure is based on the relevance judgments of the documents of the result list. Formally, the *nDCG* is computed at position p as:

$$DCG_p = rel_1 + \sum_{i=2}^{p} \frac{rel_i}{log_2 i} \quad (10)$$

$$nDCG_p = \frac{DCG_p}{IDCG_p} \quad (11)$$

where rel_i are the document relevance scores from reviews and $IDCG_p$ is the ideal *DCG*. The ideal *DCG* is the *DCG* values when sorting the documents using the relevance judgments.

We set the experiment posing each query and randomly selecting a relevant document. Then using the selected result, we perform re-ranking in the result list. The re-ranking performance is measured using *nDCG* for the initial ranking and after the re-ranking step.

5.2 Results

From Equations 1, 2 and Section 4, we created Tables 1, 2 and 3 as well as corresponding graphical representations in Figures 2 to 10 for the domain of the cars for the years 2007, 2008 and 2009. The experiments depict that *RRR*1 weighting scheme has superior performance from the formula that uses the Poisson discrete distribution. This does not mean necessarily that Poisson distribution is always worse but probably that the proper choice of λ has to be carefully performed (perhaps using extra experiments) tailored to the dataset used and also accompanied with classification techniques. The *RRR*3 results presented, are produced by weighting the entities with both the *RRR*1 and *RRR*2 weights. It is clear, as we have assumed, that the re-ranking network performs better than all the other simple ranking methods.

Table 1: Average of nDCG@10 for years 2007, 2008 and 2009.

Method	nDCG
BM25	0.879
BM25 + RRR1	0.895
BM25 + RRR2	0.879
BM25 + RRR3	0.893
BM25 + Network	0.915

Table 2: Average of nDCG@20 for years 2007, 2008 and 2009.

Method	nDCG
BM25	0.879
BM25 + RRR1	0.894
BM25 + RRR2	0.878
BM25 + RRR3	0.893
BM25 + Network	0.913

Table 3: Average of nDCG@50 for years 2007, 2008 and 2009.

Method	nDCG
BM25	0.880
BM25 + RRR1	0.892
BM25 + RRR2	0.879
BM25 + RRR3	0.890
BM25 + Network	0.905

In the following Figures 2 to 10, we present the comparisons in the performance of the simple BM25 model with the weighting scheme *RRR*1 as well as the re-ranking network. We can observe that for all 3 years, the re-ranking model and the weighting scheme outperforms clearly the classic BM25 model.

6 GENERAL CONCLUSIONS AND FUTURE WORK

In this work we have presented entity ranking techniques using opinions expressed in users' reviews.

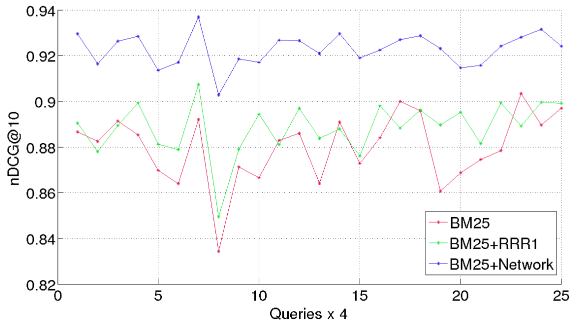

Figure 2: Measurements of the nDCG@10 for year 2007.

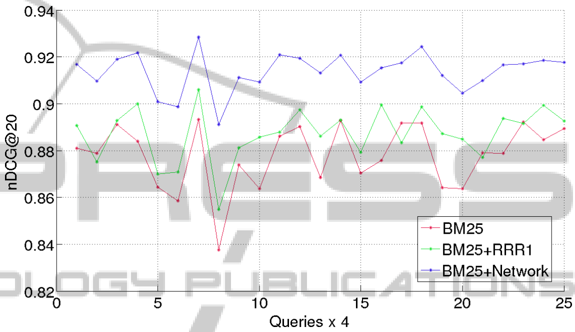

Figure 3: Measurements of the nDCG@20 for year 2007.

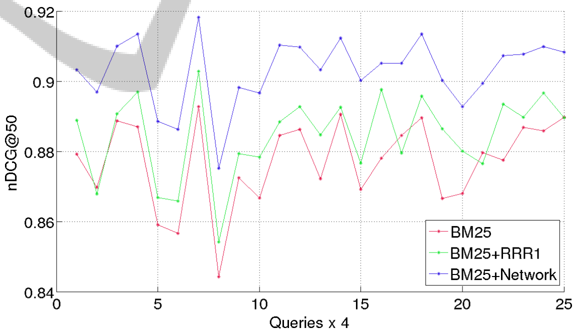

Figure 4: Measurements of the nDCG@50 for year 2007.

We examined three weighting schemas that incorporate clustering information for coping with the obstacle of the entity ranking problem. Implementing the aspects expansion, we aim in more meaningful queries incorporating additional semantic information; by that we address the problem of the initial limited collection of corresponding queries. In addition, we propose a probabilistic network scheme that employs a topic identification method so as to modify ranking of results based on users personalization. Based on this technique, our proposed system rearranges the results; also without need of storing the navigation history of each specific user, we achieve a significant improvement on the ranking of the results in terms of user preferences. The evaluation of our proposed implemented methods was

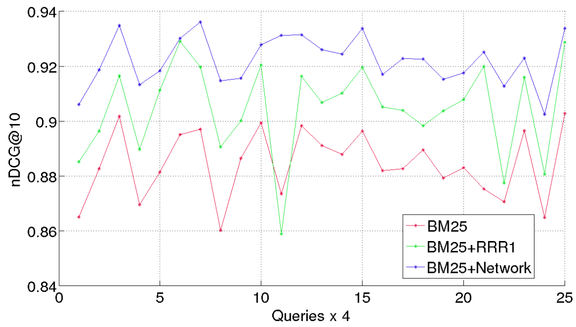

Figure 5: Measurements of the nDCG@10 for year 2008.

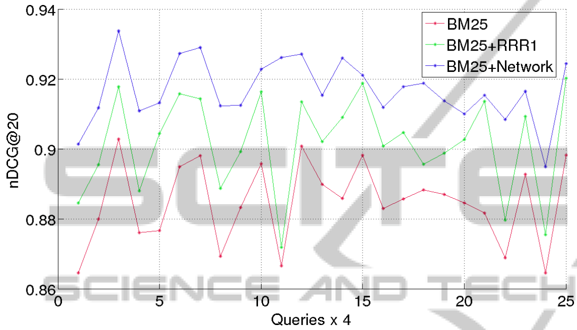

Figure 6: Measurements of the nDCG@20 for year 2008.

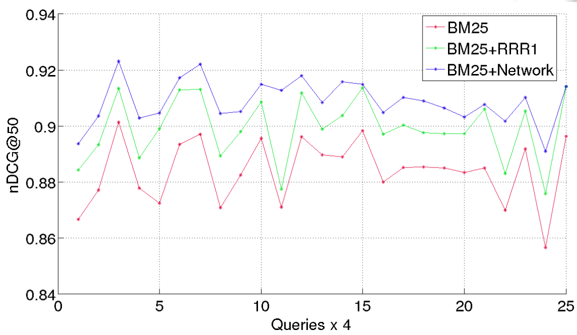

Figure 7: Measurements of the nDCG@50 for year 2008.

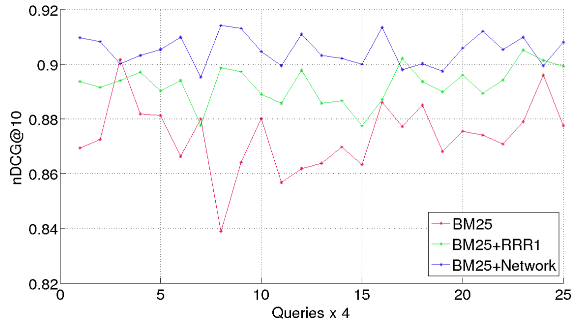

Figure 8: Measurements of the nDCG@10 for year 2009.

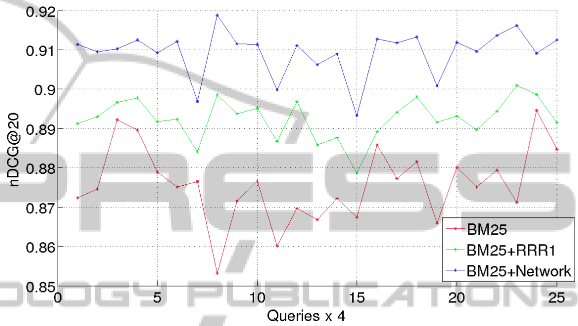

Figure 9: Measurements of the nDCG@20 for year 2009.

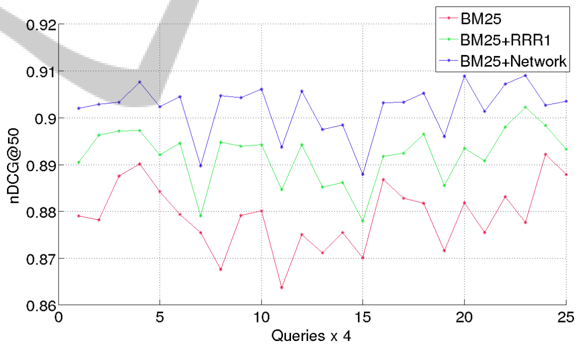

Figure 10: Measurements of the nDCG@50 for year 2009.

examined through experiments with the OpinRank Dataset where we observed an improved retrieval performance to current ranking methods.

Regarding future work, there is the attractive issue of studying an information retrieval model that favors texts (and correspondingly entities) with positive opinions on aspects while imposes penalties for negative opinions. Furthermore, another interesting point is the incorporation of other knowledge bases besides WordNet, such as YAGO and Wikipedia so as to further enhance our inference network by embedding in it additional semantic concepts. Further study and work is imminent on the Poisson weight. Ranking systems based on reviews provide us with the ability to pre-decide for all possible queries form. Since given aspects are extracted, inserted queries are known a priori. Using this knowledge, we can examine additional characteristics out of them and as a result to incorporate new weighting information retrieval techniques in such ranking systems.

REFERENCES

Abdo, A., Leclere, V., Jacques, P., Salim, N., and Pupin, M. (2014). Prediction of new bioactive molecules using a bayesian belief network. In *Journal of Chemical Information and Modeling, Volume 54, Issue 1, pp. 30-36.*

Acid, S., de Campos, L. M., Fernandez-Luna, J. M., and Huete, J. F. (2003). An information retrieval model based on simple bayesian networks. In *International*

Journal of Intelligent Systems, Volume 18, pp. 251-265.

Amati, G. and van Rijsbergen, C. J. (2002). Probabilistic models of information retrieval based on measuring the divergence from randomness. In *ACM Transactions on Information Systems (TOIS), Volume 20, Number 4, pp. 357-389*.

Antoniou, D., Plegas, Y., Tsakalidis A., Tzimas, G. and Viennas, E. (2012). Dynamic Refinement of Search Engines Results Utilizing the User Intervention. In *Journal of Systems and Software, Volume 85, pp. 1577-1587*.

Brandt, C., Joachims, T., Yue, Y., and Bank, J. (2011). Dynamic ranked retrieval. In *WSDM, pp. 247-256*.

Dave, K., Lawrence, S., and Pennock, D. M. (2003). Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In *International Conference on World Wide Web (WWW), pp. 519-528*.

Fang, H. and Zhai, C. (2007). Probabilistic models for expert finding. In *European Conference on IR Research (ECIR), pp. 418-430*.

Ganesan, K. and Zhai, C. (2012). Opinion-based entity ranking. In *Information Retrieval (IR), Volume 15, Issue 2, pp. 116-150*.

Kalervo Jarvelin, J. K. (2000). Ir evaluation methods for retrieving highly relevant documents. In *SIGIR, pp. 41-48*.

Lee, J.-W., Kim, H.-J., and Lee, S.-G. (2011). Exploiting taxonomic knowledge for personalized search: A bayesian belief network-based approach. In *Journal of Information Science and Engineering (JISE), Volume 27, pp. 1413-1433*.

Liu, B. (2012). *Sentiment Analysis and Opinion Mining*. Morgan and Claypool Publishers.

Liu, T.-Y. (2011). *Learning to Rank for Information Retrieval*. Springer.

Lu, Y., Zhai, C., and Sundaresan, N. (2009). Rated aspect summarization of short comments. In *International Conference on World Wide Web (WWW), pp. 131-140*.

Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. In *Nature Neuroscience, Volume 9, pp. 1432-1438*.

Makris, C., Plegas, Y., Tzimas, G., and Viennas, E. (2013). Serfsin: Search engines results' refinement using a sense-driven inference network. In *WEBIST, pp. 222-232*.

Makris, C. and Panagopoulos, P. (2014). Improving opinion-based entity ranking. In *WEBIST, pp. 223-230*.

Meng, W., Yu, C. T., and Liu, K.-L. (2002). Building efficient and effective metasearch engines. In *ACM Computing Surveys, Volume 34, Issue 1, pp. 48-89*.

Nasukawa, T. and Yi, J. (2003). Sentiment analysis: Capturing favorability using natural language processing. In *International Conference on Knowledge Capture (K-CAP), pp. 70-77*.

Niedermayer, I. S. P. D. (2008). An introduction to bayesian networks and their contemporary applications. In *Springer Studies in Computational Intelligence, pp. 117-130*.

Pang, B. and Lee, L. (2004). A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Annual Meeting of the Association for Computational Linguistics (ACL), pp. 271-278*.

Pang, B. and Lee, L. (2005). Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.

Prabowo, R. and Thelwall, M. (2009). Sentiment analysis: A combined approach. In *Journal of Informetrics (JOI), Volume 3, Issue 2, pp. 143-157*.

Robertson, S. E. and Zaragoza, H. (2009). The probabilistic relevance framework: Bm25 and beyond. In *Foundations and Trends in Information Retrieval, Volume 3, Issue 4, pp. 333-389*.

Teevan, J. B. (2001). Improving information retrieval with textual analysis: Bayesian models and beyond. In *Masters Thesis, MIT Press*.

Turney, P. D. and Littman, M. L. (2003). Measuring praise and criticism: Inference of semantic orientation from association. In *ACM Transactions on Information Systems (TOIS), Volume 21, Issue 4, pp. 315-346*.

Turtle, H. R. (1991). Inference networks for document retrieval. In *Doctoral Dissertation*.

Wang, H., Lu, Y., and Zhai, C. (2010). Latent aspect rating analysis on review text data: A rating regression approach. In *SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 783-792*.

Zhai, C. and Lafferty, J. D. (2001). A study of smoothing methods for language models applied to ad hoc information retrieval. In *SIGIR, pp. 334342*.