Contextual eVSM: A Content-Based Context-Aware Recommendation Framework Based on Distributional Semantics

Cataldo Musto, Giovanni Semeraro, Pasquale Lops, and Marco de Gemmis

Department of Computer Science
University of Bari Aldo Moro, Italy
{cataldo.musto,giovanni.semeraro,pasquale.lops,marco.degemmis}@uniba.it

Abstract. In several domains contextual information plays a key role in the recommendation task, since factors such as user location, time of the day, user mood, weather, etc., clearly affect user perception for a particular item. However, traditional recommendation approaches do not take into account contextual information, and this can limit the goodness of the suggestions. In this paper we extend the enhanced Vector Space Model (eVSM) framework in order to model contextual information as well. Specifically, we propose two different context-aware approaches: in the first one we adapt the *microprofiling* technique, already evaluated in collaborative filtering, to content-based recommendations. Next, we define a contextual modeling technique based on distributional semantics: it builds a context-aware user profile that merges user preferences with a semantic vector space representation of the context itself. In the experimental evaluation we carried out an extensive series of tests in order to determine the best-performing configuration among the proposed ones. We also evaluated Contextual eVSM against a state of the art dataset, and it emerged that our framework overcomes all the baselines in most of the experimental settings.

Keywords: Context-aware Recommendations, Filtering, User Modeling, Content-based Recommenders.

1 Introduction

Recommender Systems (RSs) are tools that can help users in *'making order'* in the plethora of information today available on the Web, by providing them with personalized suggestions about items that are supposed to be of interest according to their preferences or their information needs [12].

The process RSs deal with is very natural and common, since people get advice all the time, especially when they have to choose among several alternatives and the knowledge to critically discern them is not enough: *what music should I listen to? what movie should I watch? what restaurant should I choose?* A list of examples could be infinite. However, in all the abovementioned scenarios there is an aspect that plays a key role in order to determine which one is the best suggestion to provide: *the context*.

C. Huemer and P. Lops (Eds.): EC-Web 2013, LNBIP 152, pp. 125–136, 2013.

Do you need music to better concentrate or to play during a party? Do you want a movie for a funny night or for a romantic one? It is clear that, in most of the cases, the context actually influences the process of generating the recommendations. As a consequence, an effective recommendation algorithm should take into account as much contextual information as possible: location, time of the day, companion, task to be carried out, user mood, etc However, traditional recommendation approaches do not model any contextual information: collaborative filtering [14] produces recommendations by just exploiting similarities among user behavioral patterns (e.g. buying, clicking, rating, etc.). In the same way, content-based filtering [15] typically generates the recommendations by comparing the textual features describing the items with those stored in a user profile built upon the items the user enjoyed in the past. The need for recommendation approaches able to manage contextual information is taken for granted. Consequently, several context-aware recommendation algorithm recently emerged [1]. Techniques for generating context-aware recommendations can be broadly split into three categories: *pre-filtering, post-filtering and contextual modeling.* The basic idea behind pre-filtering is to generate context-aware recommendations by exploiting only the subset of the information (e.g. ratings) expressed by the users under some specific context. For example, if a user needs suggestions on the best music for a party, the algorithm builds the recommendation set by filtering out all the preferences expressed in contexts different from the target one. On the other side, post-filtering generates the recommendations by exploiting all the available data, next it uses contextual information to filter out the items that do not match some contextual constraints. For example, a post-filtering recommendation algorithm could filter the restaurants too far from the current location of the user. Finally, contextual modeling implants information about the context in the algorithm itselft, thus influencing the step of building user profiles as well as that of generating recommendations.

The main contribution of this paper is Contextual eVSM, a framework for content-based context-aware recommendations. The main building block of the framework is the enhanced Vector Space Model (eVSM) [17], an adapation of the vector space model to the requirements of content-based recommender systems (CBRS), boosted by distributional semantics [24] and quantum negation [25]. In this work the eVSM has been further extended to make it context-aware: this has been done by introducing a pre-filtering approach as well as a novel contextual modeling technique. As pre-filtering approach we adapted to CBRS the well-known *microprofiling* [5] technique. Specifically, we split user ratings according to the contextual situation the preference is expressed in, and we exploited this information to build several context-aware *(micro)* profiles that are used to generate context-aware recommendations. As contextual modeling approach, we introduced a novel technique that exploits distributional semantics to represent the context as a vector and merges it with a vector space representation of user preferences, thus building a context-aware user profile able to provide users with context-aware recommendations.

The paper is organized as follows. In Section 2 the most relevant work in the area of context-aware recommendation are sketched. Next, in Section 3, we focus the attention on Contextual eVSM: we describe the main building blocks of the eVSM, then we show how the framework is made context-aware by introducing pre-filtering and contextual modeling techniques. Details about the experimental settings are provided in Section 4: we exploit a state-of-the-art dataset to evaluate the goodness of our framework against other relevant work in the area. Finally, Section 5 contains conclusions and future directions of this research.

2 Related Work

Even if there exists a very vast literature on RSs [21], the research about context-aware RSs (CARS) is relatively new. One of the first attempts towards the construction of recommendation algorithms able to manage contextual information has been carried out by Herlocker and Kostan, that proposed in [11] a technique that adapted the recommendation list to the specific task of the user. This work represents one of the first evidences towards the goodness of the insights behind CARS, since 88% of the users involved in the user study designed by the authors preferred the context-aware version of the system. Almost in parallel, Adomavicius and Tuzhlin proposed in [3] a *multi-dimensional model* able to extend the classical user-item matrix in order to store additional contextual information. This research line has been further extended in [2], where the authors applied a *reduction-based* approach that reduces the dimensionality of the original matrix. In our experimental evaluation we exploited the same dataset and the same experimental settings used in that work, in order to guarantee comparable experimental results. The insight of reducing the complexity of the recommendation task by building a smaller matrix has been followed by Karatzoglou et al. [13], that proposed a framework for *multiverse recommendations* based on the tensor factorization. The most recent trends in the area of CARS have been discussed in the recent series of CARS[1] and CAMRa[2] workshopsas well as in a recent survey [1]. The abovementioned classification between *pre-filtering, post-filtering and contextual modeling* has been proposed by Adomavicius and Tuzhilin [4]. As stated above, pre-filtering algorithms filter the set of available data in order to exploit only those that are relevant for a certain contextual scenario. A typical approach is the micro-profiling, discussed in [5]. Similarly, Baltrunas and Ricci introduced the technique of *item splitting*, where each item is split into several fictitious items based on the different contexts in which these items can be consumed. However, pre-filtering is very prone to suffer of the sparsity problem, since it is likely that only a little subset of data is available for a certain context. In order to handle this issue, in [26] the authors introduced the concept of context relaxation for pre-filtering algorithms. In the experimental evaluation

[1] http://cars-workshop.org/
[2] http://camrachallenge.com/

they showed an improvement of the performance of their CARS. A broad comparison between pre and post-filtering techniques is provided by Panniello et al. [20]. The empirical results showed that post-filtering generally performs better than pre-filtering. However, the authors stated that it is not possible to clearly determine the best performing technique and the choice really depends on the application domain. Our contextual modeling approach got inspiration from the *weighted post-filtering* proposed by Panniello, since we used the vector space representation of the context as a weighting factor that is merged with a vector space representation of user preferences. According to the presented literature, the novelty of our work lies in the following aspects:

- Most of the proposed techniques focus on the enrichment of collaborative filtering approaches to model and store information about context. Differently, in this paper we propose an extension of a content-based recommendation framework. Up to our knowledge, this is one of the first attempts towards this direction. In [19] the authors exploited the data stored in DbPedia[3], the RDF mapping of Wikipedia, to provide context-aware movie recommendation for a mobile application, while in [8] the authors use contextual information to improve the performances of a content-based news recommender systems. Beyond these attempts, the area of context-aware content-based recommender sysems has not been properly investigated, yet;
- Our approach exploits distributional models (DMs) [24] to build a semantic vector space representation of the context. DMs state that the meaning of the terms can be inferred in a totally unsupervised way by just analyzing their usage patterns in a specific language, with the insight that terms that are usually used together (e.g. beer, wine, etc.) are supposed to share a similar meaning. According to this insight, we decided to exploit distributional models to build a semantic vector space representation of the contextas well. Specifically, we assumed that the context could be represented as a vector obtained by combining the semantic representation of the terms used to describe items labeled as relevant in that context. Similarly, in [9] the authors exploit the distributional hypothesis to calculate similarities between different contexts and use this information to relax contextual pre-filtering constraints. However, differently from this work, we used the distributional semantics as a *weighting factor* of a contextual modeling technique.

3 Contextual eVSM

The eVSM [17] is a content-based recommendation framework based on vector space model (VSM) [23].

3.1 Basis of eVSM

The whole framework is built upon the following building blocks:

[3] `http://dbpedia.org/`

- The VSM is the core of the framework. *Items* as well as *user profiles* are represented as vectors in a vector space. However, since VSM does not provide any semantic modeling of the information, distributional models are exploited to build a lightweight *semantic* representation, according to the co-occurrences of the terms within the corpus;
- Techniques based on distributional models are not scalable (e.g. LSA [10]). In order to guarantee the scalability required by CBRS, distributional semantics has been coupled with an incremental and effective dimensionality reduction technique called Random Indexing [22], that has been used to reduce the dimension of the vector space;
- Since VSM cannot model any negative evidence, a quantum negation operator, proposed by Widdows [25], has been integrated in the framework.

Thanks to the combination of distributional models, Random Indexing and quantum negation, it is possible to represent items as points in a (semantic) vector space built in a incremental and scalable way. Similarly, a semantic user profile, able to model also negative evidences (e.g. information about items the user disliked), can be learned. Specifically, let I a set of items split into I_u^+ and I_u^- (items the user liked and items the user disliked, respectively), let $d_1..d_n \in I$ be a set of already rated items, let $r(u, d_i)$ $(i = 1..n)$ the rating given by the user u to the item d_i, it is possible to define two different user profiling approaches.

In the first one, denoted as Weighted Random Indexing (WRI), the user profile is a vector that combines in a weighted way the vector space representation of the items the user liked in the past.

$$\boldsymbol{WRI(u)} = \sum_{i=1}^{|I_u^+|} \boldsymbol{d_i} * \frac{\boldsymbol{r(u, d_i)}}{\boldsymbol{MAX}} \tag{1}$$

Where MAX is the maximum rating. Next, the Weighted Quantum Negation (WQN) profile models into a single vector $\boldsymbol{WQN(u)}$ the information coming from $\boldsymbol{WRI(u)}$ with that coming from $\boldsymbol{WRI_{neg}(u)}$, a vector space representation of the items the user disliked:

$$\boldsymbol{WRI_{neg}(u)} = \sum_{i=1}^{|I_u^-|} \boldsymbol{d_i} * \frac{\boldsymbol{MAX - r(u, d_i)}}{\boldsymbol{MAX}} \tag{2}$$

Under a geometrical point of view, the user profile $\boldsymbol{WQN(u)}$ represents the projection of $\boldsymbol{WRI(u)}$ on the subspace orthogonal to those generated by $\boldsymbol{WRI_{neg}(u)}$ [7].

Finally, given a vector space representation of user preferences (WRI or WQN), the recommendation set is built by exploiting cosine similarity: specifically, the items with the highest cosine similarity are returned as recommendations. Even if a complete description of the eVSM framework is out of the scope of this paper, it is worth to note that the framework has been already evaluated in several experimental settings [17,18], where the effectiveness of the approach was always confirmed. In the next section we will evaluate the framework in the task of providing users with contextual recommendations.

3.2 Introducing Context into eVSM

The concept of *context* has been studied in multiple disciplines, and each one tends to take its own view of it. As stated by Bazire et Batillon [6], it is not possible to provide a unique universally shared definition: in the area of personalization and recommender systems, for example, a rich overview of the definitions as well as the scope of this multifaceted concept is contained in [1]. However, the definition of the concept of *context* is out of the scope of this paper. For the sake of simplicity we can consider the context *as a set of (external) factors able to influence user perception of the utility of a certain item.* Several aspects fall into this definition: the task to be accomplished, the company, the location, the mood, the weather and so on. Formally, we can define the *context* as a set of contextual variables $C = \{c_1, c_2 \dots c_n\}$. Each contextual variable c_i has its own domain $dom(c_i)$. Typically, $dom(c_i)$ is categorical. Formally, $dom(c_i) = \{v_1, v_2 \dots v_m\}$, where v_j is one of the m values allowed for the variable c_i. For example, if we consider as contextual variable the task to be accomplished, $dom(task) = \{studying, running, dancing \dots\}$. Clearly, many variables are not categorical: user location, for example, can be defined through GPS coordinates. However, in this work we just focused on those variables that can be modeled through a set of categorical values.

Pre-filtering: as pre-filtering approach, we adapted to CBRS the microprofiling technique proposed by Baltrunas and Amatriain [5]. The insight behind microprofiling is that the complete user profile, containing all the information about the preferences of the target user, can be split in several (micro) profiles containing only the information that the user expressed under a specific contextual situation. Intuitively, if the target user needs to receive suggestions about music to play during a party, it makes sense to build the recommendation set by taking into account only the preferences she expressed in that context. Formally, given a set of n contextual variables, each of which can assume m different values, the user profile (WRI or WQN) is split into at most $m \times n$ smaller microprofiles, according to the available ratings. The rating function is split as well, since user preferences can change in different contextual situations. Let $r(u, d_i, c_i, v_j)$ a contextual rating function that models the rating of user u on item d_i under the context v_j, where $v_j \in dom(c_i)$. We can define the set $I_u^+(c_i, v_j)$ as the set of the items the user likes in a specific context. Given these definition, we can define a contextual WRI profile for user U in the context v_j as:

$$\boldsymbol{preWRI(u, c_i, v_j)} = \sum_{i=1}^{|I_u^+(c_i,v_j)|} \boldsymbol{d}_i * r(u, d_i, c_i, v_j) \tag{3}$$

Due to space limitations the formula for building the negative counterpart $\boldsymbol{preWRI_{neg}(u, c_i, v_j)}$ is not provided, but it can be easily obtained from the previous one. Identically, $\boldsymbol{preWQN(u, c_i, v_j)}$ is obtained by combining both positive and negative microprofiles through quantum negation. As for uncontextual recommendations, given the vector representing a microprofile of the target user, the recommendation set is built by just calculating the cosine similarity between the profile and all the avaiable items. Since the profile is context-aware, the recommendations become context-awareas well.

Contextual Modeling: the insight behind microprofiling is very intuitive and easy to implement. However, it suffers from a clear issue: by splitting the whole user profile into several smaller profiles, it is likely that the available data are not enough to properly model user preferences. It's not by chance that several work in the state of the art [9,26] already tried to make the exact pre-filtering much more flexible and able to exploit data coming from other (similar) contextual situations. As a consequence, we introduced a novel contextual modeling approach that considers the context as a *weighting factor* that just influences the recommendation score for a certain item. Our insight is to combine the uncontextual vector space representation of user preferences $\boldsymbol{WRI(u)}$ or $\boldsymbol{WQN(u)}$ with a vector space representation of the *context* itselft. As vector space representation of the context we used $\boldsymbol{preWRI(u, c_i, v_j)}$, since it models the information coming from the items the user like in that specific context. Next, the contextual user profile is a linear combination of both vectors:

$$\boldsymbol{contextWRI(u, c_i, v_j)} = \alpha * \boldsymbol{WRI(u)} + (1 - \alpha) * \boldsymbol{preWRI(u, c_i, v_j)} \quad (4)$$

Intuitively, if the user didn't express any preference in that specific context the right part of the formula will be 0, so she will receive uncontextual recommendations. That makes sense, since we can state that it is a good choice to provide uncontextual recommendations if we don't have any evidence about user preferences in that context. Otherwise, the formula gives a greater weight to those preferences expressed in the target context, according to the weight α. As for pre-filtering, the negative counterpart $\boldsymbol{contextWQN(u, c_i, v_j)}$ or can be easily obtained, so it is omitted. Finally, given a contextual profile, we use the cosine similarity to extract the context-aware recommendations.

4 Experimental Evaluation

The goal of the experimental session was to evaluate the performances of Contextual eVSM in terms of predictive accuracy. Specifically, we designed two different experiments: in the first one we compared the effectiveness of contextual approaches with respect to their uncontextual counterparts, next, we compared our framework with another relevant state of the art approach. In order to obtain comparable experimental results, we adopted the same dataset as well as the same experimental design proposed by Adomavicius et al. [2].

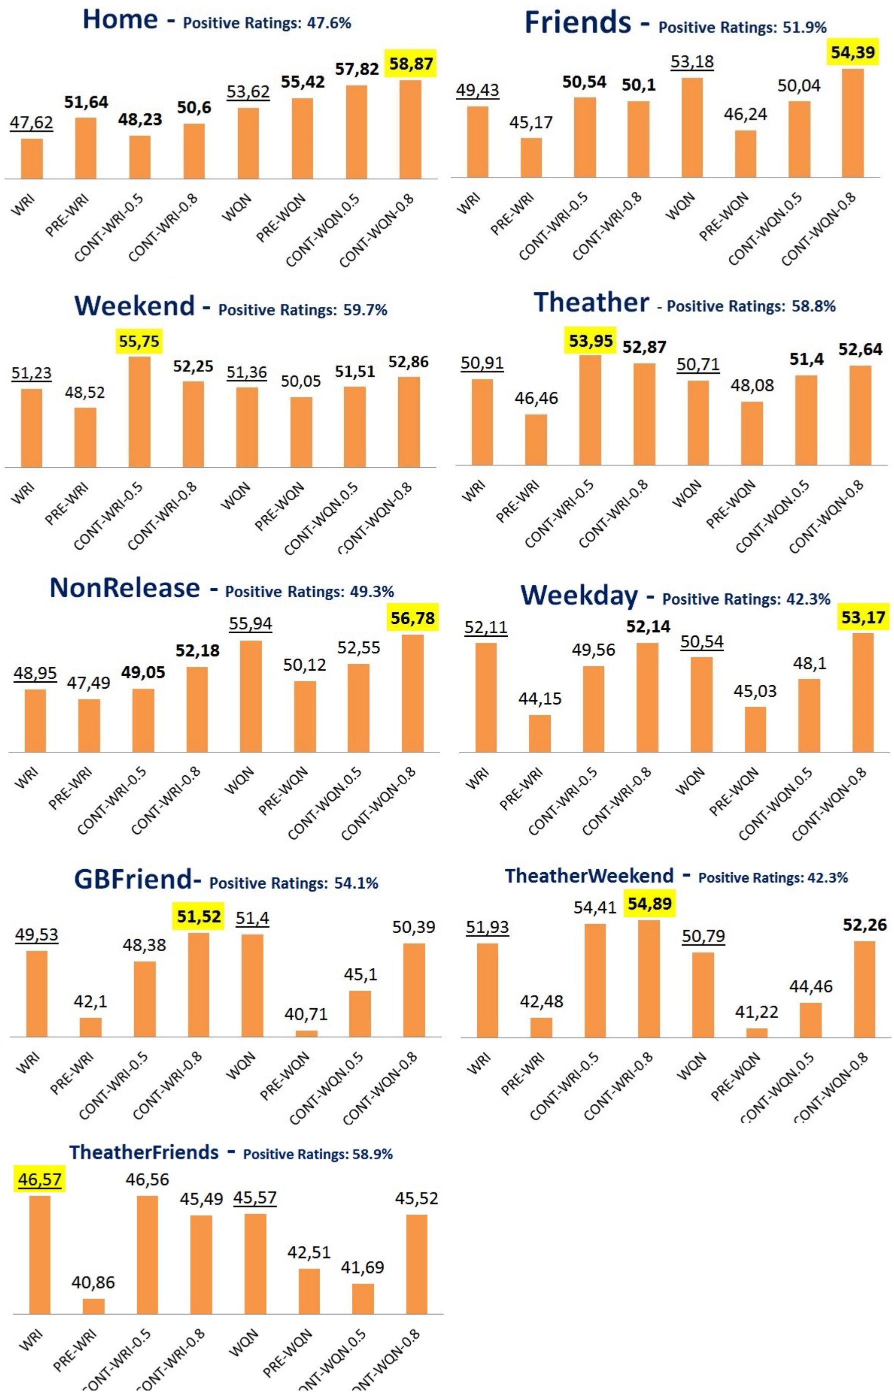

Fig. 1. Results of the experiments, split all over the contextual segments

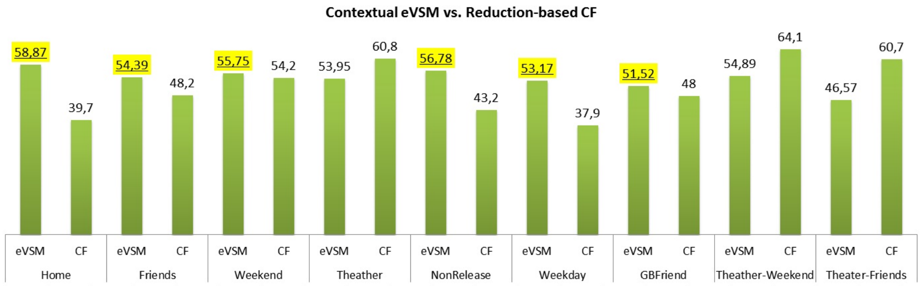

Fig. 2. Comparison with state of the art

In that papers the authors evaluated their context-aware recommender in the scenario of movie recommendation. They used a dataset crawled from IMDB[4], containing 1755 ratings coming from 117 users under different contextual situations. Specifically, four different categorical contextual variables were defined: TIME (weekday, weekend), PLACE (theather, home), COMPANION (alone, friends, boy/girlfriend, family) and MOVIE-RELATED (release week, non release week). The complete dataset has been further processed, as in [2], in order to filter all the ratings coming from users that didn't rated at least 10 movies. The final dataset contained 1457 ratings coming from 62 users on 202 movies. Since our contextual recommendation framework is a CBRS, we also gathered textual content from Wikipedia, by mapping the title of the movie with the title of the Wikipedia page. For each movie we extracted textual information about the plot, the abstract, the genre, the title, the director and the actors. According to the experimental protocol proposed by Adomavicius, the complete dataset was split into several overlapping subsets, called *contextual segments*. Each contextual segment modeled the ratings provided by the users under a specifical context, in order to evaluate the ability of the approach of providing users with good suggestions in specific contextual settings. The contextual segments containing less than 145 ratings (10% of the dataset) were filtered out. To sum up, our algorithms were evaluated against nine different contextual segments: HOME (727 ratings), FRIENDS (565 ratings), NON-RELEASE (551 ratings), WEEKEND (538 ratings), WEEKDAY (340 ratings), GBFRIEND (319 ratings), THEATHER-WEEKEND (301 ratings), THEATHER-FRIENDS (274 ratings). As experimental protocol we adopted the *bootstrapping* method [16]: for each contextual segment, 500 random re-samples were performed. In each sample 29/30th of the data were used as training and 1/30th as test. Each movie was rated on a 13-point discrete scale. All the ratings above 9 were considered as positive. Finally, we used Precision, Recall and F1-measure to evaluate the accuracy of the recommendation sets.We considered WRI and WQN as uncontextual baselines, and we compared them with both pre-filtering (PRE-WRI and PRE-WQN) and contextual modeling (CONTEXT-WRI and CONTEXT-WQN) configurations. In contextual modeling we also evaluated two different values of α, that is to say, 0.5 and 0.8.

[4] http://www.imdb.com/

To sum up, for each contextual segment 8 different configurations were evaluated. The results of the experimental sessions, in terms of F1 measure, are plotted in the next image. For each contextual segment, we underlined both the uncontextual baselines (WRI and WQN). The contextual configurations that overcame the baseline were put in bold, while the best-performing configuration was highlighted in yellow. The first outcome of the experimental evaluation is that the pre-filtering technique based on microprofiling does not improve the predictive accuracy of context-aware recommendations, since only in one segment out of nine (Home), both PRE-WRI and PRE-WQN got an improvement with respect to WRI and WQN. On the other side, it clearly emerges that our novel contextual modeling technique based on distributional semantics overcomes the baseline in 8 out of 9 segments with at least one setting. This outcome confirms those emerged from Adomavicius' experiment, since their context-aware recommender improved the F1 measure in 8 out of 9 contextual segmentsas well. Furthermore, results show that the configurations with α set to 0.8 generally got an higher F1 measure with respect to those with α=0.5. This suggests that user profiles should be modeled by giving a greater weight to user preferences, and by using contextual information just to slightly influence the recommendation score calculated by eVSM. Generally speaking, configurations with α=0.8 got the best F1 in 6 out of 8 segments. Another interesting outcome emerged by analyzing the relationship between the best-performing setting and dataset balance in terms of positive and negative ratings. Indeed, if a lot of negative evidence is available, results show that the configurations exploiting quantum negation overcome those that model only positive preferences; when the amount of positive ratings is under 52% the setting that obtains the best results is always the CONT-WQN-0.8. In all the other cases, the configurations without negation overcome those that model negative preferences. The usefulness of modeling negative evidences through our quantum negation operator further confirms the outcomes already discussed in [17] for uncontextual recommendations. In our second experimental setting we compared our best-performing configurations with the best-performing configuration coming out from Adomavicius' experiments. For the sake of clarity, it is necessary to underline that the results are just *partially comparable*: even if our work shares the same dataset as well the same experimental protocol, it is not possible to ensure that the generated samples are actually the same. However, the bigger the number of iterations, the bigger the likelihood that the results can be considered as comparable.

The comparison between contextual eVSM and the reduction-based approach proposed by Adomavicius et al. is provided in Figure 2. A quick analysis of the plot provides other interesting outcomes, since the results show that our approach clearly overcomes the state of the art algorithm in 6 out of 9 contextual segments. Even if the experiment has not been completed with a statistical test, it is likely that the difference between the algorithms is significant for most of the settings, since in 3 segments the gaps is over 10% in terms of F1-measure. This important result further confirms the goodness of the insights behind the contextual modeling approach integrated into eVSM.

5 Conclusions and Future Work

In this paper we proposed the CONTEXTUAL eVSM, a context-aware content-based recommendation framework based on VSM. Specifically, we investigated two different techniques to incorporate contextual information into CBRS: in the first one we adapted the well-known microprofiling approach to the CBRS scenario, while in the second one we introduced a novel contextual modeling approach that exploits *distributional semantics* to build a vector space representation of the context that is combined with a vector representing the preferences of the target user, in order to make the recommendation process context-aware. In the experimental evaluation the proposed approaches were evaluated against a state of the art dataset in order to determinate the best performing configuration, and it emerged that the approach based on distributional semantics can overcome both a non-contextual baseline as well as a state of the art algorithm for context-aware collaborative recommendation. In the future, we will continue through this preliminary experimental session by evaluating more values for the parameter α and by designing a statistical test to validate the outcomes presented in this paper. Furthermore, since CBRS heavily rely on textual content, we will investigate the integration of the information coming from Open Knowledge Sources, such as Wikipedia or the Linked Open Data cloud. Finally, we are going to plan a user study in order to analyze the impact of our recommendation framework on real users, in terms of predictive accuracy as well as user-centered metrics, such as novelty, diversity, serendipity and so on.

Acknowledgments. This work fullfils the research objectives of the project PON 01 00850 ASK-Health (Advanced System for the interpretation and sharing of knowledge in health care) funded by the Italian Ministry of Universty and Research (MIUR)

References

1. Adomavicius, G., Mobasher, B., Ricci, F., Tuzhilin, A.: Context-aware recommender systems. AI Magazine 32(3), 67–80 (2011)
2. Adomavicius, G., Sankaranarayanan, R., Sen, S., Tuzhilin, A.: Incorporating contextual information in recommender systems using a multidimensional approach. ACM Trans. Inf. Syst. 23(1), 103–145 (2005)
3. Adomavicius, G., Tuzhilin, A.: Multidimensional recommender systems: A data warehousing approach. In: Fiege, L., Mühl, G., Wilhelm, U.G. (eds.) WELCOM 2001. LNCS, vol. 2232, pp. 180–192. Springer, Heidelberg (2001)
4. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Recommender Systems Handbook, pp. 217–253. Springer (2011)
5. Baltrunas, L., Amatriain, X.: Towards Time-Dependant Recommendation based on Implicit Feedback. In: Context-aware Recommender Systems (CAMRa) Workshop at ACM Recsys (2009)
6. Bazire, M., Brézillon, P.: Understanding context before using it. In: Dey, A.K., Kokinov, B., Leake, D.B., Turner, R. (eds.) CONTEXT 2005. LNCS (LNAI), vol. 3554, pp. 29–40. Springer, Heidelberg (2005)

7. Birkhoff, G., Von Neumann, J.: The logic of quantum mechanics. The Annals of Mathematics 37, 823–843 (1936)
8. Cantador, I., Bellogín, A., Castells, P.: News@hand: A Semantic Web Approach to Recommending News. In: Nejdl, W., Kay, J., Pu, P., Herder, E. (eds.) AH 2008. LNCS, vol. 5149, pp. 279–283. Springer, Heidelberg (2008)
9. Codina, V., Ricci, F., Ceccaroni, L.: Semantically-enhanced pre-filtering for context-aware recommender systems. In: Proceedings of the 3rd Workshop on Context-awareness in Retrieval and Recommendation, pp. 15–18 (2013)
10. Deerwester, S.C., Dumais, S.T., Landauer, T.K., Furnas, G.W., Harshman, R.A.: Indexing by latent semantic analysis. Journal of the American Society of Information Science 41, 391–407 (1990)
11. Herlocker, J.L., Konstan, J.A.: Content-independent task-focused recommendation. IEEE Internet Computing 5(6), 40–47 (2001)
12. Jannach, D., Zanker, M., Felfernig, A., Friedrich, G.: Recommender systems: an introduction. Cambridge University Press (2010)
13. Karatzoglou, A., Amatriain, X., Baltrunas, L., Oliver, N.: Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In: Proceedings of ACM RecSys, pp. 79–86 (2010)
14. Linden, G., Smith, B., York, J.: Amazon.com Recommendations: Item-to-Item Collaborative Filtering. IEEE Internet Computing 7(1), 76–80 (2003)
15. Lops, P., de Gemmis, M., Semeraro, G.: Content-based recommender systems: State of the art and trends. In: Recommender Systems Handbook, pp. 73–105. Springer (2011)
16. Mitchell, T.M.: Machine Learning. McGraw-Hill Series in Computer Science. WCB/McGraw-Hill, Boston (1997)
17. Musto, C., Semeraro, G., Lops, P., de Gemmis, M.: Random indexing and negative user preferences for enhancing content-based recommender systems. In: Huemer, C., Setzer, T. (eds.) EC-Web 2011. LNBIP, vol. 85, pp. 270–281. Springer, Heidelberg (2011)
18. Musto, C., Semeraro, G., Lops, P., de Gemmis, M., Narducci, F.: Leveraging social media sources to generate personalized music playlists. In: Huemer, C., Lops, P. (eds.) EC-Web 2012. LNBIP, vol. 123, pp. 112–123. Springer, Heidelberg (2012)
19. Ostuni, V.C., Di Noia, T., Mirizzi, R., Romito, D., Di Sciascio, E.: Cinemappy: a context-aware mobile app for movie recommendations boosted by dbpedia. In: SeRSy. CEUR-WS, vol. 919, pp. 37–48 (2012)
20. Panniello, U., Gorgoglione, M.: Incorporating context into recommender systems: an empirical comparison of context-based approaches. Electronic Commerce Research 12(1), 1–30 (2012)
21. Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.): Recommender Systems Handbook. Springer, New York (2011)
22. Sahlgren, M.: An introduction to random indexing. In: Methods and Applications of Semantic Indexing Workshop, TKE 2005 (2005)
23. Salton, G.M., Wong, A.K.C., Yang, C.-S.: A Vector Space Model for Automatic Indexing. Comm. of the ACM 18(11), 613–620 (1975)
24. Turney, P.D., Pantel, P.: From frequency to meaning: Vector space models of semantics. J. Artif. Intell. Res. (JAIR) 37, 141–188 (2010)
25. Widdows, D.: Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In: ACL, pp. 136–143 (2003)
26. Zheng, Y., Burke, R., Mobasher, B.: Differential context relaxation for context-aware travel recommendation. In: Huemer, C., Lops, P. (eds.) EC-Web 2012. LNBIP, vol. 123, pp. 88–99. Springer, Heidelberg (2012)