

Combining Distributional Semantics and Entity Linking for Context-aware Content-based Recommendation

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

Department of Computer Science
University of Bari Aldo Moro, Italy

{cataldo.musto,giovanni.semeraro,pasquale.lops,marco.degemmis}@uniba.it

Abstract. The effectiveness of content-based recommendation strategies tremendously depends on the representation formalism adopted to model both items and user profiles. As a consequence, techniques for semantic content representation emerged thanks to their ability to filter out the noise and to face with the issues typical of keyword-based representations. This article presents Contextual eVSM (C-eVSM), a content-based context-aware recommendation framework that adopts a novel semantic representation based on distributional models and entity linking techniques. Our strategy is based on two insights: first, entity linking can identify the most relevant concepts mentioned in the text and can easily map them with structured information sources, easily triggering some inference and reasoning on user preferences, while distributional models can provide a lightweight semantics representation based on term co-occurrences that can bring out latent relationships between concepts by just analyzing their usage patterns in large corpora of data.

The resulting framework is fully domain-independent and shows better performance than state-of-the-art algorithms in several experimental settings, confirming the validity of content-based approaches and paving the way for several future research directions.

1 Introduction

Recommender Systems (RS) can support users in (real-time) decision-making by providing them with personalized access to digital content and services. However, classical personalization strategies may not be enough for some scenarios, since it is clear that people choices are also influenced by the *context* in which they have to be made. For example, the mood and the company could direct the choice of the movie to be watched. Thus, it is acknowledged that an effective recommendation algorithm can't ignore contextual information sources such as location, mood, task, company and so on, since all these factors clearly influence the perceived usefulness of a recommendation.

As a consequence, several techniques for generating context-aware recommendations recently came into play. State of the art approaches [1] are split in three main categories: *Pre-filtering* assumes that contextual information is used

to filter out irrelevant ratings *before* they are used for computing recommendations with standard methods. *Post-filtering* assumes that contextual information is used *after* the standard non-contextual recommendation methods are applied to the recommendation data. Finally, *contextual modeling* assumes that contextual information is used *inside* the recommendation algorithm together with the user and item data.

Regardless of the contextualization strategy, most of the current literature investigated context-aware collaborative filtering algorithms, especially those based on matrix factorization techniques [9]. On the other side, context-aware content-based strategies [10] did not receive the same attention, even if content-based approaches may be helpful to overcome the *new item* problem and to avoid scalability issues. Furthermore, differently from matrix factorization techniques which make the recommendation process similar to a *black box*, content-based profiles are typically more transparent and human-readable [15], thus a content-based algorithm can easily *explain* its recommendations. Finally, content-based strategies can also accelerate serendipitous encounters by triggering some reasoning and inference on the features stored in the user profile in order to introduce novel and unexpected concepts [15].

In this paper we present an extension of CONTEXTUAL EVSM [14], a content-based recommendation framework that adopts a *post-filtering* strategy to provide users with context-aware suggestions. The main contribution of this paper is the adoption of a novel representation based on both distributional semantics models [18] and entity linking techniques. The use of distributional models can provide items (and profiles) with a lightweight semantics representation based on term usages and concept co-occurrences in large corpus of data, while entity linking can be helpful to extract and identify the most relevant concepts in textual descriptions, in order to give them more emphasis when user preferences have to be modeled. Furthermore, entity linking algorithms may act as a bridge to connect free text with structured information sources, such as those coming from the Linked Open Data (LOD) cloud, thus making simpler to enable reasoning and inferences about user preferences.

The paper is organized as follows. Section 2 presents an overview of the most relevant work in the area of context-aware recommender systems. In Section 3 we focus the attention on CONTEXTUAL EVSM, by describing the novel semantic representation as well as the profiling strategies and recommendation step. Section 4 describes the experimental protocol and discusses the outcomes of the evaluation of the framework. Finally, Section 5 contains conclusions and sketches future directions of this research.

2 Related Work

The area of context-aware RSs (CARS) is quite recent, and has been fostered by the recent series of CARS¹ and CAMRa² workshops. The first proposal of a

¹ <http://cars-workshop.org/>

² <http://camrachallenge.com/>

context-aware recommendation algorithm dates back to the work by Herlocker and Konstan [8], who proposed to adapt the recommendation list to the specific task of the user. The most recent trends in the area of CARS were discussed in a recent survey [1]. The aforementioned classification in *pre-filtering*, *post-filtering* and *contextual-modeling* was introduced by Adomavicius and Tuzhilin [3].

Typically, the approaches for providing context-aware recommendations are based on the manipulation of the user \times item matrix: in [2], the authors propose a *multi-dimensional model* able to enrich that matrix with contextual information. In order to compare the results of our approach, in the experimental evaluation we used the same dataset and settings proposed in that paper. Similarly, Karatzoglou et al. [9], followed this research line by proposing a framework, called *multiverse recommendation*, in which different types of context are considered as additional dimensions in the representation of the data as a tensor. The factorization of this tensor leads to a compact model of the data which can be used to provide context-aware recommendations.

Our main contribution is the definition of a context-aware content-based recommendation framework. A similar attempt is described in [4], in which contextual information is used to improve the performance of a content-based news recommender system. The approach to context-aware recommendation proposed in this paper is inspired to the *weighted post-filtering* introduced in [16], since we used the vector space representation of the context as a weighting factor, that is combined with a vector space representation of user preferences. With respect to the current literature, the novelty of our approach lies in the fact that we exploit distributional models to build a semantic vector space representation of the context. The main insight behind distributional models [18] is that the semantics of a term can be inferred in a totally unsupervised way by just analyzing its usage patterns in a large corpus of textual data. According to this insight, our proposal is centered around the construction of a semantic vector space representation of the context, which leverages the usage of the terms adopted to label the items that are relevant under specific contextual constraints. Even if the effectiveness of these models has been acknowledged for several tasks [6], there is only a little evidence about their performance in content-based RSs [13]. Similarly to our work, Codina et al. [5] proposed a pre-filtering approach based on the intuition that ratings acquired in contextual situations similar to the target one may be used to predict missing ratings in that target contextual situation. The main differences of our approach compared to [5] are the use of the distributional semantics as a *weighting factor* to influence the final recommendation score, and the use of a content-based strategy for representing contexts.

3 Contextual eVSM

Contextual eVSM (C-eVSM) is a context-aware content-based recommendation framework. It extends the enhanced Vector Space Model (eVSM) [12] by taking into account contextual information as well. The recommendation process follows the workflow depicted in Figure 1. The only prerequisite of the algorithm is that

each item to be recommended is provided with some textual content (the plot of a movie or the the content of a news, generic metadata, etc.).

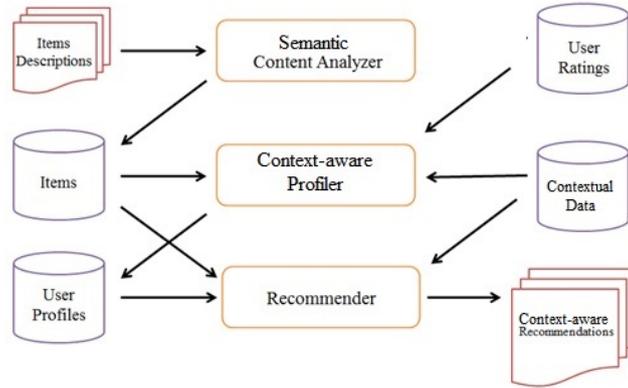


Fig. 1. Contextual eVSM workflow

The first step is performed by a SEMANTIC CONTENT ANALYZER, which processes textual descriptions through a pipeline of natural language processing (NLP) and entity linking algorithms, in order to identify the most relevant concepts found in the text. Next, by exploiting distributional semantic models, each item is represented as a (semantic) vector of co-occurrences of the concepts and their (latent) connections learned from the text. Afterwards, a CONTEXT-AWARE PROFILER builds a model of user preferences by collecting the semantic features describing the items liked in the past and combines them with contextual data. For example, if a user usually watches romantic movies on Friday night, the contextual profiler will store her preferences along with the information about the specific context the preference has been expressed. Finally, given a specific contextual setting (e.g., movie for the weekend), a RECOMMENDER module computes the relevance of an unseen item for a target user by matching the features stored in the contextual user profile with those describing the available items, with the assumption that a larger feature overlap corresponds to a greater degree of interest. In the following sections the whole pipeline will be thoroughly described.

3.1 Content Representation

The effectiveness of content-based recommendation strategies tremendously depends on the ability of the algorithm to filter out the noise from textual descriptions. As a consequence, in C-eVSM we paid attention to define a pipeline of state-of-the-art NLP and representation techniques able to produce a richer

and fine-grained semantic content representation. In our approach, each textual description has been processed through a sequence of *entity linking* algorithms. Specifically, we chose DBpedia Spotlight³, Wikipedia Miner⁴ and Tag.me⁵. The goal of these algorithms is to analyze the content and to identify and extract the most relevant concepts mentioned in the text. As an example, the resulting representation obtained for the movie Matrix is provided in Figure 2. The text processed by the algorithms is the plot of the movie gathered from Wikipedia.



Fig. 2. Entity-based Representation for the movie Matrix

It immediately emerges that such a representation, beyond being more transparent and human-readable, automatically incorporates stopwords removal, bigrams recognition as well as entities recognition and word sense disambiguation. Furthermore, it is worth to note that each concept extracted from the text is mapped to a univocal reference (in our case, a Wikipedia page). Given that each concept is mapped to a Wikipedia page, we decided to enrich the representation by performing a simple browsing on the Wikipedia categories tree. Specifically, for each concept extracted from the text the ancestor categories were included. In our case, given the concept *The Wachowskis*, the representation was enriched by adding as extra features concepts such as *Writers from Chicago* and *American Directors*, thus extending the representation with other relevant features that may be of interest for the target user.

However, even a fine-grained representation based on Wikipedia concepts does not take into account any information about the *semantics* of the terms. In order to tackle this issue, we adopted DISTRIBUTIONAL MODELS (DMs) to provide each item with a semantic vector space representation. DMs are based on a simple insight: as humans infer the meaning of a word by understanding the contexts in which that word is typically used, distributional models get information about the meaning of a word by analyzing its usage in large corpora of textual documents. This means that it is possible to infer the semantics of a term (e.g. *beer*) by analyzing the meaning of the other terms it typically co-occurs with (*wine*, *glass*, etc.) [17]. In the same way, the correlation between different

³ <http://dbpedia-spotlight.github.io/demo/>

⁴ <http://wikipedia-miner.cms.waikato.ac.nz/>

⁵ <http://tagme.di.unipi.it/>

terms (e.g., *beer* and *party*) can be inferred by analyzing the similarity between the contexts in which they are used. These approaches rely on the distributional hypothesis [7], according to which *Words that occur in the same contexts tend to have similar meanings*. This means that words are semantically similar to the extent that they share contexts. DMs represent information about terms usage in a term-context matrix (Figure 3). The advantage is that the context is very flexible, and can be adapted to the specific granularity level of the representation required by the application: for example, given a word, its context could be either a single word it co-occurs with, or a sliding window of terms that surrounds it, or a sentence, or yet the whole document.

		c1	c2	c3	c4	c5	c6	c7	c8
beer		✓		✓	✓				✓
glass		✓		✓			✓		✓
wine		✓			✓				✓
spoon			✓				✓	✓	

Fig. 3. An example of term-context matrix

In our setting, we exploited distributional models to build a semantic representation of each item. First, each Wikipedia concept found in the text is represented as a vector according to the co-occurrences with the other Wikipedia concepts. Next, we easily represented each item as a (semantic) vector by combining the vector space representation of all the Wikipedia concepts that describe it. The combination is obtained as a simple sum of the vectors, weighted through the relevance of each Wikipedia concept in the description. In this way, we combined the advantages of a fine-grained representation based on entity linking algorithms with the lightweight semantics provided by the adoption of distributional models, thus encoding semantic information without the need of implementing a complex word-sense disambiguation strategy.

3.2 Contextual Profiling and Recommendation

In C-eVSM we defined the *context* as a set of variables $C = \{c_1, c_2 \dots c_n\}$. Each contextual variable c_k has its own domain $dom(c_k)$, which is typically categorical. Formally, $dom(c_k) = \{v_1, v_2 \dots v_m\}$, where v_j is one of the m values allowed for the variable c_k . For example, if we consider as contextual variable the *task* to be accomplished, $dom(task) = \{studying, running, dancing \dots\}$.

According to this definition, we defined two different contextual profiling models, which extend those already introduced in [12]. Both models are based

on a simple insight: as DMs learn a vector-space representation of the items according to the co-occurrences between the terms that describe them, a semantic vector space representation of the *context* can be built according to the co-occurrences between the concepts that more frequently describe the items labeled as relevant in that specific contextual settings. In other terms, our assumption is that there exists a set of terms that is likely to be more descriptive of items relevant in a certain context. For example, if a user is looking for a restaurant for a romantic night, it is likely that restaurant descriptions containing terms such as *candlelight* or *sea view*, are more suitable and thus relevant in that specific context. Hence, we decided to deeply analyze the usage patterns of terms describing items relevant in different contextual situations, in order to learn a representation of the context based on those terms more relevant in that contextual setting. Up to our knowledge this is a novel contribution in the area of context-aware content-based recommender systems.

Formally, given a user u and contextual variable c_k which assumes values v_j (e.g. *task = running*), the contextual user profile can be defined as a linear combination, tuned by a parameter α , of $\mathbf{WRI}(u)$ and $\mathbf{context}(u, c_k, v_j)$:

$$\mathbf{C} - \mathbf{WRI}(u, c_k, v_j) = \alpha * \mathbf{WRI}(u) + \quad (1)$$

$$(1 - \alpha) * \mathbf{context}(u, c_k, v_j) \quad (2)$$

$\mathbf{WRI}(u)$ is an uncontextual representation of user preferences, defined as the sum over the vector space representation of the items the user liked in the past, labeled as I^+ , weighted with the normalized rating $r(u, i)$:

$$\mathbf{WRI}(u) = \sum_{i=1}^{|I^+|} \mathbf{d}_i * \frac{r(u, i)}{MAX} \quad (3)$$

and $\mathbf{context}(u, c_k, v_j)$ is the vector space representation of the context, defined as the weighted sum over the items labeled as relevant under that specific contextual setting:

$$\mathbf{context}(u, c_k, v_j) = \sum_{i=1}^{|I_u^+(c_k, v_j)|} \mathbf{d}_i * \frac{r(u, d_i, c_k, v_j)}{MAX} \quad (4)$$

Next, we defined a second contextual profiling model, called C-WQN, which encodes the information coming from the items the user disliked and adopts a QUANTUM NEGATION operator to combine positive and negative profile vectors in a single uniform representation. The discussion about the negation operator is out of the scope of this paper. For further reading refer to [13], which contains a discussion about the application of the operator in content-based recommender systems, or refer to the original Widdows' paper [19]. Finally, given a contextual setting and a semantic vector space representation of both items and user

profiles, C-eVSM generates the suggestions by calculating the *cosine similarity* between the vector representing user preferences in that specific context and the vector representing the item. The recommended items may be the top-N or those whose similarity overcomes a certain threshold.

4 Experimental Evaluation

In order to validate the performance of our framework, we carried out an extensive experimental session. The experiments had a twofold goal:

1. Evaluation of the predictive accuracy of the contextual eVSM model with respect to both a non-contextual baseline and a contextual baseline based on a simple keyword-based representation;
2. Comparison of the predictive accuracy of the contextual eVSM model with respect to state of the art algorithms.

In order to compare the accuracy of our framework with respect to state of the art algorithms, we adopted the same dataset and experimental design proposed by Adomavicius et al. [2]. In that paper, the authors evaluated their context-aware recommender system in a movie recommendation scenario. They used a dataset crawled from IMDB⁶, containing 1,755 ratings expressed by 117 users under different contextual situations. Specifically, four different categorical contextual variables were defined: TIME (weekday, weekend), PLACE (theater, 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], to filter all the ratings coming from users that did not rate at least 10 movies. The final dataset contained 1,457 ratings coming from 62 users on 202 movies. Given that our framework needs some textual content to feed the recommendation algorithm, we gathered textual information from Wikipedia. For each movie we crawled the plot, the abstract, the genre, the title, the director and the actors. Textual content was preprocessed by just filtering out stopwords according to a common stopwords list for English language.

As in the experimental protocol proposed in [2], we split the dataset into several overlapping subsets, called *contextual segments*. Each contextual segment modeled the ratings provided by the users under a specific context, in order to evaluate the ability of the approach to provide users with accurate suggestions in specific contextual settings. The contextual segments containing less than 145 ratings (10% of the dataset) were filtered out. To sum up, our algorithm was 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), THEATER-WEEKEND (301 ratings), THEATER-FRIENDS (274 ratings). In all the experiments, we adopted the *bootstrapping* method [11] as experimental protocol: for each contextual segment, 500 random re-samples were performed. In each sample 29/30th of data

⁶ <http://www.imdb.com/>

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.

4.1 Discussion of the Results

In the first experiment we investigated the effectiveness of our novel representation based on distributional semantics and entity linking against both a non-contextual baseline and a contextual variant of C-eVSM based on just distributional semantics (without entity linking). We considered WRI and WQN as non-contextual baselines, and we compared them with contextual eVSM profiles (C-WRI and C-WQN) configurations. In this preliminary evaluation we set α to 0.2, 0.5 and 0.8, without implementing a specific optimization strategy. We labeled as *recommended* all the items whose cosine similarity overcame 0.6. To sum up, for each contextual segment 8 different configurations were evaluated. Results were validated using a paired t-test ($p \ll 0.05$). Results reporting F1 measure are provided in Table 1.

The main outcome is that our novel representation based on both distributional semantics and entity linking improves the predictive with respect to both baselines. Indeed, if we compare the best performing contextual configuration (C-WRI and C-WQN) with the non-contextual counterparts (WRI and WQN), regardless of the adopted representation, it emerges that Contextual eVSM overcomes the baseline in all contextual segments. This improves Adomavicius' results, since their context-aware recommender overcame the baseline in 8 out of 9 contextual segments. Furthermore, results show that the configurations with $\alpha = 0.8$ generally outperform those with $\alpha=0.5$ and $\alpha=0.2$ in terms of F1. This suggests that user preferences, regardless the contextual situation, still play a key role, while the context has to be used to slightly influence the recommendation score computed by eVSM. As regards content representation, results show that the combined use of entity linking and distributional semantics can provide a richer and meaningful representation. Indeed, by comparing entity-based to keyword-based results, it emerges an overall increase of +4% on average (ranging from +1.34% to +6.42%) on 8 segments out of 9. By analyzing the single configurations, it clearly emerges that the use of entity linking algorithms can provide a more precise representation since the accuracy is improved in 72% of the comparisons (65 out of 90) and in 58% with a statistically significant gap (52 out of 90). Generally speaking, the representation combining both DMs and entity linking got the best results in 7 segments out of 9. It is also worth to note that in 5 segments out of 7 the configurations using Quantum Negation got the best results. This suggests that our negation operator that combines into a single representation both positive and negative preferences can provide the best results. Next, we compared the previous results with the best-performing ones of the experiments reported in [2].

The comparison between contextual eVSM using the entity linking algorithm based on Wikipedia and the reduction-based approach proposed by Adomavicius et al. is provided in Figure 4. Results show that our approach outperforms the state of the art algorithm in 7 out of 9 contextual segments. It is necessary a

further investigation to understand why our approach is outperformed for the segments THEATER-WEEKEND and THEATER-FRIENDS. Even if the experiment has not been completed with a statistical test, it is likely that the difference between the algorithms is significant since in 4 segments the gap is over 10% in terms of F1-measure. This important result further confirms the validity of the CONTEXTUAL EVSM approach, especially when compared to a context-aware algorithm based on *collaborative filtering*.

	HOME (avg. + 3.77%)		FRIENDS (avg. + 6.42%)		WEEKEND (avg. + 6.03%)	
	Keyword	Entities	Keyword	Entities	Keyword	Entities
WRI	47.62	56.13(*)	49.43	56.17(*)	51.23	58.12(*)
C-WRI-0.2	44.56	46.62	44.91	49.89(*)	55.1	57.44
C-WRI-0.5	48.23	56.38(*)	50.54	56.24(*)	55.75	58.36
C-WRI-0.8	50.61	56.74(*)	50.11	55.68(*)	52.25	58.75(*)
WQN	53.62	57.53(*)	53.18	58.25(*)	51.36	+57.63(*)
C-WQN-0.2	53.37	54.81	45.93	52.82(*)	48.36	54.44(*)
C-WQN-0.5	57.82	61.96(*)	50.04	57.20(*)	51.51	59.49(*)
C-WQN-0.8	58.91	61.3	54.39	58.37(*)	52.86	60.39(*)
	THEATER (avg. + 6.49%)		NONRELEASE (avg. + 1.51%)		WEEKDAY (avg. -2.75%)	
	Keyword	Entities	Keyword	Entities	Keyword	Entities
WRI	50.91	56.34(*)	48.95	52.16(*)	52.11	46.4
C-WRI-0.2	52.79	55.82(*)	48.24	44.99	47.12	46.42
C-WRI-0.5	53.95	57.21(*)	49.05	53.29(*)	49.56	46.77
C-WRI-0.8	52.87	58.11(*)	52.18	53.34	52.14	46.34
WQN	50.71	56.45(*)	55.94	54.94	50.54	47.32
C-WQN-0.2	46.65	54.72(*)	48.67	45.416	44.71	46.2
C-WQN-0.5	51.4	57.81(*)	52.55	57.01(*)	48.1	48.02
C-WQN-0.8	52.64	62.16(*)	56.78	56.24	53.17	47.57
	GBFRIENDS (avg. + 1.34%)		THEATER-WEEKEND (avg. + 3.17%)		THEATER-FRIENDS (avg. + 5.22%)	
	Keyword	Entities	Keyword	Entities	Keyword	Entities
WRI	49.53	50.31	51.93	51.35	47.57	46.66
C-WRI-0.2	47.28	48.19	51.64	54.64(*)	46.68	55.35(*)
C-WRI-0.5	48.38	50.72	54.41	54.95	46.56	55.96(*)
C-WRI-0.8	51.5	49.96	54.89	52.38	45.49	46.63
WQN	51.4	49.13	50.79	49.1	48.7	48.63
C-WQN-0.2	43.25	48.44(*)	41.31	54.79(*)	39.36	55.37(*)
C-WQN-0.5	45.1	50.43(*)	44.46	52.85(*)	41.69	55.92(*)
C-WQN-0.8	50.39	48.42	52.26	49.77	45.52	48.93(*)

Table 1. Results of Experiment 1. Configurations that outperform the keyword-based baseline in terms of F1 in bold. Statistical differences highlighted with (*). The best-performing configuration with a grey background.

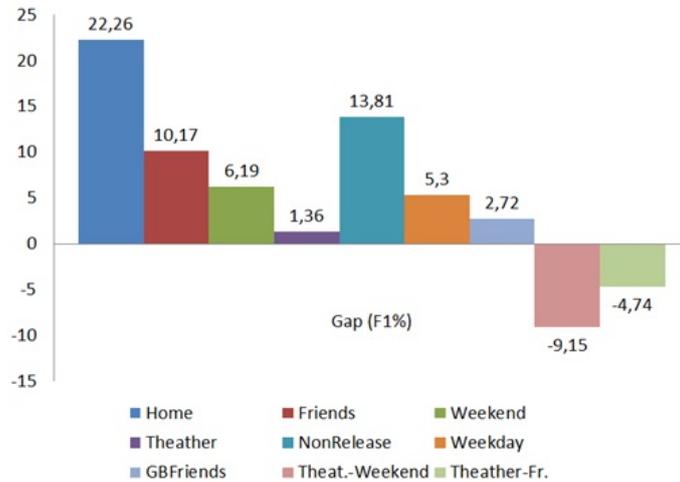


Fig. 4. Results of Experiment 2. The plot shows the gap between CONTEXTUAL EVSM and CF algorithm in terms of F1-measure.

5 Conclusions and Future Work

In this paper we proposed contextual eVSM, a context-aware content-based recommendation framework based on VSM. The original model has been extended by introducing a technique that exploits distributional semantics to build a semantic vector space representation of the context which is used to slightly influence a non-contextual recommendation score. Our novel context-aware technique has been evaluated against several non-contextual baselines, and our approach significantly improved the predictive accuracy of the baselines in most of the experiments. As a future work we will investigate the integration of information coming from the datasets available in the Linked Open Data (LOD) cloud. Indeed, our novel Wikipedia-based representation can be easily mapped to the information stored in the LOD cloud. We will also investigate possible improvements of the current experimental results by implementing some strategies for optimizing parameters in our algorithm, such as α in Equation 2. Finally, we are going to design a user study to investigate the impact of our recommendation framework on real users, in terms of predictive accuracy as well as user-centered metrics, such as novelty, diversity, and serendipity.

Acknowledgments. This work fulfils the research objectives of the projects PON 01 00850 ASK-Health (Advanced System for the interpretation and sharing of knowledge in health care) and PON 02 00563 3470993 project "VINCENTE - A Virtual collective INtelligenCe ENvironment to develop sustainable Technol-

ogy Entrepreneurship ecosystems” funded by the Italian Ministry of University and Research (MIUR)

References

1. G. Adomavicius, B. Mobasher, F. Ricci, and A. Tuzhilin. Context-aware recommender systems. *AI Magazine*, 32(3):67–80, 2011.
2. G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.*, 23(1):103–145, 2005.
3. G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In *Recommender Systems Handbook*, pages 217–253. Springer, 2011.
4. I. Cantador, A. Bellogín, and P. Castells. News@hand: A Semantic Web Approach to Recommending News. In *Adaptive Hypermedia and Adaptive Web-Based Systems*, volume 5149 of *LNCS*, pages 279–283. Springer, 2008.
5. V. Codina, F. Ricci, and L. Ceccaroni. Semantically-enhanced pre-filtering for context-aware recommender systems. In *Proceedings of the 3rd Workshop on Context-awareness in Retrieval and Recommendation*, pages 15–18, 2013.
6. T. Cohen and D. Widdows. Empirical distributional semantics: Methods and biomedical applications. *Journal of Biomedical Informatics*, 42(2):390, 2009.
7. Z. Harris. Mathematical structures of language. John Wiley & Sons, 1968.
8. J. Herlocker and J. Konstan. Content-independent task-focused recommendation. *IEEE Internet Computing*, 5(6):40–47, 2001.
9. A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of RecSys 2010*, pages 79–86. ACM, 2010.
10. P. Lops, M. de Gemmis, and G. Semeraro. Content-based recommender systems: State of the art and trends. In *Recommender Systems Handbook*, pages 73–105. Springer, 2011.
11. T. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
12. C. Musto. Enhanced vector space models for content-based recommender systems. In *Proceedings of RecSys 2010*, pages 361–364. ACM, 2010.
13. C. Musto, G. Semeraro, P. Lops, and M. de Gemmis. Random indexing and negative user preferences for enhancing content-based recommender systems. In *EC-Web 2011*, volume 85 of *LNBIP*, pages 270–281. Springer, 2011.
14. C. Musto, G. Semeraro, P. Lops, and M. de Gemmis. Contextual eVSM: A content-based context-aware recommendation framework based on distributional semantics. In *EC-Web*, volume 152 of *LNBIP*, pages 125–136. Springer, 2013.
15. F. Narducci, C. Musto, G. Semeraro, P. Lops, and M. de Gemmis. Leveraging encyclopedic knowledge for transparent and serendipitous user profiles. In *UMAP 2013*, volume 7899 of *LNCS*, pages 350–352. Springer, 2013.
16. U. Panniello and M. Gorgoglione. Incorporating context into recommender systems: an empirical comparison of context-based approaches. *Electronic Commerce Research*, 12(1):1–30, 2012.
17. H. Rubenstein and J. B. Goodenough. Contextual correlates of synonymy. *Commun. ACM*, 8(10):627–633, October 1965.
18. P. Turney and P. Pantel. From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res. (JAIR)*, 37:141–188, 2010.
19. D. Widdows. Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *ACL*, pages 136–143, 2003.