

A semantic approach to remove incoherent items from a user profile and improve the accuracy of a recommender system

Roberto Saia¹ · Ludovico Boratto¹ · Salvatore Carta¹

Received: 29 July 2015 / Revised: 29 March 2016 / Accepted: 31 March 2016 /

Published online: 25 April 2016

© Springer Science+Business Media New York 2016

Abstract Recommender systems usually suggest items by exploiting all the previous interactions of the users with a system (e.g., in order to decide the movies to recommend to a user, all the movies she previously purchased are considered). This canonical approach sometimes could lead to wrong results due to several factors, such as a change in user preferences over time, or the use of her account by third parties. This kind of incoherence in the user profiles defines a lower bound on the error the recommender systems may achieve when they generate suggestions for a user, an aspect known in literature as *magic barrier*. This paper proposes a novel dynamic coherence-based approach to define the user profile used in the recommendation process. The main aim is to identify and remove, from the previously evaluated items, those not semantically adherent to the others, in order to make a user profile as close as possible to the user's real preferences, solving the aforementioned problems. Moreover, reshaping the user profile in such a way leads to great advantages in terms of computational complexity, since the number of items considered during the recommendation process is highly reduced. The performed experiments show the effectiveness of our approach to remove the incoherent items from a user profile, increasing the recommendation accuracy.

This work is partially funded by Regione Sardegna under project NOMAD (Next generation Open Mobile Apps Development), through PIA - Pacchetti Integrati di Agevolazione "Industria Artigianato e Servizi" (annualità 2013), and by MIUR PRIN 2010-11 under project "Security Horizons".

✉ Ludovico Boratto
ludovico.boratto@unica.it

Roberto Saia
roberto.saia@unica.it

Salvatore Carta
salvatore@unica.it

¹ Dipartimento di Matematica e Informatica, Università di Cagliari, Via Ospedale 72 - 09124 Cagliari, Italy

Keywords User profiling · Semantic analysis · Magic barrier · Accuracy

1 Introduction

The rapid growth of the number of companies that sell goods through the Word Wide Web generates an enormous amount of valuable information, which can be exploited to improve the quality and efficiency of the sales criteria (Schafer et al. 1999). Because of the widely-known information overload problem, it became necessary to deal with the large amounts of data available on the Web (Wei et al. 2014). *Recommender Systems* (RS) (Ricci et al. 2011) represent an effective response to this problem, by filtering the huge amount of information about the customers in order to get useful elements to produce suggestions to them (Armano and Vargiu 2010; Addis et al. 2010; Vargiu et al. 2013). The denomination RS denotes a set of software tools and techniques providing to a user suggestions for items, where the term *item* is used to indicate what the system recommends to the users. In this paper we address one of the most important aspects related to the recommender systems, i.e., *how to represent a user profile, so that it only contains accurate information about a user, allowing a system to generate effective recommendations*.

The main motivation behind this work is that most of the solutions regarding the *user profiling* in the recommender systems context involve the interpretation of the whole set of items previously evaluated by a user, in order to measure their similarity with those that she did not consider yet, and recommend the most similar ones. Indeed, the recommendation process is usually based on the principle that user preferences remain unchanged over time and this can be true in many cases, but it is not the norm due to the existence of temporal dynamics in their preferences (Lam et al. 1996; Li et al. 2007; Widyanoro et al. 2001). Therefore, a static approach to user profiling can lead toward wrong results due to various factors, such as a simple change of preferences over time or the temporary use of their own account by other people.

Several works (Amatriain et al. 2009; Hill et al. 1995) have shown that some user ratings might be considered as outliers, due to the fact that the same user may rate the same item with different ratings, at different moments in time. This is a well-known problem, which in the literature is defined as *magic barrier* (Bellogín et al. 2014; Herlocker et al. 2004; Said et al. 2012), a term used to identify the point where, due to the noise in the data, the performance and accuracy of an algorithm cannot be further improved. After the magic barrier has been reached, any improvement in terms of accuracy might mean an overfitting instead of a performance enhancement. Thus, the primary aim of the approach we introduce is to measure the similarity between a single item and the others in the user profile, in order to improve the recommendation process by discarding the items that are highly dissimilar with the others in the user profile. In the literature, the coherence of an item with respect to the user profile is usually measured as the variance in the feature space that defines the items, typically based on the ratings given by the users (Bellogín et al. 2014). This is done by employing several metrics, such as the entropy, the mean value, or the standard deviation. Differently from the approaches at the state of the art, in this paper we consider the semantic distance between the concepts expressed by each item in a user profile, and the concepts expressed by the other ones.

This way to proceed presents a twofold advantage: firstly, it allows us to evaluate the coherence of an item in a more extensive way (i.e., by employing semantic concepts) w.r.t.

a limited mathematical approach; secondly, it reduces the cause of the *magic barrier* problem. This happens because the assumption of the *magic barrier* problem is the presence of incoherent items in the user profiles. Considering that our approach removes them, keeping in the user profiles only those items that are coherent with each other, we can consider any observed improvement as real, instead that a mere side effect (i.e., an overfitting).

To perform the task of removing semantically incoherent items from a user profile, we introduce the *Dynamic Coherence-Based Modeling* (DCBM), an algorithm based on the concept of *Minimum Global Coherence* (MGC), a metric that allows us to measure the semantic similarity between a single item and the others within the user profile. Moreover, the algorithm takes into account two factors, i.e., the position of each item in the chronology of the user choices, and the distance from the *mean value* of the *global similarity* (as “global” we mean all the items in a user profile). These metrics allow us to remove in a selective way any item that could make the user profiles non-adherent to the real preferences of the users. The main idea is that the more information in the user profile is coherent, the more the recommendations based on this profile will be reliable. Through our approach, the evaluation process of the items coherence has been moved from a domain based on rigorous mathematical criteria (i.e., variance of the user ratings in the feature space), to a new semantic domain, which presents a considerable advantage in terms of evaluation flexibility. An important aspect related with the proposed approach is its ability to operate in any domain characterized by a textual description of the items, even in absence of a user rating about them (e.g., implicit user preferences collected from the users’ browsing sessions).

In order to validate the capability of our approach to produce accurate user profiles, we are going to include the DCBM algorithm into state-of-the-art recommender systems (i.e., SVD (Koren and Bell 2015) and a User-Based Collaborative Filtering approach (Ning et al. 2015)) and evaluate the accuracy of the performed recommendations. Since the task of a recommender system that predicts the interest of the users for the items relies on the information included in a user profile, more accurate user profiles lead to an improved accuracy of the whole recommender system. Experimental results show its capability to remove the incoherent items from the user profiles, increasing the accuracy of the recommendations and reducing the computational complexity of the system, also in the context of the non-semantic approaches as those taken into account (i.e., SVD and User-Based Collaborative Filtering approaches).

The contributions of the paper are summarized as follows:

- we introduce a novel algorithm to remove incoherent items from a user profile, with the aim to improve the recommendation accuracy;
- we integrate our algorithm into state-of-the-art recommender systems, in order to improve their effectiveness and validate our proposal;
- we performed experiments on two real-world datasets and two state-of-the-art recommender systems, which compare the accuracy of a recommender system before and after the use of our approach.

This paper is based on the work presented in (Saia et al. 2014), which was completely rewritten and extended in the following ways: (i) we provide a formal notation and definition of the problem we tackle in this work, (ii) we extend the presentation of the background concepts behind this proposal, (iii) we provide more details on the proposed approach, (iv) we compare our proposal against two state-of-the-art recommender systems and two datasets.

The rest of the paper is organized as follows: Section 2 presents related work on user profiling; in Section 3 we introduce the background on the concepts and the problem handled by our proposal; Section 4 contains the formal definition of our problem, the details of the DCBM algorithm, and its integration into a recommender system; Section 5 presents the experimental framework used to evaluate our approach; Section 6 contains conclusions and future work.

2 Related work

In order to lead the potential buyers toward a number of well-targeted suggestions, related to the large amount of goods or services available today through the electronic commerce circuits, a *Recommender System* plays a determinant role, since it is able to investigate on the user preferences, suggesting them only potentially interesting items. In order to identify them, a recommender system has to *predict* that an item is worth recommending (Ricci et al. 2011).

The first class of systems that was developed employs the so-called *Collaborative Filtering* approach (Hill et al. 1995; Karypis 2001; Linden et al. 2003; Sarwar et al. 2001), which is based on the assumption that users have similar preferences on a item, if they already have rated other similar items (Su and Khoshgoftaar 2009). *Content-based recommender systems*, instead, suggest to users items whose content is similar to that of the items they previously evaluated (Lops et al. 2011; Pazzani and Billsus 2007). The early systems used relatively simple retrieval models, such as the Vector Space Model, with the basic TF-IDF weighting (which is presented in detail in the next section). Examples of systems that employ this type of content filtering are (Balabanovic and Shoham 1997; Billsus and Pazzani 1999; Lieberman 1995; Pazzani et al. 1996). Due to the fact that the approach based on a simple bag of words is not able to perform a semantic disambiguation of the words in an item description, content-based recommender systems evolved and started employing external sources of knowledge (e.g., ontologies) and semantic analysis tools, to improve their accuracy (Capelle et al. 2012; Capelle et al. 2013).

When producing personalized recommendations to users, the first requirement is to understand the needs of the users and, according to them, to build a user profile that models these needs. There are several approaches to create user profiles: some of them focus on *short-term* user profiles that capture features of the user's current search context (Budzik and Hammond 2000; Finkelstein et al. 2002; Shen et al. 2005), while others accommodate *long-term* profiles that capture the user preferences over a long period of time (Asnicar and Tasso 1997; Chirita et al. 2005; Ma et al. 2007). As shown in (Widyantoro et al. 2001), compared with the *short-term* user profiles, the use of a *long-term* user profile generally produces more reliable results, at least when the user preferences are fairly stable over a long time period. Otherwise, we need a specific strategy able to manage the changes in the user profile that do not reflect the preferences of the user and that represent a form of "noise".

The most common strategies to get useful information to build the user profiles are two, i.e., explicit or implicit. Explicit profiling strategies directly request to the users different forms of preference information, from categorical preferences (Chirita et al. 2005; Ma et al. 2007) to simple result ratings (Asnicar and Tasso 1997). Instead, implicit profiling strategies attempt to infer this information by analyzing the users behavior, and without a direct interaction with them while they perform actions in a website (Chirita et al. 2005; Liu et al. 2002; Pretschner and Gauch 1999).

However, the strategy usually adopted is the implicit one, where the user preferences are inferred without a direct interaction with her. This is the common approach because the explicit strategy presents some problems, such as those related with the *privacy* aspects (many users do not like to reveal information about their preferences), and those related with the form filling process (many users do not like to spend their time for this activity, and it is proved that the accuracy of the information depends on the time needed to provide them).

The implicit approach usually requires *long-term* user profiles, where the information about the preferences is considered over an extended period of time. However, there are some implicit approaches that involve a short-term profiling, related to the particular context in which the system operates (Shen et al. 2005).

Regardless of the type of profiling that is adopted (e.g., *long-term* or *short-term*), there is a common problem that may affect the goodness of the obtained results, i.e., the capability of the information stored in the user profile to lead toward reliable recommendations. In order to face the problem of dealing with unreliable information in a user profile, the state of art proposes different strategies. Several approaches, such as (Lam et al. 1996), take advantage from the Bayesian analysis of the user provided relevance feedback, in order to detect non-stationary user interests. Also exploiting the feedback information provided by the users, other approaches such as (Widyantoro et al. 2001) make use of a *tree-descriptor* model to detect shifts in user interests. Another technique exploits the knowledge captured in an ontology (Schickel-Zuber and Faltings 2006) to obtain the same result, but in this case it is necessary that the users express their preferences about items through an explicit rating.

There are also other different strategies that try to improve the accuracy of the information in the profiles by collecting the implicit feedbacks of the users during their natural interactions with the system (reading-time, saving, etc.) (Kelly and Teevan 2003). However, it should be pointed out that most of the strategies used in this area are effective only in specific contexts, such as for instance (Zeb and Fasli 2011), where a novel approach to automatically model the user profile, according to the change of her preferences, is designed for the articles recommendation context.

With regard to the analysis of information related to user profiles and items, there are several ways to operate and most of them work by using the *bag-of-words* model, an approach where the words are processed without taking into account the correlation between the terms (Lam et al. 1996; Widyantoro et al. 2001). This trivial way to manage the information usually does not lead toward good results, and just for this reason there are some more sophisticated alternatives, such as the semantic analysis of the content in order to model the preferences of a user (Pedersen et al. 2004). In (Stilo and Velardi 2014a, b, 2015), the problem of modeling semantically correlated items was tackled, but the authors consider a temporal correlation and not the one between the items and a user profile.

It should be noted that there is a common issue that afflicts the recommendation approaches, related to the concept of item incoherence. This is a problem that in the literature is identified as *magic barrier* (Herlocker et al. 2004), a term used to define the theoretical boundary for the level of optimization that can be achieved by a recommendation algorithm on transactional data (Said et al. 2012). The evaluation models assume as a ground truth that the transactions made in the past by the users, and stored in their profiles, are free of noise. This is a concept that has been explored in (Amatriain et al. 2009; Amatriain et al. 2009), where a study aimed to capture the noise in a service that operates in a synthetic environment was performed. It should be noted that this is an aspect that, in the context of the recommender systems, was mentioned for the first time in 1995, in a work aimed at discussing the concept of reliability of users in terms of rating coherence (Hill et al. 1995).

Our approach differs from the others in the literature, since it does not need to focus on a specific type of profile (i.e., *short-term* or *long-term*), but it can operate with any type of data that contains a textual description, overcoming the limitation introduced by the magic barrier from a novel perspective, represented by the semantic analysis of the items.

3 Background

In this section we provide some details concerning two key concepts involved in the context of this work, i.e., the spatial representation of a text document based on the *Vector Space Model*, and the functionalities offered by the *WordNet* environment.

3.1 Vector space model

Many content-based recommender systems use relatively simple retrieval models (Lops et al. 2011), such as the *Vector Space Model* (VSM), with the basic TF-IDF weighting. VSM is a spatial representation of text documents, where each document is represented by a vector in a n -dimensional space, and each dimension is related to a term from the overall vocabulary of a specific document collection. In other words, every document is represented as a vector of term weights, where the weight indicates the degree of association between the document and the term. Let $D = \{d_1, d_2, \dots, d_N\}$ indicate a set of documents, and $d_j = \{t_1, t_2, \dots, t_N\}$, $t \in T$ be the set of terms in a document. The dictionary T is obtained by applying some standard Natural Language Processing (NLP) operations, such as tokenization, *stop-words* removal and stemming, and every document d_j is represented as a vector in a n -dimensional vector space, so $d_j = \{w_{1j}, w_{2j}, \dots, w_{nj}\}$ where w_{nj} represents the weight for term t_n in document d_j . The major problems during the document representation with the VSM are the weighting of the terms and the evaluation of the similarity of the vectors. The most commonly used way to estimate the term weighting is based on TF-IDF, a trivial approach that uses empirical observations of the documents' terms (Salton et al. 1975).

3.2 WordNet environment

Due to the fact that an approach based on a simple *bag of words* is not able to perform a semantic disambiguation of the words in an item description, also motivated by the fact that exploiting a taxonomy for categorization purposes is an approach recognized in the literature (Addis et al. 2010), and by the fact that a semantic analysis is useful to improve the accuracy of a classification (Armano et al. 2011a, b), in order to perform the similarity measures used in this work we decided to exploit the functionalities offered by the WordNet environment. WordNet is a large lexical database of English, where *nouns*, *verbs*, *adjectives*, and *adverbs* are grouped into sets of cognitive synonyms (synsets), each expressing a distinct concept. Synsets are interlinked by means of conceptual-semantic and lexical relations. Wordnet currently contains about 155,287 words, organized into 117,659 synsets for a total of 206,941 word-sense pairs (Fellbaum 1998).

3.2.1 WordNet structure

The main relation among words in WordNet is the synonymy and, in order to represent these relations, the dictionary is based on *synsets*, i.e., unordered sets of grouped words that denote

the same concept and are interchangeable in many contexts. Each synset is linked to other synsets through a small number of *conceptual relations*. Word forms with several distinct meanings are represented by as many distinct synsets, so that each form-meaning pair in WordNet is unique (e.g., the *fly* noun and the *fly* verb belong to two distinct synsets). Most of the WordNet relations connect words that belong to the same part-of-speech (POS). There are four POS: *nouns*, *verbs*, *adjectives*, and *adverbs*. Both nouns and verbs are organized into precise hierarchies, defined by hypernym or *is-a* relationships. For example, the first sense of the word *radio* would have the following hypernym hierarchy, where the words at the same level are synonyms of each other: as shown in the following, some sense of *radio* is synonymous with some other senses of *radiocommunication* or *wireless*, and so on.

1. POS=*noun*

- (a) *radio, radiocommunication, wireless (medium for communication)*
- (b) *radio receiver, receiving set, radio set, radio, tuner, wireless (an electronic receiver that detects and demodulates and amplifies transmitted signals)*
- (c) *radio, wireless (a communication system based on broadcasting electromagnetic waves)*

2. POS=*verb*

- (a) *radio (transmit messages via radio waves)*

Each synset has a unique index and shares its properties, such as a gloss or dictionary definition.

In the case of *nouns* and *verbs* (the organization of adjectives and adverbs is slightly different) the WordNet hierarchies are organized into several base types (25 primitive groups for the nouns and 15 for the verbs), and all primitive groups ultimately go up to an abstract root node. As we can imagine, the network of nouns is far deeper than that of the other *parts-of-speech*. The verbs instead present a more bushy structure, and the adjectives are distributed into many clusters, as well as the adverbs, since these last are defined in terms of the adjectives (i.e., they are derived from adjectives and thus inherit the structure from them). Due to the similarity measure chosen for our work, we consider only the *nouns* and the *verbs*, exploiting the state-of-art semantic-based approach to item recommendation based on the WordNet synsets (Pedersen et al. 2004), to evaluate the semantic similarity between the items stored in the user profiles.

4 Item recommendation with incoherent items removal

The definition of the problem handled by our proposal, the notation used in the problem statement, the details about the implementation of the proposed algorithm, and its integration on a recommender system, are described in the following (Sections 4.1 and 4.2).

4.1 Problem definition

Here, after introducing the adopted notation, we define the problem handled by our proposal.

Definition 1 (User preferences) We are given a set of users $U = \{u_1, \dots, u_N\}$, a set of items $I = \{i_1, \dots, i_M\}$, and a set V of values used to express the user preferences (e.g.,

$V = [1, 5]$ or $V = \{like, dislike\}$). The set of all possible preferences expressed by the users is a ternary relation $P \subseteq U \times I \times V$. We denote as $P_+ \subseteq P$ the subset of preferences with a positive value (i.e., $P_+ = \{(u, i, v) \in P \mid v \geq \bar{v} \vee v = like\}$), where \bar{v} indicates the mean value (in the previous example, $\bar{v} = 3$).

Definition 2 (User items) Given the set of positive preferences P_+ , we denote as $I_+ = \{i \in I \mid \exists(u, i, v) \in P_+\}$ the set of items for which there is a positive preference, and as $I_u = \{i \in I \mid \exists(u, i, v) \in P_+ \wedge u \in U\}$ the set of items a user u likes.

Definition 3 (Semantic item description) Let $BoW = \{t_1, \dots, t_W\}$ be the bag of words used to describe the items in I ; we denote as d_i the binary vector used to describe each item $i \in I$ (each vector is such that $|d_i| = |BoW|$). We define as $S = \{s_1, \dots, s_W\}$ the set of synsets associated to BoW (that is, for each term used to describe an item, we consider its associated synset), and as sd_i the semantic description of i . The set of semantic descriptions is denoted as $D = \{sd_1, \dots, sd_M\}$ (note that we have a semantic description for each item, so $|D| = |I|$). The approach used to extract sd_i from d_i is described in detail in Section 4.2.1.

Definition 4 (Semantic user model) Given the set of positively evaluated items by a user I_u , we define a *semantic user model* M_u as the set of synsets in the semantic descriptions of the items in I_u . More formally, $M_u = \{s_w \mid s_w \in sd_m \wedge i_m \in I_u, \forall i_m \in I_u\}$.

Definition 5 (Item coherence) An item $i \in I_u$ is *coherent* with the rest of the items in the user profile I_u , if the similarity between the semantic description sd_i of the item and the union of the semantic descriptions of the rest of the items (i.e., $M_u \setminus sd_i$) is higher than a threshold value.

Problem 1 Given a set of items I_u that a user likes, our objective is to extract a set $\overline{I}_u \subseteq I_u$, such that each item $i \in \overline{I}_u$ is *coherent* with the others.

4.2 Our approach

As already highlighted during the description of the limits that affect the user profiling activity, individual profiles need to be as adherent as possible to the real preferences of the users, because they are exploited to predict their future interests. For this reason, in this section we propose a novel approach defined *Dynamic Coherence-Based Modeling* (DCBM) that allows us to find and remove the incoherent items from user profiles, regardless of the chosen profiling method. The implementation on a recommender system of the DCBM is described in the following subsections.

4.2.1 Data preprocessing

Before comparing the similarity between the items in a user profile, we need to follow several preprocessing steps. The first step is to detect the correct *part of speech* (POS) for each word in the text; in order to perform this task, we have used the *Stanford Log-linear Part-Of-Speech Tagger* (Toutanova et al. 2003). In the second step, we remove punctuation marks and *stop words*, i.e., the insignificant words (such as adjectives, conjunctions, etc.) that represent noise in the semantic analysis. Several stop-words lists can be found on the Internet, and in this work we have used a list of 429 stop words made available with the

Onix Text Retrieval Toolkit.¹ In the last step, after we have determined the lemma of each word using the Java API implementation for WordNet Searching JAWS,² we perform the so-called word sense disambiguation, a process where the correct sense of each word is determined, which permits us to evaluate the semantic similarity in a precise way. The best sense of each word in a sentence (i.e., the selection of the real meaning of a word in the context where it is used) was found through the Java implementation of the adapted Lesk algorithm provided by the *Denmark Technical University* similarity application (Salton et al. 1975). All the collected synsets form the set $S = \{s_1, \dots, s_W\}$ defined in Section 4.1. The output of this step is the semantic disambiguation of the textual description of each item $i \in I$, which is stored in a binary vector sd_i ; each element of the vector $sd_i[w]$ is 1 if the corresponding synset appears in the item description, and 0 otherwise.

4.2.2 Semantic similarity

The most used semantic similarity measures are five, and are those defined by Leacock and Chodorow (1998), Jiang and Conrath (1997), Resnik (1995), Lin (1998), and Wu and Palmer (1994). Each of them evaluates the semantic similarity between two WordNet synsets, and we calculate the semantic similarity by using *Wu and Palmer's* measure, a method based on the path length between a pair of concepts (WordNet synsets), which in the literature is considered to be the most accurate when generating the similarities (Dennai and Benslimane 2013; Capelle et al. 2012).

Given a set X of i WordNet synsets x_1, x_2, \dots, x_i that are related to an item description, and a set Y of j WordNet synsets y_1, y_2, \dots, y_j related to another item description, a set Q , which contains all the possible pairs between the synsets in the set X and the synsets in the set Y , is defined as in (1).

$$Q = (\langle x_1, y_1 \rangle, \langle x_1, y_2 \rangle, \dots, \langle x_i, y_j \rangle) \forall x \in X, y \in Y \quad (1)$$

In the next step, a subset Z of the pairs in Q (i.e., $Z \subseteq Q$) that have at least an element with the same POS is created (2).

$$Z = \{\langle x_i, y_j \rangle \mid POS(x_i) = POS(y_j)\} \quad (2)$$

The metric measures the similarity between concepts in an ontology (in our case it is WordNet), as shown in (3).

$$sim_{WP}(x, y) = \frac{2 \cdot A}{B + C + (2 \cdot A)} \quad (3)$$

Assuming that the *Least Common Subsumer* (LCS) of two concepts x and y is *the most specific concept that is an ancestor of both x and y* , where the concept tree is defined by the *is-a* relation, in (3) we have that $A = \text{depth}(\text{LCS}(x, y))$, $B = \text{length}(x, \text{LCS}(x, y))$, $C = \text{length}(y, \text{LCS}(x, y))$. We can note that $B + C$ represents the path length from x and y , while A indicates the global depth of the path in the taxonomy.

In the example of Fig. 1, v_4 is the parent (and also ancestor) of v_5 , while v_2 is an ancestor of both v_5 and v_3 . In this case, the LCS of v_3 and v_5 is v_2 , since it is the most specific concept that is an ancestor of both v_3 and v_5 . Note that while v_1 is a common subsumer of both v_3 and v_5 , it is not the least, since there is still a child of v_1 (in this case it is v_2), which

¹<http://www.lextek.com/manuals/onix/stopwords.html>

²<http://lyle.smu.edu/tspell/jaws/index.html>




Fig. 1 WordNet relationships tree

is also a common subsumer of both v_5 and v_3 . v_4 is not the least common subsumer since it is not an ancestor of v_3 .

In order to calculate the Wu and Palmer similarity between v_3 and v_5 , we first determine that the least common subsumer of v_3 and v_5 is v_2 . Next, we determine that the length of the path from v_3 to v_2 is 1, that the length of the path from v_5 to v_2 is 2, and that the depth of v_2 is 1 (i.e., the distance from v_2 to the root vertex v_1). Now we can determine the similarity between the synsets v_3 and v_5 , as shown in (4).

$$\text{sim}_{WP}(v_3, v_5) = \frac{2 \cdot 1}{2 + 1 + (2 \cdot 1)} = 0.40 \quad (4)$$

The similarity between two items is defined as the sum of the similarity score of all the item pairs, divided by its cardinality (the subset Z of WordNet synsets with a common part-of-speech), as shown in (5).

$$\text{sim}_{WP}(X, Y) = \frac{\sum_{(x,y) \in Z} \text{sim}_{WP}(x, y)}{|Z|} \quad (5)$$

This similarity metric is employed by our algorithm to compute the coherence of an item with the rest of the semantic user profile.

4.2.3 Dynamic coherence-based modeling

For the purpose of being able to make effective recommendations to users, their profiles need to store only the descriptions of the items that really reflect their preferences.

In order to identify which items positively evaluated by a user ($i \in I_u$) do not reflect her preference, representing for instance the result of past wrong choices or the use by third parties of her account, the Dynamic Coherence-Based Modeling (DCBM) algorithm measures the *Minimum Global Coherence* (MGC) of each single item description with the set of the other items present in her profile. In other words, through MGC, the most dissimilar item with respect to the other items is identified.

The Wu and Palmer similarity metric previously presented can be used to calculate the *MGC*, as shown in (6) (sd_i denotes the semantic description of an item i , and $M_u \setminus sd_i$ indicates the semantic user model from which the synsets in sd_i have been removed).

$$MGC = \arg \min_{i \in I_u} (sim_{WP}(sd_i, M_u \setminus sd_i)) \quad (6)$$

The basic idea is to isolate each individual item i in a user profile, semantically described by sd_i , and then measure its similarity with respect to the other items (i.e., the merging of the synsets of the rest of the items), in order to obtain a measure of its coherence within the overall context of the entire profile.

In other words, in order to identify the most distant element from the general context of the evaluated items, we are exploiting a basic principle of the differential calculus, because the *MGC* value shown upon is nothing other than the *maximum negative slope*, which is calculated by finding the ratio between the changing on y axis and the changing on x axis. This is demonstrated in Theorem 1.

Theorem 1 *The Minimum Global Coherence coefficient corresponds to the maximum negative slope.*

Proof Placing on the x axis the user iterations in a chronological order, and on the y axis the corresponding values of *GS* (Global Similarity) calculated as $sim_{WP}(sd_i, M_u \setminus sd_i)$, $\forall i \in I_u$, we can trivially calculate the slope value (denoted by the letter m), as shown in (7).

$$m = \frac{\Delta y}{\Delta x} = \frac{f(x + \Delta x) - f(x)}{\Delta x} \quad (7)$$

The mathematics of differential calculus defines the slope of a curve at a point as the slope of the tangent line at that point. Since we are working with a series of points, the slope may be calculated not at a single point but between two points. Considering that for each current user iteration Δx is always equal to 1 (in fact, for N user iterations we have that $1 - 0 = 1$, $2 - 1 = 1$, ..., $N - (N - 1) = 1$), the slope value m will always be equal to $f(x + \Delta x) - f(x)$. As (8) shows, where $sim_{WP}(I_u)$ denotes $sim_{WP}(sd_i, M_u \setminus sd_i)$, $\forall i \in I_u$, the maximum negative slope corresponds to the value of *MGC*.

$$\min \left(\frac{\Delta y}{\Delta x} \right) = \min \left(\frac{sim_{WP}(I_u)}{1} \right) = MGC \quad (8)$$

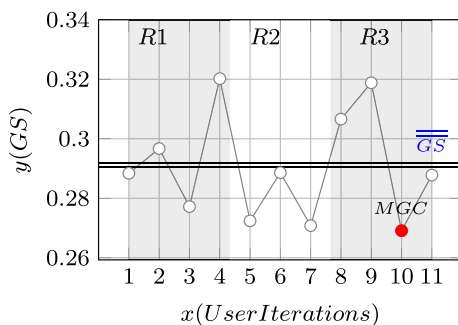
□

In Fig. 2, which displays the data reported in Table 1, we can see what we just said in a graphical way.

In order to avoid the removal of an item that might correspond to a recent change in the preferences of the user or an item not semantically distant enough from the context of the remaining items, the DCBM algorithm considers an item as incoherent and removes it, only if it meets the following conditions:

1. it is located in the first part of the user iteration history. Based on this first requirement, an item is considered far from the user's preferences only when it goes up in the first part of the iterations. This condition is checked thanks to a parameter r , taken as input by the algorithm, which defines the *removal area*, i.e., the percentage of a user profile where an item can be removed. Note that $0 \leq r \leq 1$, so in the example in Fig. 2,

Fig. 2 The maximum negative slope corresponds to the value of MGC



$r = \frac{2}{3} = 0.66$ (i.e., the element related to MGC value is located in the region $R3$, so it does not meet this first requirement);

- the value of MGC must be higher than the *mean value* of the *global similarity*.

Regarding the first requirement, it should be noted that the regions extension is strongly related both to the type of items and to their frequency of fruition, so it depends on the operative scenario. With respect to the second requirement, we prevent the removal of items when their semantic distance with the remaining items is lower than the mean value. For this reason, we first calculate the value of the mean similarity in the context of the user profile, then we define a threshold value that determines when an item must be considered incoherent with respect to the current context. Equation (9) measures the mean similarity, denoted by \overline{GS} , by calculating the average of the *Global Similarity* (GS) values, which are obtained as $sim_{WP}(sd_i, M_u \setminus sd_i), \forall i \in I_u$.

$$\overline{GS} = \frac{1}{|I_u|} \cdot \sum_{i \in I_u} (sim_{WP}(sd_i, M_u \setminus sd_i)) \quad (9)$$

where $|I_u|$ represents the total number of items stored in the profile (in the case of sample data shown in Table 1, the $\overline{GS} = 0.2906$). Once this average value is obtained, we can

Table 1 User profile sample data

| x | y | m |
|-----|--------|---------|
| 1 | 0.2884 | +0.2884 |
| 2 | 0.2967 | +0.0083 |
| 3 | 0.2772 | −0.0195 |
| 4 | 0.3202 | +0.0430 |
| 5 | 0.2724 | −0.0478 |
| 6 | 0.2886 | +0.0162 |
| 7 | 0.2708 | −0.0178 |
| 8 | 0.3066 | +0.0358 |
| 9 | 0.3188 | +0.0122 |
| 10 | 0.2691 | −0.0497 |
| 11 | 0.2878 | +0.0187 |

proceed to define the condition ρ to be used to decide when an item has to be removed (1) or not (0), based on the average value \overline{GS} (as shown in 10).

$$\rho = \begin{cases} 1, & \text{if } MGC < \overline{GS} \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

Based on the above considerations, we can now define Algorithm 1, used to remove the semantically incoherent items from a user profile. The algorithm requires as input the set I_u (i.e., the user profile), and a removal area r used to define in which part of the profile an item can be removed. In step 3 we extract the set of synsets M_u (Definition 4) from the description of the items in the user profile I_u (Definition 2). Steps 4–6 compute the similarity between each couple of synsets that belong to the user profile. In step 7, the average of the similarities is computed, so that in steps 8–15 we can evaluate if an item has to be removed from a user profile or not. In particular, once an item i is removed from a profile in step 12, its associated similarity s is removed from the list S (step 13), so that MGC in step 9 can be set as the minimum similarity value after the item removal. In step 16, the algorithm returns the user profile I_u without the removed items.

Algorithm 1 DCBM Algorithm

Require: I_u =set of items in the user profile, r =removal area

```

1: procedure PROCESS( $Y$ )
2:    $N = |I_u|$ 
3:    $M_u = GetSynsets(I_u)$ 
4:   for each Pair  $p=(sd_i, M_u \setminus sd_i)$  in  $I_u$  do
5:      $S \leftarrow sim_{WP}(p)$ 
6:   end for
7:    $a = Average(S)$ 
8:   for each  $s$  in  $S$  do
9:      $MGC = Min(S)$ 
10:     $i = index(MGC)$ 
11:    if  $i < r * size(I_u)$  AND  $MGC < a$  then
12:      Remove( $i$ )
13:      Remove( $s$ )
14:    end if
15:  end for
16:  Return  $I_u$ 
17: end procedure

```

4.2.4 Item recommendation

After the user profile has been processed by Algorithm 1, this step computes the similarity with all the items not evaluated, and recommends to a user a subset of those with the highest similarity. An interesting aspect to consider is that, thanks to our proposal, a user profile in which both the preferences in terms of ratings and in terms of synsets would be available to a recommender system. This would make possible the generation of recommendations with both collaborative filtering and content-based approaches. In this study we consider the two main subclasses of collaborative filtering approaches (i.e., the neighborhood approach and the latent factor models), since collaborative filtering approaches are known to be the most accurate. In particular, we consider SVD (Koren and Bell 2015) for the latent factor model-based approaches and a user-based approach (Ning et al. 2015) for the neighborhood subclass. This will allow us to evaluate our approach both in a scenario in which the feature space is reduced (SVD), and in which the recommendations are instead produced by considering all the items a user evaluated (user-based approach).

4.2.5 Summary

We can summarize the implementation process of the DCBM algorithm on a recommender system in the following four steps:

1. **Data Preprocessing:** preprocessing of the textual description of all items, in order to remove the useless elements and the items with a rating lower than the average;
2. **Semantic Similarity:** WordNet features are used to retrieve, from the preprocessed text, all the possible pairs between the WordNet synsets in the text of the items not evaluated and the synsets in the text of the user profile, keeping only the pairs that have at least an element with the same *part-of-speech*, for which we measure the semantic similarity according to the *Wu and Palmer* metric;
3. **Dynamic Coherence-Based Modeling:** the items dissimilar from the average preferences of a user are identified by measuring the Minimum Global Coherence (MGC). Moreover, in accordance with certain criteria, the items that are more semantically distant from the context of a user's real preferences are removed from the user profile;
4. **Item Recommendation:** the user profile without the incoherent items is employed to filter the items a user has not yet evaluated, in order to select the items to recommend.

5 Experimental framework

This section presents the framework used to evaluate our proposal. The strategy that drove our experiments is first described; then the dataset used and the preprocessing made on the data are introduced; after, the metrics used for the evaluation are presented; the last part of the section presents the obtained results.

5.1 Experimental strategy

The experimental environment for this work is based on the Java language, with the support of Java API implementation for WordNet Searching (JAWS) previously mentioned. In order to perform the evaluation, we used two real-world datasets widely-employed in the recommender systems literature, i.e., Yahoo! Webscope (R4) and Movielens 10M.

The recommender systems used to test the effectiveness of the DCBM algorithms are SVD and a classic User-Based Nearest Neighbors Collaborative Filtering approach. As previously highlighted, this will allow us to evaluate our proposal on recommender systems in which the dimensionality is reduced and in which the feature space is processed as it is.

These Mahout framework was used to implement the recommender systems. In addition to the training set, the framework requires two parameters for SVD: the number of target features and the number of training steps to run (parameter λ). The value of the first parameter can be chosen arbitrarily, on the basis of the number of features that the SVD should target; in our experiments the performance differences measured between our approach and SVD are almost the same by varying this value;³ we have set it to 19 for the Yahoo dataset and to 18 for the Movielens dataset (i.e., as the number of genres of each dataset); the second parameter, λ , will instead be tested through a set of experiments.

³The analysis has been omitted, since it did not show significant and interesting results, and in order to facilitate the reading of the paper.

For the user-based approach, it is necessary to set the number of neighbors to consider when predicting the ratings; also this parameter will be set thanks to a specific set of experiments. The distance function chosen for the neighborhood approach is the Pearson correlation, one of the most common measure of correlation to evaluate the similarity between two users (Herlocker et al. 1999).

In order to validate our proposal, a comparative analysis has been performed, by considering the values obtained by a recommender system both in a scenario in which the user profile is processed with the DCBM algorithm and in a scenario in which the profile is processed as it is. The comparisons have been made by measuring both the Root Mean Squared Error (RMSE) and the Average Difference (AD).

RMSE was chosen as a metric to compare the algorithms because, as the organizers of the Netflix prize highlight,⁴ it is well-known and widely used, it allows to evaluate a system through a single number, and it emphasizes the presence of large errors (both false positives and false negatives). AD was measured since it is being recently considered by largely employed frameworks (such as Apache's Mahout) to facilitate the interpretability of the results (Owen et al. 2011).

Three sets of experiments have been performed:

1. **Analysis of the removed items.** We analyze the amount of items removed for each user, in order to analyze the impact of the DCBM algorithm on a user profile.
2. **Recommendation accuracy measurement.** We validate the capability of our approach to remove incoherent items, by analyzing the accuracy of the generated recommendations with and without the incoherent items in the user profiles.
3. **Performance analysis.** We consider the amount of time it takes to process an item, when analyzing if it should be removed or not. This allows us to evaluate our algorithm considering its performance, and to analyze, once an item gets removed, how much time can be saved by using our proposal.

5.2 Datasets

In order to evaluate our strategy, we perform a series of experiments on two different real-world datasets, extracted by two standard benchmarks for recommender systems: Yahoo! Webscope R4⁵ and Movielens 10M.⁶

Yahoo! Webscope (R4) This dataset contains a large amount of data related to user preferences expressed on the Yahoo! Movies community that are rated on the base of two different scales, from 1 to 13 and from 1 to 5 (we use the latter). The training data is composed by 7,642 users ($|U|$), 11,915 movies/items ($|I|$), and 211,231 ratings ($|P|$). All the users in the training set have rated at least 10 items and all items are rated by at least one user. The test data is composed by 2,309 users, 2,380 items, and 10,136 ratings. There are no test users/items that do not also appear in the training data. All the users in the test set have rated at least one item and all items have been rated by at least one user. The items are classified in 19 different classes (genres), and it should be noted that an item may be classified with

⁴<http://www.netflixprize.com/faq>

⁵<http://webscope.sandbox.yahoo.com>

⁶<http://grouplens.org/datasets/movielens/>

multiple classes. The information of this dataset, training and test data, are sorted in chronological order, and it should be also noted that the test data were gathered chronologically after the training data.

Movielens 10M The second dataset used in this work is composed by 71,567 users ($|U|$), 10,681 movies/items ($|I|$), and 10,000,054 ratings ($|P|$). It was extracted at random from MovieLens (a movie recommendation website). All the users in the dataset had rated at least 20 movies, and each user is represented by a unique ID. The ratings of the items are based on a 5-star scale, with half-star increments. In this dataset the items are classified in 18 different classes (movie genres), and also in this case each item may be classified with multiple classes (genres). Since the Movielens 10M dataset does not contain any textual description of the items, to obtain this information we used a file provided by the Webscope (R4) dataset, which contains a mapping from the movie IDs used in the dataset to the corresponding movie IDs and titles used in the MovieLens dataset. In this dataset the users were selected at random for the inclusion in the training and test data.

5.3 Metrics

The accuracy of the predicted ratings was measured through the Root Mean Squared Error (RMSE) and the Average Difference (AD). Both metrics consider the test set and the predicted ratings, by comparing each rating r_{ui} , given by a user u for an item i and available in the test set, with the rating p_{ui} predicted by a recommender system. The formulas are shown below, where n is the number of ratings in the test set:

$$RMSE = \sqrt{\frac{\sum_{i=0}^n (r_{ui} - p_{ui})^2}{n}}$$

$$AD = \frac{\sum_{i=0}^n |r_{ui} - p_{ui}|}{n}$$

5.4 Experimental results

This section presents the results for each set of experiments presented in the strategy.

5.4.1 Analysis of the removed items

Figure 3 reports the number of removed items and the amount of users for which that amount of items were removed, in the Webscope dataset (Fig. 3a) and in the Movielens dataset (Fig. 3b). Both figures show a power law distribution, which indicates that for the vast majority of the users just a small number of items gets removed. However, there are cases in which hundreds of items are removed from a single user profile and get classified by our approach as noise.

In the next set of experiments we evaluate the capability of our approach to generate accurate recommendations by employing these profiles without incoherent items.

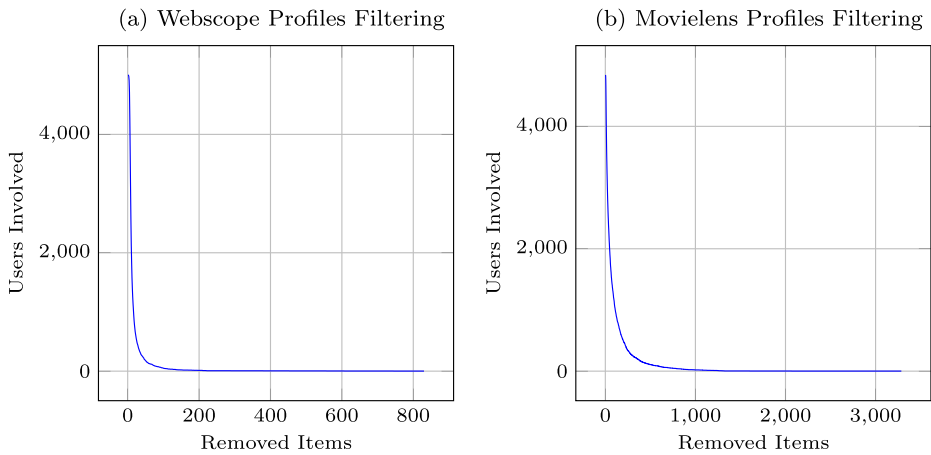


Fig. 3 Removed items

5.4.2 Recommendation accuracy measurement

Figure 4 shows the results obtained for the SVD recommender system and the two employed metrics, in both the Webscope (Fig. 4a-b) and Movielens (Fig. 4c-d) datasets. The results of the Webscope dataset show that the best results can be obtained with $\lambda = 0.141$, $RMSE = 1.0124$, $AD = 0.7788$ for the unfiltered dataset and $\lambda = 0.171$, $RMSE = 1.0128$, $AD = 0.7301$ for the data filtered with the DCBM algorithm. Regarding the results with the Movielens data, the highest accuracy is obtained with $\lambda = 0.141$, $RMSE = 0.96$, $AD = 0.7788$ for the unfiltered dataset and $\lambda = 0.091$, $RMSE = 0.9751$, $AD = 0.755$ for the data filtered with the DCBM algorithm. We also performed independent-samples two-tailed Student's t-tests, which highlighted that there is no statistical difference between the results ($p > 0.05$). Indeed, we can notice that the RMSE of the filtered data is slightly higher, while the AD is slightly lower. This means that, when the feature space is reduced by the recommendation approach, we can reduce the complexity of the system by processing less items while keeping the same accuracy.

In Fig. 5 we analyzed the other scenario, in which the feature space processed by the recommender system is the entire user profile, by considering the nearest-neighbors user-based filtering. Figure 5 shows the results with the employed metrics, in both the Webscope (Fig. 5a-b) and Movielens (Fig. 5c-d) datasets. The results of the Webscope dataset show that the best results can be obtained with $neighbors = 8$, $RMSE = 1.058$, $AD = 0.7408$ for the unfiltered dataset and $neighbors = 7$, $RMSE = 0.8785$, $AD = 0.5877$ for the data filtered with the DCBM algorithm. Regarding the results with the Movielens data, the highest accuracy is obtained with $neighbors = 98$, $RMSE = 1.0485$, $AD = 0.8255$ for the unfiltered dataset and $neighbors = 3$, $RMSE = 1.0035$, $AD = 0.7421$ for the data filtered with the DCBM algorithm. These results show that when the entire feature space is considered (as the collaborative filtering neighborhood-based approach tested in this set of experiments, or the content-based systems), the recommendation accuracy can be significantly improved ($p < 0.05$) by removing the incoherent items from the user profiles.

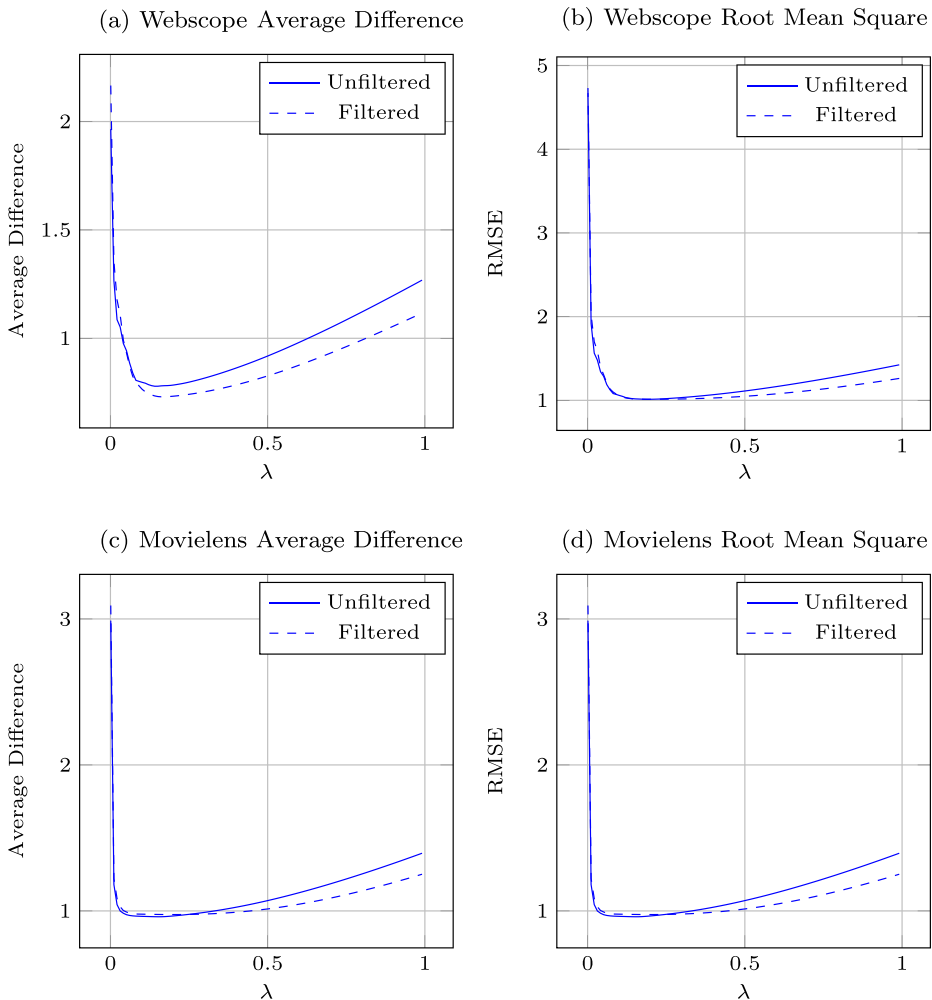


Fig. 4 SVD recommendation accuracy

5.4.3 Performance analysis

Figure 6 shows the average amount of time in seconds it takes for the DCBM algorithm to process an item and to decide if it should be removed or not, by considering the first 1000 user profiles, in order to achieve a statistic relevance. The dashed line in each subfigure indicates the average time considering all the user profiles. In the Webscope dataset (Fig. 6a) the average time to process an item is 0.21 seconds, while for the Movielens dataset (Fig. 6b) it is 0.36 seconds. This means that the algorithm presents a very good performance considering the semantic process behind the DCBM algorithm. Moreover, this quick process to decide if an item should be removed or not leads to great improvements when the recommendations are processed, both in terms of number of items to process and in terms of accuracy when the whole feature space is considered.

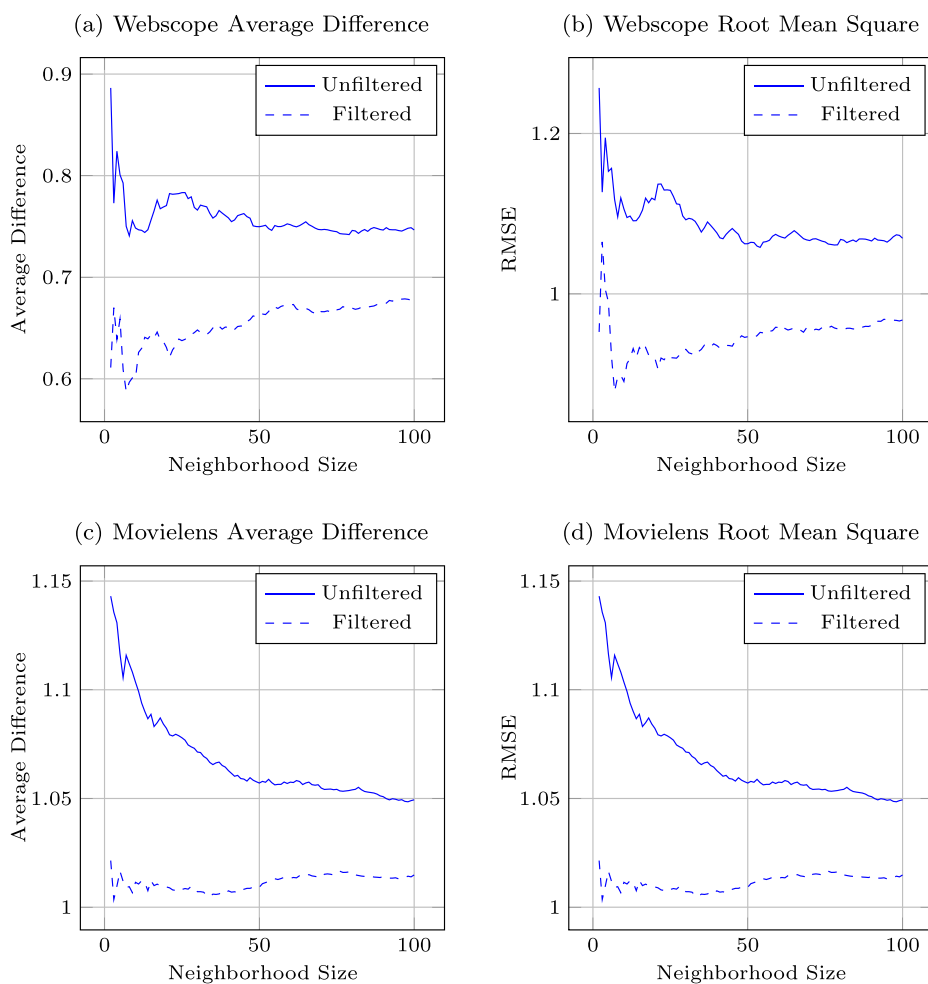


Fig. 5 User-based approach recommendation accuracy

5.4.4 Discussion

We verified that the proposed approach is actually able to remove the incoherent items from the user profiles, as highlighted in Fig. 3, which shows that there are many cases in which tens of items are removed from a single user profile. Considering that the removed items are not useful to model the real preferences of the users (i.e., they have been classified as incoherent by our approach), we obtain a considerable reduction of the computational load, without any side effect.

The main considerations that arise from the observation of the experimental results are the following: when we adopt a recommendation approach at the state of the art that does not modify the feature space, the accuracy improves significantly, as it happens with the user-based approach, whose results are shown in Fig. 5; otherwise, when we use an approach that reduces this space, as happens by using SVD, the accuracy remains unaltered, as we can observe in Fig. 4.

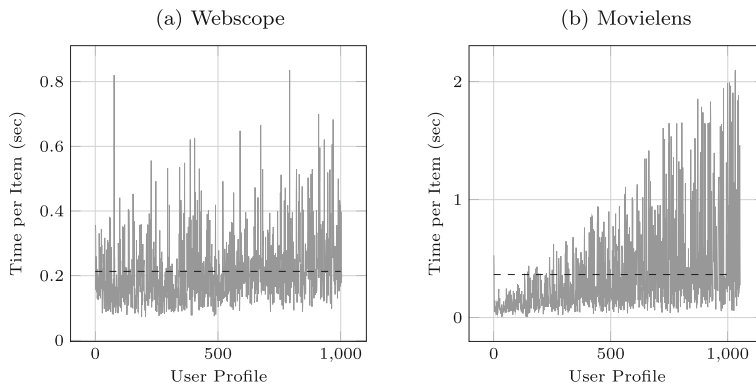


Fig. 6 Computation time

In this last case, the advantages that come from the use of our approach are to be found in the reduced number of items taken into account in the recommendation process. In fact, considering that the removal of the items from the user profiles does not require much time, as shown in Fig. 6, we have a considerable reduction of the time needed to perform the recommendation process, because it operates by analyzing a number of items that is fewer than the initial one, as reported in Fig. 3.

If on the one hand the proposed approach conducts toward more accurate recommendations, on the other hand it reduces the number of items in the user profiles, thus the computational complexity. About this aspect, it should be observed that the computational load related to the comparisons performed to detect the *MGC* does not involve the recommendation process, since it happens in another moment (i.e., when a user evaluates a new item). Even if our approach has a high computational complexity, it is meant to run in background. Indeed, when a new item is evaluated by the user, it would not be removed (even if incoherent) since it is recent, so its similarity with the rest of the user profile can be processed in background. Therefore, the recommendation task, which is the one that has to run fast because it operates online, can benefit from the proposed DCBM algorithm and from the removal of the incoherent items.

In conclusion, our approach improves the performance of a recommender system, both in terms of accuracy and of reduction of the computational time. As shown in the results of the experiments, type, and entity of the improvements is related with the recommendation strategy.

6 Conclusions and future work

In this paper we proposed a novel approach able to improve the quality of the user profiling, which takes into account the items related to a user, with the aim of removing those that do not reflect her real preferences. This is useful in many contexts, such as when the system does not allow the users to express their preferences, or when the users do not to make use of this option.

Through it we achieved a twofold result: firstly, we moved the evaluation process of the items coherence from a domain based on strict mathematical criteria (i.e., variance of the user's ratings in the feature space) to a more flexible semantic domain. Considering that

the removal of all incoherent items from the user profiles leads us toward a considerable reduction of the *magic barrier* problem, the second important result is given by the fact that we can consider each measured improvement as real, instead than a mere overfitting side effect. The experimental results show that our approach is able to reshape the user profiles in a coherent way.

A further possible extension might involve the use of a large amounts of data also related to contexts different from each other as, for example, a sales platform that gives access to very heterogeneous goods, in which we could operate in order to discover and process the semantic interconnections between different classes of items and methods, to evaluate their semantic coherence during the user profiling activity.

References

- Addis, A., Armano, G., Giuliani, A., & Vargiu, E. (2010). A recommender system based on a generic contextual advertising approach. In *Proceedings of the 15th IEEE symposium on computers and communications, ISCC 2010* (pp. 859–861). Riccione: IEEE.
- Addis, A., Armano, G., & Vargiu, E. (2010). Assessing progressive filtering to perform hierarchical text categorization in presence of input imbalance. In A.L.N. Fred, & J. Filipe (Eds.), *KDIR 2010 - Proceedings of the international conference on knowledge discovery and information retrieval* (pp. 14–23). Valencia: SciTePress.
- Amatriain, X., Pujol, J.M., & Oliver, N. (2009). I like it... I like it not: Evaluating user ratings noise in recommender systems, In Houben, G., McCalla, G.I., Pianesi, F., & Zancanaro, M. (Eds.) *Proceedings of the 17th international conference, user modeling, adaptation, and personalization UMAP, 2009, formerly UM and AH*, (Vol. 5535 pp. 247–258). Trento: Springer. Lecture Notes in Computer Science.
- Amatriain, X., Pujol, J.M., Tintarev, N., & Oliver, N. (2009). Rate it again: increasing recommendation accuracy by user re-rating. In L.D. Bergman, A. Tuzhilin, R.D. Burke, A. Felfernig, & L. Schmidt-Thieme (Eds.), *Proceedings of the 2009 ACM Conference on Recommender Systems, RecSys 2009* (pp. 173–180). New York: ACM.
- Armano, G., Giuliani, A., & Vargiu, E. (2011a). Semantic enrichment of contextual advertising by using concepts, In J. Filipe, & A.L.N. Fred (Eds.), *KDIR 2011 - Proceedings of the International Conference on Knowledge Discovery and Information Retrieval* (pp. 232–237). Paris: SciTePress.
- Armano, G., Giuliani, A., & Vargiu, E. (2011b). Studying the impact of text summarization on contextual advertising, In F. Morvan, A.M. Tjoa, & R. Wagner (Eds.), *2011 database and expert systems applications, DEXA, International workshops* (pp. 172–176). Toulouse: IEEE Computer Society.
- Armano, G., & Vargiu, E. (2010). A unifying view of contextual advertising and recommender systems, In A.L.N. Fred, & J. Filipe (Eds.) *KDIR 2010 - Proceedings of the international conference on knowledge discovery and information retrieval* (pp. 463–466). Valencia: SciTePress.
- Asnicar, F.A., & Tasso, C. (1997). ifweb: a prototype of user model-based intelligent agent for document filtering and navigation in the world wide web. In *Proceedings of workshop adaptive systems and user modeling on the world wide web'at 6th international conference on user modeling, UM97, Chia Laguna* (pp. 3–11). Sardinia.
- Balabanovic, M., & Shoham, Y. (1997). Content-based, collaborative recommendation. *Communications of the ACM*, 40(3), 66–72. doi:10.1145/245108.245124.
- Bellogin, A., Said, A., & de Vries, A.P. (2014). The magic barrier of recommender systems - no magic, just ratings, In V. Dimitrova, T. Kuflik, D. Chin, F. Ricci, P. Dolog, & G. Houben (Eds.), *Proceedings of the 22nd international conference, User modeling, adaptation, and personalization UMAP 2014*, (Vol. 8538 pp. 25–36). Aalborg: Springer. Lecture Notes in Computer Science.
- Billsus, D., & Pazzani, M.J. (1999). A hybrid user model for news story classification. In *Proceedings of the seventh international conference on user modeling, UM '99* (pp. 99–108). Secaucus: Springer-Verlag New York, Inc. <http://dl.acm.org/citation.cfm?id=317328.317338>.
- Budzik, J., & Hammond, K.J. (2000). *User interactions with everyday applications as context for just-in-time information access*. New York: ACM.
- Capelle, M., Frasinca, F., Moerland, M., & Hogenboom, F. (2012). Semantics-based news recommendation. In *Proceedings of the 2Nd international conference on web intelligence, mining and semantics, WIMS '12* (pp. 27:1–27:9). New York: ACM.

- Capelle, M., Hogenboom, F., Hogenboom, A., & Frasinicar, F. (2013). Semantic news recommendation using wordnet and bing similarities. In *Proceedings of the 28th annual ACM symposium on applied computing, SAC '13* (pp. 296–302). New York: ACM.
- Chirita, P.A., Nejdl, W., Paiu, R., & Kohlschütter, C. (2005). Using odp metadata to personalize search. In *Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval, SIGIR '05* (pp. 178–185). New York: ACM.
- Dennai, A., & Benslimane, S.M. (2013). Toward an update of a similarity measurement for a better calculation of the semantic distance between ontology concepts. In *The second international conference on informatics engineering & information science (ICIEIS2013)* (pp. 197–207): The Society of Digital Information and Wireless Communication.
- Fellbaum, C. (1998). *WordNet: An electronic lexical database*: Bradford Books.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., & Ruppín, E. (2002). Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20(1), 116–131.
- Herlocker, J.L., Konstan, J.A., Borchers, A., & Riedl, J. (1999). An algorithmic framework for performing collaborative filtering. In *SIGIR* (pp. 230–237): ACM.
- Herlocker, J.L., Konstan, J.A., Terveen, L.G., & Riedl, J. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1), 5–53.
- Hill, W.C., Stead, L., Rosenstein, M., & Furnas, G.W. (1995). Recommending and evaluating choices in a virtual community of use. In I.R. Katz, R.L. Mack, L. Marks, M.B. Rosson, & J. Nielsen (Eds.), *CHI '95 conference proceedings Human factors in computing systems* (pp. 194–201). Denver: ACM/Addison-Wesley.
- Jiang, J.J., & Conrath, D.W. (1997). Semantic similarity based on corpus statistics and lexical taxonomy arXiv preprint cmp-lg/9709008.
- Karypis, G. (2001). Evaluation of item-based top-n recommendation algorithms. In *Proceedings of the 2001 ACM CIKM international conference on information and knowledge management* (pp. 247–254). Atlanta: ACM, doi: [10.1145/502585.502627](https://doi.org/10.1145/502585.502627).
- Kelly, D., & Teevan, J. (2003). Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37(2), 18–28.
- Koren, Y., & Bell, R.M. (2015). Advances in collaborative filtering. In F. Ricci, L. Rokach & B. Shapira (Eds.), *Recommender systems handbook* (pp. 77–118): Springer, doi: [10.1007/978-1-4899-7637-6_3](https://doi.org/10.1007/978-1-4899-7637-6_3).
- Lam, W., Mukhopadhyay, S., Mostafa, J., & Palakal, M.J. (1996). Detection of shifts in user interests for personalized information filtering (pp. 317–325): SIGIR.
- Leacock, C., & Chodorow, M. (1998). Combining local context and wordnet similarity for word sense identification. In C. Fellbaum (Ed.) *WordNet: An electronic lexical database* (pp. 305–332): MIT Press.
- Li, L., Yang, Z., Wang, B., & Kitsuregawa, M. (2007). Dynamic adaptation strategies for long-term and short-term user profile to personalize search. In G. Dong, X. Lin, W. Wang, Y. Yang, & J.X. Yu (Eds.), *Advances in data and web management, joint 9th asia-pacific web conference, APWeb 2007, and 8th International Conference, on Web-Age Information Management, WAIM 2007*, (Vol. 4505 pp. 228–240). Huang Shan: Springer. Proceedings, Lecture Notes in Computer Science.
- Lieberman, H. (1995). Letizia: An agent that assists web browsing. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'95* (pp. 924–929). San Francisco: Morgan Kaufmann Publishers Inc. <http://dl.acm.org/citation.cfm?id=1625855.1625975>.
- Lin, D. (1998). An information-theoretic definition of similarity. In Shavlik, J.W. (Ed.) *Proceedings of the fifteenth international conference on machine learning (ICML 1998)* (pp. 296–304). Madison: Morgan Kaufmann.
- Linden, G., Smith, B., & York, J. (2003). Industry report: Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Distributed Systems Online*, 4(1). <http://dsonline.computer.org/0301/dwpl1ind.htm>.
- Liu, F., Yu, C., & Meng, W. (2002). Personalized web search by mapping user queries to categories. In *Proceedings of the Eleventh International Conference on Information and Knowledge Management, CIKM '02* (pp. 558–565). New York: ACM.
- Lops, P., de Gemmis, M., & Semeraro, G. (2011). Content-based recommender systems: State of the art and trends. In F. Ricci, L. Rokach, B. Shapira & P.B. Kantor (Eds.), *Recommender systems handbook* (pp. 73–105). Springer.
- Ma, Z., Pant, G., & Sheng, O.R.L. (2007). Interest-based personalized search. *ACM Transactions on Information Systems*, 25(1).
- Ning, X., Desrosiers, C., & Karypis, G. (2015). A comprehensive survey of neighborhood-based recommendation methods. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender systems handbook* (pp. 37–76). Springer, doi:[10.1007/978-1-4899-7637-6_2](https://doi.org/10.1007/978-1-4899-7637-6_2).

- Owen, S., Anil, R., Dunning, T., & Friedman, E. (2011). *Mahout in action*. Greenwich: Manning Publications Co.
- Pazzani, M., Muramatsu, J., & Billsus, D. (1996). Syskill & #38; webert: Identifying interesting web sites. In *Proceedings of the thirteenth national conference on artificial intelligence - volume 1, AAAI'96* (pp. 54–61): AAAI Press. <http://dl.acm.org/citation.cfm?id=1892875.1892883>.
- Pazzani, M.J., & Billsus, D. (2007). Content-based recommendation systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The Adaptive Web* (pp. 325–341). Berlin: Springer. <http://dl.acm.org/citation.cfm?id=1768197.1768209>.
- Pedersen, T., Patwardhan, S., & Michelizzi, J. (2004). Wordnet similarity: Measuring the relatedness of concepts. In *Demonstration Papers at HLT-NAACL 2004, HLT-NAACL-Demonstrations '04* (pp. 38–41). Stroudsburg: Association for Computational Linguistics.
- Pretschner, A., & Gauch, S. (1999). Ontology based personalized search. In *11th IEEE international conference on tools with artificial intelligence, ICTAI '99* (pp. 391–398). Chicago: IEEE Computer Society. doi:10.1109/TAI.1999.809829.
- Resnik, P. (1995). Using information content to evaluate semantic similarity in a taxonomy. In *Proceedings of the 14th international joint conference on artificial intelligence - Volume 1, IJCAI'95* (pp. 448–453). San Francisco: Morgan Kaufmann Publishers Inc.
- Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to recommender systems handbook. In F. Ricci, L. Rokach, B. Shapira, & P.B. Kantor (Eds.), *Recommender systems handbook* (pp. 1–35). Springer.
- Saia, R., Boratto, L., & Carta, S. (2014). Semantic coherence-based user profile modeling in the recommender systems context. In *Proceedings of the 6th international conference on knowledge discovery and information retrieval, KDIR 2014* (pp. 154–161). Rome: SciTePress.
- Said, A., Jain, B.J., Narr, S., & Plumbaum, T. (2012). Users and noise: The magic barrier of recommender systems, In J. Masthoff, B. Mobasher, M.C. Desmarais, & R. Nkambou (Eds.), *Proceedings of the 20th international conference, user modeling, adaptation, and personalization - UMAP 2012*, (Vol. 7379 pp. 237–248). Montreal: Springer. Lecture Notes in Computer Science.
- Said, A., Jain, B.J., Narr, S., Plumbaum, T., Albayrak, S., & Scheel, C. (2012). Estimating the magic barrier of recommender systems: a user study. In W.R. Hersh, J. Callan, Y. Maarek, & M. Sanderson (Eds.), *The 35th international ACM SIGIR conference on research and development in Information Retrieval, SIGIR '12* (pp. 1061–1062). Portland: ACM.
- Salton, G., Wong, A., & Yang, C.S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18(11), 613–620.
- Sarwar, B.M., Karypis, G., Konstan, J.A., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms, In V.Y. Shen, N. Saito, M.R. Lyu, & M.E. Zurko (Eds.), *Proceedings of the tenth international world wide web conference, WWW 10* (pp. 285–295). Hong Kong: ACM. doi:10.1145/371920.372071.
- Schafer, J.B., Konstan, J.A., & Riedl, J. (1999). Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic commerce* (pp. 158–166).
- Schickel-Zuber, V., & Faltings, B. (2006). Inferring user's preferences using ontologies. In *Proceedings, the twenty-first national conference on artificial intelligence and the eighteenth innovative applications of artificial intelligence conference* (pp. 1413–1418). Boston: AAAI Press.
- Shen, X., Tan, B., & Zhai, C. (2005). Implicit user modeling for personalized search. In O. Herzog, H.J. Schek, N. Fuhr, A. Chowdhury, & W. Teiken (Eds.), *Proceedings of the 2005 ACM CIKM international conference on information and knowledge management* (pp. 824–831). Bremen: ACM.
- Stilo, G., & Velardi, P. (2014a). Temporal semantics: Time-varying hashtag sense clustering. In *Knowledge engineering and knowledge management, Lecture Notes in Computer Science*, (Vol. 8876 pp. 563–578): Springer International Publishing.
- Stilo, G., & Velardi, P. (2014b). Time makes sense: Event discovery in twitter using temporal similarity. In *Proceedings of the 2014 IEEE/WIC/ACM international joint conferences on web intelligence (WI) and intelligent agent technologies (IAT) - Volume 02, WI-IAT '14* (pp. 186–193). Washington: IEEE Computer Society. doi:10.1109/WI-IAT.2014.97.
- Stilo, G., & Velardi, P. (2015). Efficient temporal mining of micro-blog texts and its application to event discovery. *Data Mining and Knowledge Discovery*. doi:10.1007/s10618-015-0412-3.
- Su, X., & Khoshgoftaar, T.M. (2009). A survey of collaborative filtering techniques. *Advances Artificial Intelligence*. doi:10.1155/2009/421425.
- Toutanova, K., Klein, D., Manning, C.D., & Singer, Y. (2003). *Feature-rich part-of-speech tagging with a cyclic dependency network*. Stroudsburg: Association for Computational Linguistics.
- Vargiu, E., Giuliani, A., & Armano, G. (2013). Improving contextual advertising by adopting collaborative filtering. *ACM Transactions on the Web*, 7(3), 13:1–13:22.

- Wei, C., Khoury, R., & Fong, S. (2014). Recommendation systems for web 2.0 marketing. In K. Yada (Ed.), *Data mining for service, studies in big data*, (Vol. 3 pp. 171–196). Berlin: Springer.
- Widyantoro, D.H., Ioerger, T.R., & Yen, J. (2001). Learning user interest dynamics with a three-descriptor representation. *JASIST*, 52(3), 212–225.
- Wu, Z., & Palmer, M. (1994). *Verbs semantics and lexical selection*. Stroudsburg: Association for Computational Linguistics.
- Zeb, M., & Fasli, M. (2011). Adaptive user profiling for deviating user interests. In *Computer science and electronic engineering conference (CEECE)*. 3rd (pp. 65–70). doi:[10.1109/CEECE.2011.5995827](https://doi.org/10.1109/CEECE.2011.5995827).