Expert Systems with Applications 40 (2013) 4370–4379

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

A topic-based recommender system for electronic marketplace platforms

Konstantinos Christidis *, Gregoris Mentzas

School of Electrical and Computer Engineering, Information Management Unit, National Technical University of Athens, Athens, Greece

A R T I C L E I N F O

Keywords:
Probabilistic topic models
Recommender systems
E-commerce
Electronic marketplaces

A B S T R A C T

A large number of items are described, and subsequently bought and sold every day in auction marketplaces across the web. The amount of information and the number of available items makes finding what to buy as well as describing an item to sell, a challenge for the participants. In this paper we consider two functions of electronic marketplaces. First, we address the recommendation of related items for users browsing the items offered in a marketplace. Second, in order to support potential sellers we propose the recommendation of relevant items and terms which can be used to describe an item to be sold in the marketplace. The contribution of this paper lies in the proposal of an innovative system that exploits the hidden topics of unstructured information found in the e-marketplace in order to support these functions. We propose a three-step process in which a probabilistic topic modelling approach is used in order to uncover latent topics that provide the basis for item and term similarity calculation for the corresponding recommendations. We present the design of our system and perform evaluations of the quality of the extracted topics as well as of the recommender system using real life scenarios using data extracted from a widely used auction marketplace. The evaluations demonstrate the perceived usefulness of our approach.

1. Introduction

An electronic marketplace is a virtual market where buyers and sellers meet just like in a traditional market. Internet-based electronic marketplaces in general employ information technology to match buyers and sellers with effectiveness and lower transaction costs, leading overall to more efficient markets (Bakos, 1998).

The behaviour of buyers, sellers, and intermediaries is motivated by their desire to maximise their private utility while also leading to an efficient allocation of productive resources. Even though electronic auction marketplaces deliver a low friction environment for purchases, there are two challenges in their functioning.

The first challenge refers to the information overload for the buyers caused by the large number of items that are available for sale in the marketplaces (Schwartz, 2005). The number of available products, even in narrow categories, exceeds by far the volume of products available in real life stores or even in fixed price electronic commerce sites. The information in the description of each item is in many cases lengthy and dense, demanding a significant amount of attention by the potential buyer. The aggregated volume of descriptions can be safely assumed to demand a large number of hours of careful studying by the buyer in order to find products that she is actually interested in. Frequently the buyer cannot easily identify items that are the most suitable for him and as a result she either gives up or settles for something that is not fulfilling her expectations. Moreover, lacking an overall picture of the available interesting items, she can also fail to choose a successful bidding strategy for buying a product on better terms.

The second challenge is the lack of understanding of the market competition by the sellers. The convenience of using an electronic marketplace has created a new breed of non-professional sellers that sell second-hand items. In this environment, professionals from different areas, or even nations or with different backgrounds, can compete. This generates a largely unexplored and less than welcoming landscape for potential sellers. Newcomers in electronic marketplaces do not have a clear image of what they are competing against, and therefore they are not supported on the tasks they need to perform in the marketplace: the choice of their selling strategy as well as providing a proper description of their product.

In order to address these challenges in electronic marketplaces, the use of recommender systems has been proposed, using both collaborative filtering and content-based approaches (Huang, Zeng, & Chen, 2007; Xu, Zhang, Pan, & Yang, 2005). However, existing implementations fail to utilize unstructured content in order to support the functions of recommending similar items and identifying important terms that should be included in the description of an item to be sold.

* Corresponding author. Tel.: +30 6946624882.
E-mail addresses: kchrist@mail.ntua.gr (K. Christidis), gmentzas@mail.ntua.gr (G. Mentzas).

0957-4174/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2013.01.014

We consider two functions that should be supported in electronic marketplaces. First to support the potential buyers, we consider the suggestion of related items for the user browsing the items offered in a marketplace. Second, in order to support the sellers we propose the suggestion of relevant items as well as the recommendation of terms for someone who describes an item to be sold in the marketplace. Our contribution in this paper is the proposal of an innovative recommender system that exploits the hidden topics of unstructured information found in the e-marketplace in order to support these functions. To this end we propose a three-step process in which a probabilistic topic modelling (Blei, Ng, & Jordan, 2003) approach is used in order to uncover latent topics, that provide the basis for item and term similarity calculation for the corresponding recommendations.

The remainder of this work is organised as follows. In the next section background information in the areas of electronic marketplaces, recommender systems and probabilistic topic models is provided. Section three is devoted to the description of our approach for building topic-based recommender systems in electronic marketplaces. In section four we illustrate the system architecture and its main functions. Section five outlines the elements of a case study, including the real marketplace environment, the topics extracted the system usage and evaluation. We then discuss the related work in the area and finally provide conclusions and directions for further work.

2. Background

2.1. Electronic marketplaces

E-commerce consists of the presentation of goods, attraction of customers and interaction with them, electronic purchases, subsequent support and online communication with suppliers. An electronic marketplace can be seen as a medium that assigns different roles within a community, primarily buyers and suppliers, but also other roles like logistics service providers, banks, and other intermediaries. These media facilitate the exchange of information, goods, services, and payments while also provide a relevant infrastructure (Grieger, 2003).

Markets match demand and supply and this process of matching buyers' demand with sellers' product offerings has three main components: determining product offerings, search, and price discovery. Auctions are a way of matching buyers and sellers in marketplaces. Electronic auctions involve mainly three parties, namely the auctioneer, the buyer (customer) and the seller (supplier). The behaviour of buyers, sellers, and intermediaries is motivated by their desire to maximise their private utility (Bakos, 1998). Contrary to fixed price electronic shops, in electronic marketplaces the prices change depending on the relation between offer and demand at any given point in time.

In marketplaces the dynamic pricing may differentiate as it takes one of four forms (Abramson & Means, 2001). There are four possible ways of dynamic pricing with regard to the number of the participants, one buyer and one seller, one buyer and many possible sellers, one seller and many possible buyers, many sellers and many buyers. In more detail, auctions are categorized as follows:

- *English auctions* in which a seller handles a number of bids from the possible buyers. Products have a starting price where the bidding starts and an instant-buy price for which the buyer can immediately buy the product.
- *Haggle* where one buyer and one seller bargain to reach an agreed upon price.
- *Bidding process* where a potential buyer describes on what terms s/he will buy from any seller and the sellers bid.
- *Exchange* where multiple buyers negotiate with multiple sellers.

In this work we assume the commonly used English Auction model, which allows sellers to choose the highest bidder from a number of potential buyers, and we refer to that simply as "auction". However the techniques described here can also be helpful in all other types of auctions since matching between buyers and sellers is a common need.

A number of benefits arise from using an electronic marketplace for the sellers, the buyers and the auctioneers. Sellers can have increased revenues since they can sell to a larger number of customers and they can shorten the time required for completing sells. Additionally sellers have the potential for revenue growth and market share expansion by finding new partners or better sources for supplies and getting to market earlier and faster (Sharifi, Kehoe, & Hopkins, 2006). More money can be gained by the sellers as the commission to an electronic marketplace is much lower than the commissions to intermediaries or physical auction fees. When in need for cash, sellers can quickly liquidate their stock. Finally they can build a relationship with the customers gaining loyalty by providing quality services and learning their preferences. On the other side, buyers can find unique items and collectibles for their taste. The participation in electronic auctions can be exciting and entertaining. They can also conveniently and anonymously participate in marketplaces. Eventually lower buyer search costs in electronic marketplaces promote price competition among sellers, while sellers in electronic markets thus want increasingly to differentiate their products (Bakos, 1998). Auctioneers get a new role since as intermediaries they have to perform functions that include matching buyers and sellers, providing product information to buyers and marketing information to sellers, aggregating information goods, integrating the components of consumer processes, managing physical deliveries and payments, providing trust relationships and ensuring the integrity of the markets (Bakos, 1998).

These benefits can explain the sharp increase in the use of auction marketplaces in the last years. EBay,[1] the largest auction site has over 14 million auctions hosted at any given time. A large number of other sites have emerged, such as eBid,[2] Online Auction[3] and Overstock.[4]

In order to cope with the large numbers of items and actors, electronic marketplaces commonly use taxonomies. Large initiatives have been set up to define ontologies as a means for mediating e-commerce (Fensel et al., 2001). However these taxonomies are still mostly dispersed and custom for each instance. In most cases, there is no consensus on the products making up a domain, how to describe them, and their proper product catalogue structures. The participation of large numbers of people and the subsequent information overload has created challenges to customers selecting products for online purchases and to sellers attempting to identify customers' preferences efficiently (Huang, Chung, & Chen, 2004).

2.2. Recommender systems

Recommender systems were created out of the user needs to handle the increasing volume of information that is available in the web. They are addressing, in the most general expression, the problem of estimating the utility or the ratings of items that have not yet been seen by the user (Adomavicius & Tuzhilin, 2005). Since people cannot possibly read all user reviews and all offers in the duration of their lifetime, recommender systems utilize techniques borrowed from statistics and artificial intelligence in order to predict user preferences. Business interest in this area

[1] http://www.ebay.com.
[2] http://www.ebid.net.
[3] http://www.onlineauction.com.
[4] http://www.overstock.com.

started from early adopters in the electronic business domain, such as Amazon, and is rapidly growing ever since (Linden, Smith, & York, 2003).

To address this problem different recommendation techniques have been proposed in the literature:

- *Content-based* where the content of each item is analysed and mapped against the user's past preferences in order to predict her future ratings.
- *Collaborative filtering* where the behaviour of a number of users is analysed. It is assumed that similarly behaving users in the past will continue to behave in a similar way in the future.
- *Knowledge-based* where humans discover the factors that affect user preferences (Burke, 2000).
- *Case-based* which treat the objects to be recommended as cases and use the recall of examples as the fundamental problem-solving process.
- *Hybrid recommenders* that combine the above methods (Burke, 2002).

Collaborative filtering techniques demonstrate very good results in a wide range of applications (Huang et al., 2007). However a number of limitations appear as portrayed in previous works (Cho, Kim, & Kim, 2002; Sarwar, Karypis, Konstan, & Riedl, 2000). In some cases there is not enough data available for an adequate number of items in order to predict the ratings of the users. Moreover, the processing power required for analysing a large set of users, items and ratings is unsuitable for effective recommendations. Finally the recommendations offered by the system can have questionable quality.

To address these issues, a number of approaches have been proposed including matrix factorization (Koren, Bell, & Volinsky, 2009). Moreover, machine learning techniques such as dimensionality reduction, generative models, spreading activation and link analysis have been evaluated demonstrating improvements over the previous approaches (Huang et al., 2007). Additionally, another recently popular approach in recommender systems utilizes statistically uncovered latent topics in user behaviour and content in order to facilitate the recommendation process (see for example (Hofmann, 2004) and (Krestel, Fankhauser, & Nejdl, 2009)).

2.3. Probabilistic topic models

A simple way to model text based content, is that of recording the words contained in the text and projecting them in the related term vector space (Salton, Wong, & Yang, 1975). This approach is intuitive, has been popular for years and was the basis for methods such as tf∗idf.

Latent Semantic Analysis (LSA) and probabilistic Latent Semantic Analysis (pLSA) (Hofmann, 1999) have been proposed to extend this model. Latent Semantic Analysis (LSA) was the first technique to analyse documents and the words that they contain in order to generate a set of concepts that relate to both of them. Probabilistic Latent Semantic Analysis, as proposed by Hofmann (1999), is an evolution of the previous model that incorporated a probabilistic foundation. Probabilistic topic models and specifically Latent Dirichlet Allocation (Blei et al., 2003), have come as an extension of pLSA and LSA.

In a probabilistic topic model, we identify a set of latent variables that can best explain observed data - for example, observed words in document. An illustration of the topic modelling approach, similar to the one found in the work of Steyvers and Griffiths (2007), can be found in Fig. 1.

Topics 1 and 3 are thematically related to jigsaw puzzles and used books respectively and contain different distributions over words. Topic 2 is related to dispatch details. Different documents can be produced by picking words from topics, depending on the topics found in each document as well as the words found in each topic. In this simplistic example the documents that could have been item descriptions were generated by sampling from the indicated topics, while the arrow weight demonstrates the probability that a random word in the document was generated by the given topic.

The basis of topic models is the idea that a document is a mixture of topics. Each word w_i in a document (where the index refers to the ith word token) is generated by first sampling a topic from the topic distribution, then choosing a word from the topic-word distribution. We write $P(z_i = j)$ as the probability that the jth topic was sampled for the ith word token and $P(w_i|z_i = j)$ as the probability of word w_i given topic j. The model specifies the following distribution over words in a document (1) when T is the number of topics.

$$p(w_i) = \sum_{j=1}^{T} P(w_i|z_i = j)P(z_i = j) \tag{1}$$

Latent Dirichlet Allocation is another technique that is based on probabilistic principles and overcomes problems previously encountered in pLSA. Topic mixture weights are not individually calculated for each document, but are treated as a k-parameter hidden random variable (where k is the number of topics).

The graphical model in Fig. 2, which is adapted from (Blei et al., 2003), illustrates in plate notation the generative model that describes Latent Dirichlet Allocation: z and d variables identify topics and documents respectively, while $\vartheta(d)$ is the distribution over topics for a document d and $\varphi(z)$ is the distribution over words for a topic z. These distributions can be used to generate documents in the form of a collection of words (w). D is the number of documents, T is the number of topics in the corpus and N_d the topics found in each document. Hyperparameters α and β identify the Dirichlet priors of the above multinomial distributions respectively. These hyperparameters can be changed in order to control the smoothing of the distributions that control how dense topics and documents are.

Instead of directly estimating the two required distributions, ϑ and φ, it is advisable to estimate directly the posterior distribution over z (assignment of word tokens to topics). A Gibbs sampler, which is a special case of a Monte Carlo Markov Chain, is commonly used for this approximation of $p(z)$, which subsequently is used to estimate φ and θ. Iterative evaluation of (2), after a burn-in period, leads to a sampling convergence to an estimate of z. Then using (3) the topic-word and document-topic distributions can be calculated. C^{WT} and C^{DT} are matrices of counts: contains the number of times word w is assigned to topic j, not including the current instance i and contains the number of times topic j is assigned to some word token in document d, not including the current instance i.

$$P(z_i = j|z_{-i}, w_i, d_i, \cdot) \propto \frac{C^{WT}_{w_i j} + \beta}{\sum_{w=1}^{W} C^{WT}_{wj} + W\beta} \frac{C^{DT}_{d_i j} + \alpha}{\sum_{t=1}^{T} C^{DT}_{d_i j} + T\alpha} \tag{2}$$

$$\phi_i'^{(j)} = \frac{C^{WT}_{ij} + \beta}{\sum_{k=1}^{W} C^{WT}_{kj} + W\beta} \theta_j'^{(d)} \frac{C^{DT}_{dj} + \alpha}{\sum_{k=1}^{T} C^{DT}_{d_i k} + T\alpha} \tag{3}$$

The method presents a number of advantages: it is naturally generalised to new documents and the parameters needed are not growing with the size of the training corpus. As the model is defined, there is no mutual exclusivity that restricts words to be part of one topic only. Statistical inference is used in order to approximate the underlying model, which is most probable to have generated these data. Afterwards, this model can be used in applications, for

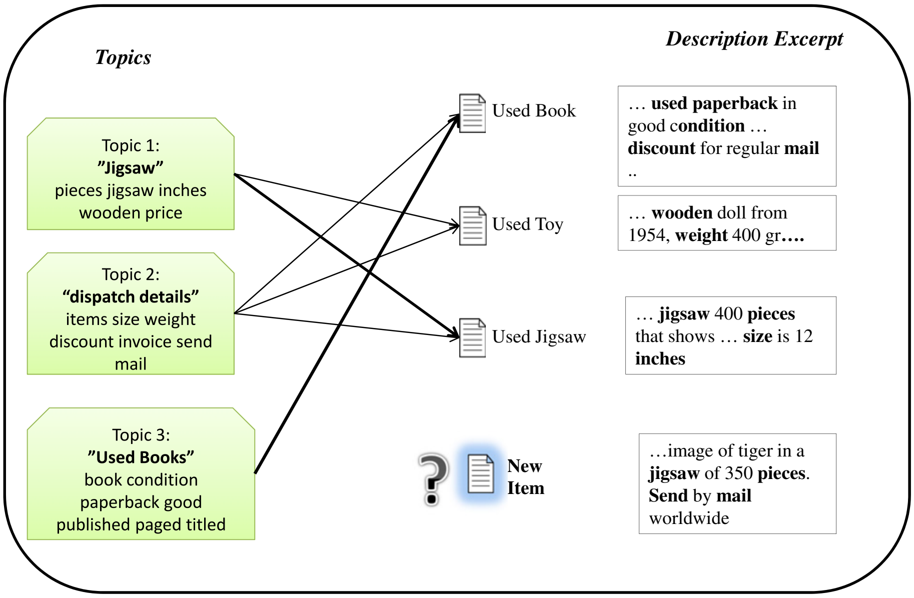

Fig. 1. Topic modelling.

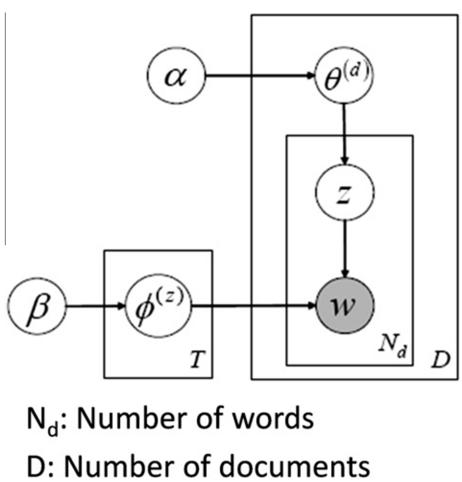

Fig. 2. Latent Dirichlet Allocation plate notation.

example to infer the distribution that could have generated a new, previously unseen, item.

It is important to note that the topics generated by this method are not claimed to be more than latent multinomial variables, nevertheless they are capturing the probability distribution of words. Additionally these distributions are exchangeable, i.e. after retraining the document no assumptions can be made to relate topics from the previous model with topics from the current model.

In this work we use a variant of the model that places a Dirichlet prior on both distributions of topics over documents and of words over topics as applied in previous works (Griffiths & Steyvers, 2004).

3. Our approach

We consider an auction site, where items are traded between buyers and sellers. To this end, the seller is expected to describe the item she expects to sell and define a starting price as well as an instant-buy price. On the other hand, the buyer is browsing between items looking for something that will fulfill her needs.

To support this activity we first use the existing item descriptions in order to extract the latent topic models and pre-calculate the similarities between items and terms, where terms are the words encountered in the descriptions. Then we utilize these similarities in our recommender system for: (1) recommending relevant items to buyers and (2) recommending relevant items and terms to sellers. Usage of topic models allows our system to take into account similarities discovered in the latent topics.

3.1. Topic model and calculation of similarities

The first part of the process is that of extracting the probabilistic topic model and calculating numerical similarities that will be subsequently used for recommendations.

Before extracting the topic model, words contained in the item descriptions are filtered. No special consideration is made about the position and the order of words (bag-of-words assumption).

To train a topic model, we use Latent Dirichlet Allocation (Blei et al., 2003). This training generates the estimated probability distributions of words in topics and topics in documents. These multinomial probability estimations are used as descriptions of each item and each topic respectively. Using the distributions we proceed to calculate a number of similarities in order to provide the recommendation functionality.

First, to calculate the similarity between items, the topic distribution of each item is used. The text of the item description is found relevant to a number of topics and to a certain degree for each topic. To calculate the similarity between the two items we use cosine similarity. For this we calculate the cosine between the vectors that represent the topic distributions of each item, as in (4).

$$item\_sim_{ij}^{topic} = \frac{\sum_{\text{topics that relate with both items}} S_{\text{topic\_item\_i}} S_{\text{topic\_item\_j}}}{\sum_{\text{topics of item i}} (S_{\text{topic\_item\_i}})^2 \sum_{\text{topics of item j}} (S_{\text{topic\_item\_j}})^2} \tag{4}$$

Second, we need to calculate the topic-based similarity between the text of an item description and the terms that have been encountered in the data set. The text is described using topics as above. Similarly the distribution of words in topics is used in order to generate a vector of topics where each word is found. Then cosine similarity is used for calculating the similarity between words and items (5).

$$term\,sim_{il}^{topic} = \frac{\sum_{\text{topics with word l}} S_{\text{topic\_item\_i}} S_{\text{topic\_word\_l}}}{\sum_{\text{topics of item i}} (S_{\text{topic\_item\_i}})^2 \sum_{\text{topics with word l}} (S_{\text{topic\_word\_l}})^2} \tag{5}$$

An overview of the suggested process is illustrated as a series of steps in the flowchart in Fig. 3.

The similarities between items and terms calculated above are used as a basis for providing recommendations in the electronic market environment.

3.2. Buyer item recommendations

A user and potential buyer at an electronic marketplace can reach possible items in a number of different ways, i.e. by looking for items in the specific categories or browsing through the latest items. When she reaches an item of interest, we consider that she can access its characteristics, and specifically the item description as well as the instant-buy and bidding price.

We isolate this part of the process where a user browses an item and we extend it, using a topic-based recommender system. The description of the item is analysed and the similarity of this distribution to the existing items is computed, as in (4). This similarity is combined with the tf∗idf similarity as in (6) using a mixing parameter μ. The tf∗idf similarity is calculated using the vectors of each item, based on the actual words contained in the item description. This similarity is highlighted and used in combination with topics in order to reflect in the system the increased importance of the fact that specific words are found in two items.

$$item\_sim_{ij} = (1-\mu)\cdot item\_sim_{ij}^{tfidf} + \mu \cdot item\_sim_{ij}^{topic} \tag{6}$$

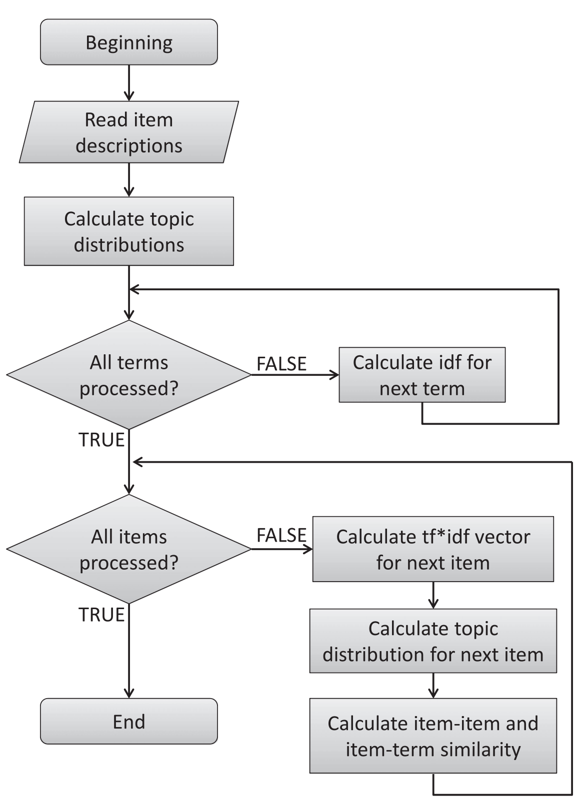

Fig. 3. Flowchart of topic analysis.

Using (6), a list of the most similar items is generated and is provided to the buyer as a recommendation list.

3.3. Seller text and item recommendations

A seller in a marketplace is a person or a company that intends to sell an item she possesses. She aspires to get the highest price possible for her item by finding a person that will appreciate its unique characteristics. Therefore, two needs emerge: the need for clear and full description of the item and the need for appropriate pricing.

In a marketplace both the buyers and the sellers seek to minimize the transaction costs, generally defined as the costs of searching for the right alternative, and negotiating and enforcing a contract (García-Dastugue & Lambert, 2003). In item descriptions important information about the item may be missing. Prospective buyers can lose interest and skip to the next seller if they cannot find the information they require. To address this issue, we provide the most relevant terms not contained in the current description. As the seller types in the description of the item, the text is analysed and a topic distribution is inferred. These topics are then used in order to locate terms that are highly used in similar items. To this end, we can use (5) to compute the similarity between items and words in the dataset. We assume that these terms can reflect important pieces of information that can be found in similar and probably competitive item descriptions, and therefore could be taken into account by the seller to complement her text.

Additionally, even as the marketplace is using auctions, the starting price as well as the instant-buy price has to be set. These prices reflect the seller's view of how valuable her item is in the marketplace. Correct positioning in the marketplace can ensure that the seller will get the highest price for her item. To address this need, while the seller types her item description the system finds similar items and presents their description and their current bidding and selling price. To locate similar items, Eq. (6) is used. This information can be useful for the seller in order to properly price her items.

3.4. System description

In the following lines we will describe the system that, as a part of an electronic market, provides the recommendations previously described.

Technically, the system is based on a relational database and is deployed in an application server. It provides two main functionalities: supporting buyers in locating items of interest and supporting the sellers in describing and pricing the items they intend to sell. To support the required recommendations, respective services have been deployed. In addition to the recommendation services, another service for topic extraction and similarity calculation is required.

In Fig. 4 an overview of the system can be found. Item descriptions found in the electronic marketplace are analysed in order to generate tf∗idf similarities and to extract a probabilistic topic model. This probabilistic topic model is stored as item to topic distributions together with topic to word distributions. These distributions are in turn used to generate topic-based similarities between items and terms. When sellers and buyers browse the marketplace, item and term recommendations are produced based on the topic and tf∗idf similarities. The specifics of these calculations can be found in the previous section.

The proposed system is general and can be applied within any auction-based electronic market. In the next section we describe how the system functions in a specific auction electronic marketplace.

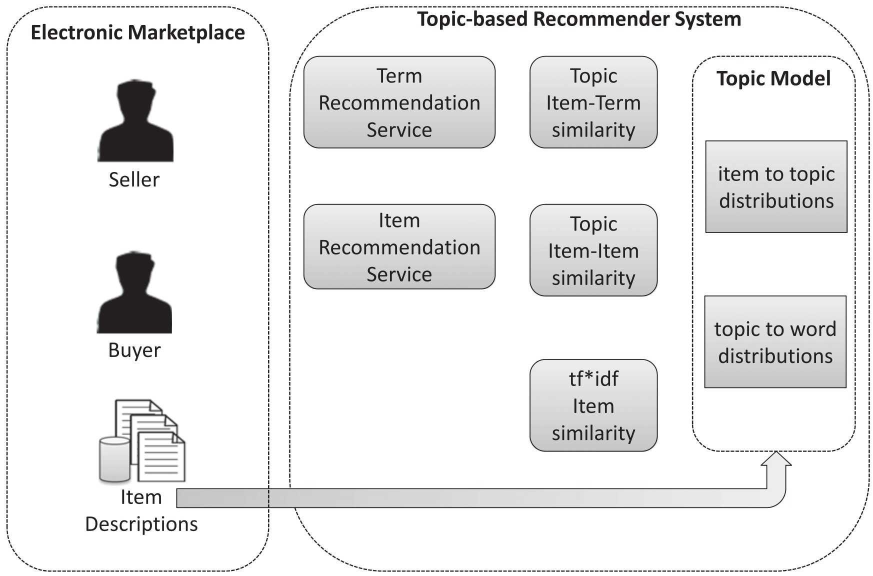

Fig. 4. System architecture.

4. Case study

In this section we describe how the topic-based recommender we propose can be integrated in an actual auction electronic marketplace. We describe the case study, the topic model extracted, a usage scenario and the evaluation performed.

4.1. EBid: An auction electronic marketplace

EBid is a large widely known auction electronic marketplace. As of June 2012 the site contained more than 1950,000 auctions while approximately 50,000 new items were added every 24 h and the value of the listings exceeded 500 million Euros.

We have retrieved 65,205 items from the marketplace's web site between the 20 and the 25th of May, 2012. The descriptions of the items as well as the current prices and instant-buy prices have been stored. These auctions were the open auctions at the time of crawling; no consideration has been made on the category of the listings or any other parameters.

4.2. Topic analysis

Latent Dirichlet Allocation has been performed on the descriptions of the items. To this end, the Mallet (McCallum, 2002) language processing toolkit has been used. Before applying the method, stop-words such as "the", "when" and "he" were removed. Stemming was not used in this experiment n as it was seen to negatively affect domain specific words such as brand names and terminology.

The number of topics after several experiments has been chosen to be 500. A number of 2000 iterations and a burn-in period of 500 were considered enough to provide stable converging topics. The model extracted contains the probabilistic relations between terms and topics, as well as between topics and documents.

An example of representative topics together with the most probable words for each topic can be found in Table 1.

4.3. Usage scenario

In this section we illustrate the functionality provided by the system in a fictional usage scenario. This scenario includes the offering and selling of a product in an auction marketplace. To better demonstrate the usage of a recommender system in an electronic marketplace we have reconstructed two web pages re-using the user interface elements of the real eBid website. We present the system using two users, Sally and Bob, a seller and a buyer respectively.

Sally has an old digital camera. After buying a brand new dSLR model last week, she decides to get rid of her old model. Most of her friends are photography enthusiasts, and her family and neighbours are not interested in an old compact camera. She decides to visit the eBid auction marketplace in order to find potential buyers for the camera.

As she types the words *camera* and the brand of the product (*canon*, see Fig. 5), the block on the right of the page is updated. Similar terms are generated by the system that can assist her in describing her device. These words in our case include *screen*, *lens* and *charger*. These terms help her realise that she needs to include descriptions of these elements in the offer. She writes a full description of the lens and the screen of the camera, as well as the fact that a charger is included. She finishes typing this description and the similar items block on the right part of the web page is also updated. Similar items from eBid are retrieved. She can compare the starting prices and the instant-buy price: similar items are sold for about $200 but they do not include a charger. Therefore she gives the opportunity for instant buying at $210 and she submits the offer.

Bob, in the meantime, browses the eBid website for a used camera for his 15 year old son. He browses through a Canon model, and the system automatically retrieves similar offers on the right part of the screen (Fig. 6). Sally's camera is suggested. Bob considers it is a bargain since a charger is included and buys it instantly.

Table 1
Topics extracted from EBid.

Topic ID	Most probable topic words
6	Light, night, day, sides, led, secret
40	ISBN, book,dog,publisher,978,product,pages,english,weight
28	Items,postage,standard,delivery,vintage,postcards,unposted,posted
52	Battery, charger, batteries, ion
68	Comics, comic, bagged, major, promotional, free, copy
95	Vinyl, producer, written, genre, art, condition, Garfunkel

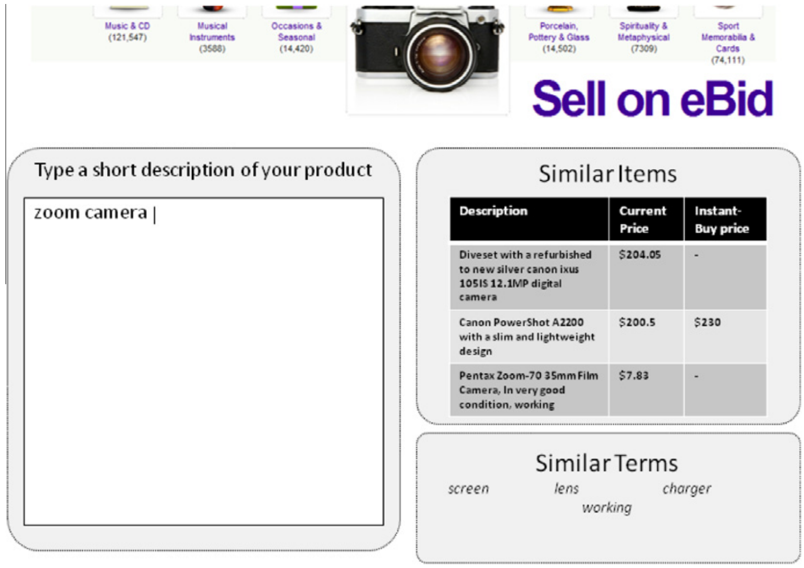

Fig. 5. Seller user interface.

In this scenario we see that the system helps matching buyers and sellers in two ways. First, the participants utilize item similarity for browsing through related items. Second the sellers use the term suggestion as well as the similar item recommendation in order to provide richer descriptions of the goods they provide, as well as competitive pricing.

5. Evaluation

In this section two evaluation methods are presented. In the first method we evaluate the quality of the topics extracted using metrics presented by Wallach et al (Wallach, Murray, Salakhutdinov, & Mimno, 2009). Additionally, we evaluate the usage of the actual recommender system that is built on the topic model.

An approach to evaluating the quality of the topic model produced is that of assessing the probability of heldout documents given a trained model. This number reflects the probability that the extracted model could have generated items that have been held out from the training dataset. We use the left-to-right evaluation probability estimation (Wallach et al., 2009) for assessing the held out probability. In Table 2 we present cross-validated (5-fold cross-validation) held out probabilities evaluated by a left-to-right evaluator as implemented in the Mallet toolkit for a varying number of topics, where the highest log likelihood per token demonstrates a better model. To be able to compare with competing models we need to divide the likelihood to the total number of tokens encountered in the dataset.

Results in Table 2 demonstrate that our choice of a number of 500 topics is reasonable. Additionally, in comparison with previous results we can conclude that the model extracted by the item descriptions is stable and can be used as a basis for providing recommendations.

In the second evaluation method we have randomly chosen 15 scenarios of using the recommender system. The scenarios involve the activity of sellers and buyers in the electronic marketplace, during which they are provided with item and term recommenda-

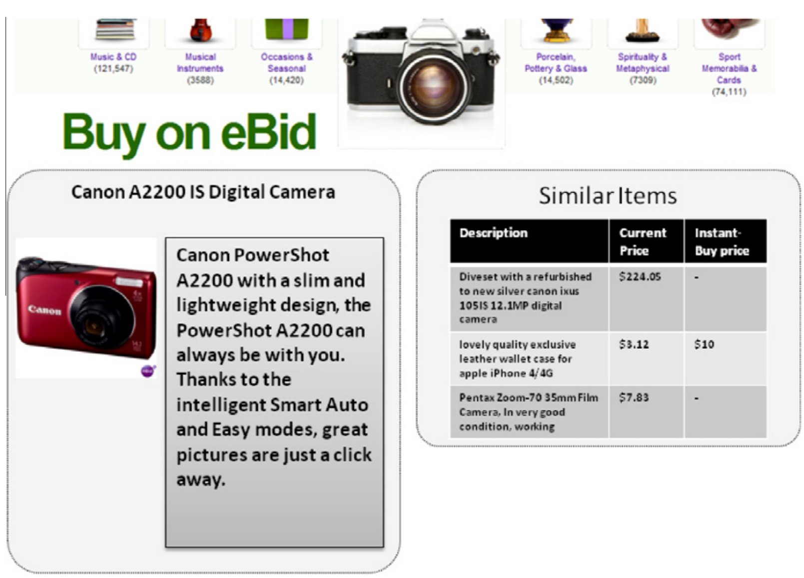

Fig. 6. Buyer user interface.

Table 2
Left-to-right evaluation results.

Number of topics	Held-out probability (average)	
	Log likelihood	Log likelihood/token
300	−8.596.153	−59,700
500	−8.560.038	−59,449
1000	−8.638.102	−59,992
2000	−8.833.686	−61,350

tions. These scenarios have been clearly described in a questionnaire that has been handed out to 32 persons. The participants of this study were residents of Greece out of which 11 have a bachelor's degree, 15 have a master's degree or equivalent and 6 are post-doctoral researchers. The participants are almost equally distributed by gender as 17 are female and 15 male. They were asked to rate how often they use the internet; visit e-commerce sites and electronic marketplace sites. Their responses were ranked in a Likert scale from 1 (never) to 5 (very often) and the results are 4.84 (very often) for internet usage, 3.84 (often) for e-commerce sites and 3.20 (occasionally) for e-marketplaces. The participants' domain of occupation is computer (21), architecture (2), finances (2), electrical engineering (3), humanities (3) and teaching (1). The ages of the participants range from 22 to 51 years while the average is 31.5 years. The study took place during June 2012 and lasted for 2 weeks.

The participants used a web interface to fill the questionnaire and evaluate the usefulness of the recommendations using a Likert scale, where 1 stands for not helpful at all while 5 stands for very helpful. In Fig. 7 we can see the number of each scenario and the perceived usefulness of the provided recommendations. The results of the user evaluation demonstrate that the item and term recommendations generated by the topic-based recommended system are generally perceived as mainly useful and very useful.

6. Related work

Recommender systems have changed from novelties to widely used tools in the e-commerce domain (Huang et al., 2004; Schafer, Konstan, & Riedi, 1999), and specifically in the areas of auctions and electronic commerce.

Some tasks presented in this paper are covered in the literature by agent-based systems, where computational models simulate the actions and interactions of autonomous agents, while other work uses mathematical and machine learning models.

A methodology is suggested for employing the use of an intelligent agent in a negotiation process, and its application is demonstrated in the case of negotiation between companies and clients (Liang, Huang, Tseng, Lin, & Tseng, 2012). The experimental evaluation confirms the advantages of the use of such systems. In another work by Lee (Lee, 2004) agents are used in order to support two functions in the electronic markets: (1) interactive generation of item recommendations and (2) automatic price negotiation. The results presented demonstrate the effectiveness of the system suggested as well as its extended capabilities. Gregg and Walczak (2006) suggest a system based on agents for informing the participants for the history of auctions and prices. Then, this data are used for improving the decision making process by people who take part in auctions. In another work, a utility model is proposed based on satisfying fuzzy constraints to satisfy a solution that is fair for both parties (Li & Wang, 2007). The model is utilizing priorities between fuzzy constraints in order to identify how mutual concessions have to be done when this is necessary while also incorporating the notion of a negotiation argument into the evaluation model allowing the agents can sometimes reach agreements that would otherwise be impossible. Rosaci and Sarné (2012) suggest a system with multiple software agents named AR-SEC, where each device exploited by a customer is associated with a device agent that autonomously monitors his/her behaviour. Furthermore, each customer is associated with a customer agent that collects in a global profile the information provided by his/her device agents and each e-commerce web site is associated with a seller agent. Based on the similarity between the global profiles, customers are partitioned in clusters, where each one is managed by a counsellor agent. These counsellor agents collaborate with the seller agent in order to generate recommendations. This fully decentralized architecture introduces a reduction of the time costs.

Additionally in the literature there have been developed a number of methods that lead to models and knowledge bases that can support recommender systems. Kim and Ahn (2008) applied GA K-means, to a real-world online shopping market segmentation case. This work has contributed in the development of a tool for pre-processing required for recommender systems. Ovsjanikov and Chen (2010) suggest a framework for extending and utilizing a topic model. This extended topic model can utilize unstable binary observations and is positively evaluated in a large datasets of e-commerce sites. Chen and Cheng (2010) describe a supervised learning method for realising semantic types of each term in product document titles. The semantic types are then used in order to improve the relevance of search results. Our methods have signif-

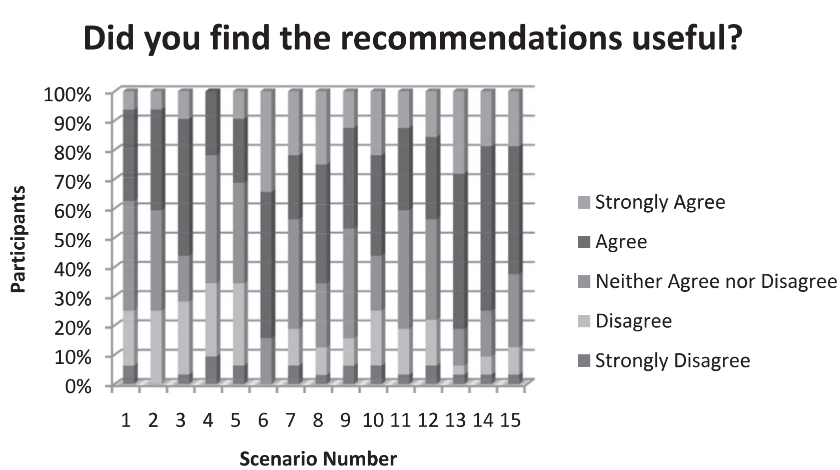

Fig. 7. Recommender system user evaluation.

icant improvement in search result precision in real world document collection and query collections.

In their general form, recommender systems in electronic commerce are designed in order to cluster consumers and products and to produce some similarity rules between them. In another work (Wenxing, Weng, Xie, & Maoqing, 2009) a conceptual framework is proposed for decision support in an online shopping mall, in which a search engine is used externally and a recommender system internally. Wang and Wu (2009) suggest a technique of direct decision support based on a mathematical model that describes the characteristics of the buyers and the profits of the suppliers. This model was developed in order the right product to be proposed to the right person offering the maximum profit for the company. In another work (Castro-Schez, Miguel, Vallejo, & López-López, 2011) the authors describe a product choice service that returns results grouped according to their similarity with the user. They also suggest a learning service for providing the users information about the relations between the characteristics of the items in a category.

Some methods of extracting utility have been developed according to the Multiple Attribute Utility Theory (MAUT) for representing the preference of the deciding person. Huang (2011) studies if these techniques that are based in utility surpass content-based methods for real time recommendations. To confirm these results, experimental evaluations have been performed in two different e-commerce environments. The results demonstrate that the performance of the proposed method depends on the context the recommendation is provided. Wang and Wu (2010) build a functional unit of recommender system for e-commerce that take into account the marketing strategies and the complexity of the items in the user interface. They also propose a technique similar to collaborative filtering with clique effects (CECF) for predicting user preferences while also demonstrating system performance in experimental evaluation.

The related work performed in the area does not consider getting insight from the large volume of unstructured content that is stored in the databases of electronic markets. It is also important that, to the best of our knowledge, there is no publication that describes the suggestion of terms for describing an item offered in a marketplace. In this work we proposed a probabilistic topic model for supporting a recommender system. This approach takes into account information in the unstructured content in order to support the users in identifying interesting items as well as in writing item descriptions.

7. Conclusions and further work

In this work we address the problem of supporting two functions in an electronic marketplace, the user's browsing of items and the seller describing and pricing an item. We propose the suggestion of related items for the user browsing the offers and the provision of relevant terms and items to the seller. We developed a topic-based recommendation system that supports these functions using latent topic analysis. We have implemented the system and applied the algorithms in the data extracted from a widely used auction marketplace. We evaluated both the topic models generated by the data set and the recommender system. The results of this evaluation lead to the conclusion that the methodology we propose can lead to a stable and helpful recommender system in an electronic marketplace.

Further work in this area could be expected in an extended evaluation of the performance of this system in both functions. Additionally, alternative methods such as labelled LDA (Ramage, Hall, Nallapati, & Manning, 2009), that combine structured data, and dynamic topic models (Blei & Lafferty, 2006), that take time into consideration, can be explored. The results of these methods can be interesting since data from an electronic marketplace contain structured information (categories, pricing) as well as change over time.

References

Abramson, M. A., & Means, G. (2001). E-Government 2001. Rowman & Littlefield.

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering, 17*(6), 734–749.

Bakos, Y. (1998). The emerging role of electronic marketplaces on the internet. *Communications of the ACM, 41*(8), 35–42.

Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd international conference on machine learning* (pp. 113–120).

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research, 3*, 993–1022.

Burke, R. (2000). Knowledge-based recommender systems. *Encyclopedia of Library and Information Systems, 69*(Suppl. 32), 175–186.

Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction, 12*(4), 331–370.

Castro-Schez, J. J., Miguel, R., Vallejo, D., & López-López, L. M. (2011). A highly adaptive recommender system based on fuzzy logic for B2C e-commerce portals. *Expert Systems with Applications, 38*(3), 2441–2454. http://dx.doi.org/10.1016/j.eswa.2010.08.033.

Chen, C.-W., & Cheng, P.-J. (2010). Title-based product search – exemplified in a Chinese e-commerce portal. In P.-J. Cheng, M.-Y. Kan, W. Lam, & P. Nakov (Eds.). *Information retrieval technology* (Vol. 6458, pp. 25–36). Springer Berlin/Heidelberg. Retrieved from <http://www.springerlink.com/content/2265532k557h8q07/abstract/>.

Cho, Y. H., Kim, J. K., & Kim, S. H. (2002). A personalized recommender system based on web usage mining and decision tree induction. *Expert Systems with Applications, 23*(3), 329–342.

Fensel, D., McGuiness, D. L., Schulten, E., Ng, W. K., Lim, G. P., & Yan, G. (2001). Ontologies and electronic commerce. *IEEE Intelligent Systems, 16*(1), 8–14.

García-Dastugue, S. J., & Lambert, D. M. (2003). Internet-enabled coordination in the supply chain. *Industrial Marketing Management, 32*(3), 251–263. http://dx.doi.org/10.1016/S0019-8501(02)00269-9.

Gregg, D. G., & Walczak, S. (2006). Auction advisor: An agent-based online-auction decision support system. *Decision Support Systems, 41*(2), 449–471. http://dx.doi.org/10.1016/j.dss.2004.07.007.

Grieger, M. (2003). Electronic marketplaces: A literature review and a call for supply chain management research. *European Journal of Operational Research, 144*(2), 280–294. http://dx.doi.org/10.1016/S0377-2217(02)00394-6.

Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. National Academy of Sciences. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=387300>.

Hofmann, T. (1999). Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 50–57). Berkeley, California, United State: ACM. doi:10.1145/312624.312649.

Hofmann, T. (2004). Latent semantic models for collaborative filtering. *ACM Transactions on Information Systems, 22*(1), 89–115. http://dx.doi.org/10.1145/963770.963774.

Huang, S. (2011). Designing utility-based recommender systems for e-commerce. Evaluation of preference-elicitation methods. *Electronic Commerce Research and Applications, 10*(4), 398–407. http://dx.doi.org/10.1016/j.elerap. 2010.11.003.

Huang, Z., Chung, W., & Chen, H. (2004). A graph model for e-commerce recommender systems. *Journal of the American Society for Information Science and Technology, 55*(3), 259–274. http://dx.doi.org/10.1002/asi.10372.

Huang, Z., Zeng, D., & Chen, H. (2007). A comparison of collaborative-filtering recommendation algorithms for e-commerce. *IEEE Intelligent Systems, 22*(5), 68–78.

Kim, K., & Ahn, H. (2008). A recommender system using GA K-means clustering in an online shopping market. *Expert Systems with Applications, 34*(2), 1200–1209.

Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer, 42*(8), 30–37. http://dx.doi.org/10.1109/MC.2009.263.

Krestel, R., Fankhauser, P., & Nejdl, W. (2009). Latent dirichlet allocation for tag recommendation. In *Proceedings of the third ACM conference on recommender systems* (pp. 61–68).

Lee, W.-P. (2004). Towards agent-based decision making in the electronic marketplace: Interactive recommendation and automated negotiation. *Expert Systems with Applications, 27*(4), 665–679. http://dx.doi.org/10.1016/j.eswa.2004.07.001.

Li, H., & Wang, H. (2007). A multi-agent-based model for a negotiation support system in electronic commerce. *Enterprise Information Systems, 1*(4), 457–472. http://dx.doi.org/10.1080/17517570701633267.

Liang, W.-Y., Huang, C.-C., Tseng, T.-L., Lin, Y.-C., & Tseng, J. (2012). The evaluation of intelligent agent performance – an example of B2C e-commerce negotiation.

Computer Standards and Interfaces, 34(5), 439–446. http://dx.doi.org/10.1016/j.csi.2012.02.003 (Bill).

Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE, 7*(1), 76–80. doi:10.1109/MIC.2003.1167344.

McCallum, A. K. (2002). MALLET: A machine learning for language toolkit. <http://www.mallet.cs.umass.edu>.

Ovsjanikov, M., & Chen, Y. (2010). Topic modeling for personalized recommendation of volatile items. *Machine Learning and Knowledge Discovery in Databases*, 483–498.

Ramage, D., Hall, D., Nallapati, R., & Manning, C. D. (2009). 0024 – Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora. In *Proceedings of the 2009 conference on empirical methods in natural language processing: Volume 1–Volume 1* (pp. 248–256).

Rosaci, D., & Sarné, G. (2012). A multi-agent recommender system for supporting device adaptivity in e-commerce. *Journal of Intelligent Information Systems, 38*(2), 393–418. http://dx.doi.org/10.1007/s10844-011-0160-9.

Salton, G., Wong, A., & Yang, C. S. (1975). A vector space model for automatic indexing. *Communications of the ACM, 18*(11), 620.

Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2000). Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 2nd ACM conference on electronic commerce* (pp. 158–167). Minneapolis, Minnesota, United States: ACM. doi:10.1145/352871.352887.

Schafer, J. B., Konstan, J., & Riedi, J. (1999). Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on electronic commerce* (pp. 158–166).

Schwartz, B. (2005). The paradox of choice: Why more is less. Harper Perennial.

Sharifi, H., Kehoe, D. F., & Hopkins, J. (2006). A classification and selection model of e-marketplaces for better alignment of supply chains. *Journal of Enterprise Information Management, 19*(5), 483–503. http://dx.doi.org/10.1108/17410390610703639.

Steyvers, M., & Griffiths, T. (2007). Probabilistic topic models. *Handbook of Latent Semantic Analysis*, 424–440.

Wallach, H. M., Murray, I., Salakhutdinov, R., & Mimno, D. (2009). Evaluation methods for topic models. In *Proceedings of the 26th annual international conference on machine learning* (pp. 1105–1112).

Wang, H.-F., & Wu, C.-T. (2009). A mathematical model for product selection strategies in a recommender system. *Expert Systems with Applications, 36*(3, Part 2), 7299–7308. http://dx.doi.org/10.1016/j.eswa.2008.09.006.

Wang, H. F., & Wu, C. T. (2010). A strategy-oriented operation module for recommender systems in e-commerce. *Computers and Operations Research*.

Wenxing, H., Weng, Y., Xie, L., & Maoqing, L. (2009). Design and implementation of web-based DSS for online shopping mall. In *IEEE international conference on control and automation, ICCA 2009* (pp. 1308–1313). Presented at the *IEEE international conference on control and automation, ICCA 2009*. http://dx.doi.org/10.1109/ICCA.2009.5410285.

Xu, B., Zhang, M., Pan, Z., & Yang, H. (2005). Content-based recommendation in e-commerce. In O. Gervasi, M. L. Gavrilova, V. Kumar, A. Laganà, H. P. Lee, Y. Mun, & C. J. K. Tan (Eds.). *Computational Science and Its Applications – ICCSA 2005* (Vol. 3481, pp. 946–955). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from <http://www.springerlink.com/content/kgj5lum7pvm8ep8p/>.