An Experiment on Recommender Systems for SME Online Shops

Henrik Stormer
Nicolas Werro
Daniel Risch

Department of Informatics

University of Fribourg

Fribourg, Switzerland

[henrik.stormer|nicolas.werro|daniel.risch]@unifr.ch

Abstract

Recommender systems are often used in electronic shops in order to suggest similar or related products, potentially interesting products for a given customer or a set of products for a marketing campaign. Most recommender systems use the collaborative filtering method in order to provide the personalization information. The collaborative filtering method is a very efficient and convenient way of achieving personalization as there is no need to introduce semantic information about the products or to manually link products and users together. However the collaborative filtering technique does need a dense matrix in order to return pertinent recommendations. This paper proposes a way of combining several types of information in order to improve the density of the input matrix. The presented solution focuses on small and medium-sized online shops that can benefit from the presented results when they want to implement a recommender system in their application

Keywords

Information technology, information attributes, information processing, user behaviour.

INTRODUCTION

In the past years, the number of personalization applications has strongly increased, especially in the field of electronic commerce where personalization becomes an important success factor (Manber 2000, Schubert 2002). The term personalization means the filtering of information for each particular person in order to provide the customers a customized or personalized interaction with a company's products, services, web site and employees (Deitel 2001). The personalization concept is a fundamental requirement for online shops. In contrast to traditional shops, electronic shops cannot provide the personal contact and the individual consultation which are important means of the customer relationship management. Hopefully, online shops may take advantage of the personalization mechanisms which can, at least partially, compensate the weaknesses of the virtual contact and help to efficiently manage the customer relationships.

Recommender systems are often used in electronic shops in order to suggest similar related or potentially interesting products for a given customer or a set of products for a marketing campaign. Most recommender systems use the collaborative filtering method in order to provide the personalized information. The starting point for collaborative filtering is an m-by-n-matrix (called rating matrix) with m referring to customers (rows) and n referring to products (columns). By using different techniques, the similarities between the products (item-based technique) or between the users (user-based technique) are calculated.

The information used to fill the rating matrix can either be gained explicitly or implicitly. Explicit information is entered by the customer directly whereas implicit information is retrieved from the user's interaction with the shop. Explicit information includes product ratings given by the customer, implicit information includes the orders and the clickstream analysis. The collaborative filtering method is a very efficient and convenient way of achieving personalization as there is no need to introduce semantic information about the products or to manually link products and users together. The customers' interactions with the shop is the only required information, however, the collaborative filtering technique does need a dense rating matrix in order to return pertinent recommendations.

The requirement of having a dense rating matrix to use the collaborative filtering method is problematic for small and medium-sized online shop systems which cannot gather enough information about their customers. In order to enlighten this problem this paper presents an experiment which was done for gaining implicit and explicit information. The aim of this paper is to compare the rating matrix density of different data sources and then, to combine explicit and implicit information in order to improve the rating matrix density. The combination of different users' information allows small and medium-sized online shops to significantly increase the rating matrix density and, therefore, the quality of the recommendations.

The remainder of the paper is divided into four sections and is structured as follows: Section 2 provides a deeper insight into the field of recommender systems. The experiment and its results are presented in Section 3. Finally Section 4 gives the conclusion and an outlook.

RECOMMENDER SYSTEMS

The purpose of recommender systems is to recommend products according to the users preferences. Well known applications of recommender systems can be found in the field of books (Linden 2003), music (McCarthy 1998, Chao 2005) and movies (Ling 2005). The broad area of recommender systems has been introduced in the mid-1990s by some early papers on collaborative filtering (Resnick 1994, Shardanand 1995). Meanwhile the term recommender system is more common because it does comprise content-based filtering, collaborative filtering as well as hybrid approaches.

Recommender System Classification

Recommender systems can be classified in three groups based on the approach used to generate the recommendations (Adomavicius 2005):

- Content-based filtering approach
- Collaborative filtering approach
- Hybrid approach

For the content-based filtering approach attributes are assigned to each product. By using information retrieval techniques on those attributes it is possible to derive the similarity between the products, so that two products with common attributes have a grade of similarity (Basu 1998). The advantage of content-based filtering is the possibility of precisely defining relations between products, namely for cross or up-selling. However this advantage comes up at a high price. On the one hand, this approach requires the manual definition of a great number of additional information, e.g. keywords and attributes for each product. This information should be permanently up-to-date. On the other hand, the content-based filtering uses complicated data mining techniques to generate the personalized information.

In contrast to content-based filtering, the collaborative filtering approach only needs information about the user interaction and transaction such as products ratings, orders or clickstream information in order to provide recommendations. This information is continuously provided by the users when browsing the websites, buying or rating products. Another major difference is that the collaborative filtering approach is based on customer context information. So the strength of this approach is its full automation and its user-based semantic. However this approach requires a certain amount of data in order to provide valuable results, i.e. the number of customers and more important the quantity of users' transactions (often called the cold start problem and the first-rater problem).

The third class of recommender systems uses a hybrid approach which is a combination of the content-based and the collaborative filtering (Burke 2002). This approach combines the advantages of having a precise description of the relationships between the objects based on the keywords and on the users' interactions. This allows pertinent recommendations from the beginning with a continuous improvement over time by gathering and using more and more users' information.

In this paper, we focus on collaborative filtering, which is the most suited approach for small and medium-sized online shop systems, and will depict a way of minimizing the cold start and the first-rater problem.

User-based and Item-based Collaborative Filtering

The collaborative filtering approach can be implemented using user-based or item-based methods. Both take as input the rating matrix with the customers in the row dimension and the products in the column dimension. This two-dimensional matrix represents the relationships between users and products either based on product ratings, purchased products or clickstream data. If product ratings are considered, each element at the intersection of a product and a customer will contain a value between -1 and +1 representing the judgement of the customer for the product where -1 denotes a strong dislike and +1 a strong affection. Figure 1 shows an example of a rating matrix.

	DVD Lost in Translation	DVD The Bourne Supremacy	DVD The Life Aquatic
Mr. Smith	$\emptyset$	$+0.5$	$\emptyset$
Mr. Johnson	-0.5	-1	$+0,5$
Mrs. Miller	$+1$	$+0.5$	$\emptyset$

Figure 1: Rating matrix example

In the example, Mr. Miller is a big fan of the product *DVD Lost in Translation* because he rated it with the highest value (+1). However, Mr. Johnson doesn't like the product and therefore rated it low (-0.5).

The same principle applies for the orders and the clickstream information with each cell containing a value between 0 and +1. Thereby +1 denotes a purchase of a product and 0 stands for a product that has not been bought yet. In the case of clickstream data the values between 0 and +1 inform how often a user has visited a webpage containing a particular product.

When applying the user-based method, in a first step, similarities between users are calculated. This calculation can be achieved applying different mathematical formulas. In this paper, the similarity between the users is assessed using the cosine method (Resnick 1994). Once the similarities between all users have been calculated a new matrix with the customers on both dimensions and the similarities as entries is returned (see Figure 2).

	Mr. Smith	Mr. Johnson	Mrs. Miller
Mr. Smith	$+1$	-1	$+1$
Mr. Johnson	-1	$+1$	-0.8
Mrs. Miller	$+1$	$+1$	-0.8

Figure 2: Similarities between customers

Based on this matrix, it is possible for each user to extract the group of most similar users (nearest neighbours). This group is then used in a second step to derive the product recommendations. The principle of the recommendation is pretty obvious; if Mr. Smith is very similar to Mrs. Miller and if Mrs. Miller strongly likes a product that Mr. Smith hasn't bought yet, the chance that Mr. Smith also likes this product is rather high. The user-based method returns personalized recommendations as each user receives propositions based on his profile (see Figure 3). In our example, it is likely that Mr. Smith is fond of the product *DVD Lost in Translation* because he is very similar to Mrs. Miller who has rated this product high.

Figure 3: Personalized recommendations based on the user's profile

In contrast to the user-based method, the item-based method directly derives the similarities between the products. Once again, several mathematical approaches can be used to calculate these similarities. In this paper, the methodology of Deshpande and Karypis (2004) has been chosen. This methodology calculates the probability that a product X will be bought if a product Y has already been bought. This represents a not personalized recommendation as every user viewing a given product will get the same recommendations. The item-based method is often referred under the motto 'Customers who bought this item also bought the following items'.

A deeper introduction to the common algorithms used for user-based and item-based collaborative filtering as well as an analysis can be found in the paper of Sarwar et al. (2200).

RECOMMENDER EXPERIMENT

The goal of the recommender experiment was to retrieve implicit as well as explicit information in a real case scenario. Explicit information was entered by the customer directly by means of a common 5 star rating. Implicit information was inferred from the behaviour of the customer on the online shop, i.e. the orderings and the clickstream data.

Experimental Setup

The experiment was started in November 2005 by setting up an online shop containing 149 movie DVDs. All DVDs contained movies released within the last 7 years. Nearly 200 students were asked to join the experiment and 83 agreed to participate. The experiment was divided in two parts:

- In the first part, all students had to virtually buy some DVDs that they already own. This part was done to gain order and clickstream information (see Figure 4).
- In the second part, each student had to rate 5 products that were presented to him. If the student knew the movie, he was asked to evaluate it with grade ranging from one to five stars.

In the first part of the experiment, 462 products were bought by the 83 participating students. An average student bought then 5.57 products. Altogether, 109 different products were ordered, meaning that 40 were not sold at all (nearly 28.8%). During the second part of the experiment, 99 ratings for 58 different products have been submitted. For calculating recommendations, user-based and item-based collaborative filtering algorithms have been implemented.

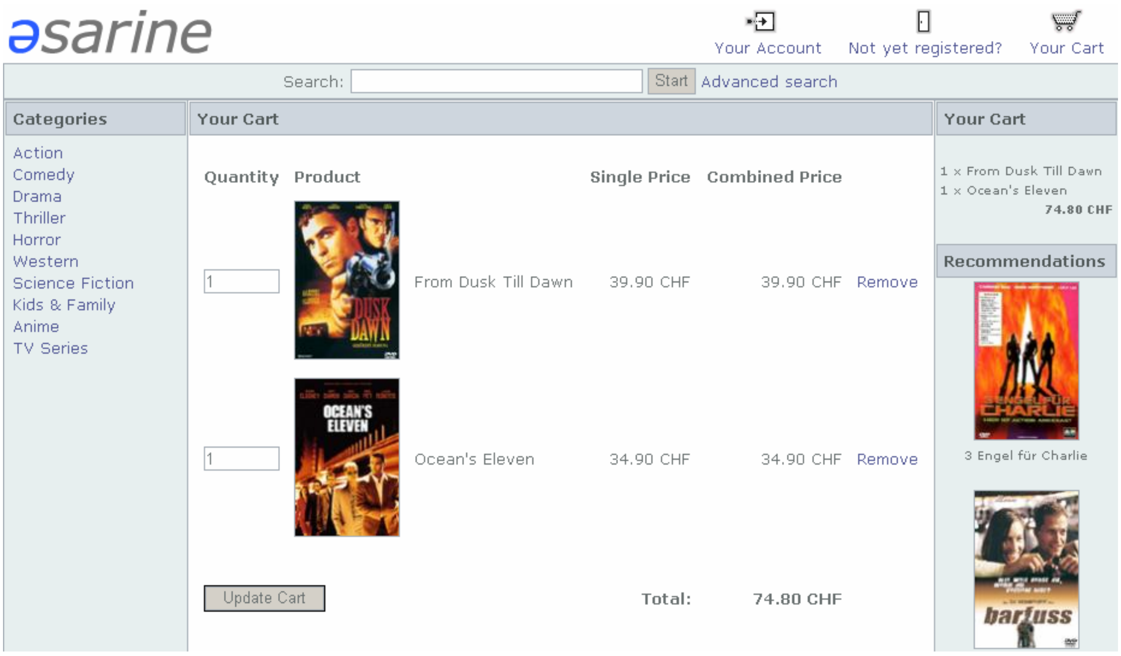

Figure 4: The order and clickstream information is gained by browsing and buying DVDs

Matrix Density with a Single Information Source

With the 149 available products and the 83 participating students, the rating matrix R contains 149 x 83 = 12367 cells. As already mentioned, three different information sources have been utilized in order to fill the rating matrix:

- The products ratings assigned by the customers
- The bought items derived from the orders
- The clickstream information

The products ratings are the best source of information an online shop can obtain because they reflect the final judgement of a customer for a given product. Unfortunately this information is rare as customers normally rate only the products they have bought and only a rather small group of users does use this functionality. The 99 ratings of the experiment using the five star model were defined into the rating matrix as follows: -1 for one star, -0.5 for two stars, 0 for three stars, +0.5 for four stars and +1 for five stars. Using this information, not even 1% of all elements of the rating matrix are filled. By using the products ratings information the matrix density is much too low to enable valuable recommendations.

The orders information is also a good way of deriving the customers preferences. This information is easily accessible and is much more dense than the product ratings. However people sometimes buy products they actually do not like. This is the case for movies they have not seen yet and think they will like it. Therefore an explicit rating is more reliable than order information. In the rating matrix, the 462 products bought by the 83 students were marked by a +1. In the experiment, each student ordered an average of 5.57 products out of the 149 available. Using the bought items information results in a matrix density of 3.74%. This matrix density is a realistic value for small and medium-sized online shop systems. This value allows the calculation of valid recommendations, however a higher density would enable a better matching between the customers and more accurate results.

The clickstream data is the last kind of information which can lead to the definition of the users taste. This is the most substantial but also the most doubtful information source. There is however a correlation between the visited webpages and the user's interest, especially if the user visited several times the same product's page. In order to reflect this information into the rating matrix, each time a customer visits a product's page the corresponding cell is increased by 0.1. By using the clickstream information the rating matrix achieves a density of 6.2%.

Matrix Density with combined Information Sources

The problematic of small and medium-sized shop systems is the opposition between the quality of the sources and their matrix density. In order to overcome this problem, a combination of different sources can be achieved in order to use the best available information and to improve the matrix density.

The combined rating matrix can be obtained by following the rules:

1. If the customer has rated the product, the rating information is used. As the rating is the most valuable information, it surpasses the orders and the clickstream information. The definition of this value is defined in the last subsection. If the customer did not rate the product, we proceeded with (2).

2. If the customer has bought the product a 0.8 value is set. The orders are the second best indicators, they outmatch the clickstream information. Note that a 0.8 value is set instead of a +1 value of the last subsection. This is done to include the little uncertainty of the orders information. If the customer did not bought the product, we proceeded with (3).

3. If the customer has visited the product page a value between 0.1 and 0.6 is set following the rule of the last subsection. The clickstream information is the less pertinent information but is useful if no ratings or orders are available. If the customer has not visited the product's page the matrix cell is not initialized.

By combining the three available information sources, it is possible to achieve a matrix density of 6.8%. This combination does not only improve the matrix density of 0.6 over the clickstream information but also provide a much better information quality. Explicit ratings have a much better impact than implicit ones (Herlocker 2004). Therefore, we use explicit ratings when available. By applying also implicit ratings, the resulting rating matrix becomes more dense.

A lot of other researchers rely on rating matrices that have a much higher density. Examples are the MovieLens database (Ling 2005) or the widely used EachMovie dataset (Pannock 2000, Domingos 2003), containing 2.8 million ratings from over 70.000 users. This leads to an average of 40 ratings per user. This is more than 4 times better then our best combination. Additionally, all ratings from the EachMovie dataset are explicit.

The proposed combination of the information sources allows at the same time to improve the rating matrix density (i.e. avoid the cold start problem and the first-rater problem) and to improve the quality of the quality of the information inside the matrix. This approach could allow small and medium-sized online shop systems to fully take advantage of the collaborative filtering approach.

CONCLUSION AND OUTLOOK

This paper showed how different implicit and explicit information could be combined to enhance the rating matrix. The combination was done using real case data gained from an experiment. The results of this paper could be used by other online shop vendors aiming to implement a recommender system in their application.

We focused on small and medium-sized online shops. These systems often cannot afford to use the content-based filtering approach due to the high investment to maintain the product related information. On the other hand, by using a collaborative filtering approach they face the matrix density problem. The combination of implicit and explicit information as done in this paper offers a simple and efficient way to improve the recommendation results.

The presented approach could be extended by the following steps:

- This paper concentrated on collaborative filtering. By using a hybrid approach, the results could be improved. However, content-based filtering approaches are much more complicated to implement. Additionally, all hybrid approaches benefit from this work because the collaborative filtering part is improved.

- The obvious way of improving the recommendations is done by using a better algorithm. In the experiment, we used standard algorithms and concentrated on the input matrix. There exist numerous approaches that have been proven to be better than the standard ones. This could be another way to improve the results.

- To gain better results, the customers should be asked to give explicit ratings. To encourage them, they could get benefits like coupons or discounts.

- Sometimes, managers of small and medium sized companies know their customers quite well. By providing a way to explicitly rate certain products for them the rating matrix density could be improved.

For the future, we are interested in comparing our results with other real online shop data. Therefore, we plan to ask other online shops to provide us their (anonymous) data. Another interesting research direction is the usage of fuzzy technology (Werro 2005a, Werro2005b) to create a hybrid approach where the fuzziness is used to overcome the problem of maintaining the product related information.

REFERENCES

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*, vol. 17, no. 6, pp. 734-749.

Basu, C., Hirsh, H., Cohen, W. 1998, 'Recommendation as classification: using social and content-based information in recommendation', *Proceedings of the 1998 workshop on recommender systems*, Menlo Park, CA, pp. 11-15.

Burke, R. 2002, 'Hybrid Recommender Systems, Survey and Experiments', *User Modeling and User Adapted Interaction*, no. 12-4, pp. 331-370.

Chao, D. L., Balthrop, J., Forrest, S. 2005, 'Adaptive Radio: Achieving consensus using negative preferences', *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*, ACM Press, pp. 120-123.

Deitel, H. M., Deitel, P. J., Steinbuhler, K. 2001, 'E-Business and E-Commerce for Managers', Prentice Hall.

Deshpande, M., Karypis, G. 2004, 'Item-Based Top-N Recommendation Algorithms', *ACM Transactions on Information Systems*, vol. 22, no. 1.

Domingos, P., Richardson, M. 2003, 'Mining the network value of customers', *Proceedings of the 7th International Conference on Knowledge Discovery and Data Mining*, ACM, pp. 57-66.

Herlocker, J. L., Konstan, J. A., Terveen, L. G., Riedl, J. T. 2004, 'Evaluating Collaborative Filtering Recommender Systems', *ACM Transactions on Information Systems*, vol. 22, no. 1.

Linden, G., Smith, B., York, J. 2003, 'Amazon.com Recommendations', *IEEE Internet Computing*, vol. 3, no. 2.

Ling, K., Beenen, G., Ludford, P., Wang, X., Chang, K., Cosley, D., Frankowski, D., Terveen, L., Rashid, A. M., Resnick, P., Kraut, R. 2005, 'Using social psychology to motivate contributions to online communities', *Journal of Computer-Mediated Communication*, vol. 10, no. 4.

Manber, U., Patel, A.,Robison, J. 2000, 'The Business of Personalization: Experience with Personalization of Yahoo!', *Communications of the ACM*, vol. 43, no. 8.

McCarthy, J. F., Anagnost, T. D. 1998, 'MusicFX: An arbiter of group preferences for computer supported collaborative workouts', *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work*, ACM Press, pp. 363-372.

Pennock, D. M., Horvitz, E., Lawrence, S., Giles, C. L. 2000, 'Collaborative filtering by personality diagnosis: A hybrid memory- and model-based approach', *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 473-480.

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. 1994, 'GroupLens: An Open Architecture for Collaborative Filtering Netnews', *Proceeding of Computer Supported Cooperative Work Conference*.

Resnick, P., Varian, H. R. 1997, 'Recommender Systems', *Communications of the ACM*, vol. 40, no. 3, pp. 56-58.

Sarwar, B., Karypis, G., Konstan, J., Riedl, J. 2000, 'Analysis of recommendation algorithms for e-commerce', *Proceedings of the Electronic Commerce Conference*.

Schubert, P., Koch, M. 2002, 'The Power of Personalization: Customer Collaboration and Virtual Communities', *Proceedings of the Eighth Americas Conference on Information Systems (AMCIS)*.

Shardanand, U., Maes, P. 1995, 'Social information filtering: algorithms for automating word of mouth', *Proceedings of conference on human factors in computing systems*, Denver, CO, pp. 210-217.

Werro, N., Stormer, H., Meier, A. 2005, 'Personalized Discount - A Fuzzy Logic Approach', *Proceedings of the 5th IFIP International Conference on eBusiness, eCommerce and eGovernment*, Poznan.

Werro, N., Meier, A., Mezger, C. 2005, 'Concept and Implementation of a Fuzzy Classification Query Language', *Proceedings of the International Conference on Data Mining*, World Congress in Applied Computing.