

Use of Radio Frequency Identification for Targeted Advertising: A Collaborative Filtering Approach Using Bayesian Networks

Esma Nur Cinicioglu, Prakash P. Shenoy, and Canan Kocabasoglu

School of Business, University of Kansas, 1300 Sunnyside Avenue,
Summerfield Hall, Lawrence, KS 66045-7585, USA
{esmanur, pshenoy, canan}@ku.edu

Abstract. This article discusses a potential application of radio frequency identification (RFID) and collaborative filtering for targeted advertising in grocery stores. Every day hundreds of items in grocery stores are marked down for promotional purposes. Whether these promotions are effective or not depends primarily on whether the customers are aware of them or not, and secondarily whether the customers are interested in the products or not. Currently, the companies are incapable of influencing the customers' decision-making process while they are shopping. However, the capabilities of RFID technology enable us to transfer the recommendation systems of e-commerce to grocery stores. In our model, using RFID technology, we get real time information about the products placed in the cart during the shopping process. Based on that information we inform the customer about those promotions in which the customer is likely to be interested in. The selection of the product advertised is a dynamic decision making process since it is based on the information of the products placed inside the cart while customer is shopping. Collaborative filtering will be used for the identification of the advertised product and Bayesian networks will be used for the application of collaborative filtering. We are assuming a scenario where all products have RFID tags, and grocery carts are equipped with RFID readers and screens that would display the relevant promotions.

Keywords: RFID, targeted advertising, Bayesian networks, learning Bayesian networks, collaborative filtering.

1 Introduction

Love it or hate it, grocery shopping occupies a significant amount of time of your life. It may seem as a straightforward task—all you need is just a shopping list. However, almost 60% of household supermarket purchases are unplanned and the result of in store decisions [Inman and Winer, 1999]. Even having a shopping list is sometimes not enough. The huge variety of products offered turns the grocery stores into labyrinths, so you have to be cautious not to get lost between the aisles as you search for the products on your list. By the time you find the item you are looking for, you may be overwhelmed to see how many different brands offer the same item.

Grocery basket selection can be thought as a reflection of customers' needs. Ideally, the products selected should represent the results of a comparison made by the customer based on the price and quality aspects of the products. Considering the nature of a simple grocery-shopping trip described earlier, a careful selection of products requires the devotion of a significant amount of time and energy on the customers' side. On the contrary, modern life imposes time constraints on the customers, which make them unwilling to spend any more time for grocery shopping than is necessary. As a result, the explosion of the size of product assortments (more than 100,000 references in a large hypermarket) no longer allows for a clear identification of differences in quality and prices inside the product mix [Bruno and Pache, 2005].

The situation on the retailers' side is also not very promising. The competition between the grocery stores is increasing every day, forcing the retailers to find new ways to influence the purchase decisions of the customers. Today a huge variety of methods to track and analyze the customers' behavior in e-commerce systems is available. For instance, amazon.com makes real-time recommendations (Customers who bought this item also bought...) to its customers based on the information of the products that have been put in a shopping cart or reviewed by the customer. However, in traditional retail stores, such systems are not used, and, therefore, the customer's behavior is considered as a black box" [Decker, 2003].

As a way to affect the consumers' purchase decisions and to introduce new products, the shelf configurations of the stores are periodically rearranged. Although this might help the retailers to find the optimal allocation of the products, it bothers the customers for not being able to find the products they are looking for.

Another way to influence the customers' purchase decisions is to do promotions. Every day hundreds of items inside the grocery stores are advertised as an effort to trigger the demand of customers for those products on promotion. Whether these special offers will become a subject of interest to customers primarily depends on whether the customers are aware of them or not. Studies suggest that more than half of the shoppers who purchased an item that was on sale were unaware that the price was reduced [Mittal, 1994]. To inform the customers about your ongoing promotions you may increase the rate of your advertisements, which will increase your costs significantly.

Advertisements and promotions are two effective ways to influence sales. However, an advertisement of a promotion will be more successful, if the promotion is particularly advertised to those shoppers who are likely to be interested in the offer. Clearly, the purchase of a product on promotion by an informed customer does not only show the success of the promotion; it is also a valid indicator of customers' interest in that particular product. Thus, the success of a promotion secondarily depends on whether a customer is interested in that product or not.

In order to understand the underlying patterns of customers' purchase decisions, most grocery stores identify its customers through customer loyalty cards via which they keep track of the products purchased by the customer. Based on this information they tailor promotions to individual customers by giving discount coupons at the checkout. However, these promotions happen after the shopping is over which tremendously reduces the impact of the promotion on the sale. The companies are mostly incapable of influencing the customers' decision making process when they

are shopping since the data about the customers' shopping behavior is only available after the decisions are made, i.e., after the shopping is over.

As a way to interact with the customer during the shopping, grocery stores install kiosks from which customers can get information about the ongoing promotions and the products displayed. However, stress and time pressure potentially force a customer to fully concentrate on the original task where the customer is not willing or able to learn the operation of a complex shopping support system [Schneider, 2004].

Having considered all of this, the e-commerce seem to have a huge advantage over traditional grocery shopping because of their capability to make targeted advertising at the same time as the consumer is shopping. Inspired by the real-time recommendation systems of e-commerce we should be looking for ways to transfer the methods of e-commerce systems to the current state of grocery shopping. The capability of RFID technology to identify individual products and collect real time data about the customer behavior inside a store makes a new model for the traditional grocery shopping feasible.

An outline of the remainder of the paper is as follows. In Section 2, we introduce the RFID technology and describe its capabilities in the domain of operations management. In section 3, we discuss 'collaborative filtering' as a way to identify our recommendations based on the customers' preferences and describe our proposed model for the grocery stores. This model promises to enable the grocery stores to make real time targeted advertising. In order to illustrate the working mechanism of our model we are going to use the data set available for the Netflix prize competition. The details of the data set used are described in section 4. Using WinMine toolkit a Bayesian net will be learned from the data set used. As the next step, using the table distributions learned with WinMine, the same BN will be built in Hugin, a commercial software, which allows us to predict the customers' preferences of the movies based on the given information. In section 5, we illustrate the use of our model via different cases. Finally, in section 6, we summarize and conclude.

2 Radio Frequency Identification

Radio Frequency Identification (RFID) is a generic term for a variety of technologies that use radio waves to automatically identify individual items" [Cavoukian, 2004]. This technology known for over 50 years, prepares to have its real bang in the business world after its potential for commercial applications has been realized. The capability of identifying individual products, ability to track the products through the processes, differentiates RFID from its preceding alternatives; but the real and huge potential of RFID systems is hidden in the massive amount of data that is captured by RFID systems.

An RFID system consists of two basic parts: a tag and a reader. Readers, depending upon design and technology used, may be a read-only or a read-write device [Finkenzeller, 1999]. They capture the information stored or gathered by the tag. The RFID tags can be either active or passive, depending whether they have their own power supply or not. Active RFID tags offer superior performance. Because they are connected to their own battery, they can be read at a much higher range-from several kilometers away. However, they are larger and more expensive. Passive tags

have no power source and no on-tag transmitter, which gives them a range of less than 10-meters and makes them sensitive to environmental constraints [Cavoukian, 2004].

Among the automatic identification systems, barcode technology has been the leader for over 20 years. Nevertheless, with the decreasing cost of the RFID tags, companies have begun to favor RFID systems over barcode technology. Although it is a fact that the reduced costs of the RFID tags have contributed a lot to the present popularity of the RFID systems, this is not the main motive why the RFID systems are preferred over barcodes. In Table 1, we illustrate the potential benefits that companies may achieve in their operation management activities by using RFID systems instead of the barcodes.

Table 1. The potential benefits of RFID systems in operations management activities

	Barcode	RFID	Potential benefit of RFIDs
Data capturing capacity	A barcode can hold only around 1000 characters of data. [Mital, 2003]	Up to 128,000 characters in an RFID chip [Mital, 2003].	The superior data capturing capacity of RFID systems offers enough room for a unique serial number, expiration date or other pertinent information [Sweeney, 2005] -This is a serious drawback of bar codes compared to the RFID systems since in RFID systems information is specific to that individual item [Cavoukian, 2004].
Cost	The barcode system is still a much cheaper identification system than the RFID technology and the experts predict that it will remain to be so.	Today Passive RFIDs sell for less than 50 cents in high volumes, and analysts predict they'll sell for five cents in high volumes by the end of this decade [Dipert, 2005].	Tags are reusable and have very long lives, so in supply chain operations where containers are continually reused, there would be no need to re-label the containers, saving on manpower and other costs associated with label production and fixing [Hopwood, 2005].
Processing times	Only one item can be read at a time because of the line of sight technology required. -The existence of dirt or dust can avoid the reading barcodes.	RFID tags can be read in harsh environments such as snow, fog, etc. with a reading distance ranging from 50 feet to 100 meters and beyond [Cavoukian, 2004]	The processing times of items increases significantly, when bar code systems are in use.
The query of components and subassemblies	Requires positioning the cases so that the labels can be read by the scanners -line of sight reading is required	Automatic check that all items from the bill of material -are received -are placed in the right location -RFID does not require positioning the cases	Convenience in order processing -helps to decrease the labor costs -reduce the order preparation times [Rutner, 2004]
A valid source of information in order preparation and processing	not applicable	How much time a worker spends on the preparation of a particular item can be measured	Management could use this data for -setting benchmarks -evaluating employees -planning labor requirements [Rutner, 2004]
Prevention of Spoilage	not applicable	Sensor-equipped tags can monitor the environment surrounding perishable items and maintain a history of environmental changes	RFID systems can be used -to detect potential spoilage conditions [Curtin, 2005] -to identify the causes of spoilage
Prevention of Theft	not applicable	The capability to locate every individual product within the inventory	Provides a tremendous opportunity for companies to prevent theft
Prevention of Shrinkage	Real time data is not available	Automatic collection of real time data	-the automatic collection of real time data prevents the shrinkage problem, and if not, makes the data available to detect the cause of shrinkage -better replenishments decisions can be made since accurate data are readily available with RFID [Lee, 2004]
Prevention of Stockouts	Captures information on how much is sold form each product	Captures information about the real time data of the current inventory (how much is sold, how much is missing)	The ability of RFID systems to prevent and detect when theft and/or shrinkage is present, makes the data more accurate thus preventing the occurrence of stockouts

The use of RFID systems in commercial applications is an emerging trend and RFID is ready to place itself as the dominant technology used in real world applications. However, the importance of RFID technology is not just limited by the convenience it provides. More importantly, RFID systems create massive amounts of data, which gives the ability to track and trace materials at the case-level within the supply chain, and at the item level from manufacturing to post sales. Therefore, the real question to be answered is: How can we transform this massive amount of data into managerially useful information?

3 Collaborative Filtering

As mentioned earlier the real potential of RFID systems is hidden in the massive amount of data collected through RFID. This application is a perfect illustration of that. We can use the RFID technology for getting real time information about the consumer behavior as they are shopping and that may enable us to inform the customer about the promotions in store in which the customer is likely to be interested. Using RFID we can get information about the products a customer is placing in his shopping basket, and using collaborative filtering we can advertise those products on promotion which the customer is more likely to be interested based on what is already in the customer's shopping basket.

Collaborative filtering, first introduced by Resnick *et al.* (1994), is defined as predicting preferences of an active user given a database of preferences of other users [Mild, 2002]. Depending on the technology used, recommendation systems are classified in two classes, content-based filtering (CBF) and collaborative filtering (CF). Content-based methods make recommendations by analyzing the description of the items that have been rated by the user and the description of items to be recommended [Pazzani, 1999]. The main difference between collaborative filtering and content-based filtering is that CF does not rely on the content descriptions of the items, but depends purely on preferences expressed by a set of users [Yu *et al.*, 2004]. Since collaborative filtering does not depend on error-prone machine analysis of content, it has significant advantages over traditional content-based filtering (ability to filter any type of content, etc.) [Herlocker *et al.*, 2000].

Popescul *et al.* [2001] describe a unified collaborative and content-based system. de Campos *et al.* [2006] describe a Bayesian network model for hybrid collaborative and content-based filtering. Adomavicius and Tuzhilin [2005] survey the field of collaborative, content-based, and hybrid recommender systems. Zhang and Callan [2001] developed an algorithm for setting dissemination thresholds while filtering documents. Linden *et al.* [2003] compare traditional collaborative filtering, cluster models and search based models with their method item-item collaborative filtering.

In e-commerce, collaborative filtering is widely used as a tool for targeted advertising. Using the capabilities of RFID, we might be able to transfer this method to traditional retail stores and base the advertisements on real-time data.

The technique used in collaborative filtering is based either on explicit or implicit voting. The data sets in explicit voting contain users explicit preference ratings for products. Implicit voting refers to interpreting user behavior or selections to impute a vote or preference [Breese *et al.*, 1998]. Our case is an example of implicit voting,

since our model will use binary choice data that identifies whether a product is placed in cart or not.

The model we are proposing is as follows. All products in a grocery store are equipped with RFID tags. The carts in the grocery store are equipped with RFID scanners, which are utilized to collect information about the products that are placed in a customers' cart. In each cart, there is also a screen where the promotions are displayed. The basic idea of our model is to inform the customer about those products on promotion that the customer is likely to be interested in buying based on the products already in the cart. The selection of the product advertised is a dynamic decision making process since it is based on the information of the products placed inside the cart while customer is shopping. Collaborative filtering will be used for the identification of the advertised product and Bayesian networks will be used for the application of collaborative filtering.

At the beginning of the shopping process, there are no products in the cart. At this stage, the system can just display those products on promotion that have the highest marginal probabilities. As the customer places products in the cart, the system can display those products that the customer is likely to be interested in purchasing based on items in the cart.

4 Dataset

The proposed model above requires data captured through RFID systems for the different market baskets of the customers. Since we did not have access to a grocery basket dataset, we decided to use the publicly available Netflix prize competition¹ dataset to illustrate our application [Netflix, 2007]. While it is not quite the same, we were able to convert the Netflix dataset to a basket dataset with movies as the products instead of grocery items.

The training data set of the Netflix prize competition constitutes of 17,770 files, one per movie. Each file contains customer ID, the rating given by the customer, and the date of the rating. The ratings are on a scale from 1 to 5, 5 as being the best rating possible. For the analysis done in this paper, 1,695 movie files from this training data set have been chosen on a random basis. These separate data files are merged into a big data set where the ratings for the movies are sorted based on the customer ID and the date has been dropped out.

The goal of our model is to predict the products that the customer may be interested in based on the products that (s)he has placed in the cart. Trying to interpret the customers' behavior suggests the need for implicit voting instead of a detailed 1 to 5 rating scale. Hence, we transformed our data set into a new data set where the ratings 3, 4 and 5 are replaced by 1's as an indicator of the movies being in a user's basket. If the customer has rated the movie as 1 or 2 or has not rated the movie at all, then the movie rating is replaced with a zero, which means that it is not in the cart. Here we are assuming that the movies not rated by the customer are movies that are not in the customer's cart.

¹ The Netflix prize competition seeks to substantially improve the accuracy of predictions about how much someone is going to love a movie based on the ratings of the movies they have already seen.

In a grocery store, there are literally hundreds of thousands of different products. For the problem of finding associations between the products that are in carts, we need to aggregate the products. For example, tomato sauce may be sold in different brands, different sizes, different packaging, etc., and all of these need to be aggregated into a single product.² The problem of finding a good aggregation can be a difficult one. We need to decide on a number of aggregated products, and a technique to do the aggregation. Cluster analysis from multivariate statistics is one method that can be used for doing the aggregation. The optimal number of aggregated products is an empirical question, and an approximate number can be found by experimentation.

After transformation of the data to the desired format, the next step was to select the movies that are going to be used for creating a Bayes net. In order to select the movies from different groupings we used cluster analysis. The FASTCLUS procedure in SAS was used for cluster analysis, where we limited the maximum number of clusters obtained to thirty³. As a result, we obtained thirty different clusters and chose one movie from each cluster on a random basis. The final data set used to build the Bayes Net constitutes of the movie preferences of 65,535 users for the thirty movies selected. The set of movies selected appears in the Bayes net model shown in Figure 1 below.

Our motivation for learning a Bayes Net is to find the predictive relationships between the movies based on the movies liked or disliked by the customer. WinMine [Heckerman *et al.*, 2000], a tool developed at Microsoft Research, is used to learn a Bayes Net. Using WinMine, the data is divided into a training set and a test set. We performed a 70/30 train/test split and had 45,874 training cases and 19,661 test cases. All of the variables are used as input-output variables (both predicted and used to predict). To set the granularity of the Bayesian network learnt by WinMine, a factor called *kappa* is used, which is a number between 0 and 1. As *kappa* approaches 1, the model becomes very dense. Since our model is already quite dense, we decreased the value of *kappa* from its default value of 0.01 to 0.00001. The resulting BN is given in Figure 1.

The accuracy of the learned model on the test set is evaluated using the log score

$$\text{Score}(x_1, \dots, x_N) = \frac{\sum_{i=1}^N \log_2 p(x_i | \text{model})}{nN},$$

where n is the number of variables in X , and N is the number of cases in the test set. Our model results in a log score of -0.4169 , meaning on average, the log probability that each variable assigns to the given value in the test case, given the values of all other variables, is -0.4169 , which translates to a probability of 0.75. Using WinMine we can also compare the difference between the provided model and the marginal model. A positive difference is desired between the provided model and the marginal model, signifying that the model outperforms the marginal model on the test set. In the same way that a regression model

² In the first iteration, we selected 33 movies that had a large number of user ratings (without doing cluster analysis) and used it to learn a Bayes net. However, that was not very effective in predicting the baskets of users in the test set (lift over marginal was about 0.04364).

³ We did not attempt to determine an optimal (or an approximate) number here. We picked thirty for convenience. Since we obtained good results, we did not experiment with other numbers.

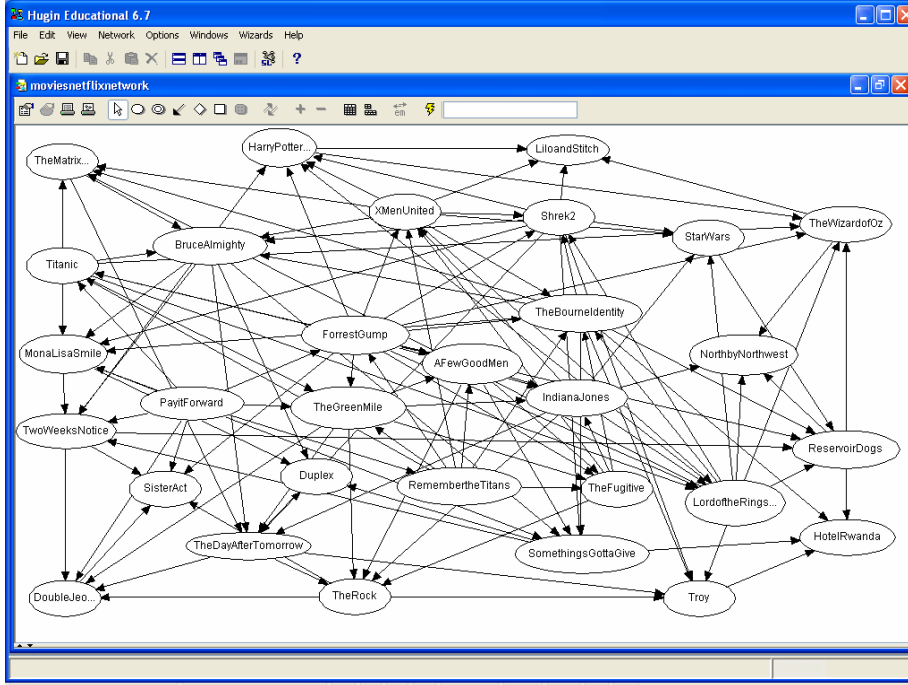


Fig. 1. A Bayes net for 30 movies from the Netflix prize dataset

is more accurate than a simple baseline model chosen in the form of a mean dependent value, the “lift over marginal” log score provides information on how well the model fits the data. The lift over marginal log score in our model is 0.1302, which suggests the performance of our model is quite good. If we ignored the products in the cart and used the marginals for prediction, the average probability of the correct prediction is 0.68 (or log score of -0.5471). Using the products in the cart, the average probability of correct prediction improves to 0.75 (or log score of -0.4169) resulting in a lift over marginal log score of $(-0.4169) - (-0.5471) = 0.1302$. There are many ways of evaluating collaborative filtering recommender systems (Herlocker *et al*, 2004), and the lift over marginal is a good conservative measure of effectiveness for our application.

5 A Case Study

In the previous section we have illustrated how a BN can be learned using the WinMine toolkit. Using the probability tables constructed by WinMine, we constructed the same Bayes Net in Hugin, a commercial software. The conditional probability table used for the movie ‘Lord of the Rings: The Two Towers’ is illustrated in Table 2 below.

Table 2. The conditional probability table for *Lord of the Rings: The Two Towers*

Lord of the Rings: The Two Towers																
Forrest Gump	0								1							
Titanic	0				1				0				1			
X-Men United	0		1		0		1		0		1		0		1	
Indiana Jones	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
0	0.88	0.76	0.53	0.21	0.83	0.59	0.37	0.16	0.80	0.52	0.33	0.13	0.70	0.41	0.27	0.10
1	0.12	0.24	0.47	0.79	0.17	0.41	0.63	0.84	0.20	0.48	0.67	0.87	0.30	0.59	0.73	0.90

The advantage of using Hugin is that we are able to enter evidence to the BN and update all probabilities accordingly using the ‘sum normal’ propagation method. In addition to that, the ‘max normal’ propagation method allows us to find states to the most probable configuration. The state of node with the most probable configuration is given the value of 100. The values for all other states are the relative values of the probability of the most probable configuration in comparison to the most probable configuration.

By using the sum-propagate normal propagation method without entering any evidence, we obtain the marginal probabilities for all the movies in the BN. The results suggest that for the state ‘1’ the movie ‘The Green Mile’ has the highest marginal probability 40.57%, and ‘Duplex’ has the lowest marginal probability 5.23%.

Suppose we want to predict whether a specific customer is going to like the movie *Forrest Gump* or not. Without having any information about the customers’ previous movie preferences the marginal probability for the state ‘1’ is 40.43% and the state for the most probable configuration is ‘0’. Suppose we get the information that the customer rented the movie *A Few Good Men* and liked it. Accordingly, the posterior marginal for *Forrest Gump* increases to 69.75%, the most likely state is still ‘0’. Next, suppose we get the information that the customer also liked *The Wizard of Oz*. The posterior marginal probability for *Forrest Gump* increases to 90.13% and the most likely state changes to ‘1’. The results for this case are summarized in Table 3 below.

As our second case, consider a scenario where we need to choose between the two movies *Mona Lisa Smile* and *Lord of the Rings: The Two Towers* to recommend to the customer. The initial most likely state is ‘0’ for both movies. Based on their marginal probabilities, which are given in Table 4 below, *Lord of the Rings: The Two Towers* should be chosen for recommendation, since it has a much higher marginal probability for the state ‘1’.

Suppose we receive information about movie preferences of the customer to whom we are going to make the recommendation. Learning that the customer liked *Pay It Forward*, *Something’s Gotta Give*, *Two Weeks Notice* and *Titanic* with the particular order given, changes the posterior marginal probabilities. Until we obtain the information that the customer liked *Something’s Gotta Give*, the marginal probabilities indicate that *Lord of the Rings* should be chosen for recommendation. After subsequent observations, *Mona Lisa Smile* takes the lead for recommendation. At the point where we learn that the customer liked *Two Weeks Notice*, the most likely state for *Mona Lisa Smile* becomes ‘1’ where for *Lord of the Rings* it is ‘0’ still. After we get the information that the customer also liked the movie *Titanic* the most likely state for both of the movies becomes ‘1’. The details of posterior marginal probabilities and the most likely states are given in Table 4 above.

Table 3. Posterior probabilities and most likely state for *Forrest Gump*

<i>Information & Rating</i>	<i>Marginal</i>	<i>Most likely state</i>
Prior	40.43%	0
<i>A Few Good Men</i> = 1	69.75%	0
<i>Wizard of Oz</i> = 1	90.13%	1

Table 4. Posterior probabilities and most likely states for *Mona Lisa Smile* and *Lord of the Rings: The Two Towers*

<i>Information & Rating</i>	<i>Mona Lisa Smile</i>		<i>Lord of the Rings: The Two Towers</i>	
	<i>Marginal</i>	<i>Most likely state</i>	<i>Marginal</i>	<i>Most likely state</i>
Prior	19.28%	0	33.85%	0
<i>Pay It Forward</i> = 1	36.86%	0	48.08%	0
<i>Something's Gotta Give</i> = 1	63.05%	0	61.37%	0
<i>Two Weeks Notice</i> = 1	72.22%	1	63.23%	0
<i>Titanic</i> = 1	76.00%	1	67.38%	1

6 Conclusions and Summary

We have proposed a system using RFID and collaborative filtering for targeting advertising in grocery stores. We have illustrated the use of such a system using the Netflix prize competition dataset.

The proposed model promises to influence the customers' decision-making process while shopping, which will increase the success of the promotions. Also, it is very important to notice that the contribution we will get through the proposed model is not just limited by the improvement of promotions. Transferring the methods of e-commerce to actual retail stores through real time data collection with RFID may give us insight about the operational problems such as the optimal placing of products inside a store. Also, many grocery stores have data on users using loyalty cards. The longitudinal information about these users can be used to further improve the effectiveness of our system. For the next stage of this research, the proposed model will be constructed using real grocery data.

Acknowledgments. We are grateful to John Charnes, Sanjay Mishra, James Lemieux, and two anonymous ECSQARU-07 reviewers for comments and suggestions for improvements.

References

1. Adomavicius, G., Tuzhilin, A.: Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Trans. on Knowledge and Data Engineering* 17(6), 734–749 (2005)

2. Breese, J.S., Heckerman, D., Kadie, C.: Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In: Proceedings of Fourteenth conference on Uncertainty in Artificial Intelligence (1998)
3. Bruno, D., Paché, G.: From Traditional Retailing to E-tailing: The death and Rebirth of the Hypermarket Format? *Frontiers of E-Business Research* (2005)
4. Cavoukian, A.: Tag, You're It: Privacy Implications of Radio Frequency Identification (RFID) Technology. Information and Privacy Commissioner (2004)
5. Curtin, J., Kaufmann, R.J., Riggins, F.J.: Making the Most Out of RFID Technology: A Research Agenda for the Study of Adaptation, Usage and Impact of RFID. *Information Technology and Management* (2007)
6. de Campos, L.M., Fernandez-Luna, J.M., Huete, J.F.: A Bayesian Network approach to Hybrid Recommending Systems. In: Proc. 11th Int. Conf. on Information Processing and Management of Uncertainty in Knowledge-based Systems (IPMU-06), pp. 2158–2165 (2006)
7. Decker, C., Kubach, U., Beigl, M.: Revealing the Retail Black Box by Interaction Sensing. In: Proceedings of ICDCS 2003, Providence, Rhode Island, pp. 328–333 (2003)
8. Dipert, B.: Reading between the Lines. Business & Company Resource Center (2005)
9. Finkenzeller, K.: RFID Handbook Radio-Frequency Identification and Applications. John Wiley & Son, New York (1999)
10. Heckerman, D., Chickering, D.M., Meek, C., Rounthwaite, R., Kadie, C.: Dependency Networks for Inference, Collaborative Filtering, and Data Visualization. *Journal of Machine Learning Research* 1, 49–75 (2000)
11. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining Collaborative Filtering Recommendations. In: Proceedings of the ACM 2000 Conference on Computer Supported Cooperative Work, pp. 241–250 (2000)
12. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ACM Trans. on Information Systems* 22(1), 5–53 (2004)
13. Hopwood, J.: What is RFID? Intellident, <http://www.intellident.co.uk>
14. Inman, J., Winer, R.: Where the Rubber Meets the Road: A Model for In-store Consumer Decision Making. *Marketing Science Institute Report*, 98–122 (1998)
15. Lee, Y.M., Cheng, F., Leung, Y.T.: Exploring the Impact of RFID on Supply Chain Dynamics. In: Proceedings of the 2004 Winter Simulation Conference (2004)
16. Linden, G., Smith, B., York, J.: Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing*, 76–80 (2003)
17. Mild, A.: An Improved Collaborative Filtering Approach for Predicting Cross-category Purchases based on Binary Market Basket Data. *Adaptive Information Systems and Modelling in Economics and Management Science* (2002)
18. Mital, T.: The Emergence of RFID Technology. ISCR Future Technology Topic Brief (2003)
19. Mittal, B.: An Integrated Framework for Relating Diverse Consumer Characteristics to Supermarket Coupon Redemption. *Journal of Marketing Research* 31(11), 533–544 (1994)
20. Netflix (2007), <http://www.netflixprize.com/>
21. Nogués, M.T.M.: Combining Machine Learning and Rule-Based Approaches in Spanish Syntactic Generation. PhD Dissertation, Institut Universitari de Lingüística Aplicada Universitat Pompeu Fabra (2006)
22. Pazzani, M.J.: A Framework for Collaborative, Content-Based and Demographic Filtering. *Artificial Intelligence Review* (1999)

23. Popescul, A., Ungar, L.H., Pennock, D.M., Lawrence, S.: Probabilistic Models for Unified Collaborative and Content-Based Recommendation in Sparse-Data Environments. In: Proceedings of 17th Conf. on Uncertainty in AI (UAI-01) (2001)
24. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J.: Grouplens: An Open Architecture for Collaborative Filtering of Netnews. In: Proceedings of the ACM Conference on Computer Supported Cooperative Work, pp. 175–186 (1994)
25. Rutner, S., Waller, M.A., Mentzer, J.T.: A Practical Look at RFID. *Supply Chain Management Review* (2004)
26. Schneider, M.: Towards a Transparent Proactive User Interface for a Shopping Assistant. In: Workshop on Multi-user and Ubiquitous User Interfaces, pp. 31–35 (2004)
27. Sweeney II, P.J.: *RFID for Dummies*. Wiley, Chichester (2005)
28. Yu, K., Schwaighofer, A., Tresp, V., Xiaowei, X., Kriegel, H.P.: Probabilistic Memory-based Collaborative Filtering. *IEEE Transaction on Knowledge and Data Engineering* 15(1), 56–69 (2004)
29. Zhang, Y., Callan, J.: Maximum Likelihood Estimation for Filtering Thresholds. In: Proc 24th Ann. Int'l ACM SIGIR Conf., pp. 294–302. ACM Press, New York (2001)