Available online at www.sciencedirect.com

SCIENCE DIRECT®

Mathematical and Computer Modelling 44 (2006) 3–15

MATHEMATICAL AND COMPUTER MODELLING

www.elsevier.com/locate/mcm

A preference scoring technique for personalized advertisements on Internet storefronts☆

Jong Woo Kim[a,*], Kyung Mi Lee[b], Michael J. Shaw[c], Hsin-Lu Chang[d], Matthew Nelson[e], Robert F. Easley[f]

[a] School of Business, Hanyang University, Seoul, Republic of Korea
[b] Department of Statistics, Chungnam National University, Teajon, Republic of Korea
[c] Department of Business Administration, University of Illinois at Urbana-Champaign, Urbana, IL, USA
[d] Department of Management Information Systems, National Chengchi University, Taipei, Taiwan
[e] College of Business, Illinois State University, Normal, IL, USA
[f] Management Department, Mendoza College of Business, University of Notre Dame, IN, USA

Received 21 October 2004; accepted 9 December 2004

Abstract

This paper describes a new personalized advertisement selection technique based on a customer's preference scores for product categories. This method performs well, despite having low data and analysis requirements, relative to other methods in use. Customer preference scores are updated based on a customer's initial profile, purchase history, and behavior in an Internet storefront, and are then used to select and display appropriate advertisements on Internet web pages when the customer visits the Internet storefront. Compared with currently available recommendation techniques such as collaborative filtering or rule-based methods, preference scoring techniques use only a single customer's data to select appropriate advertisements and do not require a learning data set, and yet have competitive performance and can reflect changes in a customers' preference. An experiment is performed to compare two alternative data storage structures, the preference table and the preference tree, with random selection and collaborative filtering.

Keywords: Internet storefront; Customization; Personalization; Recommendation techniques; Electronic commerce

1. Introduction

The Internet is making revolutionary changes in the way businesses operate. The opportunities for success are huge, and so are the risks for failure. To be a successful Internet store, it is necessary to gain customer loyalty and to turn visitors into customers [1–5]. So, one-to-one marketing (also known as database marketing or relationship marketing) is considered to be a survival strategy for Internet stores, and it is a field positioned to benefit from the

☆ This work was partially supported by the Successful Post-doctoral Fellow program of Korea Science and Engineering Foundation (KOSEF).
* Corresponding address: School of Business, Hanyang University, 17 Haengdang-dong, Seongdong-gu, Seoul, 133-791, Republic of Korea. Tel.: +82 2 2220 1067; fax: +82 2 2220 1169.
E-mail address: kjw@hanyang.ac.kr (J.W. Kim).

0895-7177/$ - see front matter
doi:10.1016/j.mcm.2004.12.011

electronic revolution and shifts in consumer sales and advertising [6,7]. Personalization introduces a fundamentally new basis for competition in the marketplace by enabling organizations to differentiate on the basis of individual customers, not on products. One-to-one marketing shifts the focus away from geography-based market share toward earning the market share of an individual customer. Information technologies provide opportunities to implement one-to-one marketing strategies more effectively and efficiently. Various information technologies are used for one-to-one marketing, which include email, dynamic HTML, intelligent agents, data mining, collaborative filtering, expert systems/rule-based systems, push, web conferencing, and web site tracking/traffic analysis [8]. There are many success stories about the applications of information technologies for one-to-one marketing [9–11].

As a means of one-to-one marketing, personalized recommendation techniques have been applied to well-known commercial sites such as Amazon, Yahoo, DEC, ZDNet, US West's, The Baan Co., Kodak, eToy, Inc., DVD Express, and Kraft Foods [8,12–14]. Personalized recommendation techniques are used to determine product recommendation or to select advertisements for display to particular customers based on demographics and the past buying behavior of individual customers [1,2,15,16]. A common feature of established personalized recommendation systems is the need for analysis over a large set of customer data, or data for a large number of customers. The methods proposed here require only a small amount of data for a single customer, and involve very simple updating rules and selection algorithms, providing a great advantage in performance and minimizing privacy concerns, which are both key issues in this market.

The development and commercialization of recommendation techniques is a growing area in business-to-consumer electronic commerce. There are several recommendation techniques such as collaborative filtering [17,18], rule-based approaches [13,19], data mining approaches [15,20], content-based prediction [21], the Proscal ideal point model [16], and the Bayesian preference model [22]. Current available recommendation tools, however, are mainly based on collaborative filtering or rule-based approaches. BroadVision's One-to-One™ System is a commercial product to support one-to-one marketing in an electronic commerce environment. The system uses a rule-based matching technique to provide proper advertisements to customers and has been applied to Internet radio, Internet television, Internet banking, and so on [13,19]. A rule-based approach can be considered as a type of customer grouping technique in which the same advertisements are displayed to customers who satisfy certain conditions. Since, even in the same group, customers may differ in their preferences, it may not be effective to show the same advertisements to all the customers satisfying some conditions. The effectiveness of rule-based approaches mainly depends on the quality of the knowledge in the rule base. It is, however, difficult to obtain valuable marketing rules from marketing experts and to validate the effectiveness of the extracted rules. To resolve the knowledge acquisition problems, data mining techniques such as decision tree induction and association rule generation techniques have been applied [15,20].

In contrast, a collaborative filtering approach selects advertisements for customers based on the opinions of others who have shown preferences similar to the customer's in the past. There are several studies and commercial products based on collaborative filtering, such as GroupLens, Jester, LikeMinds™ Personalization Server, and Net Perceptions™'s E-Commerce™ [1,5,14,23]. In this technique, a serious limitation is that a new customer has to provide preferences for a large number of items in order to view personalized advertisements. Also, for any new advertisement or product, some preference data from people in the data set is required before collaborative filtering techniques can be applied. To overcome these limitations and provide more effective advertisements based on individual preference, we propose a personalized advertisement technique which relies solely on preference scores, stored as a preference table or tree, for each customer. Using the preference score concept, we attempt to achieve true personalization without relying on customer clustering, or using similar customers' opinions. That is, personalized recommendation is achieved using only the information of each customer without reference to a customer's demographic data or other personal information. Also, the technique does not require learning data sets or learning phases and, since the preference scores are changed by customers' behaviors on Internet stores, it is possible to reflect changes in customers' preference.

The rest of this paper is organized as follows. In Section 2, we propose an architecture for providing real-time personalized advertisements and information services. Section 3 describes models and algorithms to store preference scores and to select personalized advertisements. Section 4 describes the results of an experiment to evaluate the effectiveness of the suggested techniques. Section 5 presents a comparison with other approaches and discusses preference scoring function generalization issues. Section 6 presents brief concluding remarks, including ongoing research in this area.

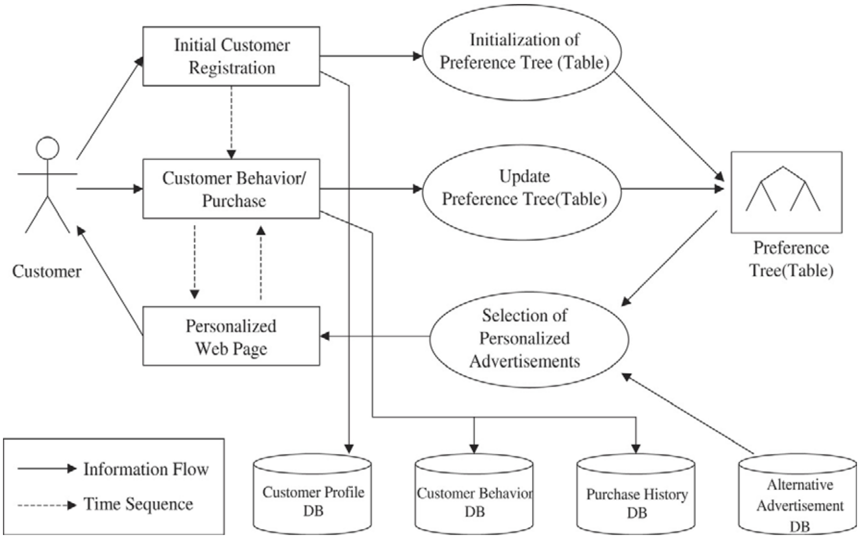

Fig. 1. Structure to provide real-time personalized advertisements.

2. Architecture for real-time personalized advertisements

To provide personalized advertisements, it is necessary to collect information regarding customers, including customer profiles (which are often written out when customers register on an Internet storefront), purchase history, and behavior data in the storefront (such as visits to a specific web page and/or clicking on advertisements). To select suitable advertisements on a real-time basis, however, due to time constraints, it is not appropriate to search all databases that contain customer-related information. To overcome this problem, we suggest a preference tree (or preference table) storage structure, which contains compressed information about customer preferences and can be accessed on a real-time basis. Though some customers may resist providing explicit category preference information to the web site on the first visit, we make the simplifying assumption that such data is available. However, this method can be adapted easily to other site architectures, for example to cases where the category preference data are not available until it can be deduced from customer browsing. This is a common limitation of recommendation techniques, and the practical result is simply that personalized recommendations cannot be made until the data becomes available.

Fig. 1 shows the system architecture used to provide real-time personalized advertisements. When a customer visits an Internet storefront, he/she first completes an initial customer profile, including basic customer information such as name, address, e-mail address, age, gender, occupation, and initial product category preferences, which is stored in a customer profile database. At the same time, a preference tree (or a preference table) is initialized that will contain personal preference information about product categories based on the initial product category preference. After registration, the customer may visit web pages in the Internet storefront, click advertisements, or purchase products. This customer behavior and purchase data are used to update records in the customer behavior and purchase history databases, respectively. At the same time, preference scores in the preference tree (or preference table) for the customer are adjusted to reflect actions taken in the Internet storefront. When it is necessary to select personalized advertisements, a set of appropriate advertisements is automatically selected from the alternative advertisement database based on preference scores in the preference tree (or preference table). Advertisements in the alternative advertisement database are managed, updated, and deleted by Internet store managers. When a new advertisement becomes available, Internet store managers add the advertisement in the alternative advertisement database with product category information.

3. Models and algorithms for providing personalized advertisements

In this section, we will describe two alternative storage structures for preference scores, preference tables and preference trees, and several algorithms for personalized advertisements selection. In order to illustrate the technique for providing personalized advertisements, we will use an imaginary Internet storefront that sells music CD albums

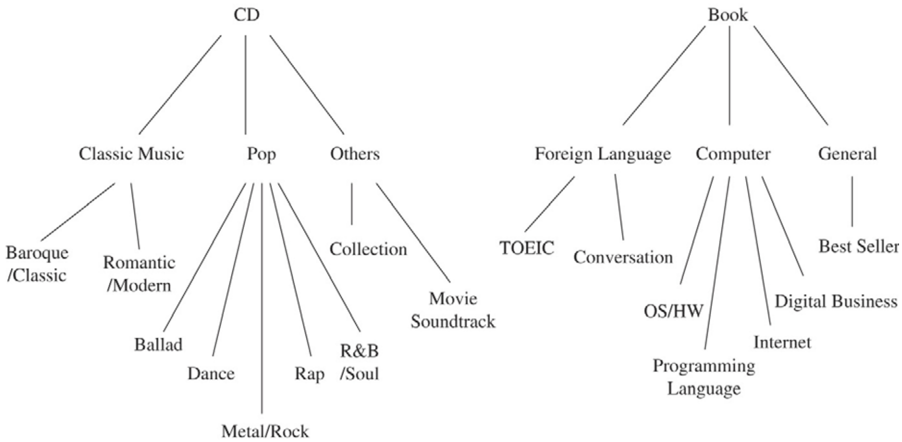

Fig. 2. Hierarchy tree of product category.

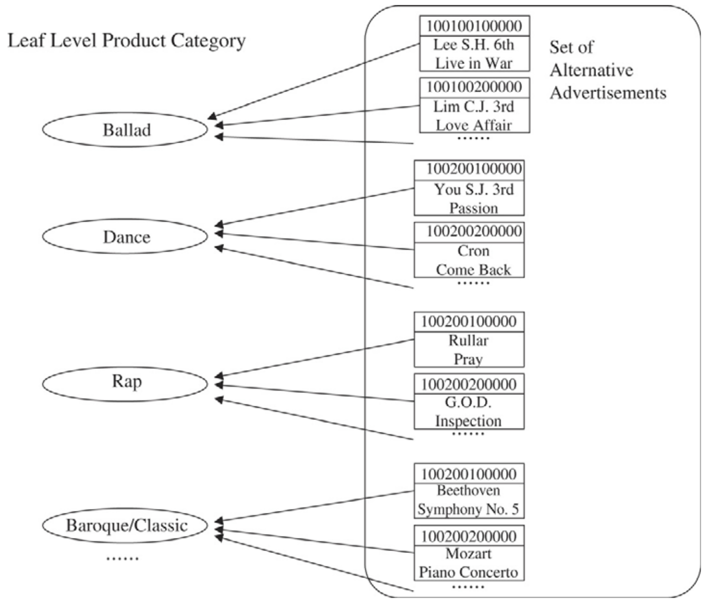

Fig. 3. Leaf nodes of hierarchy tree and product advertisements.

and books. To provide personalized advertisements, personal preference scores of a customer are managed for product categories. To do this, a hierarchy of product categories needs to be defined in advance by Internet store managers. For example, Fig. 2 shows a tree diagram depicting the hierarchy of product categories for music CD albums and books. The corresponding product advertisement alternatives (shown in Fig. 3) are linked to product categories that appear as leaf nodes in the hierarchy tree in Fig. 2.

3.1. Preference table approach

3.1.1. Basic definitions

The preference table-based approach stores customer preference scores for leaf-level product categories in table form. The preference table consists of three columns [*Customer ID* (*CID*), *Product Group ID* (*PGID*), *Preference Score* (*PS*)] which are defined in Definition 1.

Definition 1. The preference table (*PT*) is a relation such that

$$PT = \{\langle CID, PGID, PS\rangle\},$$

where *CID*, *PGID*, and *PS* are the customer identifier, the leaf-level product category identifier of the hierarchy tree, and the preference score, respectively.

Preference scores in the preference table are determined from the following three pieces of information:

1. *Initial profile* – When a customer registers in the Internet storefront, he/she is requested to choose product categories of interest. The choice of interesting product categories may be non-leaf-level product categories to simplify the registration process.
2. *Purchase* – If a customer has purchased many products that belong to a specific product category, it is reasonable to consider that he/she has a greater interest in different products within the same product category than in other product categories from which he/she has not purchased. So, the purchase history also needs to be reflected in the preference scores.
3. *Interest expression* – Behaviors that constitute interest expression at an Internet storefront include (1) visiting product description web pages, (2) printing the product description web pages, (3) clicking advertisements, and (4) putting products into their shopping cart. To make use of interest expression data, it is necessary for the Internet storefront to have monitoring functionality, and for the relative importance of different interest expressions to be assessed.

Definition 2. The preference scores in the preference table approach are defined as follows:

$$\begin{aligned} PS(i, j) &= \alpha_1 \times Profile(i, j) + \alpha_2 \times Purchase(i, j) + \alpha_3 \times Interest\_Type_1(i, j) \\ &\quad + \cdots + \alpha_{n+2} \times Interest\_Type_n(i, j) \end{aligned}$$

where $PS(i, j)$ is the preference score of customer i for the leaf-level product category j,

$Profile(i, j)$ is the profile score of customer i for the product category j, which is specified in customer i's initial profile. If customer i has specified product j as an interesting product category in his/her profile, then the values of $Profile(i, j)$ is 1, otherwise it is 0,

$Purchase(i, j)$ is the number of purchases of products that belong to product category j,

$Interest\_Type_k$ $(k = 1, \ldots, n)$ is the number of interest expressions of the kth type, and α_k $(k = 1, \ldots, n + 2)$ is the weight of each term.

In Definition 2, the weight (α_k) represents relative weights for the customer profile, the number of purchases, and the types of interest expressions. For example, we can consider that a purchase from a certain product category has more influence on preference scores than visiting a product description page in a product category. In this case the weight of a purchase (α_2) can be assigned a larger value than that of an interest expression.

3.1.2. Update of preference scores

When a customer first registers at an Internet storefront, the preference scores of the customer are based on the initial profile as follows.

For all leaf-level product categories j,

$$PS(i, j) = \begin{cases} \alpha_1 & \text{when customer } i \text{ specifies the product category } j \text{ as an interesting product} \\ & \text{category in his/her initial profile} \\ 0 & \text{otherwise.} \end{cases}$$

The preference scores are updated by purchases or interest expressions as follows.

(1) When the customer i purchases a product which is contained in the product category j,

$$PS(i, j) = PS(i, j) + \alpha_2.$$

(2) When the customer i executes an interest expression of type k in the product category j,

$$PS(i, j) = PS(i, j) + \alpha_{k+2}.$$

Table 1
Preference table, initial ($t = 0$) and after purchasing ($t = 1$) ($\alpha_1 = 2, \alpha_2 = 1$, *Customer ID* = C00001)

Product category (PGID)	Preference score (PS), $t = 0$	Preference score (PS), $t = 1$
Customer ID (CID)	C00001	C00001
Baroque/Classic	2	5
Romantic/Modern	2	2
Ballad	2	2
Dance	2	4
Metal/Rock	2	2
Rap	2	2
R&B/Soul	2	2
Collection	2	2
Movie soundtrack	2	2
TOEIC book	0	0
Conversation book	0	0
...	...	...

3.1.3. Advertisement selection

When a customer visits an Internet storefront, personalized advertisements are provided based on his/her preference scores. The advertisements that belong to product categories with high preference scores in the preference table are selected on a real-time basis. To select personalized advertisements, it is necessary to determine the number of advertisements that can be displayed at the same time and corresponding selection policies. For example, assume the number of advertisements that can be displayed concurrently is three, and we want to select two advertisements from the highest-scored product category and one from the next highest-scored product category. We also do not want to select advertisements of products that the customer has already purchased. In this case, the advertisement selection algorithm is as follows.

Selection Algorithm 1.

1. For the customer i, choose the product category j that has the highest preference score, and choose the product category k that has the next highest preference score.
2. Repeat until three advertisements are selected
   2.1 Choose two advertisements from product category j that have not been purchased by customer i.
   2.2 Choose one advertisement from product category k that has not been purchased by customer i.
   2.3 If less than three advertisements are selected, choose the two highest-scored product categories excepting product category j, and repeat steps 2.1 and 2.2.

3.1.4. A scenario using a preference table

With an example scenario, we illustrate the change of preference scores in the preference table and the selection of personalized advertisements. To simplify this example, our imaginary Internet storefront will only consider initial customer profiles and purchase history, ignoring interest expressions. That is, α_1 is 2, α_2 is 1, and $\alpha_3, \ldots, \alpha_{k+2}$ are 0, where α_1 is double α_2 so the initial customer profile is considered to be twice as important as purchases.

(1) Initial preference table

Assume that customer A visits the Internet storefront and chooses music CD albums as an interesting product category during the registration process. After registration, new records are added in the preference table for the new customer as shown in Table 1. The *Customer ID* is C00001, and preference scores of descendent product categories of music CD albums are assigned a 2, and those of other leaf-level product categories are assigned a 0.

(2) Customer A's second visit

When customer A visits the Internet storefront a second time, advertisements are selected based on the preference table. That is, two advertisements are selected from the highest product category, and one advertisement is selected from the next highest product category. Since the nine product categories have the same preference scores, a product category is randomly selected among them, and two advertisements from that product category are selected. A third advertisement from a different product category is also randomly selected among the remaining eight categories.

(3) Customer A makes a purchase

When customer A purchases three *Bach CDs*, and two *Michael Jackson CDs*, the preference table is changed as shown in the final column of Table 1. The preference score of *Classic* becomes 5, and that of *Dance* music becomes 4.

(4) Customer A's third visit

When customer A visits the Internet storefront a third time, advertisements are selected on the basis of the changed preference table. Two advertisements will be selected from the highest product category *Baroque/Classic* and one advertisement will be selected from the next highest product category *Dance*.

(5) Customer A makes another purchase

After customer A purchases an additional five classic CDs, the preference score of *Baroque/Classic* is changed to 10. When the customer visits the Internet storefront again, two advertisements of *Baroque/Classic* CDs and one advertisement of a *Dance* CD will be displayed.

3.2. Preference tree approach

3.2.1. Basic definitions

The preference table approach has a defect, in that it cannot reflect affinity levels among product categories. For example, if a customer has an interest in *Ballad* CD albums, it is reasonable to assume a further interest in *Dance* or *R&B* CD albums (which belong to the same *Pop Music* product category), rather than *Classic* albums. The preference tree approach is designed to reflect affinity among product categories.

Definition 3. The preference tree $PT'(i)$ of customer i is isomorphic to the hierarchy tree of product categories, and the set of nodes of the preference tree $PT'(i)$ is as follows:

$$N(i) = \{\langle PGID, PS\rangle\}$$

where *PGID* is a product category identifier, and *PS* is the preference score of the product category of customer i.

Definition 4. In the preference tree approach, preference scores of the leaf-level product category j and customer i are defined as in Definition 2 above. The preference scores of a non-leaf-level product category j are defined as follows:

$$PS(i, j) = Average_{k\in\{k|k \text{ is a child node of product category } j\}} PS(i, k).$$

The preference scores of leaf nodes in Definition 4 do not reflect affinity among product categories, however the adjustment of the parent category node score reflects the affinity between leaf nodes.

Definition 5. The adjusted preference scores of the leaf-level product category j and the customer i are defined as follows:

$$PS'(i, j) = \mathrm{Max}_{k\in\{k|k \text{ is } j \text{ or an ascendant node of the product category } j\}} PS(i, k).$$

3.2.2. Initializing and updating preference scores

The preference tree of a certain customer is initialized on the basis of his/her initial profile. The preference scores of the preference tree are initialized as follows.

(1) Product category nodes on the same level as the product category of greatest interest

$$PS(i, j) = \begin{cases} \alpha_1 & \text{when the customer } i \text{ specifies the product category } j \text{ as an interesting product} \\ & \text{category in his/her initial profile} \\ 0 & \text{otherwise.} \end{cases}$$

(2) Descendent nodes of the nodes in (1):

$$PS(i, k) = PS(i, j), \quad k \text{ is a descendent node of the node } j.$$

(3) Ascendant nodes of the nodes in (2):

$$PS(i, j) = Average_{k \in \{k|k \text{ is a child node of product category } j\}} PS(i, k).$$

(The *PS* values should be calculated sequentially from the bottom.)

The preference scores are updated by purchasing or interest expressions as follows.

(1) When customer i purchases a product in the leaf-level product category j,
① for the leaf node product category j,
$$PS(i, j) = PS(i, j) + \alpha_2$$
② for all ascendant nodes p of the node j,
$$PS(i, p) = Average_{c \in \{c|c \text{ is a child node of the product category } p\}} PS(i, c).$$
(The *PS* values should be calculated sequentially from the bottom.)

(2) When customer i performs a k type interest expression for the leaf-level product category j,
① for the leaf node product category j,
$$PS(i, j) = PS(i, j) + \alpha_{k+2}$$
② for all ascendant nodes p of the node j,
$$PS(i, p) = Average_{c \in \{c|c \text{ is a child node of product category } p\}} PS(i, c).$$
(The *PS* values should be calculated sequentially from bottom.)

The above update procedure does not require the update of all preference scores in the tree, rather it requires only the update of the preference scores of the related leaf node and its ascendant nodes. For example, if the preference score of *Baroque/Classic* is updated, then just two preference scores of its ascendant nodes (*Classic Music* and *CD*) need to be updated. No update of other nodes in the product category hierarchy is required. It thus satisfies requirements for handling customers' actions in Internet stores in real time.

3.2.3. Selection of advertisements

We defined the preference tree approach above using an adjusted preference score, in order to enable the consideration of affinity between categories. However, in our implementation, the preference tree stores only preference scores, not adjusted preference scores. That is because updating adjusted preference scores for every customer behavior may increase computational complexity. It is also not appropriate to calculate the adjusted preference scores of all nodes in the preference tree in order to select suitable advertisements. It is possible, however, to select advertisements efficiently using the fact that the preference score of a parent node is the average of preference scores of its child nodes.

In order to illustrate the selection algorithm, assume that the situation is the same as in the preference table case above. That is, the number of advertisements that can be displayed concurrently is three, and we want to select two advertisements from the highest product category and one from the next highest-scored product category. The advertisement selection algorithm is as follows.

Selection Algorithm 2.

1. For the customer i, choose the product category j which has the highest preference score, and choose the product category k which has the next highest preference score in the preference tree. (Note that j and k may or may not be leaf nodes.)
2. Repeat until three advertisements are selected:
   2.1. When product category j is a leaf node, choose two advertisements from product category j that have not been purchased by customer i. When the product category j is a non-leaf node, choose two advertisements from the descendent leaf product categories of the product category j that have not been purchased by customer i.
   2.2. When product category k is a leaf node, choose one advertisement from product category k that has not been purchased by customer i. When the product category k is a non-leaf node, choose one advertisement from the descendent leaf product categories of product category k that has not been purchased by customer i.
   2.3. If less than three advertisements are selected, choose the two highest-scored product categories except product category j, and repeat steps 2.1 and 2.2.

As noted above, Selection Algorithm 2 does not include adjusted preference scores. However, we can easily show that this algorithm will select product categories that are leaf nodes with the highest adjusted preference scores. This is because, in steps 1 and 2.3, we repeatedly choose product categories with the highest preference scores. If the selected product category is a leaf node, that node has the highest adjusted preference score by Definition 5. Alternatively, if the selected product category is a non-leaf node, its descendent nodes will have the highest adjusted preference scores, when compared with other nodes that are not its descendent nodes, again by Definition 5.

3.2.4. *Using preference trees for tie breaking*

Algorithm 2 results in selecting two leaf nodes that have the highest adjusted preference scores. So a leaf node that has a higher preference score cannot be selected due to its low adjusted preference score. It is also possible to use the adjusted preference scores for tie breaking in the preference table approach. That is, leaf nodes can be selected based on their preference scores at first, and when their preference scores are the same, their adjusted preference scores can be compared in order to select product categories. The following algorithm is an extension of Algorithm 1.

Selection Algorithm 1′.

1. For customer i, choose the product category j that has the highest preference score, and choose the product category k that has the next highest preference score. If two more product categories have the same highest preference scores, the adjusted preference scores of those product categories are compared to select product categories.
2. Repeat until three advertisements are selected:
   2.1. Choose two advertisements from product category j that have not been purchased by customer i.
   2.2. Choose one advertisement from product category k that has not been purchased by customer i.
   2.3. If less than three advertisements are selected, choose the two highest-scored product categories except product category j. If two more product categories have the same highest preference scores, the adjusted preference scores of those product categories are compared to select product categories. Repeat the steps 2.1 and 2.2.

3.2.5. *A scenario using a preference tree*

With the same scenario used above, we illustrate differences between the preference tree and preference table approaches in the selection of personalized advertisements. We assume that the preference score weights and other assumptions are the same as well.

(1) Initial preference tree

When customer A specifies music CD albums as an interesting product category during registration, a preference tree for the customer A is initialized. These values are shown in Table 2 ($t = 0$ column).

(2) Customer A's second visit

When customer A visits the Internet storefront a second time, advertisements are selected on the basis of the adjusted preference scores using the preference tree. Since all leaf nodes under the music CD node have the same preference score, by applying Selection Algorithm 2, two arbitrary leaf nodes under music CD are selected. Thus, two product advertisements are selected from the first leaf node and one advertisement is selected from the second.

(3) Customer A makes a purchase

When customer A purchases three *Bach* CDs and two *Michael Jackson* CDs, the preference tree is changed as shown in Table 2 ($t = 1$ column). The preference scores of *Baroque/Classic* and *Dance* nodes are updated by the purchasing behavior as follows:

$$PS(\text{Baroque/Classic}) = PS(\text{Baroque/Classic}) + 3 \times \alpha_2 = 2 + 3 = 5 \quad (\text{where } \alpha_2 = 1)$$

$$PS(\text{Dance}) = \text{PS}(\text{Dance}) + 2 \times \alpha_2 = 2 + 2 = 4 \quad (\text{where } \alpha_2 = 1).$$

Due to the update of preference scores of *Baroque/Classic* and *Dance* nodes, the preference scores of ascendant *Classic Music*, *Pop*, *CD* nodes are updated as follows:

$$PS(\text{Classic Music}) = \{\text{PS}(\text{Baroque/Classic}) + PS(\text{Romantic/Modern})\}/2 = (5 + 2)/2 = 3.5$$

$$\begin{aligned} PS(\text{Pop}) &= \{PS(\text{Ballad}) + PS(\text{Dance}) + PS(\text{Metal/Rock}) + PS(\text{Rap}) + PS(\text{R\&B/Soul})\}/5 \\ &= (2 + 4 + 2 + 2 + 2)/5 = 2.4 \end{aligned}$$

$$PS(\text{CD}) = \{\text{PS}(\text{Classic Music}) + PS(\text{Pop}) + PS(\text{Others})\}/3 = (3.5 + 2.4 + 2)/3 = 2.63.$$

Table 2
Preference tree: initial and after first and second purchases ($\alpha_1 = 2, \alpha_2 = 1$, *Customer ID* = C00001)

Product category (PGID)	Hierarchy level	Parent node	Preference score (PS), $t = 0$	Preference score (PS) $t = 1$	Preference score (PS) $t = 2$
CD	1	–	2	2.63	3.47
Book	1	–	0	0	0
Classic music	2	CD	2	3.5	6
Pop	2	CD	2	2.4	2.4
Others	2	CD	2	2	2
Foreign language book	2	Book	0	0	0
Computer book	2	Book	0	0	0
General book	2	Book	0	0	0
Baroque/Classic	3	Classic music	2	5	10
Romantic/Modern	3	Classic music	2	2	2
Ballad	3	Pop	2	2	2
Dance	3	Pop	2	4	4
Metal/Rock	3	Pop	2	2	2
Rap	3	Pop	2	2	2
R&B/Soul	3	Pop	2	2	2
Collection	3	Others	2	2	2
Movie soundtrack	3	Others	2	2	2
TOEIC book	3	Foreign language book	0	0	0
Conversation book	3	Foreign language book	0	0	0
…		…	…	…	…

(4) Customer A's third visit

When customer A visits the Internet storefront a third time, advertisements are selected on the basis of preference scores. Following Algorithm 2, two nodes that have the highest preference scores are selected, in this case *Baroque/Classic* and *Dance*, and three advertisements are selected from the two product categories.

(5) After additional purchase

After customer A purchases an additional five *Classic* CDs, the preference tree will be changed as shown in Table 2 ($t = 2$ column). When the customer visits the Internet storefront again, two advertisements are selected from *Baroque/Classic* CD advertisements and one advertisement is selected arbitrarily from the child node *Classic Music*. It is noteworthy that the selected product categories are different from those of the preference table approach, since the preference score from the *Classic* node influences its sibling nodes in the preference tree.

4. Experiment

4.1. Experimental design

We have performed a paper-based survey to evaluate the differences in effectiveness among the preference table approach, the preference tree approach, random selection, and collaborative filtering. Undergraduate university students participated in this survey. The students belonged to several departments of three universities, had similar education levels, and use the Internet frequently. The selection of students, as opposed to a random sample of customers, provides some benefits. By including students with similar educational levels and Internet experience, this sample allowed us to control significant confounding variables. However, the homogeneity of subjects may introduce some bias in the experimental results against collaborative filtering, since collaborative filtering algorithms may work better with heterogeneous sample sets. The survey yielded complete data sets for 158 subjects, after several were eliminated for incomplete data.

The survey form was designed in three parts: initial preference registration, purchase selection, and advertisement evaluation. After a brief introduction to the survey, the respondents were asked to select their preferred product categories for music CDs and books. The experiment uses the same hierarchy tree of product categories shown in Fig. 2, and weights among customer behaviors were assigned as in the scenarios in Section 3. Subsequently, the subjects selected five music CDs for purchase (though no purchase actually transpired) from 36 products: four each from nine music genres. This process was repeated for books, with five books selected for purchase from 28 products,

Table 3
Means and standard deviations ($N = 131$)

Class		Random	Collaborative filtering	Preference table (1)	Preference table (2)	Preference tree
Music CD	Mean	3.036	3.316	3.481	3.501	3.511
	Std. deviation	0.655	0.637	0.626	0.628	0.655
Book	Mean	3.224	3.336	3.466	3.496	3.466
	Std. deviation	0.529	0.519	0.529	0.595	0.614

Note: Preference table (1): Algorithm 1, Preference table (2): Algorithm 1′.

four each from seven categories. Finally, subjects were asked to evaluate 18 music CD advertisements: two from each of nine genres on the basis of their level of interest, scored on a five-point scale. This process was also repeated for books, with 14 advertisements: two from each of the seven low-level categories. To assist the subjects' purchasing and advertisement evaluation, a brief description of the products, for example in the case of CDs, title, artist, main song, price, etc., is included in the survey form along with scanned images of the CDs.

4.2. Results

To test the effectiveness of the proposed approach, three advertisements were selected for each respondent using the preference table approach (Algorithm 1 and Algorithm 1′), the preference tree approach (Algorithm 2), random selection, and collaborative filtering. Only the data required by each approach, obtained from the survey results, were used to select a set of advertisements for each respondent. Thus some of the advertisements could be selected repeatedly — there is nothing to prevent different methods from making similar recommendations for a particular respondent. This process was also repeated for books. Since collaborative filtering requires a rater data set, we randomly divide the data set into two portions: a rater set consisting of 27 records (about 20%) and a testing set consisting of 131 records. Also, collaborative filtering requires that a certain number of judgements should be available for a new subject before any advertising evaluation scores can be estimated. Given that each subject provides 18 advertising evaluation scores in music CD product categories (14 in book product categories), we randomly select five of those scores as input to the model, which then provides estimates for the other 13 (nine in book product categories). The collaborative filtering algorithm used in this experiment was a basic form that can be found in Resnick et al. [17].

The basic statistics of the test are shown in Table 3. The means in Table 3 are the averages of the respondents' evaluation scores for the three advertisements selected using each approach. The higher the average evaluation score, the better the response to the advertisements selected. In the case of the music CDs, the preference tree approach gave the best results. But in the case of books, the preference table approach (Algorithm 1′) gave the best results. T-test results are shown in Table 4. The means of the preference table approach and the preference tree approach were significantly different from random selection and collaborative filtering at a significance level of 5 per cent. Though the differences were not significant, the extended preference table algorithm (Algorithm 1′) gave better performance than the original preference table algorithm (Algorithm 1) for both of the product types. We conclude that the adjusted preference scores can be used to assist the original preference scores in the preference table approach.

5. Discussion

The main advantage of the preference scoring techniques presented here is in their limited data requirements. Most intelligent recommendation techniques require learning data sets and learning phases. Collaborative filtering requires learning data for every customer and for every new advertisement. Data mining approaches require massive amounts of learning data to extract valuable recommendation rules. The Proscal ideal point model and the Bayesian preference model also require learning data sets for parameter estimation. Preference scoring techniques also learn from an individual customer's behavior in Internet stores, but they do not use customer's demographic data for his/her personalized advertisement selection and do not require a previous separated learning phase. So, preference scoring techniques can be used in some situations where learning data is not available. The fact that the techniques perform well, despite the simple data requirements and updating procedures, makes them worth considering.

Currently available recommendation techniques are also slow to respond to changes in customer preferences. In collaborative filtering, a customer does not update prior judgments, but rates new items or, in this case, advertisements.

Table 4
Paired T-test results (P value)

Panel A: Music CDs

	Random	Collaborative filtering	Preference table (1)	Preference table (2)	Preference tree
Random	–	0.000**	0.000**	0.000**	0.000**
Collaborative filtering	–	–	0.014*	0.005*	0.004*
Preference table (1)	–	–	–	0.553	0.305
Preference table (2)	–	–	–	–	0.739
Preference tree	–	–	–	–	–

Panel B: Books

	Random	Collaborative filtering	Preference table (1)	Preference table (2)	Preference tree
Random	–	0.026*	0.000**	0.000**	0.000**
Collaborative filtering	–	–	0.016*	0.002*	0.018*
Preference table (1)	–	–	–	0.202	1.000
Preference table (2)	–	–	–	–	0.202
Preference tree	–	–	–	–	–

Note: $P^* < 0.05$, $P^{**} < 0.001$.

Correlations of these judgments with those of other customers will thus evolve slowly. Other methods perform learning in a batch mode, potentially limiting their real-time adaptability. An advantage of the proposed preference scoring methods is that they are easily updated in real-time, and any activity can affect preference ratings.

In this paper, the preference scoring function in Definition 2 is defined as a linear combination of *Profile*, *Purchase*, and $Interest\_Type_k$ functions. However, the preference scoring functions can be generalized as an arbitrary combination of functions, which are partially monotonically increasing for each of *Profile*, *Purchase*, and $Interest\_Type_k$. Future research will address finding an optimal combination of such functions. Also, the current preference scoring function structure is not weighted for recent behavior. The expansion of preference scoring function to reflect time weights and the modification of related algorithms are interesting further research issues. Another issue with preference scoring techniques involves the decision rules for advertisement selection. In this paper the decision rule in selection algorithms was to select the two highest-category and one next-highest-category advertisements. This decision rule is reasonable, but nonetheless arbitrary. So, it is desirable to develop rationales or guidelines for choosing the meta decision rules. The structure of the preference tree, in particular the number of leaf nodes per category, can affect results. The optimal design of the preference tree, as well as the optimal selection of weights used in analysis, are left for future research.

6. Conclusion

As an advertising medium, the Internet has several interesting characteristics, such as (1) interactivity, (2) rich and realistic experience, (3) aggregation of services, (4) global access, (5) targetability, (6) tracking, and (7) deliverability and flexibility [24]. Given the above characteristics, it is important to study how to maximize the effectiveness of advertisements. One-to-one marketing starts from the idea that each customer needs to be treated differently [6]. Due to the ease with which customer databases can be built, Internet storefronts have strong potential for implementing one-to-one marketing.

In this paper we propose a personalized advertisement technique based on preference scoring, as a particular implementation of one-to-one marketing. Preference scores for each customer are maintained using the customers' initial profile, purchase history, and interest expressions. Using these preference scores, personalized advertisements can be provided to customers on real-time basis. Two alternative storage structures for preference scores are proposed: preference tables and preference trees. While the preference table-based approach provides personalized advertisements in a simplified manner, it does not reflect affinity among product categories. The preference tree-based approach has the benefit of reflecting affinity among product categories, but it requires a more complex implementation and additional processing time. Both of the storage structures can reflect the changes of customers' preferences with time.

To test the effectiveness of the proposed techniques, an experiment was performed to compare the preference table-based approach, the preference tree-based approach, collaborative filtering, and random selection. The results of the experiment indicate that both of the proposed techniques provide better customer-oriented advertisement selection than pure random selection and collaborative filtering.

There are several ongoing research issues related to this study. The first category of issues relates to optimizing the preference scoring function structure in a hierarchy of product categories. The proposed techniques can also be synergistically used with other personalization techniques such as rule-based matching or collaborative filtering. So the further study of hybrid techniques is an open opportunity.

References

[1] J. Ben Schafer, Joseph A. Konstan, John Riedl, E-commerce recommendation applications, Journal of Data Mining and Knowledge Discovery 5 (1–2) (2001) 115–153.
[2] Asim Ansari, Skander Essegaier, Rajeev Kohli, Internet recommendation systems, Journal of Marketing Research 37 (August) (2000) 363–375.
[3] Richard C. Riecken, Growth in personalization and business, Communication of the ACM 43 (8) (2000) 32.
[4] Nigel Wells, Jeff Wolfers, Finace with a personalized touch, Communication of the ACM 43 (8) (2000) 31–34.
[5] LikeMinds, LikeMinds™ personalization server 3.0 technology backgrounder, White Paper, Andromedia, Inc., 1999.
[6] Don Peppers, Martha Rogers, The One to One Future: Building Relationships One Customer at a Time, Currency and Doubleday, New York, 1993.
[7] Don Peppers, Martha Rogers, Bob Dorf, The One to One Fieldbook: The Complete Toolkit for Implementing a 1to1 Marketing Program, Currency and Doubleday, New York, 1998.
[8] Cliff Allen, Deborah Kania, Yaeckel Beth, Internet World Guide to One-To-One Web Marketing, John Wiley & Sons, Inc., New York, 1998.
[9] Carol Pickering, They are watching you, Business 2.0 (February) (2000) 135–136.
[10] Jeff Sweat, Rick Whiting, Instant marketing, Informationweek August 2 (1999) 18–20.
[11] Scott Woolley, I got it cheaper than you, Forbes (November 2) (1998) 82–84.
[12] Udi Maner, Ash Patel, John Robison, Experience with personalization on Yahoo!, Communication of the ACM 34 (8) (2000) 35–39.
[13] BroadVision, http://www.broadvision.com, 2000.
[14] Net Perceptions, http://www.netperceptions.com, 2000.
[15] Jong Woo Kim, Byung Hun Lee, Michael J. Shaw, Hsin-Lu Change, Matthew Nelson, Application of decision-tree induction techniques to personalized advertisements on internet storefront, International Journal of Electronic Commerce 5 (3) (2001) spring 45–62.
[16] Robert F. Easley, Jong Woo Kim, Michael J. Shaw, Recommendation Systems Model Selection for Electronic-Commerce Personalization, Working Paper.
[17] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, John Riedl, Grouplens: An open architecture for collaborative filtering of netnews, in: Proceedings of the ACM 1994 Conference on Computer Supported Cooperative Work, ACM, New York, 1994, pp. 175–186.
[18] John S. Breese, David Heckerman, Carl Kadie, Empirical analysis of predictive algorithms for collaborative filtering, Technical Report, MSR-TR-98-12, Microsoft Research, October 1998.
[19] BroadVision, A BroadVision one-to-one White Paper, White Paper, BroadVision, 1996.
[20] Gediminas Adomavicius, Alexander Tuzhilin, Using data mining methods to build customer profiles, IEEE Computer (February) (2001) 74–82.
[21] Daniel Billsus, Michael J. Pazzani, A hybrid user model for news story classification, in: Proceedings of the Seventh International Conference on User Modeling, June 1999.
[22] Asim Ansari, Skander Essegaier, Rajeev Kohli, Internet recommendation systems, Journal of Marketing Research 37 (August) (2000) 363–375.
[23] Dhruv Gupta, Mark Digiovanni, Hiro Norita, Ken Goldberg, Jester 2.0: evaluation of a new linear time collaborative filtering algorithm, in: 22nd International ACM SIGIR Conference on Research and Development in Information Retrieval, August 1999.
[24] Chandrasekar Subramanian, Michael J. Shaw, David M. Gardner, Product marketing on the Internet, Working Paper, University of Illinois at Urbana-Champaign, 1998.