An Adaptive Personalized Recommendation Strategy Featuring Context Sensitive Content Adaptation

Zeina Chedrawy[1] and Syed Sibte Raza Abidi[1]

[1] Faculty of Computer Science, Dalhousie University, Halifax B3H 1W5, Canada
{chedrawy, sraza}@cs.dal.ca

Abstract. In this paper, we present a new approach that is a synergy of item-based Collaborative Filtering (CF) and Case Based Reasoning (CBR) for personalized recommendations. We present a two-phase strategy: in phase I, we developed a context-sensitive item-based CF method that leverages the original past recommendations of peers via ratings performed on various information items. In phase II, we further personalize the information items comprising multiple components using a CBR-based compositional adaptation technique to selectively collect the most relevant information components and combine them into one composite recommendation. In this way, our approach allows fine-grained information filtering by operating at the constituent elements of an information item as opposed to the entire information item. We show that our strategy improves the quality and relevancy of the recommendations in terms of its appropriateness to the user's needs and interests, and validated by statistical significance tests. We demonstrate the working of our strategy by recommending personalized music playlists.

1 Introduction

The volume of information over the Internet is increasing at a tremendous rate, and as a consequence the search for 'relevant' and 'useful' information is becoming proportionally difficult. Adaptive recommender systems—a class of adaptive hypermedia systems—act as mediators between information sources and information seekers [6], as they exploit the user's current specific interests and needs to (a) regulate the flow of information to users; and (b) direct users to the right information—i.e. personalized information selection and filtering [3]. Adaptive recommender systems are applied in a variety of application domains, including healthcare [15], business [14], education [16], entertainment [17], and so on.

Adaptive recommender systems use a variety of methods, spanning from adaptive hypermedia to information retrieval to machine learning to artificial intelligence, in order to determine the *relevance* and *utility* of any given information item with respect to the user's model/profile that characterizes the user's information needs, interests, preferences, demographics and consumption capacity. Functionally speaking, if an information item—which can be a document, news item, music compilation, movie, educational module, shopping list, activity plan, and so on—is deemed as relevant to the user then it is recommended to him.

In our approach to offering personalized recommendations, we extend the functionality of current adaptive recommender systems by pursuing (a) *context-sensitive information selection*; and (b) *compositional information personalization*.

Context-sensitive information selection involves the characterization and use of the context in which a recommendation is sought, processed and offered. We argue that in a collaborative filtering setting it is important to know why did a user recommend (or otherwise) a particular information item, as opposed to just tracking that the user has given a recommendation for the information item. For instance, in recommending a music CD, one user's recommendation can be due to his/her approval of the lyrics, tunes and vocals—these can be regarded as *recommendation perspectives*—whereas another user may recommend the same music CD along the recommendation perspectives of lyrics, directorship and performance. Although both users recommend the same music CD, yet the recommendation is due to different reasons or under different *contexts*. We argue that since collaborative filtering is guided by user's ratings of items it is useful to exploit the context of the rating of one user in making recommendations for other similar users. Context, for our purposes, implies a generalized set of relationships between a set of recommendation perspectives along which an information item can be rated [4]. Subsequent recommendation of the information item should be dictated by the similarity between the contexts of the recommender and the information seeker.

Compositional information personalization involves the fine-grained selection of individual components of an information item in response to a user model. An information item can be a music CD with its constituent components being the individual songs, or an E-learning module is an information item and the individual lessons are its constituent components. Compositional information personalization, therefore, involves firstly the selection of relevant information items and secondly the selection of the most salient individual components from the relevant information items [2]. A systematic compilation of these individual components, originating from different yet relevant sources, yields a fine-grained personalized information item that is much closer to the user-model as opposed to actual information items that may have some components that are not necessary of interest to the user.

We have devised a two-stage hybrid information recommendation strategy: Stage 1 uses item-based Collaborative Filtering (CF) to identify the information items that are relevant to the user-model; and Stage 2 uses compositional adaptation, in the realm of Case-Based Reasoning (CBR), to select the most salient information components (from a set of relevant information items) to dynamically compose a fine-grained personalized recommendation for the user. We extended the traditional item-based CF method [6] to establish the notion of context, whereby users are able to rate an item along multiple perspectives to elicit the basis of their recommendation. Our strategy is a hybrid of information retrieval viz. CF methods and artificial intelligence viz. CBR based compositional adaptation. It is important to note that the CF method in the first stage serves as the case retrieval mechanism vis-à-vis the CBR approach that we use to adapt the information items (or past cases) in stage 2 (Fig.1).

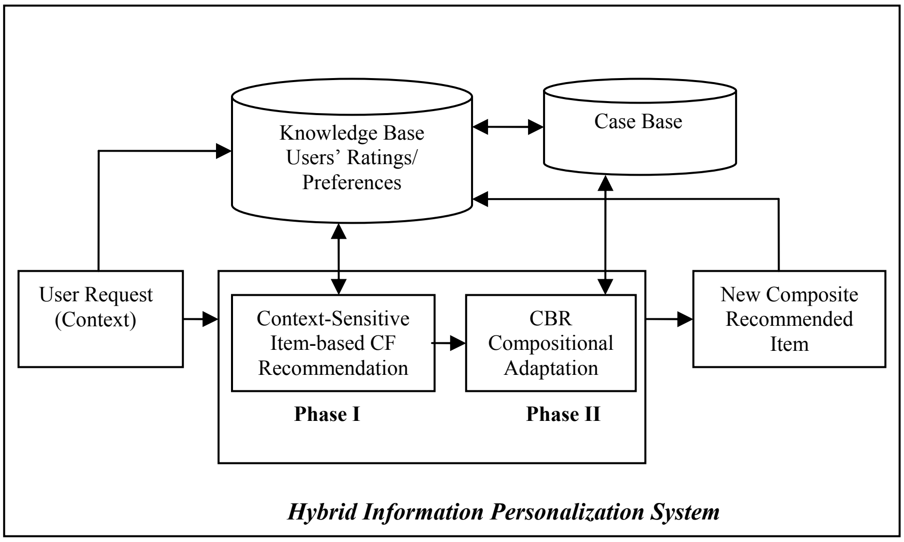

Fig. 1. PRECiSE Framework

3 Related Work

Recommender Systems (RS) are used to customize information content upon an individual's needs and interests to provide the most "*relevant*" and "*precise*" information. Since RS are grounded to solve real world problems, they have been adopted in both business and academic domains. One way to recommend objects is based on the explicit ratings of other people who have rated/purchased the objects in the past. Collaborative Filtering (CF) does this by first asking people to rate given items and recommending new items that have been highly rated by similar users. CF systems are the most widely used recommender systems [6]. Recently, integrating CBR in CF has gained a lot of attention and success [18] because the problem-solving principles of CBR make them an interesting candidate for integration with similarity based information filtering methods, such as CF methods [20]. As a matter of fact, hybrid recommendations were extended to contain knowledge-based techniques such as Case Based Reasoning for the purposes of improving the quality of recommendations and reducing the effect of the traditional CF cold-start problem. PTV [9] is a hybrid recommender system of which operates in the TV listings domain to recommend TV programs and produce personalized TV guides for active users. CoCoA [17] is a recommendation system for music compilation. It uses a case-based retrieval engine based on the CF users' ratings to propose an existing collection of sound tracks. Other hybrid systems were also developed such as Tapestry [21], GroupLens [5].

We present our adaptive recommendation system—***PRECiSE (Personalized Recommendations in a Context-Sensitive Environment***), that combines item-based CF and CBR techniques to generate fine-grained personalized information and demonstrate its working through an exemplar application for recommending a personalized music playlist.

4 Experimental Design

Our experimental dataset contains 100,000 ratings performed by 943 users on 1682 music compilations where each compilation is a collection of 10 songs. Each user in the dataset has rated a number of compilations along three pre-defined perspectives *song lyric, singer performance* and *song rhythm*. The dataset is divided into training (80%) and test sets (20%). The training set is used to identify the N (N=10) recommended items while the test set is used to measure the quality of the recommendation in terms of the F1-metric *(F_1)* and the appropriateness degree *(AD)*. The F1-metric combines recall *(R)* and precision *(P)* at similar weights as defined in [7]. Let *hits* be the total number of recommended items that were really rated by any of the users but were excluded from the training set to be part of the test set; when a set of recommended items is generated for a user, if the rated item in the test set exists in the recommended set, then a hit is recorded. Let t be the total number of users in the test set and N the total number of recommended items. Therefore,

$$R = \frac{hits}{t}; P = \frac{hits}{N}; F_1 = \frac{2 * R * P}{R + P} \quad (1)$$

The appropriateness degree *(AD)* determines the degree of similarity between a recommended item and a user profile. The calculation of the *(AD)* is detailed in Section 5.2.

Given a database of preferences that contains users' ratings on items along multiple perspectives. A user, who is seeking for advice, specifies the perspectives he is interested in and the relevance of each perspective is expressed in terms of weight values denoting the contribution of each perspective in the recommendation based on the user's interests. The system then recommends music compilations well-focused towards the user needs.

5 PRECiSE Methodology

5.1 Phase I: Item-Based CF Recommendation

The first phase of our strategy involves item-based CF recommendation to select a set of information items based on the ratings/recommendations of like-minded peers. Information filtering is pursued via collaboration based approaches, including CF, that filter information by making use of the opinions of peers with similar interests—peers critique the information items by rating them along a number of dimensions that reflect the quality and utility of the item.

We believe that the role of context in generating personalized recommendations is paramount as it determines both the relevance and usefulness of the recommended information. Note that the traditional CF approaches compare items along a single perspective—i.e. whether the item was liked/disliked by the user, without recourse to what perspectives of the item were rated and how were they rated. Original Item-based CF algorithms compare the user model of a user with the user models of other peers to recommend items that may be of interest to the user. The user model is expressed as a vector of attribute/value pairs defining items rated by the user, and the corresponding ratings' values with respect to all given perspectives. In our framework, we create a context for information filtering where a user can rate an item along multiple pre-defined perspectives. The rating on these perspectives is subsequently used to recommend the item to similar users. Hence, we extended the item-based CF method proposed by Sarwar et al. [6] so that, instead of having one similarity value for two items ***i*** and ***j***, a similarity vector of dimension ***P*** is generated (***P*** is the total number of perspectives available for a user to rate an item). The components of this vector are the individual similarities calculated based on the perspectives selected by a user. For instance, for a music playlist compilation problem the set of perspectives can be lyrics, tunes, band, vocals, direction, video, etc. Based on the user's context—i.e. a selection of perspectives—we compare the compilations that have been previously rated by other users and compute the degree of similarity between them.

Our context-based CF algorithm is described below:

1. The user chooses the relevant perspectives along which similarity between items is desired. Preferences of all users are induced through previous ratings along multiple perspectives on chosen items that are expressed as numerical scores. A separate rating matrix $M_p(u,i)$ is generated for every perspective, where $M_p(u,i)$ is the rating of user u on item i for perspective p.

2. We identify and isolate users that have rated both items and then we apply the similarity technique proposed by [6]. Let $PS_p(i,j)$ be the *Perspective Similarity* between two co-rated items i and j with respect to perspective p, and is calculated using (Eq.2).

$$PS_p(i,j) = \frac{\sum_{u \in U} (M_p(u,i) - \overline{M}_p(u))(M_p(u,j) - \overline{M}_p(u))}{\sqrt{\sum_{u \in U} (M_p(u,i) - \overline{M}_p(u))^2} \sqrt{\sum_{u \in U} (M_p(u,j) - \overline{M}_p(u))^2}} \quad \textbf{(2)}$$

where $\overline{M}_p(u)$ is the average rating of user u on all rated items. U is the set of all users in the CF knowledge base. Let W_p the weight assigned to perspective p; the Contextual Similarity $CS(i,j)$ between items i and j is then computed as:

$$CS(i,j) = \frac{\sum_{p=1}^{P} W_p * PS_p(i,j)}{\sum_{p=1}^{P} W_p} \quad \textbf{(3)}$$

When a user selects a specific context for similarity computation, the summation in (Eq.3) is achieved over the selected perspectives only.

3. Let the set I_u contains all items that have been rated by user u. For every item $i \in I_u$, we find the set of k most similar items (K_u). The set K_u excludes any item that has been rated/preferred by u and hence belong to the set I_u.

4. For every item $i \in K_u$, we compute its similarity $S\text{-}set(i,I_u)$ to the set I_u. This similarity is the sum of the similarities (calculated in *Step 2*) between all items rated by user u and item i.

$$S-set(i \in K_u, I_u) = \sum_{j \in I_u} CS(i,j) \qquad \textbf{(4)}$$

5. We sort the set K_u by the similarity *S-set(i,I_u)* in decreasing order and select the top ***N*** items. The selected ***N*** items/compilations would most likely interest the user.

5.2 Phase II: CBR-mediated Compositional Adaptation

Case adaptation in a CBR cycle allows for the adaptation of past cases (or solutions) in line with the description of the current problem [1, 11]. In our approach, we regard an information item (comprising a set of individual components) as a past case. The case selection process is achieved by the CF method in phase I of our strategy. Here, we apply a compositional case adaptation method to select the most pertinent constituent components of a composite information item [8]; provided the information components are mutually independent and compatibility between components is not an issue. The adapted case—i.e. a set of information components selected from the multiple information items—represents a fine-grained information item [12].

The basis for our adaptation strategy is defined by (a) the frequency of occurrence of a solution component in the similar cases and, (b) the degree of similarity between the user request and the retrieved case (measured in terms of the *appropriateness degree*). Our compositional adaptation of past cases is achieved as follows:

1. Given that a case is an item comprising multiple components. The similarity between a retrieved case C and a user u is calculated as shown in Eq.4. The similarity $S(u,C)$ between user u and every retrieved case C as $S(u,C)=S\text{-}set(i,I_u)$.

2. For every similar item case C_i, the Normalized Similarity (*NS)* of C_i to user u over the entire set of retrieved cases (*RC*) is calculated as follows:

$$Temp = \sum_{i=1}^{RC} 1/S(u,C_i) \qquad \textbf{(5)}$$

$$NS(u,C_i) = 1 - 1/(S(u,C_i) * Temp) \qquad \textbf{(6)}$$

3. Let $Comp_{C_i}$ be a component of the solution derived from a past case C_i, and $AD_u^{Comp_{C_i}}$ be the appropriateness degree for $Comp_{C_i}$, then

For every user u in the test set,

For i = 1 to *RC*

If $Comp_{C_i}$ exists in the solution of the similar case C_i

$$AD_u^{Comp_{C_i}} = AD_u^{Comp_{C_i}} + NS(u,C_i) \qquad \textbf{(7)}$$

In order to compute the Appropriateness Degree (*AD*) of each component, the normalized similarities of the similar cases that contain this component are added to one another.

4. We sort the distinct components of the N items by their *AD*, and select the M top components—i.e. the components that are most similar to the user. The value for M can be specified by the user.

5. The M selected components are amalgamated into one information item (or a new adapted case) that is recommended to the user.

Note that the application of the compositional adaptation method not only takes into account the global similarity between the present and past cases, but it is additionally driven by an attribute-level similarity between the current and retrieved cases.

6 Empirical Results

6.1 Performance of the Contextual Similarity in CF Recommendation

The quality of our item-based CF recommender system depends on how accurately it captures the user's preferences (ratings) as well as its ability to accurately match those preferences with similar users. Our experiments show that as we increase the number of perspectives, the F1-metric and the appropriateness degree increase and as a consequence a more personalized recommendation is provided. In fact, the appropriateness degree increases significantly by 52% when we apply 3 perspectives instead of only one perspective. Our conclusion is grounded in the observation that when the similarity is based on fewer perspectives, the similar item space which includes all items similar to the target item is large and contains items that are of potentially of little interest to the target user.

We performed statistical tests to study whether the difference between the means F1-metric for all three alternatives (1, 2, and 3 perspectives) is significant. In this regards, we used the General Linear Model approach [13] to perform the statistical test. The null hypothesis for the test is that all means are the same. Our test gave strong evidence against the null hypothesis (p-value, which denotes the probability that the null hypothesis is true, is much less than 0.001), which led us to conclude that the difference between the three F1-metric means is statistically significant.

6.2 Performance of the Compositional Adaptation Technique

The *AD* computed in (Eq.7) is used as a measure for the efficiency of our compositional adaptation technique. We calculate the *AD* of the newly adapted item which reflects the similarity degree between the adapted item and the individual user's preferences (represented by the set of items rated by the user). Basically, the *AD* of the final solution is the average sum of the *AD*s of the M (most relevant) components that constitute the final personalized information item.

In Fig.2, we note an increase in the appropriateness degree from phase I to phase II for selected users. An average percent increase of 61% was recorded for all users in the test dataset. The composite solution is more appropriate to the user's preferences than the items recommended through the CF phase; we conclude that our CBR-based compositional adaptation has provided the opportunity to pursue personalized information content more focused toward an individual user's interests.

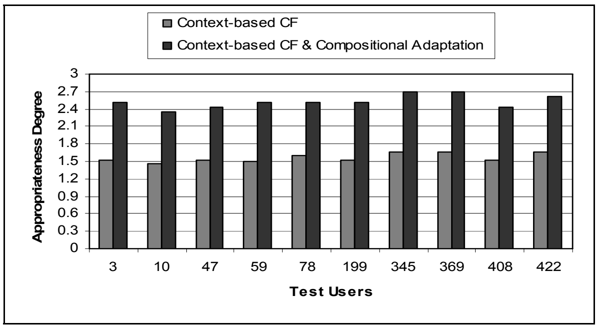

Fig. 2. Impact of the compositional adaptation on the appropriateness of the final solution

In order to show that the difference in the appropriateness degree between phase I and phase II is statistically significant, we apply the General Linear Model approach for testing whether the 2 mean values in appropriateness degree over all test users between phase I and phase II are different. The null hypothesis of the statistical test is that the 2 means are the same. Our test provided strong evidence against the null hypothesis (p-value < 0.001); we conclude that the 2 mean *AD*s are statistically different.

7 Working Example

We present a working example for recommending personalized music playlists. The user *u* has rated a set of CD compilations (see Fig. 3), where each compilation consists of 10 songs. The hit set contains those compilations that are recommended and were rated by *u* but were excluded from the training set for testing purposes. Note that users, music compilations (i.e. items) and songs (i.e. components) are represented by their IDs. The output of Phase I and Phase II is shown below.

User ID	Compilations (Information Items)				
130	1	17	42	96	159
	328	444	578	892	1095

Fig .3. Preferred music compilations for user ID '130'

- *Output from Phase I:* CF is applied with a context of 3 perspectives yields a hit set containing 6 compilations (see Fig.4). The appropriateness degree (AD) is averaged over the 10 recommendations. The F1-metric with a context (i.e. 0.34) is found to be better than without context (i.e. 0.76).

User ID	Compilations (Information Items)					F1 Metric	AD
130	47	92	295	331	332	**0.34**	**1.134**
	348	379	876	1012	1197		

Fig. 4. *N* Recommended music compilations (Phase I). Hit List is shaded grey. *N=10*

- *Output from Phase II:* We apply compositional adaptation to recommended compilations (in Fig.4). Figure 5, shows a sample of the recommended items together with their components. The components selected in phase II are shaded in figure 5, whereas figure 6 shows the final personalized recommended music playlist.

User ID	Item ID	Songs (Components)					AD
130	332	2	26	46	63	88	**1.801**
	348	22	35	104	125	187	
	379	2	20	41	110	128	
	876	20	25	125	196	198	
	1012	24	35	113	161	198	

Fig .5. A sample of the 10 recommended compilations and their constituent songs. Shaded songs are the most relevant to the user

User ID	Personalized Music Playlist				
130	**2**	**20**	**22**	**25**	**46**
	71	**104**	**125**	**187**	**198**

Fig. 6. New composite recommendation comprising the 10 most relevant songs for user ID '130'

We note an improvement in the quality of the final personalized recommendation in terms of the *F1-metric*, and the *appropriateness degree* that has increased significantly by 58.8%. The personalized music playlist, originating from different items, will be presented to the user as being most relevant to his/her interest.

8 Concluding Remarks and Future Work

In this paper, we have introduced a new personalized recommendation strategy featuring a unique hybrid of item-based Collaborative Filtering and Case Based Reasoning—i.e. a hybrid of information retrieval and artificial intelligence methods. We addressed personalization at a fine-grained level, whereby in the first stage the collaborative information filtering strategy initiates the process guided by peer based recommendations for pertinent information items. Next, in the second stage, the compositional adaptation method takes into account the degree of relevance of the retrieved information items and the weighted frequency of the recurring constituent information components in order to select the most appropriate information components. We introduced the notion of context—basically rating an item along distinct perspectives, where the more perspectives are used to rate the item the more focused and appropriate the retrieved information items would be to the user. This is an improvement from the single-dimensional binary rating scheme observed by CF systems. Our empirical results show that the usage of context as well as the compositional adaptation has provided more precise personalized recommendations in line with the user's needs. Our statistical tests also showed that the improvement of the recommendation was statistically significant.

For future work, we plan to explore quantifying the users' ratings based on the Multi-Attribute Utility Theory [19]. For instance, we evaluate initially the overall rating value on every rated item as a weighted addition of its ratings along the multiple perspectives, and finally we compute the similarity between rated items.

References

1. Aamodt, A., Plaza, E.: Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. AI Communication, IOS Press (1994)
2. Abidi, S.S.R.: Designing Adaptive Hypermedia for Internet Portals: A Personalization Strategy Featuring Case Based Reasoning with Compositional Adaptation. Lecture Notes in Artificial Intelligence 2527, Springer-Verlag, Berlin (2002)
3. Belkin, N. J., Croft, W.B.: Information Filtering and Information Retrieval; Two Sides of the Same Coin. Communications of the ACM, Vol. 35. (1992) 29-38
4. Dilley, R.: The problem of Context. Berghahn Books, New York (1999)
5. Resnick, P., N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl: GroupLens: An Open Architecture for Collaborative Filtering of Netnews, Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work: Chapel Hill, NC (1994) 175-186
6. Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Item-based Collaborative Filtering Recommendation Algorithms. Proceedings of the International WWW Conference (10). Hong Kong (2001)
7. Sarwar, B., Karypis, J., Konstan, J., Riedl, J.: Analysis of Recommendation Algorithms for E-Commerce. 2nd Conf. on Electronic Commerce (EC'00), New York (2000)
8. Wilke, W., Bergmann, R.: Techniques and Knowledge Used for Adaptation During Case-Based Problem Solving. Proceedings of the 11th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (1998)
9. Cotter, P. and B. Smyth: PTV: Intelligent Personalized TV Guides, Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence. AAAI Press, MIT Press (2000) 957-964
10. Geong, Y. Y.: An Analysis of Collaborative Filtering Systems. KMS Research Paper, School of Information, University of Texas, Austin (2003)
11. Schmidt, R., Vorobieva, O., Gierl, L.: Case-Based Adaptation Problems in Medicine. Proceedings of WM2003: Professionelles Wissensmanagement, Kollen-Verlag Bonn (2003)
12. Reyhani, N., Badie, K., Kharrat, M.: A New Approach to Compositional Adaptation Based on Optimizing the Global Distance Function and Its Application in an Intelligent Tutoring System. Proceedings of 2003 IEEE Intl. Conf. on Information Reuse and Integration, Las Vegas, USA (2003) 285-290
13. Daniel Wayne W.: Biostatistics, A Foundation for Analysis in the Health Sciences. Wiley (1987)
14. Kobsa A.: Customized Hypermedia Presentation Techniques for Improving Online Customer Relationships. Knowledge Engineering Review, Vol. 16(2). (1999) 111-155
15. Abidi, S.S.R., Chong, Y., Abidi, S.R.: Patient Empowerment Via 'Pushed' Delivery of Customized Healthcare Educational Content Over the Internet. 10th World Congress on Medical Informatics, London (2001)
16. Henze, N., Nejdl, W.: Extensible Adaptive Hypermedia Courseware: Integrating Different Courses and Web Material. In P. Brusilovsky, O Stock & C Strappavara (Eds.) Adaptive Hypermedia and Adaptive Web-based Systems, Springer-Verlag (2000) 109-120
17. Aguzzoli, S., Avesani, P., Masssa, P.: Compositional CBR via Collaborative Filtering. Proceedings of ICCBR'01 Workshop on CBR in Electronic Commerce, Vancouver, Canada (2001)
18. Burke, R.: A Case-Based Approach to Collaborative Filtering. Advances in Case-Based Reasoning, 5th European Workshop EWCBR 2000. Springer-Verlag, New York (2000)
19. Winterfeld, D. von, Edwards, W.: Decision Analysis and Behavioral Research. Cambridge, England, Cambridge University Press (1986)
20. Goker, M. H. and B. Smyth: Workshop on Case Based Reasoning and Personalization, 6th European Conference on Case Based Reasoning ECCBR, Aberdeen, Scotland (2002)