Improving Collaborative Filtering with Multimedia Indexing Techniques to create User-Adapting Web Sites

Arnd Kohrs – Bernard Merialdo
Institut EURECOM
Department of Multimedia Communications
BP 193 – 06904 Sophia-Antipolis – France
{kohrs,merialdo}@eurecom.fr

Abstract

The Internet is evolving from a static collection of hypertext, to a rich assortment of dynamic services and products targeted at millions of Internet users. For most sites it is a crucial matter to keep a close tie between the users and the site.

More and more Web sites build close relationships with their users by adapting to their needs and therefore providing a personal experience. One aspect of personalization is the recommendation and presentation of information and products so that users can access the site more efficiently. However, powerful filtering technology is required in order to identify relevant items for each user.

In this paper we describe how collaborative filtering and content-based filtering can be combined to provide better performance for filtering information. Filtering techniques of various nature are integrated in a weighed mix to achieve more robust results and to profit from automatic multimedia indexing technologies. The combined approach is evaluated in a prototype user-adapting Web site, the Active WebMuseum.

Keywords

Collaborative filtering, content-based filtering, Web museum, image indexing, user-adapting Web sites

1 Introduction

More and more information is available on the Internet through Web sites. Generally, these Web sites focus on providing information on services and products. Sometimes, they also allow to perform transactions. The site owner compiles a vast amount of information and organizes it in a structure of Web pages in a way that is most desirable for expected users. Users have been recognized as a very valuable asset to Web sites, e.g. as advertisement audience or potential customers. Therefore, Web sites try to tie their users to their services by providing more efficient (e.g., less time-consuming) access to preferred content. A single structure for all users fails to achieve broad satisfaction because users have different backgrounds and focuses. Therefore, Web sites which adapt to individual users are likely to be attractive and successful. Web sites which adapt their structure to the individual user are called *user-adapting Web sites*.

This trend of user-adapting Web sites, in contrary to the static collection of hypertext documents, necessitates new technologies and tools to adapt to users. One key technology is information filtering, which can be used identify relevant information for each user.

Information filtering techniques fall in two independent categories: content-based filtering and collaborative filtering. Content-based filtering is based on content analysis of the considered objects, e.g. term frequency for text documents, and its relation to the user's preferences. For content-based filtering it is therefore necessary that the results of content analysis and user preferences can reliably and automatically be determined. While recent research shows good results for the content-based filtering of text documents, filtering of other media, as audio and video, is hard due to the limitations of content analysis technology available. Collaborative filtering, on the other hand, does not show this limitation. In collaborative filtering, objects are selected for a particular user when they are also relevant to similar users and, in general, the content of the objects is ignored. Therefore, collaborative filtering is especially interesting for objects for which content analysis is difficult or impossible, for example non-textual. However, the performance of collaborative filtering relies on the amount of available opinions on the considered objects and it therefore performs poorly when few or no opinions are known.

In order to build better performing filtering systems both techniques can be combined. In recent research several approaches combining both techniques have been studied [3, 5, 10]. However, these approaches are limited to text documents. For objects such as images these approaches are not appropriate.

The purpose of our research is to explore the combination of content-based and collaborative filtering on media for which indexing techniques are limited. We have implemented our ideas in a prototype called the Active WebMuseum. The museum is a Web based service, that gives personalized tours through a virtual collection of art paintings. In the basic prototype, personalization was achieved through collaborative filtering only. The opinions on paintings collected through the Active WebMuseum are now used for further studies on combining content-based and collaborative filtering.

In this paper, we first describe our prototype of a user-adapting Web site, the Active WebMuseum, in more detail in section 2. Then, the content-based filtering, which is based on color and texture indexing techniques for images, is described in section 3. Collaborative filtering is explored in more detail in section 4. We then present our combined approach together with an evaluation in section 5. Finally, we conclude and describe future steps in section 6.

2 The Active WebMuseum

In an ideal world a visitor of a museum would enter a museum and then find in the first corridor exactly those items, which he would find most interesting. Given that real museums serve many people at the same time, it is not feasible to rearrange the collection for individual visitors. Often, real museums offer tours, which might be covering a particular topic or addressing a particular interest group, but having personal tours which show exactly the items of hight interest is impractical because items cannot be physically moved.

When a museum's art collection is presented through the Web, it becomes feasible to rearrange the collection for each individual visitor. Numerous museum sites already exist on the Web. They present images of arts contained in a hypertext structure, so that the navigation within a Web-based museum emulates strolling through the corridors of a real museum. Existing sites are static, which means that the hypertext structure linking the objects has been defined once for all, and is the same for all users, in the same way that the topology of buildings does not change. In contrast, our Active WebMuseum has a dynamic topology which is adapting to the museum visitor's taste and choices.[1]

2.1 The site's content model

In the Active WebMuseum project, we use filtering techniques to create a user-adapting Web site, in which the navigation structure is created for each specific user, based on predictions of what this user should prefer.

In order to allow dynamic restructuring of a museum site based on collaborative filtering, we introduced some simplifications:

- The site content is determined by a corpus of paintings. Each painting is contained in one Web page, which displays the painting and further gives additional information (title, painter and date).
- The pages are organized in categories, which can be accessed through virtual corridors.

By changing existing corridors and creating new ones, which contain references to pages showing paintings, it is possible to dynamically restructure the museum site in a way adapted to each user.

2.2 Access Paths

In this section we describe how the access to the corpus of paintings is implemented in the design of the user interface of the Active WebMuseum. In general we pursue three goals:

1. Provide typical access paths, which are present in most real museums, e.g. accessing corridors which contain arts by one chosen artist, or one epoch.
2. Provide dynamic corridors, which contain arts grouped by a chosen criteria, e.g paintings which are similar in color to a chosen reference painting.
3. Personalize the tour through the museum, e.g. showing arts, which are most relevant first to the visitor.

While the first two aspects are more or less already present in current Web-based art collections, in this paper we focus on the third point.

[1] The Active WebMuseum (accessed through `http://www.eurecom.fr/~kohrs/museum.html`) uses the collection of paintings from the *WebMuseum, Paris*(accessed through `http://metalab.unc.edu/wm/`), which has been created by Nicolas Pioch and contains roughly 1200 paintings by about 170 painters.

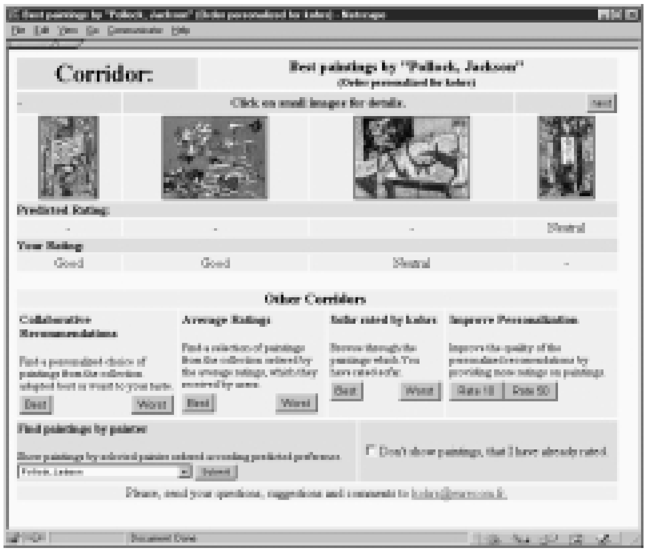

Figure 1: Browsing dynamic corridors: If the user has chosen a dynamic corridor (in the example a corridor containing paintings by Jackson Pollock), he is presented iconized paintings ordered according to his preference. From here the visitor may choose to get closer to a painting by clicking on it. Further the user may choose another access path to the collection: Paintings ordered according to personal preference, paintings ordered according to average rating by other visitors, paintings by other painters, etc. The user may also choose to further precise his user profile by providing more ratings.

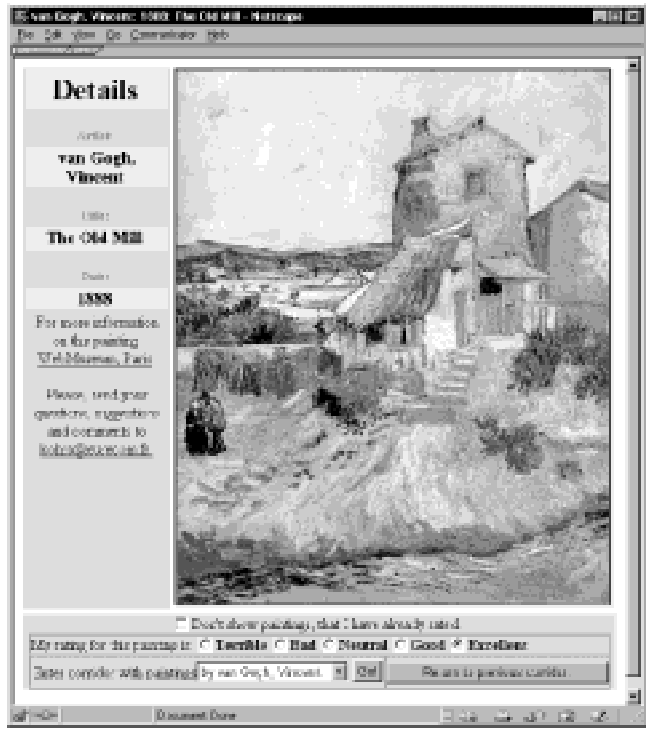

Figure 2: A single painting in detail close-up: When the user choses an iconized painting from a corridor it is presented in more detail (artist, title, creation date). If not already done, the visitor has to assign a rating to the painting expressing much he enjoys the painting. From the detailed view of a painting the visitor may return to the previous corridor, or he may choose to create a new corridor containing paintings similar to the current painting. As similarity criteria he may choose: same artist, same century, similar color or similar texture.

2.2.1 Typical Access Paths

In the Active WebMuseum we incorporate the commonly provided access path ¨by painter¨. The arts of one selected painter can be selected and are then presented in a virtual corridor as if visiting a real corridor containing all the arts of that particular painter. Further, it is possible to access the arts of one epoch (paintings which fall within one hundred years time intervals).

Other access criteria can be suggested, such as style, portraits etc., but we did not emphasize this issue since the focus of our work is on filtering techniques.

2.2.2 Dynamic Corridors

While the provisions of typical access paths provide the user with more or less the same service as real museums (arts grouped by painters and epochs and shown in corridors), it should be the purpose of a Web-based museum to provide additional functionality which is impossible in real museums.

In our case of the Active WebMuseum we provided some functionality which allows the user to dynamically create corridors containing arts which satisfy a selected criteria. The general idea is that the visitor, after viewing one painting, might want to view paintings which are similar. By choosing a similarity criteria (same painter, same century, similar color, similar texture) the user warps directly to corridors, which are dynamically created.

2.2.3 User-Adapting Corridors

As we described earlier, in the case of the Active WebMuseum, the main objects of change in order to adapt to the user are the corridors. The goal is to rearrange corridors so that museum visitors find faster paintings, which they like. The Active WebMuseum provides several dynamically created access paths to the collection of paintings:

- by color: paintings which are similar in color to the current painting.
- by texture: paintings which are similar in texture to the current painting.
- by ranking: paintings ordered according to average rank of all users.
- by recommendation: paintings ordered by personal ranking (as predicted by collaborative filtering).

Each of these access paths is mapped on a dynamic corridor, which allows the sequential browsing of appropriate images for that particular corridor.

Some corridors contain a large number of paintings so that it is very important to display the most significant paintings first. For some corridors the order is predetermined, e.g. the corridor containing paintings similar in color to a particular reference painting. On the other hand, for some corridors the order of presentation is not constrained by the paintings themselves, e.g. the corridor containing paintings by Vincent Van Gogh. Instead of using arbitrary orderings, e.g. alphabetical, the order of presentation is adapted in such cases to the visitor's preference (predicted by collaborative filtering if unknown).

When presented with preferred paintings first, the visitor can spend more time with paintings which he most enjoys.

2.3 Acquiring Preferences

In the previous section, we explained how preferences are used in order to transform corridors to personalized corridors. This approach is based on the existence of the visitor's preferences. In the Active WebMuseum, visitors can express preferences by giving symbolic ratings to paintings (*excellent*, *good*, *neutral*, *bad*, *terrible*). For historic reasons, the symbolic ratings are then mapped on numerical ratings in the interval [0..10]. For paintings which have not been rated by the visitor, the ratings are predicted using other users ratings and collaborative filtering technology.

Rating paintings should not be the prime occupation of a user. Therefore, the ratings can be conveniently provided while wandering within the museum. We noticed in initial trials, that users are hesitant in giving ratings, because giving a rating demands the inconvenience of having to make a decision. Therefore, we provided in the user-interface that

- the ratings can be provided with very little effort (one mouse click) without disrupting the users chosen tour, and
- if a painting is viewed in detail a rating is mandatory so that the visitor evaluates this painting or otherwise cannot continue his tour.

In the following sections we describe how ratings for paintings are predicted first by using content-based filtering and then collaborative filtering and finally using both in combination.

3 Content-Based Filtering

It is reasonable to expect that images with similar content will be almost equally interesting to users. The problem is that defining image content and image similarity is still an open problem. Ongoing research in Multimedia indexing is focusing on two directions:

- either each image is described by a textual caption, and captions are compared using techniques derived from document retrieval,

- or analysis and recognition techniques are applied to the image pixels to extract automatically features which are compared using some distance measure in the feature space.

We focus on the second approach, because it can be entirely automated. In our prototype, we have currently implemented two feature extraction components, derived from the work described in [13, 14]: color histograms and texture coefficients.

3.1 Color Histograms

The original paintings are available in RGB format, where each pixel is defined by the values (0-255) of the three components red, blue and green. We project these values in the HSV space (Hue, Saturation, Value) which models more accurately the human perception of colors. The HSV coefficients are quantized to yield 166 different colors. For each image, the histogram of these 166 colors is computed (proportion of pixels with a given quantized color).

To compare two images, we compute the L_1 distance (equation 1) between their color histograms:

$$\begin{aligned} h_i(j) &: \text{percentage of number of pixels} \\ & \quad \text{of painting } i \text{ with the color } j. \\ L_1(h_k, h_l) &= \sum_j |h_k(j) - h_l(j)| \\ d^{color}(p, p') &= L_1(h_p, h_{p'}) \qquad (1) \end{aligned}$$

3.2 Texture Coefficients

While color histograms do not take into account the arrangement of pixels, texture coefficients can be computed to characterize local properties of the image. We are using a wavelet decomposition using the Haar transform, by which a number of sub-images corresponding to a frequency decomposition are generated. These sub-images (see figure 3) are quantized to binary values, so that each pixel of the original image is associated with a binary vector of length 9. The histogram of these vectors (it has length 512) is the feature vector associated to the texture analysis of the image. Again, the L_1(see equation 1) distance is used to measure the texture distance between images.

3.3 Content-based Prediction

Using the individual ratings, that users assigned to paintings, and the previously described content-based distances between paintings, we measured a correlation between image distance and the difference of ratings, which the same user assigned to the paintings. The results for color histograms and texture coefficients are

(a) Original image

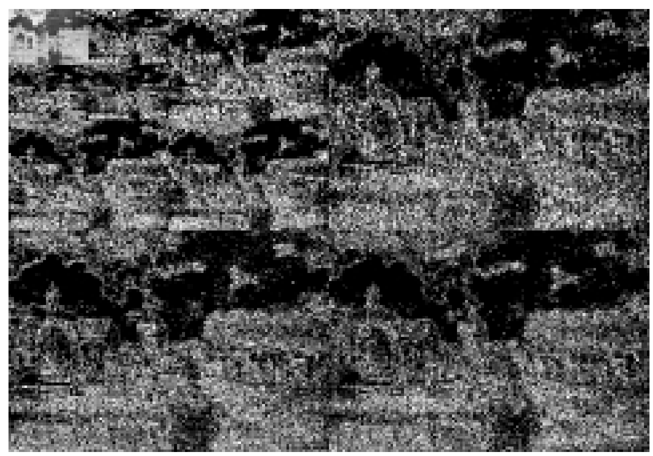

(b) Wavelet decomposition

Figure 3: In order to determine the similarity between paintings according to texture, all the painting in the database are decomposed into sub images using wavelet decomposition. From the decomposition a feature histogram is derived, which can then be compared by the use a vector metric.

plotted in figure 4. These measurements suggest, that paintings which are close in color or in texture receive in general similar ratings by the same users.

The primary goal for the Active WebMuseum is to select most relevant paintings. This is achieved through predicting the ratings, that a user would assign. Based on the findings of the previous measurements concerning the relationship between image distance and rating difference we derived a content-based prediction model. We use a linear estimator for the content-based prediction, which is illustrated in the following formula:

$$\begin{aligned} r_{u,i} &: \text{ user } u\text{'s rating for image } i \\ I_u &: \text{ Images, rated by user } u. \end{aligned}$$

Distance intervals:

$$\begin{aligned} j &= 1..n_\lambda \\ interval_1 &= [0,1), interval_2 = [1,1.5)... \end{aligned}$$

Distance classes:

$$C_j(i) = \{i' \in I_u : d^{color}(i,i') \in interval_j\}$$

Prediction for image i for user u:

$$p^{color}(u,i) = \sum_{j \in 1..n_\lambda} \lambda_j \cdot \frac{\sum_{i' \in C_j(i)} r_{u,i'}}{|C_j(i)|}$$

Expressed in words, the prediction works as follows: If a prediction is to be made for a user u and a target painting i, all the paintings previously rated by user u are grouped into distance classes ($C_j(i)$) according to color-based distance to target painting i. Each class is associated with a weight λ_j. The prediction is then the weighed sum of the mean ratings of each class. The weights λ_j are estimated through linear regression by using a priorly separated subset of the ratings.

For each content-based criteria, color histogram and texture coefficients, a predictor was created: p^{color} and $p^{texture}$. Later in this paper we describe how these predictors are used in combination with a collaborative filtering predictor in order to improve prediction results.

4 Collaborative Filtering

Collaborative filtering is a filtering technology which can be used for personalized recommendation. It is currently successfully applied to several content domains, e.g. books[2], movies[3], and even jokes.[4] Collaborative filtering gains more and more popularity in the e-commerce world, since it is an excellent technology to improve customer relations by personalizing offers and at the same

[2]Amazon (http://www.amazon.com/)
[3]Moviecritic (http://www.moviecritic.com/)
[4]Jester (http://shadow.ieor.berkeley.edu/humor/)

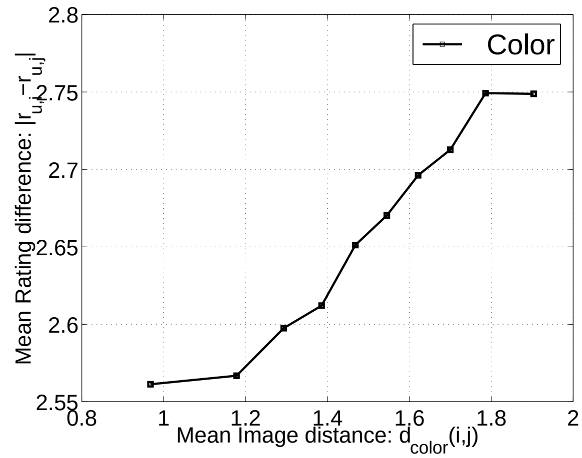

(a) Color-histogram distances

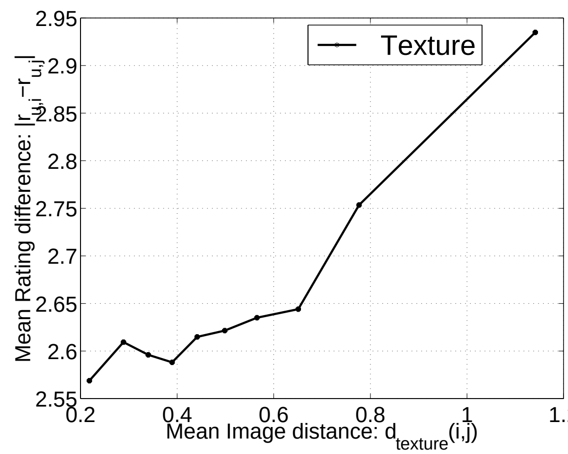

(b) Texture-coefficients distances

Figure 4: Correlation of distances between color-histograms(texture-coefficients) of paintings and rating differences: For each user and for each painting, that a user rated, all occurring distances between the paintings were collected together with the according absolute difference of ratings. The distances are then sorted and grouped. The mean distance of each group determines the values for the x-axis. For each group the mean absolute rating difference determines the y-coordinate.

time increasing sales by targeting products, information and advertisement in a personalized way. Collaborative filtering systems recommend objects for a target user based on the opinions of other users by considering how much the target user and the other users have agreed on other objects in the past. This allows this technique to be used on any type of objects and thus build a large variety of services, since collaborative filtering systems consider only human judgments on the value of objects. These judgments are usually expressed as numerical ratings, revealing the user's preferences for objects.

The importance of collaborative filtering is reflected in a growing number of research activities. One of the earliest projects is the GroupLens project [8, 9, 10], which focuses on filtering news articles from the Usenet, and recently also movie recommendation. Ringo [12, 11] was a collaborative filtering prototype for recommending music, leading to the spin-off company Firefly[5].

Most collaborative filtering systems collect the users opinions as ratings on a numerical scale, leading to a sparse matrix $rating(user, item)$ (in short $r_{u,i}$). Collaborative filtering systems then use this rating matrix in order to derive predictions. Several algorithms have been proposed on how to use the rating matrix to predict ratings [6, 12, 4].

In our Active WebMuseum we apply a commonly used algorithm, proposed in the GroupLens project and also applied in Ringo, which is based on vector correlation using the Pearson correlation coefficient. In the following we will describe the underlying formulas in more detail to make the general idea of automatically using other users as expert recommenders more understandable.

Usually, the task of a collaborative filtering system is to predict the rating of a particular target user u for an item i. The system compares user u's ratings with the ratings of all other users, who have rated the considered item i. Then a weighted average of the other users ratings is used as a prediction.

If I_u is set of items that a user u has rated then we can define the mean rating of user u as:

$$\overline{r_u} = \frac{1}{|I_u|} \sum_{i \in I_u} r_{u,i}$$

Collaborative filtering algorithms predict the ratings based on the ratings of similar users. When Pearson correlation is used, similarity is determined from the correlation of the rating vectors of target user u and the other users u':

$$\rho(u, u') = \frac{\sum_{i \in I_u \cap I_{u'}} (r_{u,i} - \overline{r_u})(r_{u',i} - \overline{r_{u'}})}{\sqrt{\left(\sum_{i \in I_u \cap I_{u'}} (r_{u,i} - \overline{r_u})^2\right)\left(\sum_{i \in I_u \cap I_{u'}} (r_{u',i} - \overline{r'_u})^2\right)}}$$

[5]Firefly (www.firefly.com) specializes in personalization and privacy on the Internet.

It can be noted that$\rho \in [-1, +1]$.

The value of ρ measures the similarity between the two users' rating vectors. A high value close to 1 signifies high similarity and a low value close to 0 signifies low correlation (not much can be deduced) and a value close to -1 signifies that users are often of opposite opinion.

The prediction formula (shown below) is based on the assumption that the prediction is a weighted average of other users' ratings. The weights refer to the amount of similarity between the user u and other users.

$$\begin{aligned} U_i &: \text{Users, who rated item } i. \\ p^{collab}(u, i) &= \overline{r_u} + k \sum_{u' \in U_i} \rho(u, u')(r_{u',i} - \overline{r_{u'}}) \\ \text{with } k &= \frac{1}{\sum_{u' \in U_i} \rho(u, u')} \end{aligned}$$

The factor k normalizes the weights.

5 Combining Collaborative and Content-based Filtering

In earlier work we discovered typical problematic cases for collaborative filtering systems [6], cases when not enough ratings are available, due to an insufficient amount of users or too few ratings per user. In the contrary, content-based schemes are less sensible to sparsity of ratings, since the performance for one user relies exclusively on his user-profile and not on the number of users in the system. However, comparative studies have shown, that collaborative filtering can outperform content-based filtering [2, 1]. Collaborative filtering should therefore be favored over content-based filtering.

In cases where collaborative filtering is limited by an insufficient amount of users and ratings, a combination of content-based and collaborative filtering should lead to better filtering performance. Besides the improvements of performance for cases of sparsity, a system which uses a combined approach can also recommend items which have not yet received any ratings e.g., new items, which is not possible for a system relying only on collaborative filtering.

In the following we present briefly recent research, which pursues the combination of content-based and collaborative filtering.

5.1 Existing Approaches

Fab[3] is an agent-based document filtering system. An agent society adapts through genetic algorithms and machine learning to topics and users. The agents are grouped into two different groups: The selection agents adapt to the preferences of a specific user and the collection agents adapt to topics. The collaborative aspect is achieved by the use of the same specialized collection agents for a group of related users and by forwarding highly rated documents to similar users. Similarity between users is determined from the keywords of preferred documents, by the use document retrieval techniques.

Sarwar [10] suggests filter-bots, specialized agents which detect features (spelling accuracy and message length) in news articles for the GroupLens Usenet filtering system. According to the feature detection the agents insert artificial ratings into the system.

The previously described projects combine collaborative and content-based filtering technology. However, a coherent method for combining collaborative filtering with content-based filtering has not been described yet. Further, both approaches were designed to operate on textual documents. For text documents the IR community has produced powerful analysis models, e.g. the Vector Space Model. For other media, such as images, it is not obvious how collaborative filtering can gain from existing content-analysis techniques.

In the following we briefly present our first approach of combining content-based and collaborative filtering. Our goal is to show that the performance of collaborative filtering is improved when combined with content-based filtering based on color and texture of paintings.

5.2 Linear Combination

For the following considerations we assume an existing collaborative filtering system, as the Active WebMuseum. The combination with content-based filtering is therefore rather an extension of collaborative filtering. As research leads to additional content-analysis tools, the extension approach should not limit the number of content-based extensions. Therefore we combine the content-based predictors with the collaborative predictor p^{collab}, as described in section 4, linearly using the following formula:

$$\begin{aligned} p^{comb}(u,i) &= \mu^{collab} \cdot p^{collab}(u,i) + \\ & \quad \mu^{color} \cdot p^{color}(u,i) + \\ & \quad \mu^{texture} \cdot p^{texture}(u,i) \\ \text{with } \sum \mu &= 1 \end{aligned}$$

The weights $\mu^{\{collab,color,texture\}}$ are estimated by the use of linear regression with a set-aside subset of the ratings.

5.3 Evaluation

After the Active WebMuseum has been online for several months, about 4000 ratings by 140 users were collected. Figure 5 depicts a histogram of the collected user ratings.

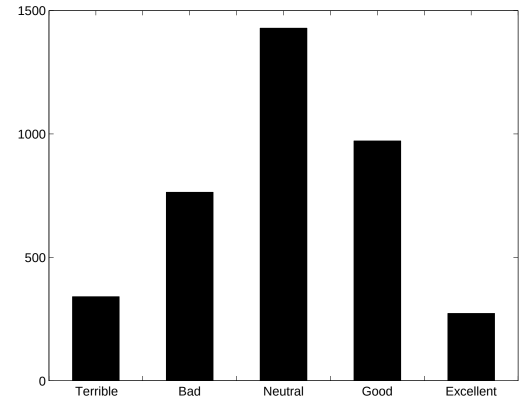

Figure 5: Distribution of the user ratings for the paintings in the Active WebMuseum.

In order to evaluate our approach, we measured the performance in terms of prediction precision for various combination configurations through off-line experiments on the collected data.

For the measurements 10 ratings for each user of a subset of 15 users were randomly separated in a test-set of ratings. Then the system was used to predict ratings in the test-set, using various prediction methods with the remaining ratings as a training set. The predictions are then compared to the original test-set to derive prediction errors.

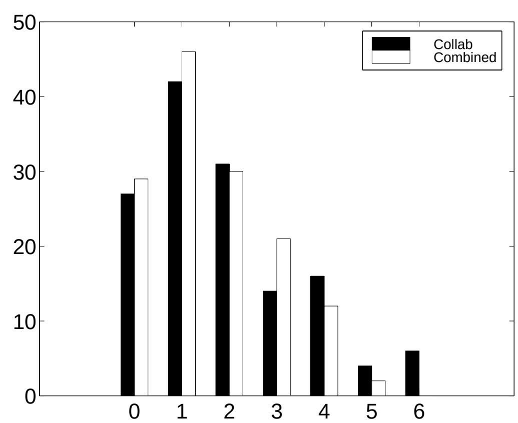

Figure 6: Histogram of the absolute prediction errors.

Figure 6 shows the histogram of the absolute prediction errors created by using only the collaborative predictor and using the combined predictor. It can be noted that while the collaborative predictor shows more frequent smaller errors, the combined predictor avoids large errors. However, it is hard to judge which one should be better. Also since the data-set is rather small the histograms of absolute errors the result changes depending on the selected test-set. In order to measure the performance of the prediction more robustly, the division into test-set and training set was repeated five times. After each run the prediction is evaluated using mean absolute error and correlation as distances between the test-set and the predicted set. These measurements were then averaged.

Table 1 lists the measured precision for the previously discussed predictors. Here, it is interesting to note the improvements of the combined approach compared to the pure collaborative approach. An improvement of mean absolute prediction error (MAE) in the combined prediction over the collaborative prediction can be identified. Further, an improvement of standard deviation of the absolute error (DEV) can be observed indicating, that the predictions are more robust using the combination, i.e. large prediction errors are likely avoided. The increase of the mean correlation (COR[6]) indicates that the overall ordering of the paintings in the test-set is more respected by the prediction when the combination is used instead of collaborative prediction by itself.

Prediction Method	**MAE**	**DEV**	**COR**
Collab	1.972	1.397	0.353
Combined	1.824	1.159	0.383

Table 1: Prediction precision of collaborative and combined predictors: The mean absolute prediction error(MAE), its standard deviation(DEV), and the correlation (COR) between the prediction and the test-set were measured. The measurements are averaged over 5 different test-sets and over all users.

Table 2 lists measurements of COR for variations of the combinations of predictors. The measurements show that the color predictor performs better than the texture predictor when used without combination. Further, the combination of the collab predictor with both the color and the texture predictor outperforms all the other combinations suggesting if another content-based predictor was added the combined predictor could be even more improved.

The observed improvements through the combination of collaborative filtering and content-based filtering, based on color and texture, suggest that the principle is valid. However, the presented combination model does not distinguish between individual users, i.e. for each user the same mix of predictors is used. For some users the content-based measures might be less appropriate as for others.

[6]For the correlation the same correlation formula as for the collaborative prediction (section 4) is used by replacing r_u with the ratings in the test-set and $r_{u'}$ with the predictions.

Prediction Method	**COR**
Texture	0.128
Color	0.178
Color & Texture	0.218
Collab & Texture	0.364
Collab & Color	0.377
Collab & Color & Texture	0.383

Table 2: Prediction precision with variation of the combination.

The goal of the prediction is to present users only relevant paintings (paintings, which they would rate highly), so that the users get the most satisfaction from visiting the site of the Active WebMuseum. Therefore, a good measure for the comparison of different prediction strategies should be focused on user satisfaction. We believe that the previously used measures which have been commonly used in the literature, are related to user satisfaction but do not focus on the goals of the users. More appropriate measures should be designed, for example another measure would monitor the ratings given by the users when they visit the site. If the system works well, then only positive ratings should be expected (except for new users). We have not yet investigated further in this direction.

6 Conclusion and Future Work

In this paper we identify filtering as a key technology for user-adapting Web sites, sites which allow users to access information more efficiently by adapting to their preferences. However, filtering is a hard problem, and cannot be addressed by one filtering technology alone. Due to limitations of both collaborative and content-based filtering, it is useful to combine these independent approaches to achieve better filtering results and therefore better user-adapting Web sites. We validate the combined use of collaborative and content-based filtering in several steps: First, we established a prototype user-adapting Web site, the Active WebMuseum, in order to apply our findings and to collect data about users. We demonstrated how filtering is used to personalize a Web museum. Second, we show for the example of art paintings how multi-media indexing technology can be used to derive content-based filtering. And third, we describe how various filtering technologies can be combined. This combination uses a weighed mix of available filtering techniques, so that the filtering result improves, even if weak techniques are included, as in our case, recommendation of paintings based on color and texture. The validity of the combination approach is supported by performance measurements.

In the future we plan to improve the combination method, as to adapt the mix of filtering techniques to the individual user. In [7] we describe how the approach by Sarwar [10] which we mentioned in section 5.1 can be extended for color and texture leading to more adaption to the individual users.

Also, in addition to texture and color, other automatic indexing techniques can be used to integrate more content-based schemes, e.g. fractal distances between images. Another important issue is a reliable measure of the systems performance. Monitoring the users' ratings during the use of an user-adapting Web site will lead to a measure, which comes closer to user satisfaction.

We believe that our work is useful for other applications which fall into the category of multi-media databases, for example online poster galleries or online music sales, given that personalization is an issue and automatic indexing technology is available.

Acknowledgments: This research was supported by a regional research grant, and by Eurecom's industrial members: Ascom, Cegetel, France Telecom, Hitachi, IBM France, Motorola, Swisscom, Texas Instruments, and Thomson CSF. The Active WebMuseum is based the collection of paintings of the *WebMuseum, Paris*, which has been created by Nicolas Pioch.

References

[1] Joshua Alspector, Alexander Kolcz, and Nachimuthu Karunanithi. Feature-based and clique-based user models for movie selection: A comparative study. *User Modeling and User Adapted Interaction*, (7):279–304, 1997.

[2] Joshua Alspector, Alexander Kolcz, and Nachimuthu Karunanithi. Comparing feature-based and clique-based user models for movie selection. In *DL '98. Proceedings of the third ACM Conference on Digital libraries*, pages 11–18, 1998.

[3] Marko Balabanovic and Yoav Shoham. Fab: Content-based, collaborative recommendation. *Communications of the ACM*, 40(3):66–72, March 1997.

[4] Jack Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Madison, WI, July 1998. Morgan Kaufmann Publisher.

[5] Joaquin Delgado, Naohiro Ishii, and Tomoki Ura. Content-based collaborative information filtering. In *Cooperative Information Agents II*, Lecture Notes in Artificial Intelligence 1435, pages 206–215. Springer Verlag Berlin Heidelberg New York, 1998.

[6] Arnd Kohrs and Bernard Merialdo. Clustering for collaborative filtering applications. In *Computational Intelligence for Modelling, Control & Automation*. IOS Press, 1999.

[7] Arnd Kohrs and Bernard Merialdo. Using color and texture indexing to improve collaborative filtering of art paintings. In *CBMI'99: European Workshop on Content-Based Multimedia Indexing*, 1999.

[8] Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordan, and John Riedl. Grouplens: Applying collaborative filtering to usenet news. *Communications of the ACM*, 40(3):77–87, March 1997.

[9] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of ACM CSCW'94 Conference on Computer-Supported Cooperative Work*, Sharing Information and Creating Meaning, pages 175–186, 1994.

[10] Badrul M. Sarwar, Joseph A. Konstan, Al Bochers, Jon Herlocker, Brad Miller, and John Riedl. Using filtering agents to improve prediction quality in the grouplens research collaborative filtering system. In *Proceedings of ACM CSCW'98 Conference on Computer-Supported Cooperative Work*, 1998.

[11] Uprendra Shardanand. Social information filtering for music recommendation. Master's thesis, MIT, 1994.

[12] Uprendra Shardanand. Social information filtering: Algorithms for automating "word of mouth". In *Proceedings of Human Factors in Computing Systems, CHI '95*, 1995.

[13] J. R. Smith and S.-F. Chang. Local color and texture extraction and spatial query. In *IEEE Proceedings International Conference on Image Processing*, September 1996. Lausane, Switzerland.

[14] J. R. Smith and S.-F. Chang. Tools and techniques for color image retrieval. In *Symposium on Electronic Imaging: Science and Technology - Storage and Retrieval for Image and Video Databases IV*, volume volume 2670, San Jose, CA, February 1996.