Dissimilarity Representation of Images for Relevance Feedback in Content-Based Image Retrieval

Giorgio Giacinto and Fabio Roli

Dept. of Electrical and Electronic Engineering - University of Cagliari
Piazza D'Armi 09123 Cagliari, Italy
e-mail {giacinto,roli}@diee.unica.it

Abstract. Relevance feedback mechanisms are adopted to refine image-based queries by asking users to mark the set of retrieved images as being relevant or not. In this paper, a relevance feedback technique based on the “dissimilarity representation” of images is proposed. Each image is represented by a vector whose components are the similarity values between the image itself and a “representation set” made up of the images retrieved so far. A relevance score is then assigned to each image according to its distances from the sets of relevant and non-relevant images. Three techniques to compute such relevance scores are described. Reported results on three image databases show that the proposed relevance feedback mechanism allows attaining large improvements in retrieval precision after each retrieval iteration. It also outperforms other techniques proposed in the literature.

1. Introduction

The vast majority of content based image retrieval (CBIR) techniques relies on the representation of images by low-level features, e.g., color, texture, shape, etc., [2,21]. Content-based queries are often expressed by visual examples in order to retrieve from the database all images that are “similar” to the examples. It is easy to see that the effectiveness of content-based image retrieval systems (CBIR) strongly depends on the choice of the set of visual features and on the choice of the “metric” used to model the user’s perception of image similarity. A number of metrics have been proposed in the literature to adequately measure (dis)similarities in a given feature space [19].

However, no matter how suitable for the task at hand the features and the similarity metric have been designed, the set of retrieved images often fits the user’s needs only partly. Typically, different users may categorise images according to different semantic criteria [1]. Thus, if we allow different users to mark the images retrieved with a given query as relevant or non-relevant, different subsets of images will be marked as relevant. Accordingly, the need for mechanisms to adapt the CBIR system response based on some feedback from the user is widely recognised.

This issue has been studied thoroughly in the text retrieval field, where the relevance feedback concept has been introduced [17]. Techniques developed for text retrieval should be suitably adapted to content based image retrieval, on account of differences in both feature number and meaning, and in similarity measures [11,15].

Relevance feedback techniques proposed in the literature involve the optimisation of one or more CBIR components, e.g., the formulation of a new query and/or the modification of the similarity metric to take into account the relevance of each feature to the user query.

Query reformulation is motivated by the observation that the image used to query the database may be placed in a region of the feature space that is "far" from the one containing images that are relevant to the user [6-7,12,15].

Other CBIR systems employ parametric similarity metrics whose parameters are computed from relevance feedback [18,20]. Rather than modifying the similarity metric, Frederix et al. proposed a transformation of the feature space so that relevant images represented in the new feature space exhibit higher similarity values [5]. A probabilistic feature relevance scheme has been proposed in [14].

Theoretical frameworks involving both the computation of a new query and the optimisation of the parameters of similarity metric have been also proposed [8,16].

In this paper, we propose a relevance feedback mechanism based on the representation of the images in the database in terms of their (dis)similarities from a *representation set*. The use of the dissimilarity representation of objects has been recently studied in the pattern recognition field [3-4]. This representation allows building pattern classifiers characterised by low error rates, as the pair-wise (Euclidean) distances between patterns of the same class are usually smaller than those of patterns of different classes.

Analogously, in a CBIR system, images represented by dissimilarities should exhibit the same property, i.e., relevant images should be characterised by smaller pair-wise distances than those between relevant and non-relevant images. In particular, the proposed technique is based on the representation of the images of the database in terms of their (dis)similarities from the set of images retrieved so far. A relevance score for each image of the dataset is then computed by taking into account the distances between the image itself and each of the relevant and non-relevant images retrieved. The relevant score is used to rank the images so that the first k are returned to the user in response of the relevance feedback.

In Section 2, the dissimilarity representation of images in the context of CBIR systems is proposed. The relevance feedback mechanism is described in Section 3, where three measures are described to rank the images according to their distances from the relevant and non-relevant images retrieved so far. Experiments with three image datasets are reported in Section 4. The reported results show that the proposed method outperforms other relevance feedback mechanisms recently described in the literature. Conclusions are drawn in Section 5.

2. Dissimilarity representation of images

Let us consider an image database whose images I are represented in a d-dimensional low-level feature space, e.g., color, texture, etc. Let us assume that a (dis)similarity metric $S(I_j,I_k)$ has been defined in such feature space. In the following, we will neither make any assumption about the feature space, nor about the similarity metric employed.

Let **Q** be the image used to query the image database. For each image **I** of the database, the value of $S(\mathbf{Q},\mathbf{I})$ is computed and the first k images, with the largest similarity values, are shown to the user (the value of the k parameter is chosen by the user). The user may either stop the search if satisfied with the results, or decide to refine the query by selecting the *relevant* images among the k image returned by the system. In this case the user provides the so-called *relevance feedback*. In order to exploit such feedback, we propose to represent each image **I** of the database in terms of its similarities with each of the k images returned by the system:

$$\mathbf{I} = (S(\mathbf{I},\mathbf{I}_1), S(\mathbf{I},\mathbf{I}_2),..., S(\mathbf{I},\mathbf{I}_k)) \quad (1)$$

where $\mathbf{I}_1$, $\mathbf{I}_2$,..., $\mathbf{I}_k$ are the k images more similar to query **Q** according to the similarity measure S. It is worth noting that also the images $\mathbf{I}_1$, $\mathbf{I}_2$,..., $\mathbf{I}_k$ are represented according to equation (1). This representation allows using the Euclidean distance measure to compute the dissimilarity between pairs of images [3-4]. The image **I** will be as much as relevant as it is *near* to the relevant images and, at the same time, *far* from the non-relevant ones. The techniques proposed to exploit this property of the dissimilarity representation will be illustrated in the following section.

It is worth noting that while the query mechanisms employed by some image databases are based on the combination of a number of similarity measures (usually related to different image representations), the following discussion is limited to a single feature space where only one similarity metric is defined. The proposed method can be easily generalised to address the case where more than one similarity measure (possibly related to different image representations) is used. However, this topic is out of the scope of the present paper and will be further discussed elsewhere.

3 Relevance Feedback

The goal of this section is to compute a "relevance" score for each image of the database by exploiting the dissimilarity representation of the images.

Let us denote with R the subset of indexes $j \in \{1,...,k\}$ related to relevant images, and NR the subset of indexes $j \in (1,...,k\}$ related to non-relevant images. Let r and nr be the cardinality of the sets R and NR, respectively.

For each image $\mathbf{I}_s$ of the database, let us compute the k distances d_{js} between the image itself and the k images $\mathbf{I}_j$, $j \in \{1,...,k\}$. Since the k images $\mathbf{I}_j$, $j \in \{1,...,k\}$, have been marked as being either relevant or non-relevant, the image $\mathbf{I}_s$ is as much as relevant as the values of d_{js}, for $j \in R$, are *small* and, at the same time, the values of d_{js}, for $j \in NR$, are *large*. In other words the k distances can be subdivided into two subsets, i.e., the subset made up of the r distances related to relevant images, and the subset made up of the nr distances related to non-relevant images.

A number of functions developed in the frameworks of the fuzzy set and clustering theories can be used to "aggregate" the r relevance distances and the nr non-relevance distances, so that the two values, one related to the degree of *relevance*, and the other related to the degree of *non-relevance*, can be associated to the image $\mathbf{I}_s$ [9-10]. In this work, we propose to use the “min” function, so that the *relevance* of each image $\mathbf{I}_s$ is computed in terms of the *nearest* relevant image and in terms of the *nearest* non-relevant image:

$$distR_s = \min_{j \in R} d_{js} \tag{2}$$

$$distNR_s = \min_{j \in NR} d_{js} \tag{3}$$

Let us linearly normalise the values of $distR_s$ and $distNR_s$ for all the images of the database, so that they take values in the range [0,1]:

$$dR_s = \frac{distR_s - \min\limits_s\left(\min\left(distR_s, distNR_s\right)\right)}{\max\limits_s\left(\max\left(distR_s, distNR_s\right)\right) - \min\limits_s\left(\min\left(distR_s, distNR_s\right)\right)} \tag{4}$$

$$dNR_s = \frac{distNR_s - \min\limits_s\left(\min\left(distR_s, distNR_s\right)\right)}{\max\limits_s\left(\max\left(distR_s, distNR_s\right)\right) - \min\limits_s\left(\min\left(distR_s, distNR_s\right)\right)} \tag{5}$$

Finally, let us denote with $\mu_R(s) = (1 - dR_s)$ the degree of *relevance* of $\mathbf{I}_s$, and with $\mu_{NR}(s) = (1 - dNR_s)$ the degree of *non-relevance* of $\mathbf{I}_s$. These values can be interpreted as the degrees of “membership” of the considered image to the relevant and non-relevant sets of images respectively. As a consequence, relevant images are those with large values of $\mu_R(s)$ and large values of $1 - \mu_{NR}(s) = \mu_{\overline{NR}}(s) = dNR_s$. In other words, $\mu_{\overline{NR}}(s)$ is a measure of how much the image $\mathbf{I}_s$ should *not* be considered as non-relevant, i.e., a measure of its relevance.

In order to assign a unique value of *relevance* to each image $\mathbf{I}_s$, the two membership values, namely $\mu_R(s)$ and $\mu_{\overline{NR}}(s)$, should be aggregated. It is easy to see that such aggregation function should be: i) monotonically increasing in $\mu_R(s)$, and ii) monotonically increasing in $\mu_{\overline{NR}}(s)$. To this end, in this paper, we consider two simple functions that take values in the range [0,1]:

$$average: \ \frac{1}{2}\left(\mu_R(s) + \mu_{\overline{NR}}(s)\right) \tag{6}$$

$$ratio: \ e^{-\frac{1-\mu_R(s)}{\mu_{\overline{NR}}(s)}} \tag{7}$$

The first function is commutative, so that the two values $\mu_R(s)$ and $\mu_{\overline{NR}}(s)$ play the same role in the final value of the degree of relevance. On the other hand, the second function is not commutative, as $\mu_{\overline{NR}}(s)$ is used to *weight* the degree of relevance measured by $\mu_R(s)$ (it is worth noting that the only purpose of the exponential is to have an aggregated value in the range [0,1]). In other words, the value of $\mu_{\overline{NR}}(s)$ is used to "support" the degree of relevance provided by $\mu_R(s)$, and is not considered a degree of relevance in itself.

Finally, the two aggregation functions, namely the *average* and the *ratio*, can be further aggregated to take advantage of their complementarity. As an example, the "probabilistic product" can be used [10]:

probabilistic product:

$$\frac{1}{2}\left(\mu_R(s)+\mu_{\overline{NR}}(s)\right)+e^{-\frac{1-\mu_R(s)}{\mu_{\overline{NR}}(s)}}-\frac{1}{2}\left(\mu_R(s)+\mu_{\overline{NR}}(s)\right)\cdot e^{-\frac{1-\mu_R(s)}{\mu_{\overline{NR}}(s)}} \quad \textbf{(8)}$$

which represents a "union" of the two functions (the "average" and the "ratio").

The relevance values computed either according to equation (6), or (7), or (8) are then used to rank the images and the first *k* are presented to the user.

To sum up, the proposed technique works as follows:

1. each image of the database is represented in terms of the distances between the image itself and the images retrieved so far;
2. the distances between each image of the database and the sets of relevant and non-relevant images are computed and a relevance score is obtained;
3. the images of the database are ranked according to this score and the first *k* are returned to the user.

4. Experimental Results

In order to test the proposed method and compare it with other methods described in the literature, three image databases have been used: the MIT database, a database contained in the UCI repository, and a subset of the Corel database. These databases are currently used for assessing and comparing relevance feedback techniques [13-16].

The MIT database (ftp://whitechapel.media.mit.edu/pub/VisTex) contains 40 texture images that have been manually classified into fifteen classes. Each of these images has been subdivided into sixteen non-overlapping images, obtaining a data set with 640 images. Sixteen Gabor filters were used to characterise these images, so that each image is represented by a 16-dimensional feature vector [14].

The database extracted from the UCI repository (ftp://ftp.ics.uci.edu/pub/machine-learning-databases/statlog/segment/) consists of 2,310 outdoor images. The images are subdivided into seven data classes. Nineteen colour and spatial features characterise each image. (Details are reported in the UCI web site).

The database extracted from the Corel collection is available at the KDD-UCI repository (http://kdd.ics.uci.edu/databases/CorelFeatures/CorelFeatures.data.html). We used a subset made up of 19513 images, manually subdivided into 43 classes. For each image, four sets of features were available at the web site. In this paper, we report the results related to the Color Moments (9 features), and the Co-occurrence Texture (16 features) feature sets [13].

For each dataset, the Euclidean distance metric was used. A linear normalisation procedure has been performed, so that each feature takes values in the range between 0 and 1.

For the MIT and UCI databases, each image is used as a query. In the case of the Corel dataset, 500 images have been randomly extracted and used as query. The top twenty nearest neighbours of each query are returned. Relevance feedback is performed by marking images belonging to the same class of the query as relevant, and all other images in the top twenty as non-relevant. This experimental set up affords an objective comparison among different methods and is currently used by many researchers [13-16].

For the sake of comparison, retrieval performances obtained with three methods recently described in the literature are also reported, namely the RFM (Relevance Feedback Method) [15], the Bayesian Query Shifting [7], and the PFRL (Probabilistic Feature Relevance Learning) [14].

RFM exploits relevance feedback by computing a new query $\mathbf{Q}_1$ according to the Rocchio formula [17]. The new query is computed as a linear combination of the original query, and the mean vectors of relevant and non-relevant images retrieved so far. The coefficients of the linear combination are usually chosen by heuristics. This query shifting formulation requires that the similarity between image vectors be measured by the cosine metric. To this end, data has been pre-processed by the so-called $tf \times idf$ normalisation, used in the information retrieval domain [15], that converts image feature vectors into weight vectors.

The Bayesian Query Shifting technique (Bayes QS) computes a new query according to the following formula derived from the Bayes decision theory:

$$\mathbf{Q}_1 = \mathbf{m}_R + \frac{s}{\|\mathbf{m}_R - \mathbf{m}_N\|}\left(1 - \frac{k_R - k_N}{\max(k_R, k_N)}\right)(\mathbf{m}_R - \mathbf{m}_N) \qquad (9)$$

where $\mathbf{m}_R$ and $\mathbf{m}_N$ are the mean vectors of relevant and non-relevant images respectively, σ is the standard deviation of the images belonging to the neighbourhood of the original query, and k_R and k_N are the number of relevant and non relevant images, respectively. More details on this method can be found in [7].

PFRL is a probabilistic feature relevance feedback method aimed at weighting each feature according to the information extracted from the relevant images. This method uses a weighted Euclidean metric to measure the similarity between images. The weights are computed according to the feature values of the relevant images retrieved. Two parameters chosen by experiments are used to optimise the performances. More details on this method can be found in [14].

As the performances of the RFM and PFRL techniques depend on the choice of the values of some parameters, the results reported hereafter are related to the best ones obtained in a number of experiments.

Experiments with the MIT data set

Figure 1 shows the average percentage retrieval precision after 10 retrievals attained by the three Dissimilarity Relevance Feedback (DRF) measures proposed in the previous section, and the three methods used for comparison purposes. After one iteration of relevance feedback (1 rf), the best performances are attained by the Bayes QS technique (91.80%), while the performances of the DRF techniques are between 89.31% (ratio) and 90.72% (probabilistic product). RFM and PFRL do not perform well, with an average precision of 84.55% and 85.47%, respectively. After two relevance feedback iterations, DRF provided the highest precision and clearly outperforms the other techniques starting from the third iteration. In addition, the proposed DRF allows further improvements, reaching the retrieval precision of 97.01% after 9 iterations with the "ratio" function. The other three techniques used for comparison does not allow to further improve the retrieval performances attained after four iterations if a larger number of iterations is performed.

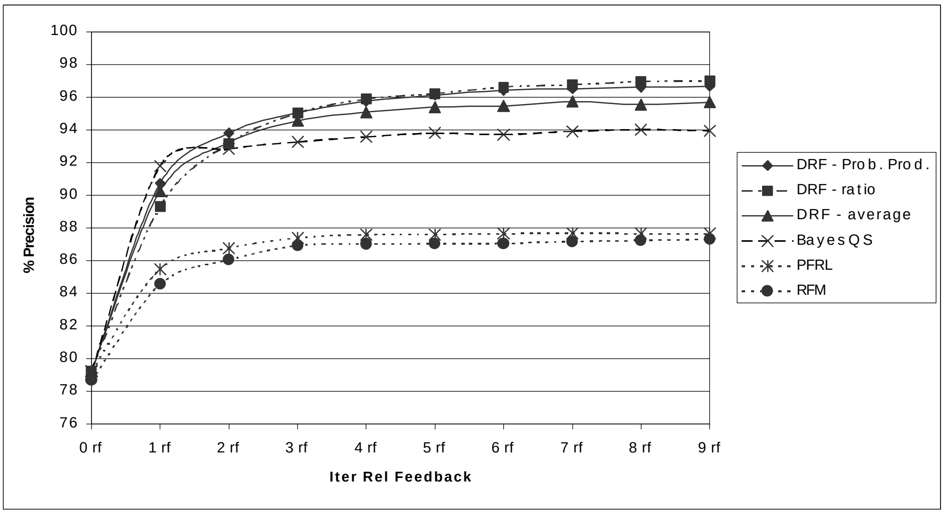

Fig. 1. Average percentage precision retrieval for the MIT data set. Nine relevance feedback iterations were performed with the considered relevance feedback techniques.

The comparison of the three DRF functions shows that the "average" measure provides lower performances than those of the "ratio" and the " probabilistic product". Thus it follows that, at least for this data set, the rationale behind the "ratio" measure seems more appropriate than the one behind the "average" measure. It is worth recalling that the non-relevant membership value is used to weight the relevant membership value in the "ratio" measure, while it is used just as a relevance membership value in the "average" measure.

Experiments with the UCI data set

For this data set, the choices of both the similarity metric and the normalisation procedure affect the precision of the first retrieval (0 rf). The results reported in Figure 2 show that the $tf \times idf$ normalisation is less suited than the linear normalisation, and the cosine metric (RFM) provides a lower precision (84.03%) than the Euclidean metric (90.21%) employed by all other techniques. After the first relevance feedback iteration, the best performance is attained by the Bayes QS (96.33%). The DRF with the "probabilistic product" function (96.23%) and the DRF with the "average" function (96.04%) provided similar performances, while the other techniques provided lower performances.

All the considered techniques allow improving the retrieval precision if the number of relevance feedback iteration is increased. The proposed DRF techniques attained the highest performances at each iteration. In particular, the highest values have been attained by the DRF with the "ratio" function and the DRF with the "probabilistic product" aggregation function, thus confirming the comments on the results related to the MIT data set.

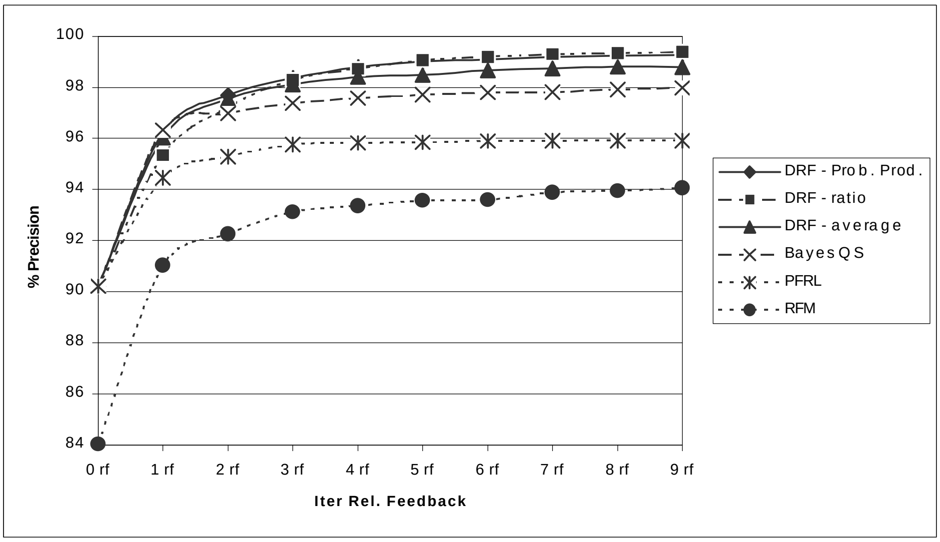

Fig. 2. Average percentage precision retrieval for UCI data set. Nine relevance feedback iterations were performed with the considered relevance feedback techniques.

Summing up, as far as the MIT and UCI data sets are concerned, the precision of the DRF after the first iteration is very close to the highest one of the Bayes QS technique. DRF clearly outperformed other techniques after two or more iterations, thus showing its validity, and the ability to exploit additional information provided by further feedbacks.

Experiments with the Corel data set

This data set allows a more thorough comparison among methods, as it is made up of a larger number of images. Figures 3 and 4 show the results with the color moments and co-occurrence texture feature sets, respectively. The retrieval performances related to the original query (0 rf) are quite low, thus showing that the chosen feature sets are not suited for the task at hand. In addition, different similarity metrics provided quite different results. In particular, the cosine metric provided better results than those of the Euclidean metric with both feature sets. The cosine metric attained a precision of 20.36% with the "color moments" feature set, and 17.89% with the "co-occurrence texture" feature set, while the Eulidean metric attained precisions of 17.30% and 14.13%, respectively.

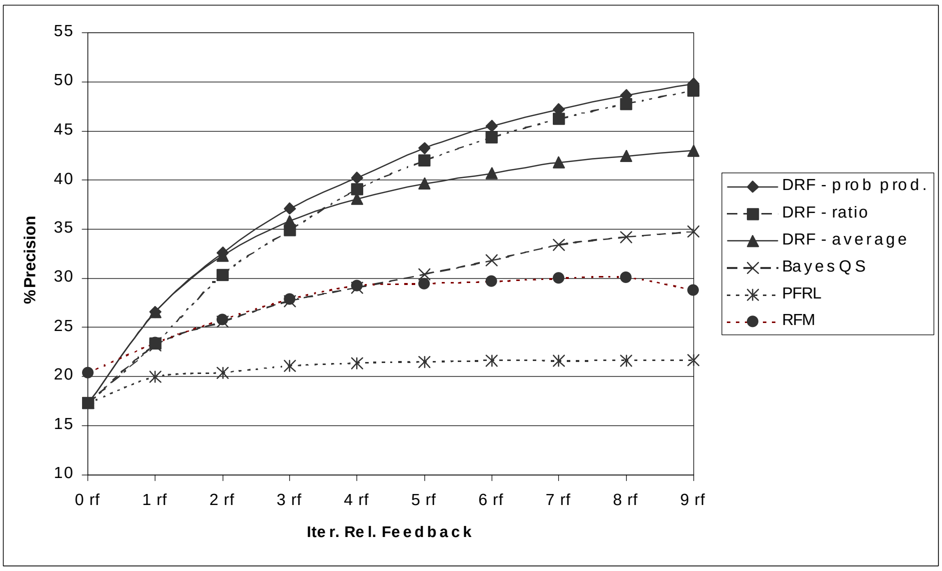

Fig. 3. Corel Data set – Color Moments feature set. Average percentage precision retrieval for Corel Data set – Color Moments feature set. Nine relevance feedback iterations were performed with the considered relevance feedback techniques.

Each of the considered relevance feedback techniques allows to improve these results. However, the DRF outperformed the other methods starting from the first iteration. The results related to the "color moments" feature set (Figure 3) show that the DRF with the "average" function and DRF with the "probabilistic product" provide higher performances than those provided by the RFM, the DRF with the "ratio" function, the Bayes QS, and the PFRL. All the techniques but the PFRL allows increasing the retrieval precision when further iterations are performed. As an example, after three iterations the DRF with the "probabilistic product" function attains a precision of 37.12%, while the DRF with the "average function" reaches 35.83%, and the DRF with the "ratio" function attains 34.91%. The precisions of the RFM and the Bayes QS are around 28%, while the precision of the PFRL is around 21%. As noted for the results of the previous data sets, the best performances after 9 iterations are attained by the DRF with the “probabilistic product” function (49.80%) and the DRF with the “ratio” function (49.10%), the DRF with the “average” function providing much smaller performances (43%).

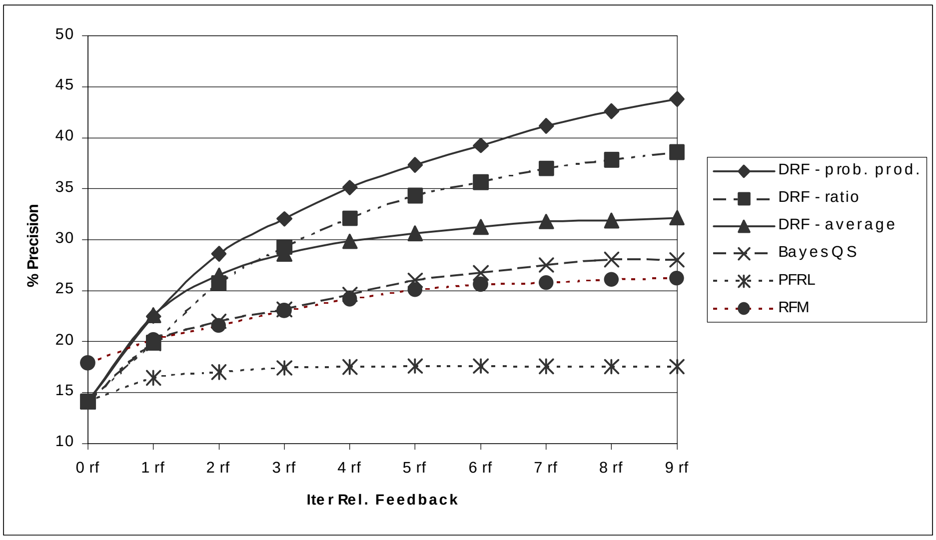

Fig. 4. Average percentage precision retrieval for Corel Data set – Co-occurrence texture feature set. Nine relevance feedback iterations were performed with the considered relevance feedback techniques.

Similar conclusions can be drawn from the results with the “co-occurrence texture” feature set. In particular, the DRF with the “probabilistic product” function attained the best performances at each iteration, while the “average” function performed better than the “ratio” function for the first three iterations. Therefore, the “probabilistic product” allows combining the strengths of the two relevance scores.

Figures 6 and 7 show the retrieval results related to the query shown in Figure 5. In particular, Figure 6 shows the images retrieved from the Corel dataset when no feedback is performed. According to the overall results presented in the previous tables, the performance is quite poor: only 4 images out of 20 match the user’s needs. After three feedbacks (Fig. 7) a large fraction of the images (16 out of 20) are relevant to the user’s needs.

Fig. 5. One of the query used in the experiments on the Corel dataset

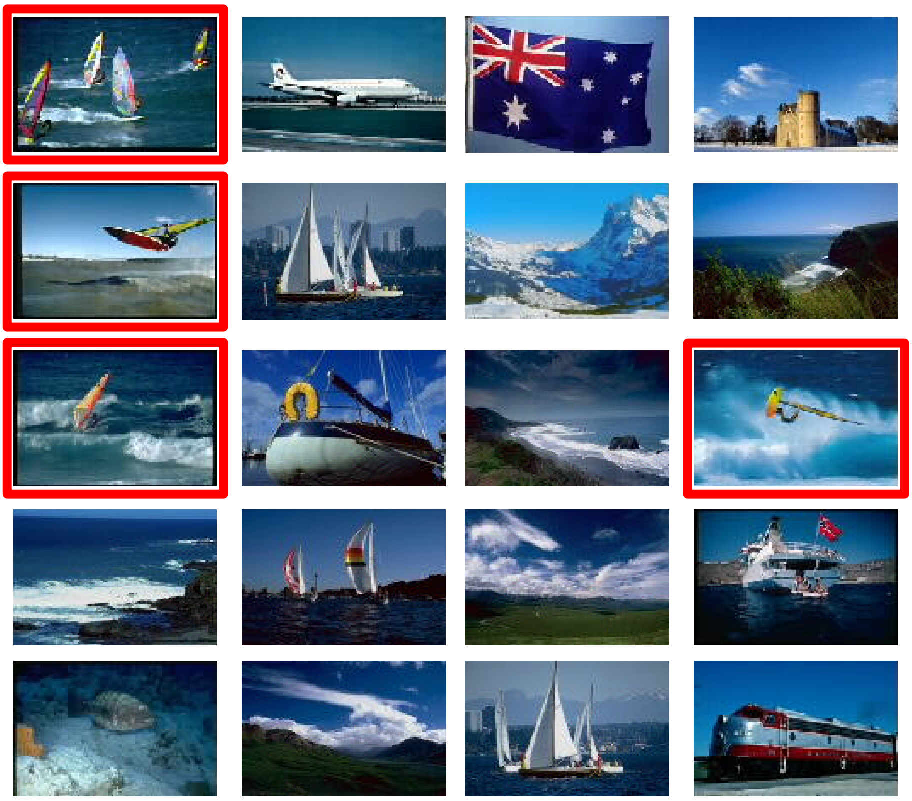

Fig. 6. Retrieval result related to the query shown in Fig. 5 with no feedback. Relevant images have a bold border.

Fig. 7. Retrieval results after three feedbacks. The number of relevant images increased significantly, and non-relevant images are ranked in

5. Conclusions

In this paper, we have presented a relevance feedback technique based on a dissimilarity representation of images. The reported results on three image databases showed the superiority of the proposed method with respect to other relevance feedback techniques, especially when a number of feedbacks are performed. The superiority of DRF was more evident in the case of the Corel image dataset, where the precision without relevance feedback is quite poor.

It is worth noting that many experiments presented in the literature on the Corel data set are based on the combination of relevance feedbacks from different feature sets, thus providing high performances. The proposed technique is also well suited for combining different feature sets, as the "dissimilarity" representation is independent on any feature-based representation. However, this topic is out of the scope of the present paper and will be discussed elsewhere.

As far as the computational complexity of the proposed technique is concerned, a large number of distances are to be computed. Nevertheless, the response time between two consecutive feedbacks was around 0.6s on the Corel dataset (made up of 19513 images), on a Celeron 450 MHz PC using the Win98 OS. This response time is far below the classic limit of 1.0s for the user's flow of thought to stay uninterrupted. Thus, despite the computational complexity of the algorithm, the response time on a not-so-fast machine can be considered acceptable for a large database. However, the response time of the implemented algorithm could be further improved by taking into account, for example, that consecutive retrievals share a number of images. As a consequence at each step only the distances related to new retrieved images should be computed, provided that distances computed in previous steps are stored.

Finally, it is worth remarking that the performances of some techniques used for comparison are heavily affected by the choice of the value of a number of parameters, while the proposed technique does not rely on parameters computed by heuristics.

Acknowledgments

The authors wish to thank Dr. Jing Peng for providing the features of the MIT data set and the C-code of the PFRL algorithm used to perform some of the reported experiments. They also thank Dr. Manuel Ortega for providing the images of the Corel Data set.

References

1. Bhanu, B., Dong, D.: Concepts Learning with Fuzzy Clustering and Relevance Feedback. In: Petra, P. (Ed.): Machine Learning and Data Mining in Pattern Recognition. LNAI 2123, Springer-Verlag, Berlin (2001) 102-116
2. Del Bimbo A.: Visual Information Retrieval. Morgan Kaufmann Pub. Inc., San Francisco, CA (1999)
3. Duin R.P.W., de Ridder D., Tax D.M.J.: Experiments with object based discriminant functions: a featureless approach to pattern recognition. Pattern Recognition Letters **18**(11-13) (1997) 1159-1166
4. Duin R.P.W., Pekalska E., de Ridder D.: Relational discriminant analysis. Pattern Recognition Letters **20**(11-13) (1999) 1175-1181
5. Frederix G., Caenen G., Pauwels E.J.: PARISS: Panoramic, Adaptive and Reconfigurable Interface for Similairty Search. Proc. of ICIP 2000 Intern. Conf. on Image Processing. WA 07.04, vol. III (2000) 222-225
6. Giacinto, G., Roli, F., Fumera, G.: Content-Based Image Retrieval with Adaptive Query Shifting. In: Petra, P. (Ed.): Machine Learning and Data Mining in Pattern Recognition. LNAI 2123, Springer-Verlag, Berlin, (2001) 337-346
7. Giacinto G., Roli F.: Query shifting based on Bayesian decision theory for content-based image retrieval. Proc. of S+SSPR2002, Canada, LNCS 2396, Springer-Verlag 607-616

8. Ishikawa Y., Subramanys R., Faloutsos C.: MindReader: Querying databases through multiple examples. In Proceedings. of the 24th VLDB Conference (1998) 433-438
9. Jain A.K., Dubes R.C.: Algorithms for Clustering Data. Prentice Hall (1988)
10. Kuncheva L.I.: Fuzzy Classifier Design. Springer-Verlag (2000)
11. McG Squire D, Müller W., Müller H., Pun T.: Content-based query of image databases: inspirations from text retrieval. Pattern Recognition Letters **21**(13-14) (2000) 1193-1198
12. Nastar C., Mitschke M., Meilhac C.: Efficient query refinement for Image Retrieval. Proc. of IEEE Conf. Computer Vision and Pattern Recognition, CA (1998) 547-552
13. Ortega M., Rui Y., Chakrabarti K., Porkaew K., Mehrotra S., Huang T.S.: Supporting ranked boolean similarity queries in MARS. IEEE Trans. on KDE **10**(6) 905-925 (1998)
14. Peng J., Bhanu B., Qing S.: Probabilistic feature relevance learning for content-based image retrieval. Computer Vision and Image Understanding **75**(1-2) (1999) 150-164
15. Rui Y., Huang T.S., Mehrotra S.: Content-based image retrieval with relevance feedback: in MARS. In Proceedings of the IEEE International Conference on Image Processing, IEEE Press (1997) 815-818
16. Rui Y., Huang T.S.: Relevance Feedback Techniques in Image retrieval. In Lew M.S. (ed.): Principles of Visual Information Retrieval. Springer-Verlag, London, (2001) 219-258
17. Salton G,, McGill M.J.: Introduction to modern information retrieval. McGraw-Hill, New York (1988)
18. Santini S., Jain R.: Integrated browsing and querying for image databases. IEEE Multimedia **7**(3) (2000) 26-39
19. Santini S., Jain R.: Similarity Measures. IEEE Trans. on Pattern Analysis and Machine Intelligence **21**(9) (1999) 871-883
20. Sclaroff S., La Cascia M., Sethi S., Taycher L.: Mix and Match Features in the ImageRover search engine. In Lew M.S. (ed.): Principles of Visual Information Retrieval. Springer-Verlag, London (2001) 219-258
21. Smeulders A.W.M., Worring M., Santini S., Gupta A., Jain R.: Content-based image retrieval at the end of the early years. IEEE Trans. on Pattern Analysis and Machine Intelligence **22**(12) (2000) 1349-1380