UNBALANCED LEARNING IN CONTENT-BASED IMAGE CLASSIFICATION AND RETRIEVAL

Luca Piras, Giorgio Giacinto

Dept. of Electrical and Electronic Eng.
University of Cagliari
09123 Piazza D'armi, Cagliari, Italy
Email: luca.piras@diee.unica.it, giacinto@diee.unica.it

ABSTRACT

Nowadays very large archives of digital images can be easily produced thanks to the availability of digital cameras as standalone devices, or embedded into a number of portable devices. Each personal computer is typically a repository for thousands of images, while the Internet can be seen as a very large repository. One of the most severe problems in the classification and retrieval of images from very large repositories is the very limited number of elements belonging to each semantic class compared to the number of images in the repository. As a consequence, an even smaller fraction of images per semantic class can be used as training set in a classification problem, or as a query in a content-based image retrieval problem. In this paper we propose a technique aimed at artificially increasing the number of examples in the training set in order to improve the learning capabilities, reducing the unbalance between the semantic class of interest, and all other images. The proposed approach is tailored to classification and relevance feedback techniques based on the Nearest-Neighbor paradigm. A number of new points in the feature space are created based on the available training patterns, so that they better represent the distribution of the semantic class of interest. These new points are created according to the k-NN paradigm, and take into account both relevant and non-relevant images with respect to the semantic class of interest. The proposed approach allows increasing the generalization capability of NN techniques, and mitigates the risk of classifier over-training on few patterns. Reported experiments show the effectiveness of the proposed technique in Content-Based Image Retrieval tasks, where the Nearest-Neighbor approach is used to exploit user's relevance feedback. The improvement in precision and recall gained in one feature space allows also to outperform the improvement in performances attained by combining different feature spaces.

***Keywords*—** Unbalanced Learning, Small Sample-Size, Artificial Pattern Injection, Image Retrieval, Image Classification

1. INTRODUCTION

As the years go by, it is ever-easier to have access to a ever-greater amount of electronically archived images. As a consequence there is an increasing need of tools for searching, classifying and retrieving them. The use of metadata associated to the images solves the problems only partly, as the process of assigning metadata to images is not trivial, is slow, and closely related to the persons who performed the task. This is especially true for retrieval tasks in very high dimensional databases, where images exhibits high variability in semantic. It turns out that the description of image content tends to be intrinsically subjective and partial, and the search for images based on keywords may fit users' needs only partially. To this end, the analysis of image content is performed not only by human experts, but also by the automatic analysis of the visual content [1, 2]. On the other hand, the automatic analysis of image content allows for the automatic propagation of labels.

The main reason for the difficulty in devising effective image retrieval and classification tools is caused by the vast amount of information conveyed by images, and the related subjectivity of the criteria to be used to assign labels to images [2, 3, 4]. In order to capture such subjectivity, image retrieval and classification tools may employ relevance feedback techniques. Recently, relevance feedback techniques have been formulated in terms of a classification problem [5]. From this perspective, the image retrieval task share similar problems with classification tasks.

One of the most severe problem in the design of the classifier, is the unbalance between the number of samples of the class the user is interested in, and all other images of the database that share some characteristics with that class. In the machine-learning literature the unbalance problem has been widely investigated and solution based either on under-sampling the majority class [6, 7], or on over-sampling the minority class have been proposed. Other authors proposed to blend techniques based on under-sampling the majority class, with a special form of over-sampling of the minority class [8]. However, it is easy to see that these solutions tend to produce a distortion of the "real" distribution of the classes. In the image retrieval domain, some solutions proposed so far involve the reduction of the set of images that are non-relevant to user's interest by the creation of bootstrap samples from the set of non-relevant images [9]. An ensemble of balanced training sets is thus created, each ensemble being made up of all the available relevant images and one of the bootstrap samples. This solution is computationally quite expensive, as an ensemble of classifiers has to be created, and the choice of the most appropriate set of parameters for the various parts of the systems is far from being a trivial task.

In this paper, we address the unbalance problem by creating new random artificial patterns of the class of interest. The new patterns take into account nearest-neighbor relations between the available samples of the class of interest. To avoid creating noisy patterns, the new patterns are constrained to lie in a region of the feature space that contains images belonging to the class of interest. The exploitation of the nearest-neighbor information is performed according to the rationale of the k-nn noise injection technique, and accounts for the manifold where relevant images are assumed to lie [10].

Recently, some other papers addressed the unbalance problem by proposing the generation of artificial patterns [11]. However, while those papers, propose the generation of synthetic images for the sake of receiving informative feedback from the user, in this paper new patterns are generated in the feature space to improve the performances of learning mechanisms. Thus, no visual representation of the patterns is provided.

This paper is organized as follows. The proposed technique for creating artificial feature vectors is illustrated in Section 2. Section 3 shows the integration of the proposed technique in the learning process whereas Section 4 describes the Nearest-neighbor technique used. Experimental results on an image dataset are reported in Section 5. Reported results show that the proposed method allows improving the performances of a relevance feedback mechanisms based on Nearest-neighbor. Conclusions are drawn in Section 6.

2. GENERATION OF ARTIFICIAL FEATURE VECTORS

The ever-increasing availability of very large archives of digital images makes harder and harder the tasks of image classification and retrieval. The main difficulty basically arises from the wide variety of semantic classes in the repositories. As a consequence, the number of images per semantic class is a tiny fraction of the whole repository, and the fraction of images available for training or querying purposes is even smaller. In addition, in order to "cover" the wide spectrum of high-level semantic concepts, features made up of a large number of components need to be extracted. In such large feature spaces, it is easy that images from different semantic classes can be considered similar each other. Thus, the learning process where the goal is to classify/retrieve images belonging to a target class, is usually unbalanced, as there are more chances of collecting images that do not belong to the target class, than chances of collecting target images [12, 13].

To provide a solution to the unbalance problem, we propose to artificially increase the number of samples of the target class, thus reducing the unbalance between the semantic class of interest, and all other images. The new patterns are not "real" new images but synthetic samples generated in the feature space according to the distribution of the available samples. We will refer to the proposed technique as "*Directed Pattern Injection*" (DPI). This technique has been inspired by the work on "Directed Noise Injection" [10], which generates noisy samples in the neighborhood of each sample point. In our proposal, we take into account the peculiar characteristics of the task at end, so that the new patterns can provide additional information that is embedded into the available samples.

The proposed DPI technique, basically creates some artificial patterns in the feature space by taking into account the K nearest images of the class of interest with respect to some reference point. Let us consider a feature space F with p components and denote with X a pattern of F. We will create a number of artificial relevant patterns U related to X as [10]:

$$U = X + Z. \tag{1}$$

where the *artificial feature vector* Z is evaluated as:

$$Z = \lambda \times \frac{1}{K} \sum_{k=1}^{K} \xi_k \left(X - Q_k \right), \tag{2}$$

where $\xi_k \sim N(0,1)$, λ is a scale factor and $Q_1, \ldots, Q_K$ are X's K nearest patterns. For each X and its K nearest patterns $Q_1, \ldots, Q_K$ it is possible to generate potentially an infinite number of artificial feature vectors Z by varying the coefficients $\xi_1, \ldots, \xi_K$. Consequently, it is possible to train a classifier with a number of samples as large as needed. The value of K, i.e. the number of nearest-neighbors considered for the creation of artificial patterns, should not be large to avoid taking into account pattern that are actually far from the reference point. According to the work by Skurichina [10] and to the results of some preliminary experiments, we suggest using a value of K equal to 2.

The proposed approach is expected to be effective when used with learning techniques based on the Nearest-Neighbor paradigm. In fact, new patterns are generated according to a linear combination of directions depending on the distribution of the K-nearest neighbors of a reference image, while the coefficients of the combination are chosen randomly from a normal distribution. Starting from the assumption that the patterns of the same class lie in a subspace of the feature space, this technique permits to identify the subspace determined by the K-nearest neighbors through the linear combination of known patterns and to generate the new ones in it. Figure 1 shows an example of how the proposed technique works, where the gray area indicates where artificial patterns can be created. Further details on the implementation of the proposed technique will be provided in the following section.

It is worth noting that the proposed *Directed Pattern Injection* technique is cheaper in terms of computation time if compared to, for example, techniques that produce new patterns along all directions (*Spherical Gaussian PI*). In fact if the patterns are added only toward "useful" directions it is possible to save memory and processing time. Another advantage of DPI is that it is not dependent on the knowledge of patterns distributions, and so it is possible to successfully apply it to any kind of data distribution.

The DPI technique depends on the values assigned to a number of free parameters, that are, the image X used as a reference, the number of nearest-neighbors considered, the number of artificial feature vectors that are generated, and the scale factor λ. These values cannot be "a priori" chosen. On the other hand, tuning all the free parameters in order to find the optimal configuration is a difficult and computationally expensive search task. We thus selected these parameters according to some heuristics that are reported in the following section.

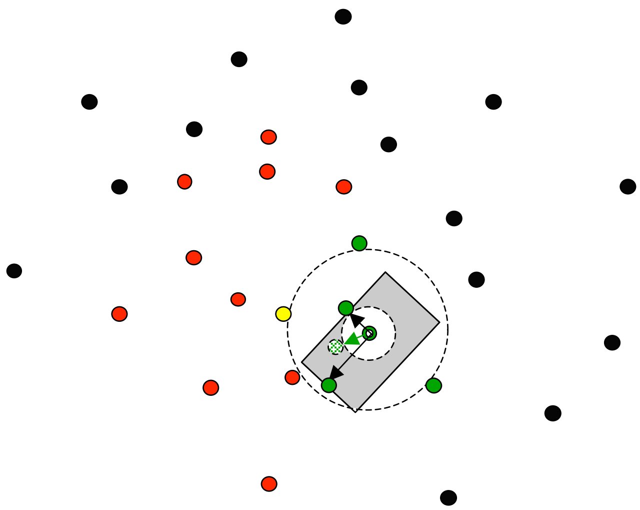

Fig. 1. Example of K-NN "synthetic" pattern generation in a two-dimensional feature space, where $K = 2$, $\xi_1 = \xi_2 = 0.5$ and $\lambda = 1$.

3. CATEGORY LEARNING WITH ARTIFICIAL FEATURE VECTORS

The DPI method is useful both for image retrieval and image classification, as the number of training patterns belonging to the class of interest can be increased to cope with unbalanced training sets. In particular, in the case of image retrieval, it can be beneficial for relevance feedback techniques, where very often non-relevant images outnumber relevant images. It is worth noting that the choice of the number of artificial patterns to be created is not a trivial task. In fact if the number of them is too large there is the concrete risk to add noise to the dataset, thus producing a distortion of the "real" distribution of images. For these reasons, we propose to constrain the generation of new patterns so that the ratio between images that belong to the class of interest, and those that do not, is constantly equal to a predefined ratio $1 : n$, where $n = 2, \dots, 5$. In cases when the ratio in the training set exceeds the above ratio, then the artificial generation of patterns is not executed at all. In addition to this choice, there are at least two other parameters that need to be adjusted: Should the values of the parameters $\xi_1, \dots, \xi_K$ be constrained? Which strategy can be used to choose the reference point X used to generate the artificial patterns?

Let us provide an answer to the first question. While the proposed approach generates new vectors only in some directions of the feature space (see Equation (2)), the generated patterns may still lie outside the region explored, i.e., the region defined by the known neighborhood of the reference point. In our opinion, it is too risky to generate pattern outside that region, as we have no information available about the distribution of images. This can happen depending on the random values of the coefficients $\xi_1, \dots, \xi_K$ in equations (2) and (1). To avoid this risk, we propose to constrain the creation of new patterns in the region delimited by the nearest and the farthest known image of interest w.r.t. the reference image used in equation (1). Figure 1 shows an example of how the proposed technique works, where the grey area indicates where artificial patterns can be created, and the two dashed circles bound the area where artificial patterns are accepted. The green ring represents the reference point, chosen as the mean vector of all the images of the target class, while the red, green, and black dots represent non-target, target and unlabeled images, respectively. The green dots pointed by the black arrows represents the two target images nearest to the reference point, while the green dot drawn with a broken line and pointed by the green arrow, represents the injected "synthetic" pattern. It is easy to see that the new pattern falls in the region containing the samples of the target class, thus increasing the density of this area according to neighborhood information.

Different choices for the reference point can be investigated. Let us recall that the directions used to generate new patterns depends on the choice of reference point, and the related nearest neighbors. In addition, the choice of the reference point may also depend on the application at hand. In an image retrieval problem the reference point could be selected as being **the pattern associated to the query image**, as the user asked for images similar to the query. While this can be a reasonable choice, it can also exhibit some drawbacks, as its representation in the low-level feature space may not reflect its representativeness w.r.t. the images the user considers as being relevant. In other words, the so-called "semantic gap" between user perception of similarity and its representation in the low-level feature space may suggest to use a different point in the feature space as a query vector.

As an alternative, we propose to use **the mean vector of all the known images of the target class** as the reference point, thus taking into account the distribution of the images of the class of interest in the feature space. This choice, with respect to the first one, takes into account all available information, and thus can be used in different application scenarios, including classification and retrieval.

Other options could be also investigated such as, the use of each known image of the target class as a reference image, and the use of a new point computed according to the known target and non-target images. The first option requires some extra parameter to be set, as the total number of artificial patterns that needs to be generated can be smaller than the number of available images, and some target patterns should not participate in the process. The second option requires defining the new point according to some heuristics, e.g., those used for performing query movement in image retrieval tasks [2, 3, 4]. These options have been tested, and the experimental results have been reported in [14]. Good results are attained by using the mean vector, and, in some cases, by using a new point computed according to the known target and non-target images.

In order to keep the system simple to implement, by reducing the number of parameters which affects the final performance of the system, we decided to use the mean vector of all images belonging to the target class as the reference point. In addition, this choice allows creating synthetic patterns that lie in the region where we actually observed images of interests.

4. NEAREST NEIGHBOR APPROACH FOR CONCEPT LEARNING

As above-mentioned, the DPI approach can be effective for techniques based on the Nearest-Neighbor paradigm. The use of the nearest-neighbor paradigm is motivated by its use in a number of different pattern recognition fields, where it is difficult to produce a high-level generalization of a class of objects, but where neighborhood information is available [15, 16]. In particular, nearest-neighbor approaches have proven to be effective in outlier detection, and one-class classification tasks [17, 18]. In this work we resort to a technique proposed in [19] where a score is assigned to each image of a database according to its distance from the nearest image belonging to the target class, and the distance from the nearest image belonging to a different class. This score is further combined to a score related to the distance of the image from the region of relevant images. The combined score is computed as follows:

$$rel(I)_{stab} = \left(\frac{n/k}{1+n/k}\right) \cdot rel_{BQS}(I) + \left(\frac{1}{1+n/k}\right) \cdot rel_{NN}(I) \tag{3}$$

where n and k are the number of non-relevant images and the whole number of images retrieved after the latter iteration, respectively. The two terms rel_{NN} and rel_{BQS} are computed as follows:

$$rel_{NN}(I) = \frac{\|I - NN^{nr}(I)\|}{\|I - NN^{r}(I)\| + \|I - NN^{nr}(I)\|} \tag{4}$$

where $NN(I)$ denotes the nearest neighbor of I, $\|\cdot\|$ is the metric defined in the feature space at hand,

$$rel_{BQS}(I) = \frac{1 - e^{1 - d_{BQS}(I) \Big/ \max_{I} d_{BQS}(I)}}{1 - e} \tag{5}$$

where d_{BQS} is the distance of image I from a modified query vector computed according to the Bayes decision theory (Bayes Query Shifting, BQS) [20]. This technique has been designed to implement relevance feedback in image retrieval, but in principle it is possible to set a threshold, and use it for image classification.

5. EXPERIMENTAL RESULTS

This section describes the experimental results related to the use of the technique proposed in Sections 2 and 3 in a Content-Based Image Retrieval problem. Experiments have been carried out using the Caltech-256 dataset obtained from the California Institute of Technology repository[1] has been used. It consists of 30607 images subdivided into 257 semantic classes [21]. Images have been represented using three different feature vectors, namely *Tamura* [22] (18 components), *Scalable Color* [23] (64 components), and the *Color and Edge Directivity Descriptor* ($Cedd$, 144 components) [24]. In order to test the performances of the proposed mechanism, 500 query images have been randomly extracted from all the 257 classes. The top twenty best scored images for each query are returned to the user. 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. It is worth noting that at each round of relevance feedback, the user is asked to mark twenty brand new images never seen before. Performances are evaluated in terms of retrieval precision and recall. Precision is measured by taking into account the top twenty best scored images, regardless they have been already labelled by the user. On the other hand, the recall takes also into account the brand new images labelled by the user at each iteration for providing the feedback to the system.

In order to choose the most suitable values of the parameters discussed in Section 3, a number of preliminary experiments have been performed. Accordingly, we computed the normalization parameter as $\lambda = \left(\xi_1^2 + \xi_2^2\right)^{-2}$, and created new "synthetic images" at each iteration only when the ratio between relevant and non-relevant patterns is smaller then 0.5, so that the final ratio between relevant and non relevant images is equal to $1 : 2$. Thus it means that artificial patterns are generated only in cases of very unbalanced training sets, and that the number of artificial patterns is kept as small as possible, so that they can provide useful information rather than adding noise.

Figures 2(a) and 2(b) show the performances in terms of precision and recall using the three feature representations. As it is expected, the best performances are attained by the $Cedd$ representation, as it allows to better capture the different semantic of the classes in the dataset. On the other hand, both the $Tamura$ and $Scalable\ Color$ representations allow capturing only partially the semantic of the classes. If we compare the retrieval results without relevance feedback, it can be seen that (Figure 2(a)) the average precision attained with the $Scalable\ Color$ representation is equal to 2.5%, while the $Tamura$ representation allows attaining a precision rate equal to 6%, and the precision attained by $Cedd$ is around 10%. A similar trend can be seen in the recall measure (Figure 2(b)), where $Scalable\ Color$ produce a 0.3% rate, the $Tamura$ representation reaches 0.5% and the $Cedd$ allows recalling 1%.

As the three feature representations provide complementary information on image semantic, we performed an experiment by combining the relevance scores computed separately for each feature representation. In particular, the images are ranked according to the average relevance score computed using each feature representation. Reported results show that modest improvements are attained in precision w.r.t. the performances attained by the $Cedd$ representation, while the improvements are

[1] http://www.vision.caltech.edu/Image_Datasets/Caltech256/

clearly remarkable if we make the comparison with the precision attained by the $Scalable\ Color$ and the $Tamura$ representations. On the other hand, if we consider the recall measure, the performance of the combination is clearly better than that attained by the $Tamura$ and $Scalable\ Color$ representations, while it is slightly worse than that attained by the $Cedd$ representation. As the combination of feature representations is not in the scope of this paper, we don't report experiments related to combination mechanisms based on weighted rules. Moreover, while weighted rules could in principle take into account the strength and weakness of each feature representation, they require a careful design in order to properly estimate the weights.

According to the results of the above experiments, we decided to use the proposed technique for artificial pattern generation in the $Cedd$ feature representation, as it provided the best performances compared to the other feature representations. Figures 2(a) and 2(b) show that the artificial generation of patterns (**DPI - Cedd**) allows improving the learning capabilities in the case of unbalanced training set. In particular, the precision attained by proposed technique on the $Cedd$ feature representation is higher than the precision attained in the $Cedd$ feature space without artificial pattern generation from the second iteration on. In addition, the precision is also higher than that attained by the combination of the three feature representations, starting from the third iteration. At the end of the ninth iteration, the improvement in precision attained by the proposed technique is nearly equal to 3.5%. If we consider the recall measure reported in Figure 2(b), we can observe a behavior similar to the one seen in the case of the precision. In particular, significative improvements can be seen since the third iteration.

Another aspect that makes the DPI an effective technique, is the computation time saving. It is worth noting that while the combination of complementary feature representation allows attaining improvements in precision, the computational overhead of the combination with respect to the most computational demanding feature space (i.e., the $Cedd$ representation) is equal to 77%. On the other hand, the overhead of the proposed DPI technique is equal to 15%.

6. CONCLUSIONS

In this paper we proposed a technique that address the unbalance problem in image classification and retrieval tasks. We proposed to artificially increase the number of samples of the class of interest by creating new samples according to nearest-neighbor information. Reported results on an image retrieval task with relevance feedback show that the proposed technique allows improving the performances both in precision and recall. In particular the improvement can be seen not only with respect to the performances attained without artificial patterns, but also with respect to performance attained by combining different feature spaces. The computation overhead of the proposed technique is small if compared to the overhead due to the combined use of different feature spaces.

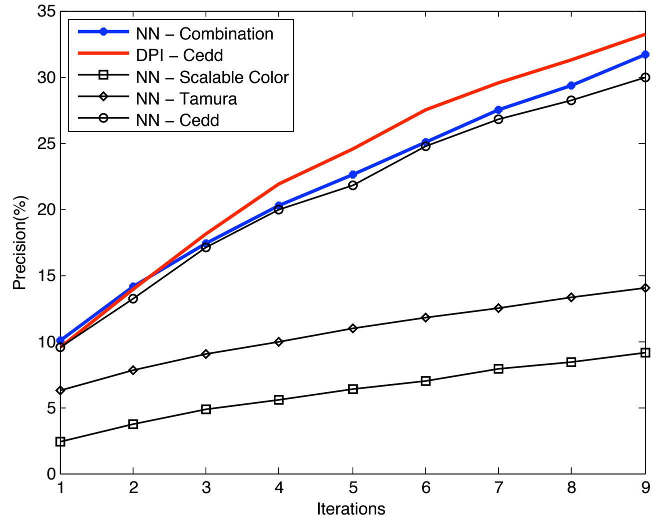

(a) Precision

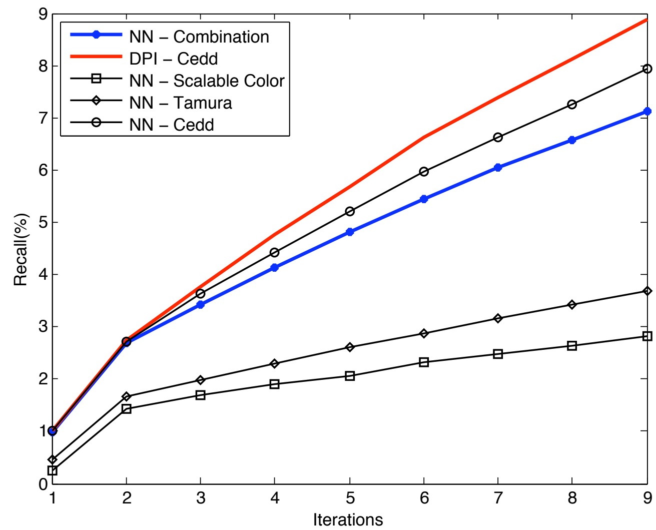

(b) Recall

Fig. 2. Experimental results.

7. REFERENCES

[1] Alberto Del Bimbo, *Visual information retrieval*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.

[2] Michael S. Lew, Nicu Sebe, Chabane Djeraba, and Ramesh Jain, "Content-based multimedia information retrieval: State of the art and challenges," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 2, no. 1, pp. 1–19, 2006.

[3] A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 12, pp. 1349–1380, 2000.

[4] X. S. Zhou and T. S. Huang, "Relevance feedback in im-

age retrieval: A comprehensive review," *ACM Multimedia Systems*, vol. 8, no. 6, pp. 536–544, 2003.

[5] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z. Wang, "Image retrieval: Ideas, influences, and trends of the new age," *ACM Computing Surveys*, vol. 40, no. 2, pp. 1–60, 2008.

[6] Miroslav Kubat and Stan Matwin, "Addressing the curse of imbalanced training sets: one-sided selection," in *Proc. 14th International Conference on Machine Learning*. 1997, pp. 179–186, Morgan Kaufmann.

[7] Nathalie Japkowicz, "The class imbalance problem: Significance and strategies," in *In Proceedings of the 2000 International Conference on Artificial Intelligence (ICAI*, 2000, pp. 111–117.

[8] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer, "Smote: Synthetic minority over-sampling technique," *J. Artif. Intell. Res. (JAIR)*, vol. 16, pp. 321–357, 2002.

[9] D. Tao, X. Tang, X. Li, and X. Wu, "Asymmetric bagging and random subspace for support vector machines-based relevance feedback in image retrieval," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 28, pp. 1088–1099, 2006.

[10] M. Skurichina, S. Raudys, and R. P. W. Duin, "K-nearest neighbors directed noise injection in multilayer perceptron training," *IEEE Trans. on Neural Networks*, vol. 11, no. 2, pp. 504–511, March 2000.

[11] B. Thomee, M.J. Huiskes, E. Bakker, and M.S. Lew, "Using an artificial imagination for texture retrieval," in *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, Dec. 2008, pp. 1–4.

[12] X. Zhou and Thomas S. Huang, "Small sample learning during multimedia retrieval using biasmap," in *Proc. IEEE Intl Conf. Computer Vision and Pattern Recognition*, 2001, vol. 1, pp. 11–17.

[13] Robert P. W. Duin, "Pattern recognition in almost empty spaces," in *WIC Winter Symposium*, Eindhoven, Netherlands, Jan. 2004.

[14] Luca Piras and Giorgio Giacinto, "K-nearest neighbors directed synthetic images injection," in *WIAMIS '10: Proceedings of the 2010 11th International Workshop on Image Analysis for Multimedia Interactive Services*, Desenzano del Garda, Italy, April 2010.

[15] David W. Aha, Dennis Kibler, and Marc K. Albert, "Instance-based learning algorithms," *Machine Learning*, vol. 6, no. 1, pp. 37–66, 1991.

[16] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, John Wiley and Sons, Inc., New York, 2001.

[17] M. M. Breunig, H. P. Kriegel, R. T. Ng, and J. Sander, "LOF: identifying density-based local outliers," *SIGMOD Rec.*, vol. 29, no. 2, pp. 93–104, 2000.

[18] D. Tax, *One-class classification*, Ph.D. thesis, Delft University of Technology, The Netherlands, 2001.

[19] Giorgio Giacinto, "A nearest-neighbor approach to relevance feedback in content based image retrieval," in *CIVR '07: Proceedings of the 6th ACM international conference on Image and video retrieval*, New York, NY, USA, 2007, pp. 456–463, ACM.

[20] G. Giacinto and F. Roli, "Bayesian relevance feedback for content-based image retrieval," *Pattern Recognition*, vol. 37, pp. 1499–1508, 2004.

[21] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," Tech. Rep. 7694, California Institute of Technology, 2007.

[22] Hideyuki Tamura, Shunji Mori, and Takashi Yamawaki, "Textural features corresponding to visual perception," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 8, no. 6, pp. 460–473, June 1978.

[23] Information technology - Multimedia content description interface - Part 3: Visual, ISO/IEC Std. 15938-3:2003, 2003.

[24] Savvas A. Chatzichristofis and Yiannis S. Boutalis, "Cedd: Color and edge directivity descriptor: A compact descriptor for image indexing and retrieval," in *ICVS*, Antonios Gasteratos, Markus Vincze, and John K. Tsotsos, Eds. 2008, vol. 5008 of *Lecture Notes in Computer Science*, pp. 312–322, Springer.