

Pseudo-Relevance Feedback Diversification of Social Image Retrieval Results

Bogdan Boteanu · Ionuț Mironică ·
Bogdan Ionescu

Received: date / Accepted: date

Abstract In this paper we introduce a novel pseudo-relevance feedback (RF) perspective to social image search results diversification. Traditional RF techniques introduce the user in the processing loop by harvesting feedback about the relevance of the query results. This information is used for recomputing a better representation of the needed data. The novelty of our work is in exploiting the automatic generation of user feedback in a completely unsupervised diversification scenario, where positive and negative examples are used to generate better representations of visual classes in the data. First, user feedback is simulated automatically by selecting positive and negative examples from the initial query results. Then, an unsupervised hierarchical clustering is used to re-group images according to their content. Diversification is finally achieved with a re-ranking approach of the previously achieved clusters. Experimental validation on real-world data from Flickr shows the benefits of this approach achieving very promising results.

Keywords Social image search results diversification · Pseudo-relevance feedback · Hierarchical clustering · Image retrieval · MediaEval benchmarking.

Bogdan Boteanu
LAPI, University “Politehnica” of Bucharest, 061071 Bucharest, Romania
E-mail: bboteanu@alpha.imag.pub.ro

Ionuț Mironică
LAPI, University “Politehnica” of Bucharest, 061071 Bucharest, Romania
E-mail: imironica@alpha.imag.pub.ro

Bogdan Ionescu
LAPI, University “Politehnica” of Bucharest, 061071 Bucharest, Romania
E-mail: bionescu@alpha.imag.pub.ro

1 Introduction

Nowadays the Internet plays a critical role in our daily life, becoming the main source of information. Besides text, most of these information can also be found in multimedia files, such as images, videos and audio. As the amount of data increases continuously, it is essential to find our needed information in an effective way. In this context, image retrieval still remains an important topic as a large amount of multimedia data is stored as images on the Internet.

An effective retrieval system should be able to summarize search results and give a global view by surfacing the results which are *relevant* and in the same time they cover different aspects, i.e., they are *diverse*, of the query. Most of the queries involve many possibilities to interpret the interrogations because they may refer to different subtopics, e.g., animals can be part of particular species, cars belong to different types and producers, points of interest can be photographed from different angles and so on. Current state-of-the art retrieval systems, e.g., Flickr¹, Google Images², focus mainly on relevance and providing real time capabilities for the retrieval process which is often in the detriment of the effectiveness and accuracy of the results. In effect, on many occasions, results tend to be unsatisfactory for the common user. For instance, results can be noisy and not relevant to what the users are searching for; or provide many similar results in terms of visual appearance, e.g., duplicate images. An example in this respect is presented in Figure 1, namely a Flickr search for “*Rialto Bridge*”. One can observe that although the search query is more or less common, provided results tend to be redundant (e.g., see images marked by the red rectangle, which are near-duplicates) and many are not relevant to the query or are ambiguous (e.g., see images marked by ‘X’ and ‘?’). A good example of relevant and diverse images are the results marked by the blue rectangles, where we can observe that the images are related to the query and are also visually different, i.e., they depict different angles, different moments of the day, etc.

By increasing the relevance of the results, one can expect to provide more accurate information to the user; whereas by widening also the pool of possible results (i.e., diversify) one can increase even more the effectiveness of the system, e.g., by responding to the needs of different users, by tackling queries with unclear information needs, by widening the pool of possible results, reducing the number/redundancy of the returned items, and so on. In general, relevance and diversity tend to be antinomic, i.e., increasing solely the relevance will tend to provide many near duplicates while too much diversification may result in losing relevant items [1]. The key of the entire diversification process is to mitigate these two components. Although relevance was studied more thoroughly in existing literature than diversification [2,3,4] and even though a considerable amount of diversification literature exists (in particular in the

¹ <https://www.flickr.com/>

² <https://www.google.com/imghp>



Fig. 1 Search for “Rialto Bridge” on Flickr: images marked by ‘X’ are not relevant to the query, images marked with ‘?’ are ambiguous, images marked by the red rectangle depict redundant results and images in blue rectangles depict examples that are relevant and diverse.

area of text retrieval, from where the concept emerged), the topic is currently still an open research issue, especially for social media [5,6,7,8].

In this paper we investigate a particular class of methods, i.e., relevance feedback (RF), which were developed initially for increasing relevance only, but show promising results also for diversification. RF take advantage directly of the human expertise in the retrieval process [9,10,11]. A general RF scenario consists of the following steps. For a certain retrieval query, the user gives his opinion by marking the results as relevant or irrelevant. Then, the system automatically computes a better representation of the information needed based on this feedback and the retrieval is further refined. This process can go on several cycles. Although effective in many scenarios, there are some drawbacks of this approach. If to consider a real-world retrieval scenario it is not realistic to assume that relevance judgments are available at the time of the retrieval and the process of completing the feedback can be a time consuming task for the user. In the same time, traditional relevance feedback harvests only relevance information, which limits its capabilities to improving only these aspects. A particular case of RF is the pseudo-relevance feedback [12,13]. Pseudo-RF represents an automatic method to simulate the users’s feedback, so that the results are improved without any interaction with the system. The main idea behind this concept is to assume that a small number of top-ranked items in the initial retrieval results are relevant, and use them to re-rank the results. Many experiments such as the results from Trecvid [14] and ImageClef [15] show improvement in retrieval system’s performances using this concept.

In this work, we introduce a new pseudo-relevance feedback approach in an unsupervised diversification scenario based on a hierarchical clustering algorithm and a re-ranking scheme. The automatic generation of user feedback, i.e., the positive and negative examples, is used to generate better representa-

tions of visual classes in the data. To disambiguate the diversification need, we focus in this paper on a specific use case, i.e., the search for images with tourist landmarks. Landmark locations are popular with social media platforms and also benefit from spatio-temporal and visual invariance which makes them suitable for benchmarking scenarios.

The remainder of the paper is organized as following. Section 2 presents the overview of the literature on diversification and relevance feedback and positions our contribution. The proposed approach is explained in Section 3. Section 4 and 5 deal with the experimental setup and validation, while Section 6 concludes the paper.

2 Previous work

The concept of search result diversification emerged naturally from the text community, which is responsible for the very first publicly available indexing systems; therefore most of the work on diversification is concentrated around text-based search [16, 17, 18]. However, with the success of the Internet and social media, the problem has been extended to media information and many diversification approaches have been introduced, e.g., use of visual descriptors [5, 8, 19], user tagging credibility information [20], video features [21].

From their early stages until present, many directions have been investigated. A class of approaches are the *election-based* techniques, where images vote for other images to capture similarity and to achieve diversification. For instance, Leuken et al. [5] propose a method where every image decides by which image (besides itself) it is best represented, accumulating a score which in the end determine chances of the image of being elected as representative. The remaining images are assigned to the representative images using some similarity rules, e.g., a Nearest Neighbor approach.

Another direction is to use *clustering techniques*, in which diversification is achieved by employing some clustering/classification algorithms; some of the commonly used being Hierarchical Clustering, k-means, Support Vector Machines (SVM), etc. For instance, Dang-Nguyen et al. [20] propose for diversification a method based on Hierarchical Clustering algorithm, followed by isolated cluster removing and re-ranking. They employ face, blur, GPS and user credibility-based pre-filtering to remove irrelevant images. Data is represented by text (TF-IDF), visual and user credibility descriptors. Another example is the approach of Sarac and Duygulu [22] which uses for relevance a binary SVM classifier and for diversification k-means clustering. The output is formed by the centroids of each class and by images which are closest to the centroids, according to a k-Nearest Neighbors approach. Data is represented by visual descriptors.

Some other approaches use *re-ranking techniques*, which attempt to compute a new ranking of the results using a different modality than the one used in the initial retrieval, which is typically text. This is often achieved using the visual similarities between retrieved images and other representative

images. For instance, Pedronette et al. [23] propose a re-ranking approach based on the k-Nearest Neighbors of a query image, which define the contextual information for the query image. Data is represented by visual and textual descriptors. Spyromitros-Xioufis et al. [24] propose a re-ranking strategy based on the scores of a utility function which jointly considers relevance and diversity. They use a supervised classification model that capture a query and application-specific notion of relevance. Ksibi et al. [25] use an 'adaptive diverse relevance ranking' which automatically predicts an effective trade-off between relevance scores and diversity scores according to the query ambiguity level. Relevance score is assessed using a random walk with restart process and diversity score is a combination of semantic and visual scores.

Apart from the "automated" techniques, more and more attention is paid to the concept of "human in the loop", where human input/feedback is included in the processing chain. In a first attempt, this was carried out by conducting user studies on the systems' results. As this approach is far from being able to perform in real time, other perspectives had to be considered. One lead is to take advantage of the potential of *crowdsourcing* platforms [26] in which humans (i.e., users around the world) act like a computational machine that can be accessed via a computer interface. Issues like validity and reliability are present here and opened to further investigation, since the involvement of untrained people in high complexity tasks may not always bring good results. Another perspective is the use of *relevance feedback*, in which a part of the information is assessed by the user and further provided to the system. This concept has proven to increase retrieval accuracy and gives more personalized results for the user. The relevance feedback problem can be formulated either as a two class classification of the negative and positive samples; or as an one class classification problem, i.e., separate positive samples by negative samples. Using a graphic interface, a human assessor usually annotates a part of the retrieved images as relevant or irrelevant, according to a relevance criteria. The number of the images to be annotated is much lower than the number of the results. These images are further processed so that the result is a re-ranked list of relevant results. Basically, the algorithm learns from the provided data and tries to refine the whole list of results in order to obtain better precision in terms of relevance. A part of these results can be further evaluated again by the user in a second iteration. The process can be repeated until there are no big variations in terms of precision or until a minimum desired degree of performance is achieved.

One of the earliest and most successful RF algorithms is the Rocchio's algorithm [27] (which is still used at the present time). Using the set of relevant and irrelevant documents selected from the current user relevance feedback window, the Rocchio's algorithm modifies the features of the initial query by adding the features of positive examples and subtracting the features of negative examples to the original feature. Another relevant approach is the Relevance Feature Estimation (RFE) algorithm [28]. It assumes that for a given query, according to the user's subjective judgment, some specific features may be more important than others. A re-weighting strategy is adopted

which analyzes the relevant objects in order to understand which dimensions are more important than others in determining “what makes a result relevant”. Features with higher variance with respect to the relevant queries lead to lower importance factors than elements with reduced variation. Machine learning techniques can be used also to this task. These approaches involve a training step using the data collected from the user. Then all the results are ranked according to the classifiers’s confidence level [29], or classified as relevant or irrelevant depending on some output functions [30]. Some of the most successful techniques use Support Vector Machines [29], Nearest Neighbor approaches [5], classification trees, e.g., use of Random Forests [31]; or boosting techniques, e.g., AdaBoost [30].

Although they provide good results, the main disadvantage of RF is in the need of human expertise, which may not be available at any time and is also time consuming, due to the manual annotation process. A good tradeoff between performance and time are the pseudo-relevance feedback techniques, where the user is substituted by automatic methods of selecting the positive and negative examples. Some of the first approaches in this respect were text-based [32], where first top-ranked documents are the positive examples and the last documents are the negative ones. Positive examples and negative examples can be used as a training set for machine-learning techniques [33,34] and for multimedia retrieval [35]. This perspective seems very promising especially in the context of the current state of the art retrieval systems, which are able to provide an initial high relevance of the results.

In this paper, we exploit a novel perspective of the diversification via the use of pseudo-RF and clustering/re-ranking. Traditional RF introduces the user in the loop by harvesting feedback about the relevance of the search results. This information is used as ground truth for recomputing a better representation of the data needed. We propose an alternative pseudo-relevance solution for rendering this process completely automatic, while still maintaining a high relevance. First, user feedback is simulated automatically by selecting positive and negative examples from the initial query results. Then, a completely automatic unsupervised hierarchical clustering is used to re-group images according to their content. Diversification is finally achieved with a re-ranking approach of the previously achieved clusters. Experimentation on real-wold Flickr data shows the benefits of this approach which despite its complete automation is able to outperform or provide comparable results to user-based relevance feedback approaches as well as other consecrated diversification schemes.

3 Proposed approach

The proposed approach operates on top of an existing retrieval system and works as a re-ranking step that refines the initial query results. It involves the following steps:

1. *examples selection*: positive and negative examples are selected from the query results. The total number of the selected examples is split in equal sub-intervals (see Section 3.1);
2. *clustering*: for each sub-interval, an unsupervised classification step is used to cluster these examples (see Section 3.2);
3. *merging*: because different sub-intervals may contain similar clusters, clusters are further merged using a folding approach. For this step, we adapted the approach in [5] (see Section 3.3);
4. *pruning*: each resulting cluster is further evaluated based on the number of the relevant and irrelevant images within. This step will ensure the relevance of the refined images (see Section 3.4);
5. *diversification*: following the initial ranking of the retrieved images, cluster images are progressively selected to form the refined diversified query results (see Section 3.5).

Each of the processing steps is detailed in the following sections.

3.1 Selection of positive and negative examples

The first step of the proposed approach consists of selecting a number of positive and negative query examples. Instead of using a classic relevance feedback strategy where the user is supposed to provide these examples, we use a pseudo-relevance feedback approach [12]. In general, current retrieval technology is capable of providing high quality results in terms of relevance. Therefore, we can assume that most of the very first returned results, tend to be relevant for the query. In contrast, the very last of the results are highly likely to be noisy and irrelevant. For instance, in the case of Flickr’s state-of-the-art retrieval system, results in [36, 37] show that, in average, among the first 50 returned images, at least 37 images are relevant to the query, i.e., 75.37% — estimate obtained for 549 location related queries where the query was formulated as keywords using the location’s name.

Inline with this, we retain the first N_p images from the initial N , ranked images as positive examples and the last N_n images as negative examples (we denote this as *experimental assumption 1*, i.e., the use of first images and last images as positive and negative examples, respectively, to improve the overall system’s performance — see Figure 2). This leads to a total number of N_t examples (i.e., $N_p + N_n$) that constitutes the ground truth. To address the borderline case when the number of returned images is lower than N_t , we adopt the following approximations where the positive examples are a fraction of the number of initial ranking images:

$$N'_p = \frac{N}{N_t} \cdot N_p, \quad N'_n = N - N'_p \quad (1)$$

The immediate advantage of this strategy is in the complete automation of the relevance feedback process. No real user interaction is actually required,

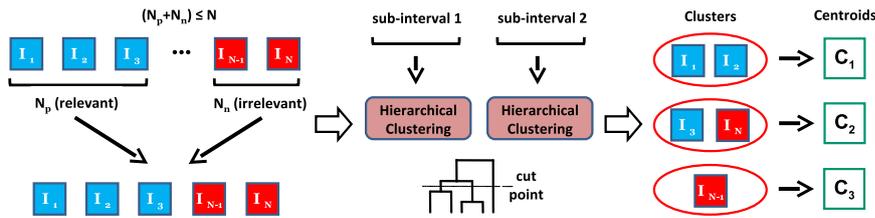


Fig. 2 Example selection and clustering: selection of positive and negative examples (N_p and N_n), N is the total number of returned images; followed by sub-intervals clustering and centroid representation.

which reduces significantly the processing time as well as the need for conducting complex user studies. Moreover, this step has impact over the diversification part of the approach, because improving the relevance will also improve diversification.

3.2 Clustering

Equipped with the ground truth, we use a clustering strategy to group images with similar content. We select a Hierarchical Clustering (HC) [38] scheme that proved highly effective in various diversification scenarios [36, 39]. The HC scheme uses the "bottom up" approach (i.e., agglomerative). The clustering algorithm works as follows. First, each of the images is assigned to an individual cluster. Then, pairs of new formed clusters are merged according to a similarity measure (i.e., distance metric). The process is repeated until it ends with a single, overall, cluster. Besides its low complexity, HC has the advantage of providing a dendrogram of classes by grouping images iteratively based on a certain distance metric. This allows for adapting the number of output classes to the target scenario based on the selection of a cutting point in the dendrogram.

We apply HC only to the selected positive and negative examples, according to the following strategy. All other unselected images are ignored in the following steps. First, we apply a non-overlapping sliding window approach by dividing the selected examples into small sub-intervals of size s (see Figure 2). Preliminary experiments [40] showed that HC separates the data more precisely if we have a small number of images as input, e.g., tens rather than having hundreds. The clustering of less images is better in terms of visual information, i.e., the obtained clusters tend to contain only similar images and images from one cluster are different to images from another cluster (we denote this as *experimental assumption 2*, i.e., the use of sub-interval division improves the overall system's performance).

Second, as shown in Figure 2, the actual HC is applied independently to each of the sub-intervals and the cutoff points of the resulting dendrograms are chosen based on the inconsistency coefficient threshold N_c .

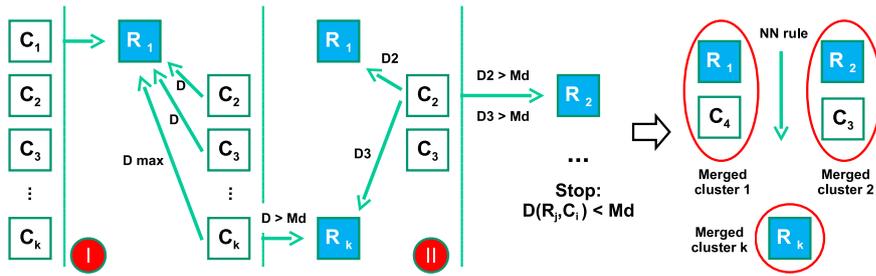


Fig. 3 Merging similar image clusters: C_i is the centroid of cluster i , R_j is the centroid selected as representative, D represents the Euclidean distance between the centroids of the clusters and Md represents the average distance computed between all images and the average image. Centroids are merged according to their similarity.

3.3 Merging

The next step is to merge similar image clusters. Because of the sub-division applied, some resulting clusters may be similar, e.g., images of one cluster formed in the sub-interval W_i may be visually similar to other images in another cluster from sub-interval W_j ($i \neq j$). The key is to merge these clusters into a single one, representing the sub-interval W_{ij} .

To do so, we adapted the folding approach in [5] which diversifies a list of retrieved images by first constructing a set S of representative images and then by assigning the rest of the images to images of set S . The construction of the representative images is done dynamically, starting from the first retrieved image which is by default assigned to set S . Next retrieved image is considered representative if it is sufficiently dissimilar to all the images in this set, i.e., the distance is greater than the mean of distances to the mean image of entire set of all images. In this case, the image is assigned to set S and the process is repeated until there are no representative images left. Each of the remaining images is assigned to a representative image using the Nearest Neighbor approach. This last step represents the construction of clusters around the representative images, so that similar images are grouped in the same cluster.

We apply the same principle for our merging problem. Instead of using a ranked list of images, we use the centroids of each cluster (computed by averaging all images) obtained after HC (see Figure 2). After running the folding approach in [5] we end up with similar centroids grouped together, while the (sufficiently) different centroids will remain in different clusters. Then, all the images corresponding to each centroid are assigned to their new cluster (see Figure 3). Although cluster similarity highly depends on the size of the clusters, this assumption will hold for this case of sufficiently small clusters that compose a small sub-interval. The reason we choose this approach is the power of discrimination of the method. Despite its high complexity, the method is able to separate well the classes [5] and thus to produce very good results.

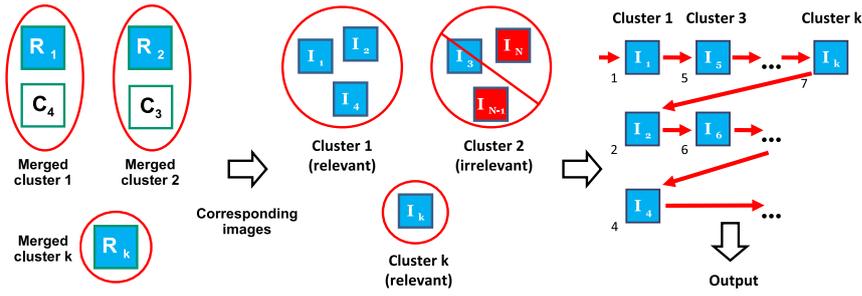


Fig. 4 Cluster pruning and diversification.

Therefore, we use Hierarchical Clustering to cluster a large number of images and then a more complex method with high power of discrimination applied on a reduced number of images. This leads to a good combination in terms of performance and resources consumed.

3.4 Pruning

In the following, we adopt a supplementary pruning step. A cluster is declared irrelevant if it contains only negative examples or if the number of negative examples is higher than of the positive ones, namely:

$$N_n^{(i)} \geq \frac{N^{(i)}}{2} \quad (2)$$

where $N_n^{(i)}$ is the number of negative examples in cluster i and $N^{(i)}$ is the total number of examples in cluster i (we denote this as *experimental assumption 3*, i.e., the adoption of un-relevant clusters improves the overall system's performance — see Figure 4). This assumption is based on the fact that cluster images are supposed to be similar with each other. Therefore, if more than half of the present examples are negative, there is a high probability that all the images are in fact negative examples and were assigned wrongly to the positive category. Experimental results show that this is a good approximation (see Section 5.2). The advantage of this step consists in removing images with a high probability of being un-relevant, and consequently in improving relevance and diversification.

3.5 Diversification

The final step is the actual diversification of the results, which is a round-robin approach. To improve also the relevance, we take into account the initial ranking of the results, as the first retrieved images have a higher probability to be relevant than the last ones. To enforce the diversity, we restrict the output

to contain at least one image from each HC generated cluster. The algorithm is the following.

First, for each of the HC output relevant clusters (the clusters declared as irrelevant are discarded from diversification), the images are sorted according to their initial ranking, so that the first image in a cluster is the one which has the highest rank in the initial retrieval results. Then, considering the order described above and starting with the first cluster, i.e., the class labeled as the first one by the HC scheme, we select as output each cluster first ranked image. This leads to k images, where k is the total number of cluster. The process is repeated iteratively, and clusters are covered again by selecting the second ranked images, third ranked and so on. If in a certain cluster, the number of images is lower than the number of the current iteration (e.g., in the third iteration we attempt to select images from a cluster that has only two images), then that particular cluster is disregarded in the current and further iterations. Finally, the process is repeated until the desired number of images is achieved (see Figure 4).

4 Experimental Setup

4.1 Datasets

To test our approach, we selected two publicly available social image retrieval diversification datasets, namely Div150Multi [41], and Div150Cred [36]:

- *Div150Multi dataset*: contains 292 location related queries (e.g., museums, bridges, parks, monuments) with up to 300 photos per query and associated metadata retrieved from Flickr using Flickr’s default “relevance” algorithm (a total of 86,769 images). Images are annotated for both relevance and diversity by human assessors. In particular, for diversity, images are clustered into similar appearance classes. This dataset proposes a development set containing 153 of the queries (45,575 images) and an evaluation set of 139 queries (41,394 images). The evaluation data contains results obtained with two kind of queries: ‘single-topic’ — only the name of a location is used; and ‘multi-topic’ — combined queries related to events and states associated with locations, e.g., “Oktoberfest in Munich”, “Bucharest in winter”;
- *Div150Cred dataset*: similar to Div150Multi, it contains 153 location related queries with up to 300 photos per query and associated metadata retrieved from Flickr (a total of 45,375 images). Images are also annotated for both relevance and diversity by human assessors and the data is divided into a development set containing 30 of the queries (8,923 images) and an evaluation set of 123 queries (36,452 images).

To be able to compare to the benchmarking results reported on these datasets, for evaluation we used the same experimenting conditions, i.e., the method parameters were optimized on the development data and the actual evaluation is performed on the evaluation set.

4.2 Evaluation metrics

To assess performance, we compute the standard cluster recall at a cutoff at X images ($CR@X$) [6] and the precision at X images ($P@X$), given by:

$$CR@X = \frac{N}{N_{gt}}, \quad P@X = \frac{N_r}{X} \quad (3)$$

where N is the number of image clusters represented in the first X ranked images and N_{gt} is the total number of image clusters from the ground truth (N_{gt} is limited to a maximum of 25 clusters for the Div150Multi [41], and Div150Cred [36] datasets), N_r is the number of relevant images among the first X ranked results. $CR@X$ assesses how many clusters from the ground truth are represented among the top X results provided by the retrieval system. Since clusters are made up of relevant photos only, relevance of the top X results is implicitly measured by $CR@X$, along with diversity. To have a clearer view of relevance, $P@X$ measures the number of relevant photos among the top X results. To account for an overall assessment of both diversity and precision, we also report $F1@X$, i.e., the harmonic mean of $CR@X$ and $P@X$. Results are reported as overall average values over all the queries in the dataset.

4.3 Image content representation

In the clustering process images are represented by content descriptors. Although the approach is not dependent on a certain type of description scheme, the choice of the descriptors influence significantly the results and should be adapted to the specificity of the evaluation data.

Given the specificity of the task, i.e., diversifying social media contents, we tested a broad category of visual and text-based descriptors, which are known to perform well in image retrieval tasks:

- *global color naming histogram* (CN, 11 values): maps colors to 11 universal color names: “black”, “blue”, “brown”, “grey”, “green”, “orange”, “pink”, “purple”, “red”, “white”, and “yellow” [42];
- *global Histogram of Oriented Gradients* (HoG, 81 values): represents the HoG feature computed on 3 by 3 image regions [43];
- *global color moments* computed on the HSV Color Space (CM, 9 values): represent the first three central moments of an image color distribution: mean, standard deviation and skewness [44];
- *global Locally Binary Patterns* (LBP, 16 values) computed on gray scale representation of the image [45];
- *global Color Structure Descriptor* (CSD, 64 values): represents the MPEG-7 Color Structure Descriptor computed on the HMMD color space [46];
- *global statistics on gray level Run Length Matrix* (GLRLM, 44 dimensions): represent 11 statistics computed on gray level run-length matrices for 4 directions: Short Run Emphasis, Long Run Emphasis, Gray-Level Non-uniformity, Run Length Non-uniformity, Run Percentage, Low Gray-Level

- Run Emphasis, High Gray-Level Run Emphasis, Short Run Low Gray-Level Emphasis, Short Run High Gray-Level Emphasis, Long Run Low Gray-Level Emphasis, Long Run High Gray-Level Emphasis [47];
- global descriptor which is obtained by the concatenation of all values and *spatial pyramid representations* of these (denoted with 3x3): each of the previous descriptors is computed also locally; the image is divided into 3 by 3 non-overlapping blocks and descriptors are computed on each patch;
 - *generic convolutional neural network* based descriptors (CNN, 4,096 values): based on the reference CNN model provided along with the Caffe framework, model which is learned with the 1,000 ImageNet classes used during the ImageNet challenge. The descriptors are extracted from the last fully connected layer of the network (named fc7) [24];
 - *adapted convolutional neural network* based descriptor (CNN ad, 4,096 values): based on a CNN model obtained with an identical architecture to that of the Caffe reference model, which is learned with 1,000 tourist points of interest classes whose images were automatically collected from the Web. Similar to CNN generic, the descriptors are extracted from the last fully connected layer of the network (named fc7) [24];
 - histogram representations of *term frequency* (TF), *document frequency* (DF) and *term frequency - inverse document frequency* (TF-IDF) information computed on image metadata (descriptor average size of 500 values) [48];
 - *user image annotation credibility* information (Credibility, 9 values): an estimation of the global quality of tag-image content relationships for a user's contributions [41, 49, 50].

Descriptors were experimented individually or in combination. Fusion is carried out with an early fusion approach preceded by a max-min normalization.

4.4 Pre-filtering

To improve more the relevance of the results, we pass the initial retrieved images through several pre-filtering steps, namely:

- first, we use the Viola-Jones [51] *face detector* to filter out images with persons as the main subject. These images are in general irrelevant for the common user. The output of the filter for an image consists in a number of pairs of coordinates, indicating where it is most likely to have the face of a person. If this number is greater than a threshold T_f , then the image is considered to contain faces and it is removed;
- second, we use an *image blur detector* to remove the out of focus images. Regardless their content, severely blurred images are in general not satisfactory results for a query. We use the aggregation of 10 state-of-the-art blur indicators³, namely: Brenner's indicator, graylevel variance, normal-

³ An implementation of the used blur indicators is available at <http://www.mathworks.com/matlabcentral/fileexchange/27314-focus-measure/content/fmeasure/fmeasure.m>

ized GLV, energy of gradient, thresholded gradient, energy of Laplacian, modified Laplacian, variance of Laplacian, Tenengrad, and sum of wavelet coefficients. An image is rejected if the average of the normalized values is lower than a threshold T_b ;

- finally, in particular for these datasets, we use a *GPS-based filter*. The filter rejects the images that are positioned too far away from the query location, and therefore which cannot be relevant shots for that location. We use a tolerance radius of T_d Kms. For accurate results, distance between GPS coordinates is computed using the Harvesine formula [52]. Images for which the GPS information is not available are passed by the filters.

Limitations: The proposed filters are more or less domain restricted, in particular adapted to the tourism search scenario. However, these filters are a pre-processing step which can be turned on or off if the method is to be applied to other domains or data and not specifically a part of the algorithms.

5 Experimental results

This section presents the experimental validation results on the two datasets. We have conducted the following experiments: Section 5.1 deals with method’s parameter tuning; Section 5.2 deals with the experimentation of the adopted experimental assumptions of the proposed approach; Section 5.3 compares the proposed method to reference relevance feedback approaches from the literature; Section 5.4 situates our results in the context of the state-of-the-art diversification approaches; Section 5.5 concludes the experimentation by discussing several visual diversification examples.

5.1 Parameter tuning

The performance of the proposed approach depends on the choice of several parameters. We run the following experiments on the Div150Multi dataset.

First test consists of tuning the pre-filters to leverage the relevance. We employed filters only on *single – topic* set, as on *multi – topic* set pre-filtering doesn’t improve performance in terms of relevance. For instance, face-based filter is not employed because of the specificity the dataset, for which the majority of the photos represent people attending location-related events (e.g., carnivals, concerts) or representing crowds or groups of demonstrators etc. These photos are considered an exception, thus they are relevant, because people attending such events are natural part of the events themselves. We tested the impact of this type of pre-filtering over the performance in terms of relevance and we noticed that, instead of increasing or preserving it, it is even worse, by significantly reducing the precision. For the blur detector, we didn’t noticed any performance improvement, thus we concluded that there were no blurred photos in the dataset, and regard to the distance-based filter, the dataset was not provided with GPS coordinates, as there were not available

for the majority of the queries (e.g., “rainbow over the lake”, “animals at zoo”, “snow in the city”).

We vary T_b from 0 to 0.6 with a step of 0.1, T_f from 1 to 3 with a step of 1 and T_d from 0 to 5 with a step of 1 (see Section 4.4) so that we obtain maximum precision on the development data. Precision was calculated on the first 200 images retrieved and best performance is achieved by taking $T_b = 0$, $T_f = 1$, and $T_d = 3$, which leads to $P@200 = 0.7472$. Without pre-filtering, the precision is 0.7273, which means that we obtain an improvement of around 2 percentage points.

We also conducted some experiments to see the influence of the parameters for each filter separately, starting from the baseline case when the results are retrieved by Flickr. Thus, we varied T_b from 0 to 1 with a step of 0.1, T_f from 1 to 10 with a step of 1, and T_d from 1 to 10, with a step of 1. Results are presented in Figure 5. The most important observation is that the biggest performance improvement is given by the face-based filter, i.e., around 2 percentage points (no filtering means $T_f = Inf$, because the filter rejects the images with at least T_f faces), followed by distance-based filter - 0.5 percentage points (no filtering means $T_d = Inf$, because the filter rejects the images with distance equal or greater than T_d) and blur detector - 0.2 percentage points (no filtering means $T_b = 1$). We can notice that the overall tendency for each parameter setting is to decrease the performance in terms of precision as the values increase, because they get closer to the baseline case where no filter is employed.

The following test consists of determining the best descriptor - HC parameter combination (i.e., choice of N_p , N_n , N_c — see Section 3.1). In what concerns the descriptors used, we tested individual descriptors and most of the early fusion combinations of the implemented descriptors. For brevity reasons, we restrict the presentation to the use of only one descriptor for each modality and combinations of all of them for each type (e.g., all visual, all textual), as the other combinations didn’t lead to any more significant results. We experiment with varying N_p (number of positive examples) from 100 to 200 with a step of 50, N_n (number of negative examples) from 0 to 20 with a step of 10, and N_c (which is related to the number of diversity classes) from 0.7 to 1.3 with a step of 0.2. We use the HC’s standard Euclidean distance metric. We experiment also with various sizes of the image sub-interval, s , namely 10, 20 and 40 images, and we select the value for the best results, which is $s = 20$ images. We use this value in the next experiments. Table 1 presents the results for the optimal configurations (pre-filtering is used only on single-topic data). We report the metrics at a cutoff at 20 images which is the official metric for this dataset. Overall, the best results in terms of $F1@20$ are obtained with CNN descriptors and $N_p = 200$, $N_n = 20$, $N_c = 0.7$, which yielded $F1@20 = 0.5444$. CNN descriptors proved to be very efficient for diversification, maintaining in the same time a good performance for relevance. We use these configurations for the next experiments.

The final experiment was conducted for selecting the best distance metric - cluster merging approach for the HC. For the previously selected param-

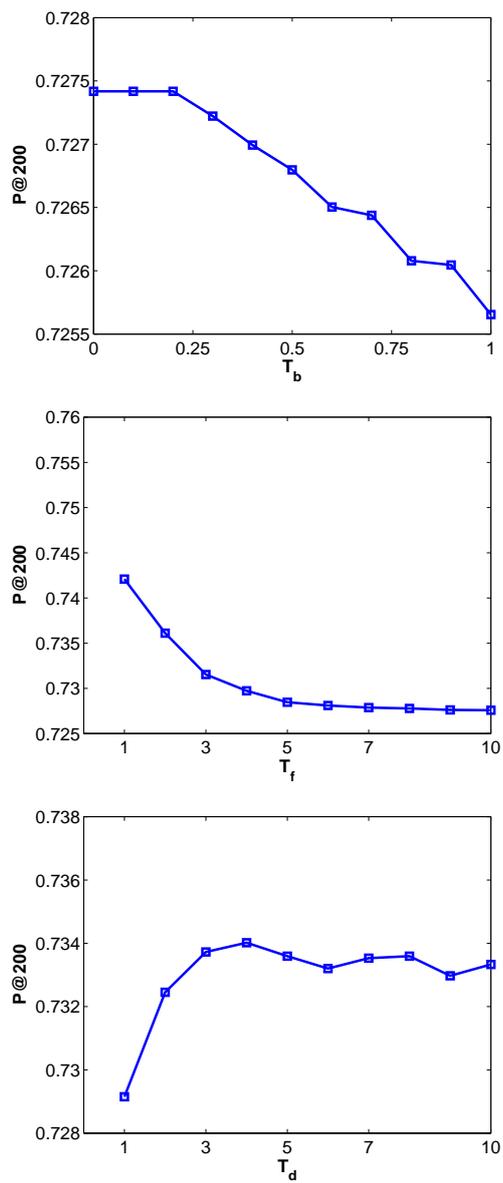


Fig. 5 Influence of the filters' parameters: T_b, T_f, T_d

eters, we experiment now with various distance metrics, namely: Euclidean, sEuclidean, cityblock, Minkowski, Chebyshev, cosine, correlation, Spearman, Hamming, and Jaccard [38]; combined with several modalities of computing distances between clusters, namely: single [53] - shortest distance of all dis-

Table 1 Diversification results for various descriptor - HC parameter combinations (best results are represented in bold).

<i>set</i>	<i>modality</i>	<i>descriptor</i>	$N_p-N_n-N_c$	$P@20$	$CR@20$	$F@20$
Single-topic	Visual	CNN ad	100-20-0.7	0.7065	0.465	0.5435
	Text	Text all	200-20-1.3	0.6819	0.4217	0.5074
	Credibility	Credibility	100-0-1.1	0.7007	0.4404	0.5247
Multi-topic	Visual	CNN all	200-10-0.7	0.7079	0.4666	0.5472
	Text	Text all	250-20-1.3	0.685	0.4365	0.5165
	Credibility	Credibility	280-20-0.9	0.685	0.4313	0.5123
Overall	Visual	CNN ad	100-10-0.7	0.7004	0.4689	0.5444
	Text	TF	200-20-1.3	0.6845	0.4209	0.5057
	Credibility	Credibility	150-20-1.3	0.6838	0.4297	0.5122

Table 2 Diversification results for various HC metrics - merging approaches combinations (best results are represented in bold).

<i>set</i>	<i>distance-merging approach</i>	$P@20$	$CR@20$	$F1@20$
Single-topic	Euclidean-single (default)	0.7065	0.465	0.5435
	Euclidean-median	0.7036	0.459	0.538
	Cosine-centroid	0.713	0.4714	0.5509
Multi-topic	Euclidean-single (default)	0.7079	0.4666	0.5472
	Euclidean-median	0.7029	0.4528	0.5348
	Cosine-single	0.7171	0.4731	0.5533
Overall	Euclidean-single (default)	0.7072	0.4658	0.5454
	Euclidean-median	0.7032	0.4558	0.5364
	Correlation-ward	0.7043	0.4708	0.5471

tances between each image of one cluster and images of another; complete [54] - furthest distance of all distances; average - unweighted average distance, the resulting distance is the average of all distances between images of clusters; centroid - centroid distance which is the distance between the centroids of two clusters; median - the resulting distance is the median distance of all distances between images of two clusters; Ward [55] - inner squared distance based on minimum variance algorithm; and weighted - weighted average distance, the resulting distance is a combination of all distances between images of each cluster. In all cases, the minimum resulting distance criteria is used to merge two clusters and form a larger one. The best combination results are reported in Table 2. Overall, the highest $F1@20$ is achieved for Correlation and Ward centroid, $F1@20 = 0.5471$.

We also experiment the influence of the parameters N_p , N_n and N_c over the performance of the proposed approach. Thus, we start from the best performance configuration presented in Table 1 and vary separately each parameter while keeping constant the other ones, i.e., we vary N_p from 10 to 280 with a step of 10, N_n from 0 to 20 with a step of 5 and N_c from 0 to 1.3 with a step of 0.1. Results are reported for $F1$ metric score and presented for each set in Figure 6. As one can see, selecting the number of classes has the biggest

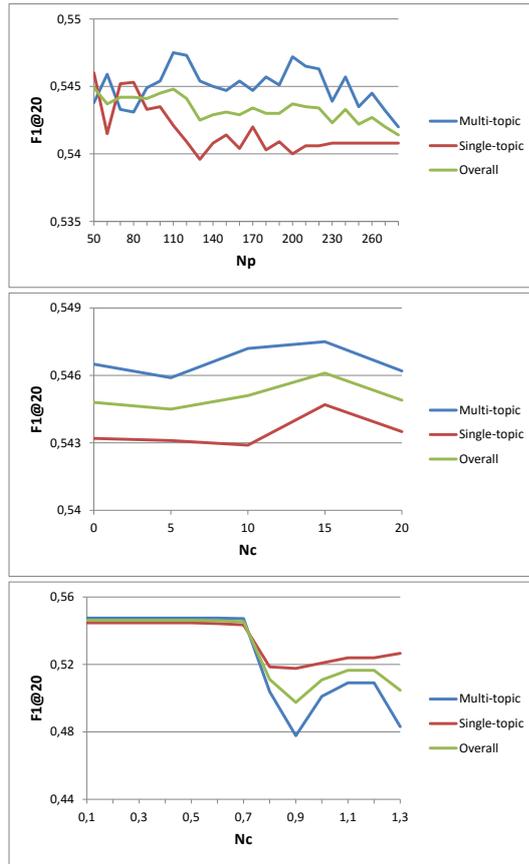


Fig. 6 Influence of the number of the proposed method's parameters: N_p, N_n, N_c

influence over the overall performance, i.e., a variation of around 6 percentage points for *Multi-topic* set and around 3 percentage points for *Single-topic* set. The performance is approximately constant for N_c from 0.1 to 0.7, then it decreases abruptly for greater values. The second most important parameter is the number of positive examples followed by the number of negative examples. Another interesting observation is related to the number of positive examples. While for the *Multi-topic* set a greater number means better performance, on the *Single-topic* set a greater number of positive examples affects negatively the overall performance. This could be an indicator that within the *Single-topic* set only the first 100 images are in general truly relevant, according to the ground truth, while for *Multi-topic* this number is higher.

5.2 Experimental assumptions

In this section we conduct several experiments in order to assess the advantages of some of the main assumptions of the proposed approach. Experiments are performed on the Div150Multi dataset which is the biggest of the two and using previous best performance parameter tuning. Of course, a proof of these experimental assumptions would require a rigorous significance test of the results. However, this preliminary validation shows promising results.

Experimental assumption 1: the adoption of the positive and negative examples (see Section 3.1). We test the impact of taking as negative examples the very last of the returned images. Therefore, for the best results, we set $N_n = 0$ (no negative examples). In this case we achieve $F1@20 = 0.5432$ on *single – topic* data and $F1@20 = 0.5465$ on *multi – topic* data. This is lower than the use of negative examples by around 0.7 percentage points in both cases. Another test is to assess the number of un-relevant images in the set of negative examples. This can be an indicator of how important is the number of negative examples. For 20 negative examples, 19.37 of them are in average un-relevant in the case of *Single – topic* set, while for *Multi – topic* set 9.95 out of 10 images are un-relevant, which means a ratio of more than 95 percent in both cases.

Experimental assumption 2: the adoption of interval subdivision (see Section 3.2). We test the usefulness of dividing the set of the retrieved images into small sub-intervals and the aggregation of similar clusters. We took into consideration two scenarios: (i) the data is split and HC is applied for each sub-interval, then we merge the similar clusters; (ii) HC is applied on all retrieved images. In the second case, we obtained $F1@20 = 0.5582$ on *single – topic* data and $F1@20 = 0.5381$ on *multi – topic* data. The adoption of the interval division leads to a slightly lower performance for *single – topic* data (e.g., around 0.5 percentage points lower $F1@20$) and a significant improvement for the *multi – topic* data, i.e., an increase of 1.5 percentage points ; which overall is better than the scenario without this step.

Experimental assumption 3: the adoption of irrelevant classes in HC (see Section 3.4). To test the usefulness of building the irrelevant classes and thus removing them from the results, we experiment by considering all the classes as relevant. This yields an $F1@20 = 0.541$ on *single – topic* data and $F1@20 = 0.546$ on *multi – topic* data, which is almost 1, respectively 0.7 percentage points lower than the result with removing irrelevant classes.

Diversification improvement: we test the benefits in terms of solely diversification of the proposed method. To do so, we consider the ideal case when the selected positive examples are relevant (selected from the relevance ground truth). This way, we enforce to see the diversification capabilities. First, we consider as baseline the case when all images are relevant and they are retrieved according to their initial rank from Flickr and no diversification is employed. The results in this case are presented in Table 3. Then, we evaluate the performance of the proposed method using the relevance ground truth,

Table 3 Performance of the proposed method in the ideal case when the positive examples are selected from the relevance ground truth (descriptors were combine using early fusion approach).

<i>set</i>	<i>modality</i>	<i>descriptor</i>	N_p - N_n - N_c	$P@20$	$CR@20$	$F@20$
Single-topic	Baseline	-	-	0.9884	0.5018	0.6
	Visual	CNN all	300-0-0.9	0.9884	0.5951	0.7294
	Text	Text all	300-0-1.3	0.9884	0.5289	0.6744
	Credibility	Credibility	300-0-1.3	0.9884	0.5615	0.7016
Multi-topic	Baseline	-	-	0.9843	0.4893	0.6294
	Visual	CNN ad	300-0-0.9	0.9843	0.6219	0.7478
	Text	Text all	300-0-0.7	0.9843	0.5694	0.7028
	Credibility	Credibility	300-0-0.7	0.9843	0.5862	0.7138
Overall	Baseline	-	-	0.9863	0.4955	0.6396
	Visual	CNN ad	300-0-0.9	0.9863	0.6031	0.7341
	Text	Text all	300-0-0.7	0.9863	0.5495	0.6885
	Credibility	Credibility	300-0-0.9	0.9863	0.5687	0.706

i.e., the positive examples are restricted to be genuine relevant. The idea behind this experiment is to compare the improvement in terms of diversification given by the proposed method when the relevance remains constant. In this case, we select all images which are relevant in their initial order and feed them to the proposed method. For the image selection we set $N_p = 300$ and $N_n = 0$, i.e., all images are positive examples, and for HC we varied N_c from 0.7 to 1.3 with a step of 0.2. In terms of modality, we employed all types of descriptors: visual, textual and credibility (see Section 4.3), or combinations of them, using early fusion. The results are presented in Table 3. By analyzing the results we can observe that the proposed method improves diversification on all sets, regardless of the modality used. Overall, in terms of $F1@20$, we achieved the lowest value of 0.6885 using text descriptors, which is still almost 4.9 percentage points higher compared to baseline, and the highest value of 0.7341 using CNN descriptors, which is almost 9.4 percentage points higher compared to baseline. This means that the proposed method achieves better diversification over the baseline, when the input is the same.

5.3 Comparison to other relevance feedback approaches

In this experiment we compare our results against other relevance feedback approaches from the literature, namely: Rocchio [27], Relevance Feature Estimation [28] (RFE) and some classification-based approaches: Support Vector Machines (SVM) [29], AdaBoost [30], Random Forests [31] and k-Nearest Neighbors [56]. User relevance feedback is simulated with the images' ground truth in a window of 10 images (which is a typical browsing scenario). We experimented with two situations: (i) feedback is simulated with the relevance ground truth (*relevance*); (ii) feedback is simulated with the diversity ground truth by selecting diverse images in the initial feedback window (*diversity*) —

Table 4 Comparison to relevance feedback approaches on Div150Multi dataset [41](best results are represented in bold).

<i>set</i>	<i>RF Approach</i>	<i>feedback</i>	<i>descriptor</i>	<i>P@20</i>	<i>CR@20</i>	<i>F@20</i>
Single-topic	Proposed	pseudo-rel.	CNN ad	0.713	0.4714	0.5509
	Rocchio [27]	relevance	Text all-CNN all	0.7906	0.4467	0.5563
	Rocchio [27]	diversity	Text all-CNN all	0.7768	0.4618	0.5656
	RFE [28]	relevance	Vis all-CNN all	0.7572	0.4262	0.5335
	RFE [28]	diversity	Text all-CNN all	0.7029	0.4569	0.5419
	SVM [29]	relevance	Text all-CNN all	0.821	0.4432	0.5603
	SVM [29]	diversity	Text all-CNN all	0.7594	0.4533	0.5513
	Random Forests [31]	relevance	LBP3x3	0.8036	0.393	0.506
	Random Forests [31]	diversity	LBP3x3	0.7413	0.4315	0.5287
	AdaBoost [30]	relevance	CNN ad	0.6812	0.4352	0.5128
	AdaBoost [30]	diversity	CNN ad	0.7022	0.4292	0.511
	kNN [56]	relevance	CNN ad	0.7906	0.4569	0.5622
	kNN [56]	diversity	CNN all	0.6935	0.4198	0.5053
Multi-topic	Proposed	pseudo-rel.	CNN all	0.7171	0.4731	0.5533
	Rocchio [27]	relevance	Text all-CNN all	0.8064	0.4471	0.5584
	Rocchio [27]	diversity	Text all-CNN all	0.7521	0.4744	0.5704
	RFE [28]	relevance	Cred-CNN all	0.7379	0.4571	0.5528
	RFE [28]	diversity	Vis all-CNN all	0.7514	0.4616	0.5574
	SVM [29]	relevance	Cred-CNN all	0.7607	0.4615	0.557
	SVM [29]	diversity	CNN all	0.9036	0.4594	0.5926
	Random Forests [31]	relevance	GLRLM	0.8086	0.4024	0.5134
	Random Forests [31]	diversity	HOG	0.7457	0.47	0.5532
	AdaBoost [30]	relevance	CNN ad	0.7279	0.4641	0.551
	AdaBoost [30]	diversity	CNN gen	0.7136	0.4564	0.54
	kNN [56]	relevance	CNN ad	0.825	0.4905	0.5976
	kNN [56]	diversity	CNN gen	0.7364	0.4475	0.5443
Overall	Proposed	pseudo-rel.	CNN ad	0.7043	0.4708	0.5471
	Rocchio [27]	relevance	Text all-CNN all	0.7986	0.4469	0.5578
	Rocchio [27]	diversity	Text all-CNN all	0.7644	0.4681	0.568
	RFE [28]	relevance	Vis all-CNN all	0.7629	0.4319	0.5369
	RFE [28]	diversity	Vis all-CNN all	0.745	0.45	0.5474
	SVM [29]	relevance	Cred-CNN all	0.7439	0.451	0.5485
	SVM [29]	diversity	CNN ad	0.9025	0.432	0.5688
	Random Forests [31]	relevance	GLRLM	0.7802	0.3902	0.4999
	Random Forests [31]	diversity	HOG	0.7432	0.4376	0.5316
	AdaBoost [30]	relevance	CNN ad	0.7047	0.4498	0.532
	AdaBoost [30]	diversity	CNN ad	0.7187	0.4393	0.5249
	kNN [56]	relevance	CNN ad	0.8079	0.4638	0.58
	kNN [56]	diversity	CNN all	0.7165	0.4276	0.5208

this should allow for more emphasis on the diversification. All the approaches were tuned to best performing parameters (parameters - descriptors). Descriptors were combined using early fusion approach.

First, we compare to the 2015 MediaEval Retrieving Diverse Social Images benchmarking [41] (Div150Multi dataset). For the proposed approach we use the combination of parameters which yielded best performance, i.e., descriptor- N_p - N_n - N_c -distance-merging approach. These are: for *single - topic* CNN ad-

Table 5 Comparison to relevance feedback approaches on Div150Cred dataset [36] (RBF - Radial Basis Function kernel; best results are represented in bold).

<i>RF approach</i>	<i>feedback</i>	<i>descriptor</i>	<i>P@20</i>	<i>CR@20</i>	<i>F1@20</i>
proposed	pseudo-rel.	Vis all-Text all-Credibility	0.7907	0.4555	0.5707
Rocchio [27]	relevance	CN	0.8549	0.3385	0.4718
Rocchio [27]	diversity	CSD	0.7126	0.3429	0.455
RFE [28]	relevance	CN	0.828	0.3239	0.4526
RFE [28]	diversity	CN	0.787	0.3561	0.4773
SVM RBF [29]	relevance	GLRLM	0.8508	0.369	0.505
SVM RBF [29]	diversity	all visual	0.75	0.4086	0.5172
AdaBoost [30]	relevance	GLRLM	0.8077	0.3666	0.4934
AdaBoost [30]	diversity	LBP	0.7463	0.3779	0.4935

100-20-0.7-cosine-centroid, for *multi – topic* CNN all-200-10-0.7-cosine-single, and *overall* CNN ad-100-10-0.7-correlation-ward. Results are presented in Table 4. The first observation is the fact that the best results are obtained for all methods using CNN descriptors or a combination of them, except for Random Forests. In terms of F1 metric score, best result ($F1@20 = 0.5976$) is achieved on *multi – topic* set, using Support Vector Machines (SVM) [29] with linear kernel, using diversity ground truth. The same method performs best on *overall* set ($F1@20 = 0.5688$) using the same type of Ground Truth. On *single – topic* set, best performance is achieved using Rocchio’s algorithm with diversity ground truth ($F1@20 = 0.5656$). Another observation is that, overall, best performance is achieved using the diversity ground truth, $F1@20 = 0.5976$. An explanation for this can be that in the first retrieved images there are many relevant and diverse images, and the SVM-RF algorithm only retrieve on the first position the same images, because they have the greatest confidence score.

In terms of relevance, the proposed method doesn’t perform best on any of the sets, and this can be a consequence to the fact that there are images which are not relevant in the first retrieved images and the only major improvement is the use of the three filters. However, in terms of diversification, the proposed method performs best on *single – topic* and *overall* sets, with $CR@20 = 0.4714$ and $CR@20 = 0.4708$, which is an improvement of almost 1 and 0.3, respectively, percentage points over the best RF method, i.e., Rocchio ($CR@20 = 0.4618$ and $CR@20 = 0.4681$ respectively). On *multi – topic* set, the proposed method is on the third place, with $CR@20 = 0.4731$, after SVM and Rocchio approaches. The good diversification is achieved due to the fact that the main steps of the proposed method consist in diversification strategies: interval division, Hierarchical Clustering, merging, and diversification.

We conclude that although the method doesn’t perform best in terms of relevance — which is somehow expected given the automated feedback assumption — it achieves good performance in terms of diversification and overall it is capable to achieve comparable/better results than other, user based, RF techniques. It has the main advantage that the user ground-truth is simulated automatically and no human presence is required. Moreover, the RF

techniques were simulated with the ideal ground truth, which gives maximum performance. However, in real case, human feedback is not always accurate and reliable due to time constraints, which gives more weight to the achieved results.

Second, we compare to the 2014 MediaEval Retrieving Diverse Social Images benchmarking [57] (Div150Cred dataset). The best performance parameter combination in this case is visual all-textual all-credibility-180-0-0.7-cityblock-single. Results are presented in Table 5. On this dataset one can observe that the use of diversified feedback instead of only relevance allows for improvement over the last one. However, regardless the use of actual image ground truth, the best traditional relevance feedback result in terms of $F1@20$ is 0.5172, achieved with SVM and Radial Basis Function (RBF) kernel. This is around 5 percentage points less than the proposed approach. Best result is also achieved in terms of visual diversification, $CR@20 = 0.4555$, which is almost 5 percentage points greater than the best performing RF technique on this dataset, SVM-RBF. In terms of relevance, the proposed method is on the 5th place, which is a good result if we take into account that it is focused mainly on the diversification improvement. As for the previous results, these results are very promising considering the fact that the proposed approach uses automatically generated feedback.

5.4 Comparison to other diversification approaches

In this experiment, we compare the results against state-of-the-art diversification approaches from the literature.

First, we compare to the 2015 MediaEval Retrieving Diverse Social Images benchmarking [41] (Div150Multi dataset). For the fairness of the comparison, we present only the approaches that use similar techniques (e.g., relevance feedback, hierarchical clustering), namely: UNED [58] — uses relevance human feedback, which is the input to several Local Logistic Regression models that estimate the probability for a certain image to belong to a relevant set; Recod [59] — uses a geographic-based, face-based, and blur-based pre-filtering strategies followed by a re-ranking and a k-Medoids-based diversification step; MIS [60] — uses face-based and distance-based filters to remove irrelevant images, followed by a distance-based re-ranking and an Agglomerative Hierarchical Clustering approach for diversification; ETH-CVL [61] — uses relevance ground truth provided for development data to train a generic SVM on CNN features for relevance scoring.

Results are presented in Table 6. Overall, the proposed approach achieves the second best result in terms of $F1@20$, after MIS [60] — $F1@20 = 0.5473$, with $F1@20 = 0.5471$. However, we achieve the best diversification performance on *single - topic* and *multi - topic* sets, with $CR@20 = 0.4714$, $CR@20 = 0.4731$ respectively. Results show that, overall, methods achieving the highest relevance are not necessarily the ones with the highest diversification, e.g., UNED [58] has $P@20 = 0.7766$ compared to $P@20 = 0.7043$

Table 6 Comparison to diversification approaches from 2015 MediaEval Retrieving Diverse Social Images benchmarking [41] (best results are represented in bold).

<i>set</i>	<i>Approach</i>	<i>pre-filtering</i>	<i>modality</i>	<i>P@20</i>	<i>CR@20</i>	<i>F1@20</i>
Single-topic	proposed	yes	CNN	0.713	0.4714	0.5509
	UNED [58]	no	visual-textual	0.7645	0.4194	0.524
	Recod [59]	yes	credibility	0.763	0.4301	0.539
	MIS [60]	yes	CNN	0.7935	0.4546	0.5595
	ETH-CVL [61]	no	visual-textual	0.842	0.442	0.5674
Multi-topic	proposed	yes	CNN	0.7171	0.4731	0.5533
	UNED [58]	no	visual-textual	0.7886	0.4491	0.5519
	Recod [59]	yes	visual	0.735	0.4221	0.5133
	MIS [60]	yes	CNN	0.7636	0.4354	0.5353
	ETH-CVL [61]	no	visual	0.6829	0.4622	0.5333
Overall	proposed	yes	CNN	0.7043	0.4708	0.5471
	UNED [58]	no	visual-textual	0.7766	0.4344	0.538
	Recod [59]	yes	credibility	0.7198	0.4309	0.5219
	MIS [60]	yes	CNN	0.7784	0.445	0.5473
	ETH-CVL [61]	no	visual	0.6853	0.4724	0.5453

achieved with the proposed approach, but their diversification is lower. In terms of modality, exploiting the CNN information allows for the best performance. The use of pre-filtering techniques do improve performance in certain cases, but is not strictly necessary, e.g., ETH-CVL [61] achieves $F1@20 = 0.5453$ without any filtering.

Second, we compare to the 2014 MediaEval Retrieving Diverse Social Images benchmarking [57] (Div150Cred dataset). For the same reason of fairness of the comparison, we present only the approaches that use similar techniques (e.g., relevance feedback, hierarchical clustering), namely: UNED [62] — uses Hierarchical Agglomerative Clustering to group together the images according to latent topics related to the images, detected with Formal Concept Analysis; Recod [63] — uses a geographic-based, and a face-based pre-filtering strategies followed by a re-ranking and a k-Medoids-based diversification step; MIS [64] — uses a distance-based reranking, followed by an Adaptive Hierarchical Clustering approach for diversification; PeRCeiVe@UNICT [65] — uses a faced-based and a GPS-based pre-filtering to remove irrelevant images, followed by Random Forest predictors to define a dissimilarity measure between images, a Cluster-filtering and a Hierarchical Clustering-based ranking.

Results are presented in Table 7. The proposed approach achieves best result in terms of $F1@20$, with $F1@20 = 0.5707$, which is around 2 percentage points higher than the method from MediaEval 2014, UNED [62] — $F1@20 = 0.5502$. We also achieve best diversification performance with $CR@20 = 0.4555$ and best relevance with $P@20 = 0.7907$.

Table 7 Comparison to diversification approaches from 2014 MediaEval Retrieving Diverse Social Images benchmarking [57] (best results are represented in bold).

Approach	pre-filtering	modality	P@20	CR@20	F1@20
proposed	yes	Vis all-Text all-Credibility	0.7907	0.4555	0.5707
UNED [62]	no	textual	0.7772	0.4343	0.5502
RECOD [63]	yes	Vis-Cred	0.7598	0.4288	0.5423
MIS [64]	no	visual-textual	0.6732	0.3982	0.4949
PeRCeiVe@UNICT [65]	yes	textual	0.7553	0.3902	0.5063

Flickr initial ranking**Proposed method****Fig. 7** Visual Example for query 'Christ Church Cathedral' in Missouri from *single – topic* set: first two lines represent first 10 retrieved images from Flickr and last two lines represent the first 10 images refined with the propose approach. The order of the images is from top to bottom and from left to right.

5.5 Visual example

The final experiment consists on discussing some visual examples of the diversification results. We present two examples, first with query 'Christ Church Cathedral' in Missouri from *single – topic* set and the second one with query 'Pingxi Sky Lantern Festival' from *multi – topic* set of the Div150Multi dataset. In both examples, we take the first 10 retrieved images, considering two cases:

Flickr initial ranking



Proposed method



Fig. 8 Visual Example for query 'Pingxi Sky Lantern Festival' from *multi – topic* set: first two lines represent first 10 retrieved images from Flickr and last two lines represent the first 10 images refined with the propose approach. The order of the images is from top to bottom and from left to right.

(i) images are retrieved in their initial order from Flickr, which is the baseline and (ii) images are refined by the proposed approach.

The first example is depicted in Figure 7. Images are retrieved from left to right and top to bottom, according to their rank. We have in the first case $P@20 = 0.4$ and $CR@20 = 0.32$, which means that relevance and diversification are low. We can also observe in the figure that the number of relevant images is low and consequently diversification is poor. With the proposed approach we obtain more relevant images and we can observe that they are also visually diverse. In this case $P@20 = 0.8$ and $CR@20 = 0.44$, which is a visible improvement over the baseline case.

For the baseline of the second example, images are relevant, but more than half of them are redundant, as it can be observed in Figure 8. With Flickr initial ranking, precision is $P@20 = 0.95$ and cluster recall $CR@20 = 0.2778$. By refining the results with the proposed method, we obtain $P@20 = 0.8$ and $CR@20 = 0.6667$, which is a significant improvement over the initial ranking.

This can also be observed in the figure, where images are more diverse in terms of visual information.

6 Conclusions

In this article we addressed the problem of social image search result diversification from the perspective of relevance feedback techniques. We proposed a novel perspective that renders the feedback process completely automatic via pseudo-relevance feedback and considers in priority the diversification, in addition to the relevance of the results. The method operates on top of an existing retrieval system.

Experimental validation on Flickr data (datasets from the 2014 and 2015 MediaEval Retrieving Diverse Social Images tasks) show the potential of this approach. It has superior or at least comparable performance in terms of visual diversification than traditional state-of-the-art relevance feedback methods, but its big advantage is the automatic process of selecting images as feedback, which doesn't require a human annotator. Moreover, the proposed approach achieves similar or better performance than other, more complex, state-of-the-art diversification approaches from the literature. It allows in particular to achieve better diversification of the results. We therefore proved the benefits of the pseudo-relevance assumption in the context of result diversification opening new perspectives for this area of research.

Future work will mainly address exploring more complex diversification scenarios, such as the ones involving more general concept queries where relevance and diversification are relative to each concept in particular. This is a scenario more closer to a general purpose retrieval system. However, the lack of experimentation data is a practical issue for tackling this.

Acknowledgements This work has been funded by the Ministry of European Funds through the Financial Agreement POSDRU 187/1.5/S/155420 PROSCIENCE, and by the ESF POSDRU/159/ 1.5/S/132395 InnoRESEARCH programme. We acknowledge also the MediaEval Benchmarking Initiative for Multimedia Evaluation (<http://www.multimediaeval.org>) for providing the data.

References

1. B. Ionescu, A. Popescu A.-L. Radu, H. Müller, "Result Diversification in Social Image Retrieval: A Benchmarking Framework," in *Multimedia Tools and Applications*, pp. 1–31, 2014.
2. A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, R. Jain, "Content-based Image Retrieval at the End of the Early Years," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1349–1380, 2000.
3. R. Datta, D. Joshi, J. Li, J. Z. Wang, "Image Retrieval: Ideas, Influences, and Trends of the New Age," *ACM Comput. Surv.*, vol. 40, no. 2, pp. 1–60, 2008.
4. R. Priyatharshini, S. Chitrakala, "Association Based Image Retrieval: A Survey," in *Mobile Communication and Power Engineering*, Springer Communications, Computer and Information Science, vol. 296, 2013, pp. 17–26.

5. R. H. van Leuken, L. Garcia, X. Olivares, R. van Zwol, "Visual Diversification of Image Search Results," in Proc. of the 18th Int. Conf. on World Wide Web, Madrid, Spain, Apr. 20-24, 2009, pp. 341–350.
6. M. L. Paramita, M. Sanderson, P. Clough, "Diversity in Photo Retrieval: Overview of the ImageCLEF Photo Task 2009," in Multilingual Information Access Evaluation II. Multimedia Experiments, 2009, pp. 45–59.
7. S. Rudinac, A. Hanjalic, M.A. Larson, "Generating Visual Summaries of Geographic Areas Using Community-Contributed Images," IEEE Trans. on Multimedia, vol. 15, no. 4, pp. 921–932, 2013.
8. B. Taneva, M. Kacimi, G. Weikum, "Gathering and Ranking Photos of Named Entities with High Precision, High Recall, and Diversity," in Proc. of the 3rd ACM International Conference on Web Search and Data Mining, New York, USA, Feb. 03-06, 2010, pp. 431–440.
9. Y. Yang, F. Nie, D. Xu, J. Luo, Y. Zhuang, Y. Pan, "A Multimedia Retrieval Framework based on Semi-Supervised Ranking and Relevance Feedback", IEEE Trans. on Pattern Analysis and Machine Intelligence, vol. 34, no. 4, pp. 723–742, 2012.
10. S. Jones, L. Shao, "Content-based Retrieval of Human Actions from Realistic Video Databases," Information Sciences, vol. 236, pp. 56–65, 2013.
11. X. Y. Wang, B. B. Zhang, H. Y. Yang, "Active SVM-based Relevance Feedback using Multiple Classifiers Ensemble and Features Reweighting," Engineering Applications of Artificial Intelligence, vol. 26, no. 1, pp. 368–381, 2013.
12. G. Cao, J. Y. Nie, J. Gao, S. Robertson, "Selecting Good Expansion Terms for Pseudo-Relevance Feedback," ACM SIGIR Conf. on Research and Development in Information Retrieval, Singapore, Jul. 20-24, 2008, pp. 243–250.
13. L. Jiang, T. Mitamura, S. I. Yu, A. G. Hauptmann, "Zero-example Event Search Using Multimodal Pseudo Relevance Feedback," in Proc. of Int. Conf. on Multimedia Retrieval, Glasgow, U.K., Apr. 01-04, 2014, pp. 297–304.
14. P. Over, G. Awad, M. Michel, J. Fiscus, G. Sanders, W. Kraaij, A.F. Smeaton, G. Queenot, "TRECVID 2013 — An Overview of the Goals, Tasks, Data, Evaluation Mechanisms and Metrics," TRECVID Workshop, NIST, USA, 2013. <http://www-nlpir.nist.gov/projects/tvpubs/tv13.papers/tv13overview.pdf>.
15. A. Garcia Seco de Herrera, J. Kalpathy-Cramer, D. Demner Fushman, S. Antani, H. Müller, "Overview of the ImageCLEF 2013 Medical Tasks," in Working Notes of Cross Language Evaluation Forum (CLEF), Valencia, Spain, Sep. 23-26, 2013.
16. X. Zhu, A. Goldberg, J. V. Gael, D. Andrzejewski, "Improving Diversity in Ranking using Absorbing Random Walks," Human Language Technologies: The Annual Conf. of the North American Chapter of the Association for Computational Linguistics, Rochester, NY, USA, Apr. 22-27, 2007, pp.97–104.
17. E. Vee, U. Srivastava, J. Shanmugasundaram, P. Bhat, S. A . Yahia, "Efficient Computation of Diverse Query Results," in Proc. of 24th IEEE Int. Conf. on Data Engineering, Cancun, Mexic, Apr. 07-12, 2008, pp. 228–236.
18. V. Dang, W. B. Croft, "Diversity by Proportionality: An Election-based Approach to Search Result Diversification," in Proc. of the 35th Int. ACM SIGIR Conf. on Research and Development in Information Retrieval, Portland, Oregon, USA, Aug. 12-16, 2012, pp. 65–74.
19. M. R. Vieira, H. L. Razente, M. C. N. Barioni, M. Hadjieleftheriou, D. Srivastava, C. Traina Jr., V. J. Tsotras, "On Query Result Diversification," in Proc. of IEEE Int. Conf. on Data Engineering, Hannover, Germany, Apr. 11-16, 2011, pp. 1163–1174.
20. D.-T. Dang-Nguyen, L. Piras, G. Giacinto, G. Boato, F. De Natale, "A Hybrid Approach for Retrieving Diverse Social Images of Landmarks," in Proc. of IEEE Int. Conf. on Multimedia and Expo (ICME), Torino, Italy, Jun. 29-Jul. 03, 2015, pp. 1–6.
21. L. Cao, R. Ji, W. Liu, Y. Gao, L. Y. Duan, C. Men, "Weakly Supervised Topic Grouping of YouTube Search Results," in Proc. of the 19th IEEE Int. Conf. on Image Processing (ICIP), Coronado Springs - Disney World Orlando, FL, USA, Sep. 30-Oct. 03, 2012, pp. 2885–2888.
22. M. I. Sarac, Pinar Duygulu, "Bilkent-RETINA at Retrieving Diverse Social Images Task of MediaEval 2014," in Proc. of the MediaEval 2014 Workshop, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.

23. D. C. G. Pedronette, R. da S. Torres, R. T. Calumby, "Using Contextual Spaces for Image Re-ranking and Rank Aggregation," in *Multimedia Tools and Applications*, vol. 69, no. 3, pp. 689–716, 2014.
24. E. Spyromitros-Xioufis, S. Papadopoulos, A. L. Gînscă, A. Popescu, I. Kompatsiaris, I. Vlahavas, "Improving Diversity in Image Search via Supervised Relevance Scoring," in *Proc. of ACM Int. Conf. on Multimedia Retrieval*, Shanghai, China, Jun. 23-26, 2015, pp. 323–330.
25. A. Ksibi, A. B. Ammar, C. B. Amar, "Adaptive Diversification for Tag-based Social Image Retrieval," *Int. Journal of Multimedia Information Retrieval*, vol. 3, no. 1, pp. 29–39, 2014.
26. M. Soleymani, M. Larson, "Crowd-sourcing for Affective Annotation of Video: Development of a Viewer-reported Boredom Corpus," *SIGIR Workshop on Crowd-sourcing for Search Evaluation*, Geneva, Switzerland, Jul. 23, 2010.
27. J. Rocchio, "Relevance Feedback in Information Retrieval," in *The Smart Retrieval System Experiments in Automatic Document Processing*, Prentice Hall, Englewood Cliffs, NJ, USA, 1971, pp. 313–323.
28. Y. Rui, T. Huang, S.-F. Chang, "Image Retrieval: Current Techniques, Promising Directions and Open Issues," in *Visual Communication and Image Representation*, vol. 10, no. 1, pp. 39–62, 1999.
29. S. Liang, Z. Sun, "Sketch Retrieval and Relevance Feedback with Biased SVM Classification," in *Pattern Recognition Letters*, vol. 29, no. 12, pp. 1733–1741, 2008.
30. J. Yu, Y. Lu, Y. Xu, N. Sebe, Q. Tian, "Integrating Relevance Feedback in Boosting for Content-based Image Retrieval," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Honolulu, Hawaii, USA, Apr. 15-20, 2007, pp. 965–968.
31. A. Bosch, A. Zisserman, X. Munoz, "Image Classification Using Random Forests and Ferns," in *Proc. of IEEE Int. Conf. on Computer Vision*, Rio de Janeiro, Brazil, Oct. 14-21, 2007, pp. 1–8, ISSN 1550-5499.
32. J. Carbonell, R. F. Y. Yang, R.D. Brown, Y. Geng, D. Lee, "Translingual Information Retrieval: A Comparative Evaluation," in *Proc. of 15th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, NAGOYA, Aichi, Japan, Aug. 23-29, 1997, pp. 708–715.
33. F. Denis, "PAC Learning from Positive Statistical Queries," in *Proc. of 9th Int. Conf. on Algorithmic Learning Theory*, Lecture Notes in Artificial Intelligence, vol. 1501, pp. 112–126, Springer - Verlag, 1998.
34. B. Liu, W.S. Lee, P.S. Yu, X. Li, "Partially Supervised Classification of Text Documents," in *Proc. of 19th International Conference on Machine Learning*, Sydney, Australia, Jul. 08-12, 2002, pp. 387–394.
35. R. Yan, A. Hauptmann, R. Jin, "Multimedia Search with Pseudo-Relevance Feedback," in *Proc. of Int. Conf. on Image and Video Retrieval (CIVR)*, Urbana, Illinois, USA, Jul. 24-25, 2003, pp. 238–247.
36. B. Ionescu, A. Popescu, M. Lupu, A. L. Gînscă, B. Boteanu, H. Müller, "Div150Cred: A Social Image Retrieval Result Diversification with User Tagging Credibility Dataset," in *ACM Multimedia Systems (MMSys)*, Portland, Oregon, USA, Mar. 18-20, 2015, pp. 207–212.
37. B. Ionescu, A.-L. Radu, M. Menéndez, H. Müller, A. Popescu, B. Loni, "Div400: A Social Image Retrieval Result Diversification Dataset," in *ACM Multimedia Systems (MMSys)*, Singapore, Mar. 19-20, 2014, pp. 29–34.
38. A. K. Jain, M. N. Murty, P. J. Flynn, "Data Clustering: a Review," in *ACM computing surveys (CSUR)*, vol. 31, no. 3, pp. 264–323, 1999.
39. D. Cai, X. He, Z. Li, W.-Y. Ma, J.-R. Wen, "Hierarchical Clustering of WWW Image Search Results using Visual, Textual and Link Information," in *Proc. of the 12th Annual ACM Int. Conf. on Multimedia*, New York, NY, USA, Oct. 10-16, 2004, pp. 952–959.
40. B. Boteanu, I. Mironică, B. Ionescu, "Hierarchical Clustering Pseudo-Relevance Feedback for Social Image Search Result Diversification," in *Proc. of 13th Int. Workshop on Content-Based Multimedia Indexing CBMI 2015*, Prague, Czech Republic, Jun. 10-12, 2015, pp. 1–6.
41. B. Ionescu, A. L. Gînscă, B. Boteanu, A. Popescu, M. Lupu, H. Müller, "Retrieving Diverse Social Images at MediaEval 2015: Challenge, Dataset and Evaluation," in *Proc. of the MediaEval 2015 Workshop*, Wurzen, Germany, Sep. 14-15, 2015, CEUR-WS.org, ISSN 1613-0073, 1436, <http://ceur-ws.org/Vol-1436/>.

42. V. de Weijer, C. Schmid, J. Verbeek, D. Larlus, "Learning Color Names for Real-world Applications," *IEEE Trans. on Image Processing*, vol. 18, no. 7, pp. 1512–1523, 2009.
43. O. Ludwig, D. Delgado, V. Goncalves, U. Nunes, "Trainable Classifier-Fusion Schemes: An Application To Pedestrian Detection," in *Proc. of the 12th Int. IEEE Conf. on Intelligent Transportation Systems*, St. Louis, Missouri, USA, Oct. 04-07, 2009, pp. 1–6.
44. M. Stricker, M. Orengo, "Similarity of Color Image," in *IS&T/SPIE's Symposium on Electronic Imaging: Science & Technology*, 1995, pp. 381–392.
45. T. Ojala, M. Pietikinen, D. Harwood, "Performance Evaluation of Texture Measures with Classification based on Kullback Discrimination of Distributions," in *Proc. of the 12th IAPR Int. Conf. on Pattern Recognition*, vol. 1, Oct. 09-13, 1994, pp. 582–585.
46. B. S. Manjunath, J. R. Ohm, V. V. Vasudevan, A. Yamada, "Color and Texture Descriptors," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 11, no. 6, pp. 703–715, 2001.
47. X. Tang, "Texture Information in Run-Length Matrices," *IEEE Trans. on Image Processing*, vol. 7, no. 11, pp. 1602–1609, 1998.
48. H. C. Wu, R. W. P. Luk, K.F. Wong, K.L. Kwok, "Interpreting TFIDF Term Weights As Making Relevance Decisions," *ACM Trans. on Information Systems*, vol. 26, no. 3, pp. 1–37, 2008.
49. A. Popescu, "CEA LIST's Participation at MediaEval 2013 Placing Task," in *Proceedings of the MediaEval 2014 Workshop*, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.
50. A. Sun, S. S. Bhowmick, "Image Tag Clarity: in Search of Visual-Representative Tags for Social Images," in *Proc. of the 1st SIGMM Workshop on Social Media*, ACM, Beijing, China, Oct. 19-24, 2009, pp. 19–26.
51. P. Viola, M. J. Jones, "Robust Real-Time Face Detection," *Int. Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
52. G. R. van Brummelen, "Heavenly Mathematics: The Forgotten Art of Spherical Trigonometry," Princeton University, 2013.
53. P. H. A. Sneath, R. R. Sokal, "Numerical Taxonomy," Freeman, London, UK, 1973.
54. B. King, "Step-wise Clustering Procedures," in *Journal of the American Statistical Association*, vol. 62, no. 317, pp. 86–101, 1967.
55. J. H. Ward Jr, "Hierarchical Grouping to Optimize an Objective Function," in *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 236–244, 1963.
56. G. Giacinto, "A Nearest-Neighbor Approach to Relevance Feedback in Content-Based Image Retrieval," in *Proc. of the 6th ACM Conf. on Image and Video Retrieval*, Amsterdam, Netherlands, Jul. 09-11, 2007, pp. 456–463.
57. B. Ionescu, A. Popescu, M. Lupu, A. L. Gincă, H. Müller, "Retrieving Diverse Social Images at MediaEval 2014: Challenge, Dataset and Evaluation," in *Proc. of the MediaEval 2014 Workshop*, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.
58. A. Castellanos, X. Benavent, A. G.-Serrano, E. de Ves, J. Cigarrn, "UNED-UV @ Retrieving Diverse Social Images Task," in *Proc. of the MediaEval 2015 Workshop*, Wurzen, Germany, Sep. 14-15, 2015, CEUR-WS.org, ISSN 1613-0073, 1436, <http://ceur-ws.org/Vol-1436/>.
59. R. T. Calumby, I. B. A. do C. Araujo, V. P. Santana, J. A. V. Munoz, O. A. B. Penatti, L. T. Li, J. Almeida, G. Chiachia, M. A. Goncalves, R. da S. Torres, "Recod @ MediaEval 2015: Diverse Social Images Retrieval," in *Proc. of the MediaEval 2015 Workshop*, Wurzen, Germany, Sep. 14-15, 2015, CEUR-WS.org, ISSN 1613-0073, 1436, <http://ceur-ws.org/Vol-1436/>.
60. M. Zaharieva, L. Diem, "MIS @ Retrieving Diverse Social Images Task 2015," in *Proc. of the MediaEval 2015 Workshop*, Wurzen, Germany, Sep. 14-15, 2015, CEUR-WS.org, ISSN 1613-0073, 1436, <http://ceur-ws.org/Vol-1436/>.
61. S. S. Ravindranath, M. Gygli, L. V. Gool, "ETH-CVL @ MediaEval 2015: Learning Objective Functions for Improved Image Retrieval," in *Proc. of the MediaEval 2015 Workshop*, Wurzen, Germany, Sep. 14-15, 2015, CEUR-WS.org, ISSN 1613-0073, 1436, <http://ceur-ws.org/Vol-1436/>.
62. A. Castellanos, A. Garcia-Serrano, J. M. C. Recuero, "UNED @ Retrieving Diverse Social Images Task," in *Proc. of the MediaEval 2014 Workshop*, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.

63. R. T. Calumby, V. P. Santana, F. S. Cordeiro, O. A. B. Penatti, L. T. Li, G. Chiachia, R. da S. Torres, "Recod @ MediaEval 2014: Diverse Social Images Retrieval," in Proc. of the MediaEval 2014 Workshop, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.
64. M. Zaharieva, P. Schwab, "A Unified Framework for Retrieving Diverse Social Images," in Proc. of the MediaEval 2014 Workshop, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.
65. C. Spampinato, S. Palazzo, "PeRCeiVe Lab@UNICT at MediaEval 2014 Diverse Images: Random Forests for Diversity-based Clustering," in Proc. of the MediaEval 2014 Workshop, Barcelona, Spain, Oct. 16-17, 2014, CEUR-WS.org, ISSN 1613-0073, 1263, <http://ceur-ws.org/Vol-1263/>.