

The medGIFT Group in ImageCLEFmed 2012

Alba G. Seco de Herrera, Dimitrios Markonis, Ivan Eggel, Henning Müller

University of Applied Sciences Western Switzerland (HES-SO)
Sierre, Switzerland
alba.garcia@hevs.ch

Abstract. This article presents the participation of the medGIFT group in ImageCLEFmed 2012. Since 2004, the group has participated in the medical image retrieval tasks of ImageCLEF each year. There are three types of tasks for ImageCLEFmed 2012: modality classification, image-based retrieval and case-based retrieval. The medGIFT group participated in all three tasks. MedGIFT is developing a system named ParaDISE (Parallel Distributed Image Search Engine), which is the successor of GIFT (GNU Image Finding Tool). The alpha version of ParaDISE was used to run most of the experiments in the competition.

Results show that our approach using Bag-of-Visual-Words (BoVW), Bag-of-Colors (BoC) and Lucene for the image captions is the best run for mixed modality classification. The same approach is also the best for the image retrieval task in terms of bpref. In the case-based retrieval task, the Lucene baseline is the best run in terms of mean average precision (MAP). We were the only group presenting mixed and visual runs in these tasks.

1 Introduction

ImageCLEF is the cross-language image retrieval track¹ of the Cross Language Evaluation Forum (CLEF). ImageCLEFmed is part of ImageCLEF focusing on medical images [1–5]. The medGIFT² research group has participated in ImageCLEFmed since 2004. MedGIFT is currently developing ParaDISE (Parallel Distributed Image Search Engine), which is the successor of GIFT³ (GNU Image Finding Tool). The alpha version of ParaDISE was used to run most of the experiments.

The Bag-of-Visual-Words (BOVW) and Bag-of-Colors (BOC) approaches were used for visual retrieval and the textual baseline is based on the open source Lucene⁴ system. In ImageCLEFmed 2011, the BoVW approach was applied using the Scale Invariant Feature Transform (SIFT) [6] but color information was not used. In 2012, there are two main novelties in our system: first, we introduce the BoC [7] descriptor that represents local image colors. We combine

¹ <http://www.imageclef.org/>

² <http://medgift.hevs.ch/>

³ <http://www.gnu.org/software/gift/>

⁴ <http://lucene.apache.org/>

BoC with BoVW and textual descriptors in order to yield better results. Second, we use training set expansion strategies since for some of the image categories only very few annotated examples were available. This significantly improves classification performance.

The widely used BoVW [8] method is applied as follows: a training set of images is chosen and a number of local descriptors are extracted from each image of this set. The descriptors are then clustered using a clustering method (such as k-means or DENCLUE). The centroids of the clusters are treated as visual words that represent the specific local patterns. Hence, we have a visual-word vocabulary describing all types of local image patterns. Local features are then also extracted from each image in the database and mapped onto the visual-word vocabulary. An image is then represented as a histograms of the visual-word occurrences, the BoVW used as feature vector in the classification task. When an image is queried, a similarity measure (such as histogram intersection) is used to compare the query image histogram and the database image histogram, providing a similarity score.

The BoC representation is analogous to the BoVW representation and to the Bag-of-Words representation of text documents. As the two methods are quite complementary and as the representations as histograms are very similar, the two approaches can easily be fused for medical image classification.

In Section 2, we describe the datasets and the techniques used. We evaluate the runs submitted to the ImageCLEFmed 2012 benchmark in Section 3. Finally, conclusions are presented in Section 4.

2 Datasets and Techniques

This section describes the basic techniques used in ImageCLEFmed 2012 by the medGIFT group. For the the majority of the runs, the alpha version of ParaDISE (Parallel Distributed Image Search Engine) developed by the medGIFT group is used. This image search engine is built on top of Hadoop [9], a MapReduce [10] paradigm of parallel computing and the Cassandra ⁵ DBMS (Database Management System). This allows for computational and storage scalability. The ParaDISE component-based architecture is designed to simplify the plugin of different image features and representations. Apart from the 3 baseline visual runs that were also given publicly to the participants, we submitted 10 runs (2 textual, 4 visual and 2 mixed runs) to the image-based retrieval task, 8 runs (2 textual, 4 visual and 2 mixed runs) to the case-based retrieval task and 8 runs (4 visual and 4 mixed runs) to the modality classification task.

2.1 Image Collection

We used the database provided by ImageCLEFmed 2012. The database contains over 300,000 images of 75'000 articles of the biomedical open access literature. This is a subset from the PubMed Central⁶ database containing over one million

⁵ <http://cassandra.apache.org/>

⁶ <http://www.ncbi.nlm.nih.gov/pmc/>

images. This set of articles contains all articles in PubMed that are open access but the exact copyright for redistribution varies among the journals. A more detailed description of the ImageCLEFmed 2012 setup is given in [4].

2.2 Textual Techniques

For text retrieval the standard settings of the Apache Lucene text retrieval system are used. The documents containing the journal texts are cleaned off their XML elements and only the remaining text is used.

Two indexations were done for ImageCLEF 2012: (1) the full text of all articles was indexed on an article basis, and (2) the captions were indexed on a figure basis. In the past it was shown that for case-based retrieval an indexation of the full text had best results whereas for image-based retrieval the caption text delivered much better results.

2.3 Visual Techniques

The baseline for the visual description of the image is the Bag-of-Visual-Words (BoVW) approach. In this approach, local SIFT descriptors are extracted from each image. Then, the descriptors of the image are quantized, assigning each one to its nearest neighbour from a fixed set of local descriptors, called “visual vocabulary”. The image is then represented by a histogram of the frequency of the “visual words”. Similarity between two images can be quantified using a distance metric to measure the distance between the two histograms.

In our runs, the SIFT implementation in the `fiji`⁷ image processing package was used for the extraction of the local descriptors as in our participation in ImageCLEFmed 2011. In order to create the visual vocabulary, our implementation of the density-based clustering algorithm DENCLUE [11] was used. The reasons for this choice are the features and the nature of the data set that needs to be clustered. The data set to be clustered is large (1,000 training images produce approximately 2'500'000 descriptors) and high dimensional (SIFT descriptors are 128-dimensional). The DENCLUE algorithm is highly efficient for clustering large-scale data sets, can detect arbitrarily shaped clusters and handles outliers and noise well. Moreover, opposed to other density-based clustering algorithms it performs well for high-dimensional data. However, when using a density-based clustering algorithm care needs to be taken for data sets containing clusters of different densities. To deal with this, the parameter ξ that controls the significance of the candidate cluster in respect to its density was set to zero. The same visual vocabulary used in last year's participation [12] was also used for the creation of the bags of visual words.

Based on the BoVW approach, the Bag-of-Colors (BoC) [7] is an image description technique introduced in [13]. Similarly to the BoVW, the technique uses a color vocabulary C previously learned on a sub set of the collection to represent the image. A color vocabulary $C = \{c_1, \dots, c_{k_c}\}$, with $c_i = (L_i, a_i, b_i) \in$

⁷ <http://fiji.sc/wiki/index.php/Fiji>

CIELab is constructed by finding the most frequently occurring colors in each image of a sub set of the collection. The modality classification training set was used for the creation of the color vocabulary. The CIE (International Commission on Illumination) 1976 L*a*b (CIELab) space was used, being a perceptually uniform color space. CIELab is a space defined by L for luminance and a, b for the color-opponent dimensions for chrominance [14, 15]. The BoC of an image I is defined as a vector $h_{BoC} = \{\bar{c}_1, \dots, \bar{c}_k\}$ such that, for each pixel $p_k \in I \forall k \in \{1, \dots, n_p\}$, with n_p being the number of pixels of the image I :

$$\bar{c}_i = \sum_{k=1}^{n_p} \sum_{j=1}^{n_p} g_j(p_k) \forall i \in \{1, \dots, k_c\}$$

where

$$g_j(p) = \begin{cases} 1 & \text{if } d_\varepsilon(p, c_j) \leq d_\varepsilon(p, c_l) \forall l \in \{1, \dots, k_c\} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

From experiments on the modality classification task of ImageCLEFmed 2012, $k_c = 100$ and $k_{vw} = 238$ were chosen as the sizes of the color and visual vocabulary, respectively. For the comparison of the images the histogram intersection [16] is used as distance measure.

2.4 Fusion Techniques

In image retrieval and classification tasks, often different features, systems and results can be combined to deliver improved results. Moreover, multiple query images can describe in more detail the visual characteristics to be retrieved. In these cases a fusion strategy needs to be used. Two main fusion strategies exist, early and late fusion. In the early fusion, the vectors of the features or the systems are merged into a single vector. For the early fusion of multiple positive and/or negative query images, Rocchio's algorithm can be used:

$$\mathbf{q}_m = \alpha \mathbf{q}_o + \beta \frac{1}{|D_r|} \sum_{\mathbf{d}_j \in D_r} \mathbf{d}_j - \gamma \frac{1}{|D_{nr}|} \sum_{\mathbf{d}_j \in D_{nr}} \mathbf{d}_j \quad (2)$$

where α, β and γ are weights, \mathbf{q}_m is the modified query, \mathbf{q}_o is the original query, D_r is the set of relevant images and D_{nr} is the set of non-relevant images. In our scenario there is a lack of non-relevant images, so only the second term of the right part of the equation is used. This algorithm can be applied only for vectors of the same feature spaces so it is not applicable for the fusion of different visual features or retrieval systems, in general.

In the late fusion, the retrieval results of the features or the systems are fused. Two main categories of late fusion techniques exist, the score-based and rank-based methods. In 2009, the ImageCLEF@ICPR fusion task was organized to compare late fusion techniques using the best ImageCLEFmed visual and textual results [17]. Studies such as [18] show that combSUM (3) and combMNZ(4), score-based methods, proposed by [19] in 1994 are robust fusion strategies. With

the data from the ImageCLEF@ICPR fusion task, combMNZ performed slightly better than combSUM but the difference was small and not statistically significant.

$$S_{\text{combSUM}}(i) = \sum_{k=1}^{N_k} S_k(i) \quad (3)$$

$$S_{\text{combMNZ}}(i) = F(i) * S_{\text{combSUM}}(i) \quad (4)$$

where $F(i)$ is the frequency of image i being returned by one input system with a non-zero score, and $S(i)$ is the score assigned to image i .

In general, rank-based fusion (RRF) worked better than score-based fusion. The reciprocal rank fusion [20] is a simple fusion method based on ranks(5).

$$RRF_{\text{score}}(d \in D) = \sum_{r \in R} \frac{1}{k + r(d)} \quad (5)$$

where, D is the set of documents retrieved, R is the set of rankings of the documents and $k = 60$.

3 Results

This section details the techniques that were used to produce the runs for ImageCLEFmed 2012 and then evaluates the runs.

3.1 Image and Case Retrieval Techniques

Two strategies were compared for the fusion of the multiple queries, early fusion and late fusion (see Section 2.4). For the fusion of visual features and the fusion of visual and textual systems, late fusion was applied. Score-based and rank-based fusion techniques were compared. To summarize, fusion was used in three cases:

- fusing multiple visual features to produce visual runs;
- fusing textual and visual runs to produce mixed runs;
- fusing vectors (early fusion) of query images which belong to the same topic or their results (late fusion).

In the image retrieval task, as the fulltext search retrieved articles instead of images, an article-to-image mapping was used. If an image was contained in multiple articles, only the article with the highest score was taken into account giving its score to the image. If multiple images were contained in the same article, all the images received the common article’s score.

The same strategy was applied in the case-based task, because the image search retrieved images instead of articles. The article received the score of the best scored image that it contained. If multiple articles contained the same image they all received the common image’s score.

Table 1 contains a summary of the techniques used, while Tables 2 and 3 show the details of the submitted runs.

Table 1. Overview of the techniques used for the retrieval.

Name	Technique
BoVW	Content-based image search using Bag-of-Visual-Words representation
BoC	Content-based image search using Bag-of-Colors representation
Full text	Lucene textual search, searching into the full text of the articles
Captions	Lucene textual search, searching into the captions of images

Table 2. Techniques used for the image-based runs.

Run ID	Queries fusion	Techniques	Techniques fusion	Run Type
ir1	Reciprocal	BoVW	n/a	Visual
ir2	combMNZ	BoVW	n/a	Visual
ir3	combMNZ	BoVW + BoC	combMNZ	Visual
ir4	Reciprocal	BoVW + BoC	Reciprocal	Visual
ir5	Rocchio	BoVW + BoC	combMNZ	Visual
ir6	Rocchio	BoVW + BoC	Reciprocal	Visual
ir7	n/a	Full text	n/a	Textual
ir8	n/a	Captions	n/a	Textual
ir9	combMNZ	BoVW + BoC + Full text + Captions	combMNZ	Mixed
ir10	Reciprocal	BoVW + BoC + Full text + Captions	Reciprocal	Mixed

Table 3. Techniques used for the case-based runs.

Run ID	Queries fusion	Techniques	Techniques fusion	Run Type
cr1	Reciprocal	BoVW	n/a	Visual
cr2	combMNZ	BoVW	n/a	Visual
cr3	combMNZ	BoVW + BoC	combMNZ	Visual
cr4	Reciprocal	BoVW + BoC	Reciprocal	Visual
cr5	Rocchio	BoVW + BoC	combMNZ	Visual
cr6	Rocchio	BoVW + BoC	Reciprocal	Visual
cr7	n/a	Full text	n/a	Textual
cr8	n/a	Captions	n/a	Textual
cr9	combMNZ	BoVW + BoC + Full text + Captions	combMNZ	Mixed
cr10	Reciprocal	BoVW + BoC + Full text + Captions	Reciprocal	Mixed

3.2 Modality Classification Techniques

Driven by the good performance of [3] in the ImageCLEFmed 2011 modality classification task, a similar approach was followed by the medGIFT group in 2012. The approach involves automatically expanding the labelled training set to improve the performance of the classification for classes that are poorly represented (e.g. the training set containing very few images of that class). To achieve this, training images are used as queries in the full 300,000 images of the ImageCLEFmed 2012 data set and the l highest ranked retrieved images are added as training images into the class of the query image.

Two methods of expanding the training set were used. In the first expansion technique, s images taken randomly of each class were used as queries. As in the original training set the number of images per class varies from 5 to 50, by choosing $s = 5, l = 20$ we can theoretically obtain a relatively balanced training set (105–150 per class) of 4,100 images. In the second, all the training images were queried, resulting in a larger non-balanced training set. E. g. for $l = 20$ an expanded training set of 21,000 images can be obtained theoretically. In practice, smaller sizes were obtained, mainly because of two reasons: retrieved images that are already contained in the training set were discarded and images retrieved multiple times by query images of different classes were discarded as well.

For a run to qualify as visual, we considered that the expanded training set used in this run needs to be created only by visual means. This means that the queries on the full data set used only visual features for the retrieval. Similarly, this was repeated using mixed (visual and textual) queries for the mixed runs. This resulted in a final number of 5 training sets (2 balanced, 2 non-balanced and the original training set).

A k - nn classifier using weighted voting was used to classify the test images. For the choice of the classifier parameters the results of [7] were taken into account and $k = 11, k = 7$ were used for the visual runs. However, since the non-balanced expanded training set was significantly larger, double the value $k = 14$ was also tested for this case. The inverse of the similarity score of the k - nn images was used to weight the voting.

Table 4 gives the details of the submitted runs.

3.3 Image Retrieval Evaluation

Table 5 displays the results of the medGIFT runs for the image retrieval task. *ir9*, a mixed approach, achieved the highest MAP, GM-MAP and bpref among all our submitted runs in the image-based retrieval task. This result highlights the potential benefits of combining textual and visual features.

For the visual runs, the baseline of BoVW, *ir1* and *ir2*, did not demonstrate good results. However, when BoVW is fused with BoC the results improve significantly. Furthermore, when using RRF for the fusion of techniques the MAP decreases, which contradicts our hypothesis.

Table 4. Modality classification runs

Run ID	Techniques	Fusion Rule	Training Set	k	Run Type
mc1	BoVW	n/a	original	11	Visual
mc2	BoVW + BoC	combMNZ	original	7	Visual
mc3	BoVW + BoC	combMNZ	visual non-balanced	7	Visual
mc4	BoVW + BoC	combMNZ	visual non-balanced	14	Visual
mc5	BoVW + BoC	combMNZ	visual balanced	7	Visual
mc6	BoVW + BoC + Captions	Reciprocal	original	7	Mixed
mc7	BoVW + BoC + Captions	Reciprocal	mixed non-balanced	7	Mixed
mc8	BoVW + BoC + Captions	Reciprocal	mixed non-balanced	14	Mixed
mc9	BoVW + BoC + Captions	Reciprocal	mixed balanced	7	Mixed

Table 5. Image retrieval results

Run ID	Run type	MAP	GM-MAP	bpref	P10	P30
best visual run		0,0101	0,0004	0,0193	0,0591	0,0439
ir1	Visual	0,0016	0	0,0048	0,0273	0,0318
ir2	Visual	0,0017	0	0,0058	0,0227	0,0318
ir3	Visual	0,0049	0,0003	0,0138	0,0364	0,0364
ir4	Visual	0,004	0,0002	0,0103	0,0227	0,0318
ir5	Visual	0,0033	0,0003	0,0133	0,0364	0,0333
ir6	Visual	0,003	0,0001	0,01	0,0273	0,0227
best textual run		0.2182	0.082	0.2173	0.3409	0.2045
ir7	Textual	0.1397	0.0436	0.1565	0.2227	0.1379
ir8	Textual	0.1562	0.0424	0.167	0.3273	0.1864
best mixed run		0.2377	0.0665	0.2542	0.3682	0.2712
ir9	Mixed	0.2005	0.0917	0.1947	0.3091	0.2
ir10	Mixed	0.1167	0.0383	0.1238	0.1864	0.1485

Our best results in early precision (P10 and P30) on the other hand were obtained with a text retrieval approach, *ir8*. This is in contradiction with past results where the best overall results were often textual whereas early precision was often better with combined or visual approaches.

3.4 Case Retrieval Evaluation

Table 6 shows the results of the medGIFT runs for the case based retrieval task. In this task, *cr8* achieved the highest MAP (0,169) among all submitted runs. It demonstrates that a Lucene baseline still obtains very good results with relatively low effort. MedGIFT was the only lab that submitted purely visual and mixed runs for the case-based task. Although the results of our visual runs are lower than textual runs, *cr9*, a mixed approach, performs better than the average of all submitted runs in this task.

Table 6. Case based retrieval results

Run ID	Run type	MAP	GM-MAP	bpref	P10	P30
cr1	Visual	0,0008	0	0	0,0038	0,0013
cr2	Visual	0,0016	0	0,0032	0,0038	0,0013
cr3	Visual	0,0302	0,001	0,0293	0,0231	0,009
cr4	Visual	0,0366	0,0014	0,0347	0,0269	0,0141
cr5	Visual	0,0007	0	0	0	0,0013
cr6	Visual	0,0008	0,0001	0,0007	0	0,0013
second best textual run		0,1508	0,0322	0,1279	0,1538	0,1167
cr7	Textual	0,169	0,0374	0,1499	0,1885	0,109
cr8	Textual	0,0696	0,0028	0,0762	0,0962	0,0615
cr9	Mixed	0,1017	0,0175	0,0857	0,1115	0,0679
cr10	Mixed	0,0514	0,009	0,0395	0,0654	0,0564

As all of the best results are from text retrieval runs it becomes clear that visual techniques need to be used in different ways than was the case for obtaining good results. Most likely a matching of image occurrences in articles in a more complex way would be necessary for obtaining good results for the case-based tasks using visual data.

3.5 Modality Classification Evaluation

Finally, Table 7 presents the classification accuracy of the submitted medGIFT runs for the modality classification task. The runs *c8*, *mc6*, *mc7* achieved the three best accuracies in the mixed run category. The visual runs achieved an average performance with the inclusion of BoC as a global descriptor to improving the classification accuracy. It can be observed that the runs *mc4*, *mc8* using the non-balanced expanded training sets and $k = 14$ are outperforming

Table 7. Modality classification results

Run ID	Visual						Mixed			
	mc1	mc2	mc3	mc4	mc5	best run	mc6	mc7	mc8	mc9
Accuracy (%)	11.1	38.1	41.8	42.2	34.2	69.7	64.2	63.6	66.2	58.8

the runs *mc6*, *mc2* that use the original training set. These runs also perform better than the runs *mc3*, *mc7* that use $k = 7$, confirming our hypothesis that using a larger k can improve results. Moreover, in experiments not submitted as official runs using larger values of k leads to a better performance, with the mixed run reaching an accuracy of 68.5% using the triple value $k = 21$.

4 Conclusions

This article describes the methods and results of the the medGIFT group for the ImageCLEF 2012 medical tasks. We submitted ten runs each for the ad-hoc image retrieval and the modality classification tasks and nine runs for the case-based retrieval task.

In ImageCLEFmed 2012 we concentrated on fusion methods of the visual and textual features and training set expansion strategies. We included the BoC approach, which significantly improves classification performance.

In the image-based retrieval task, our strategy leads to limited results obtained by our submitted runs. However, we still need to do a more detailed analysis to try to understand what is hurting performance in these runs. Our Lucene baseline achieves the best MAP for the case-based retrieval task. Finally, in the modality classification task we submitted the three runs with the best accuracies in the mixed run category and with only one group delivering better results with a visual approach.

Future work will to go beyond in the work of fusion methods, as well as further research in new visual features. Visual information can be used in better ways, for example to classify images by modality and if modality names appear in a query then filtering by this modality. If the same modalities occur in a query and an article this can also give evidence for a higher relevance.

5 Acknowledgments

The research leading to these results has received funding from the European Union’s Seventh Framework Programme under grant agreement 257528 (KHRES-MOI), 249008 (Chorus+) and 258191 (Promise).

References

1. Clough, P., Müller, H., Sanderson, M.: The CLEF cross-language image retrieval track (ImageCLEF) 2004. In Peters, C., Clough, P., Gonzalo, J., Jones, G.J.F.,

- Kluck, M., Magnini, B., eds.: Multilingual Information Access for Text, Speech and Images: Result of the fifth CLEF evaluation campaign. Volume 3491 of Lecture Notes in Computer Science (LNCS)., Bath, UK, Springer (2005) 597–613
2. Müller, H., Deselaers, T., Kim, E., Kalpathy-Cramer, J., Deserno, T.M., Clough, P., Hersh, W.: Overview of the ImageCLEFmed 2007 medical retrieval and annotation tasks. In: CLEF 2007 Proceedings. Volume 5152 of Lecture Notes in Computer Science (LNCS)., Budapest, Hungary, Springer (2008) 473–491
 3. Kalpathy-Cramer, J., Müller, H., Bedrick, S., Eggel, I., Seco de Herrera, A.G., Tsirikas, T.: The CLEF 2011 medical image retrieval and classification tasks. In: Working Notes of CLEF 2011 (Cross Language Evaluation Forum). (September 2011)
 4. Müller, H., Seco de Herrera, A.G., Kalpathy-Cramer, J., Fushman, D.D., Antani, S., Eggel, I.: Overview of the imageclef 2012 medical image retrieval and classification tasks. In: Working Notes of CLEF 2012 (Cross Language Evaluation Forum). (September 2012)
 5. Müller, H., Clough, P., Deselaers, T., Caputo, B., eds.: ImageCLEF – Experimental Evaluation in Visual Information Retrieval. Volume 32 of The Springer International Series On Information Retrieval. Springer, Berlin Heidelberg (2010)
 6. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2) (2004) 91–110
 7. Seco de Herrera, A.G., Markonits, D., Müller, H.: Bag-of-colors for biomedical document image classification. In: MICCAI workshop MCBR-CDS. (September 2012) Forthcoming
 8. Sivic, J., Zisserman, A.: Video google: A text retrieval approach to object matching in videos. In: Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2. ICCV '03, Washington, DC, USA, IEEE Computer Society (2003) 1470–1477
 9. White, T.: Hadoop: The Definitive Guide. O'Reilly Media, Inc. (2010)
 10. Dean, J., Ghemawat, S.: MapReduce: simplified data processing on large clusters. *Communications of the ACM — 50th anniversary issue* **51** (January 2008) 107–113
 11. Hinneburg, A., Keim, D.A.: An efficient approach to clustering in large multimedia databases with noise. In: Conference on Knowledge Discovery and Data Mining (KDD). Volume 5865., AAAI Press (1998) 58–65
 12. Markonits, D., Eggel, I., Seco de Herrera, A.G., Müller, H.: The medGIFT group in ImageCLEFmed 2011. In: Working Notes of CLEF 2011. (2011)
 13. Wengert, C., Douze, M., Jégou, H.: Bag-of-colors for improved image search. In: Proceedings of the 19th ACM international conference on Multimedia. MM '11, New York, NY, USA, ACM (2011) 1437–1440
 14. Sharma, G., Trussell, H.J.: Digital color imaging. *IEEE Transactions on Image Processing* **6**(7) (1997) 901–932
 15. Banu, M., Nallaperumal, K.: Analysis of color feature extraction techniques for pathology image retrieval system, IEEE (2010)
 16. Swain, M.J., Ballard, D.H.: Color indexing. *International Journal of Computer Vision* **7**(1) (1991) 11–32
 17. Müller, H., Kalpathy-Cramer, J.: The ImageCLEF medical retrieval task at icpr 2010 — information fusion to combine visual and textual information. In: Proceedings of the International Conference on Pattern Recognition (ICPR 2010). Lecture Notes in Computer Science (LNCS), Istanbul, Turkey, Springer (August 2010)
 18. Zhou, X., Deppe, A., Müller, H.: Information fusion for combining visual and textual image retrieval. In: 20th IEEE International Conference on Pattern Recognition (ICPR). (August 2010) 1590–1593

19. Fox, E.A., Shaw, J.A.: Combination of multiple searches. In: Text REtrieval Conference. (1993) 243–252
20. Cormack, G.V., Clarke, C.L.A., Büttcher, S.: Reciprocal rank fusion outperforms condorcet and individual rank learning methods. In: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, New York, NY, USA, ACM (2009) 758–759