Overview of the Second Workshop on Medical Content–Based Retrieval for Clinical Decision Support

Adrien Depeursinge[1], Hayit Greenspan[2], Tanveer Syeda–Mahmood[3], Henning Müller[1]

[1]University of Applied Sciences Western Switzerland, Sierre, Switzerland
[2]Tel Aviv University, Israel
[3]IBM Almaden Research Lab, USA
henning.mueller@hevs.ch

Abstract. The second workshop on Medical Content–Based Retrieval for Clinical Decision Support took place at the MICCAI conference in Toronto, Canada on September 22, 2011. The workshop brought together more than 40 registered researchers interested in the field of medical content–based retrieval. Eleven papers were accepted and presented at the workshop. Two invited speakers gave overviews on state–of–the–art academic research and industrial perspectives. The program was completed with a panel discussion on the role of content–based retrieval in clinical decision support. This overview introduces the main highlights and discussions in the workshop, summarizes the novelties and introduces the presented papers, which are provided in these proceedings.

1 Introduction

Medical content–based retrieval has received a large amount of research attention over the past 15 years [1, 2]. Despite large amounts of research and also the availability of benchmarking data sets for medical retrieval [3], there are still only few tests of real systems in clinical practice such as [4]. Image retrieval has always been at the crossing between several sciences, namely image analysis or computer vision, information retrieval and medical informatics. The MICCAI (Medical Image Computing for Computer–Assisted Intervention) conference is a logical target for a workshop on content–based retrieval.

A workshop titled *Content–based Image Retrieval for Biomedical Image Archives: achievements, problems and prospects* focusing on medical image retrieval, took already place in MICCAI 2007 in Brisbane, Australia. The first workshop on *Medical Content–based Retrieval for Clinical Decision Support*, took place at MICCAI 2009 in London, United Kingdom [5]. The motivation was to show that in addressing true clinical challenges, purely visual image retrieval needs to be augmented with additional information, such as text (coming from various sources including image captions or the patient record). Thus, the challenges in content–based retrieval for clinical applications are in fact multimodal in nature and require the combination of state–of–the–art tools for each modality and the fusion across the modalities to obtain optimal results.

The second workshop on Medical Content–Based Retrieval for Clinical Decision Support took place at the MICCAI conference in Toronto, Canada on September 22, 2011. The workshop brought together more than 40 registered researchers. The 2011 workshop web page was set up at[1] to advertise the event. Submissions were requested in the following principle areas of interest:

- data mining of multimodal medical data,
- machine learning of disease correlations from mining multimodal data,
- algorithms for indexing and retrieval of data from multimodal medical databases,
- disease model–building and clinical decision support systems based on multimodal analysis,
- practical applications of clinical decision support using multimodal data retrieval or analysis,
- algorithms for medical image retrieval or classification using the ImageCLEF collection.

A specific goal was also to promote the data sets created in the ImageCLEF[2] challenge [6–8]. Using standard data sets can help identifying well–performing techniques and help measuring performance improvements within and across techniques and systems. Several articles in the workshop used standard data sets, which allows to well judge the obtained performance.

In total, 17 papers were submitted to the workshop. All papers were reviewed by at least three external reviewers from the scientific committee as well as one member from the organization committee. Through this process 11 high-quality papers were selected for oral presentation at the workshop. In addition to the scientific papers, two invited speakers gave insights into their current research directions and projects. One invited speaker was from the academic field (Nicholas Ayache, INRIA, France) and one from industry (Dorin Comaniciu, Siemens research, USA). The workshop finished with a panel that discussed the role of medical content–based retrieval for clinical decision support.

The workshop led to many lively discussions on application areas, technologies (particularly the use of various types of visual words) and future ideas for the current tools. Feedback from the participants was very positive for continuing the workshop seires at future MICCAI conferences.

2 Highlights of the presentations

This section details the main highlights of the workshop by discussing the main novelties presented in the invited presentations and the scientific papers. Papers on computer–aided diagnosis for a specific medical application are presented in Section 2.2 and general management of visual information in large databases in Section 2.3.

[1] http://www.mcbr-cds.org/

[2] http://www.imageclef.org/

2.1 Main novelties presented

The workshop discussed several novel ideas in the field of medical content–based retrieval. Dominant was the application of many types of *visual words* and bag of visual words (BoVW) techniques [9–14]. Key to strong performance is the combination of feature space definition and strong machine learning tools. The optimization of BoVW remains an interesting research topic as the parameters can vary widely depending on the type of application.

The use of 3D data was introduced in the current workshop in several of the works presented [10, 15, 13, 16, 17]. This contrasts with past years, where most approaches were concentrating on retrieval in 2D images. Visual words were shown to be useful in the 3D space.

A key feature in the use of visual words is the need to define interest points or *salient points* [9, 13, 17]. Various techniques can be used — from random sampling to dense sampling (and others). The technique to give best performance needs to be tested for each application. For example, it has been shown that dense sampling is optimal if a specific subset of images is to be analyzed such as lung CTs or radiographs of various anatomic regions. This has been shown to give best results in the ImageCLEF 2009 competition [18], as well as in several of the works presented in the current workshop.

Another presented challenge is the combination of *text and visual retrieval* [15, 12, 19]. *Cross–modality retrieval* is presented in [16]. A less frequently studied subject is the *perceived visual similarity* [9]. This domain is of importance for applications that aim at the integration into clinical practice.

Several papers dealt with *efficiency* questions using a variety of approaches such as *thumbnails* [17], *inverted files* [16] and *random clustering* [14]).

In summary, the works presented covered a large variety of techniques and applications, addressing state–of–the–art topics in medical content–based retrieval. Initial work with defined clinical applications, as well as evaluation methodologies, were presented. Both are of great importance for future acceptance in the medical community.

In the years ahead, we can expect to see medical image retrieval applications integrated into viewing stations and into clinical information systems, as complementary to text analysis tools.

2.2 Computer–aided diagnosis

In [9], André et al. introduce a smart atlas for videos from probe–based confocal laser endomicroscopy (pCLE) using CBIR. The videos are transformed to 2D mosaic images built from the course of the probe [20]. pCLE mosaic images are a good candidate for CBIR, because little experience exists with this relatively novel imaging modality. Therefore, retrieving similar images with attached diagnosis can support both training and diagnosis. The bag–of–visual–words method [21] is used to extract prevailing visual concepts in a feature space spanned by the scale invariant feature transform (SIFT) descriptors [22] using a dense sampling. The similarity between pCLE videos is computed as the χ^2

distance between the histograms of the occurrences of the various visual words (VW) over the whole image. The evaluation of the retrieval performance addresses an important point of CBIR systems for clinical decision support, where retrieving visually similar images with distinct diagnoses provides important cues to assist the decision of the clinicians. Indirect and direct evaluations of the system are performed: the indirect performance is computed as the classification performance based on the pathological classes (e.g., benign, neoplastic), and a direct retrieval performance based on the perceived visual similarity. The latter, which may closely assess the needs of the clinicians, is based on a visual similarity score obtained from endoscopists using an online survey tool. Such acquired ground truth is used to estimate the interpretation difficulty as well as to learn the perceived similarity by "shortening" the distance between instances perceived as very similar. The consistency between VWs and eight visual semantic concepts defined by experts in pCLE is also verified.

Local 3D texture quantification is used by Burner et al. [10] to retrieve high–resolution computed tomography (HRCT) images of the lungs with emphysema and metastases. The system uses a region of interest (ROI) delineated by the user as a query. It then searches for HRCT images containing regions with similar texture properties. The regional texture properties are encoded using bags of VW in a feature space spanned by a multiscale 3D extension of local binary patterns (LBP) [23]. In the database, the ROIs are defined using the supervoxel algorithm [24] to divide the lung parenchyma into homogeneous regions. The distance between ROIs are computed using the diffusion distance [25] between the histograms of VW. The evaluation shows that the proposed method outperforms approaches based on a global similarity measure.

A system for the retrieval of similar ROIs in HRCT images from patients affected by interstitial lung diseases is proposed by Foncubierta et al. in [11]. VWs are used to encode the characteristics of 6 types of lung tissue expressed in a wavelet domain. The influences of two intrinsic parameters of the proposed methods are investigated. First, the importance of scaling parameter of difference of Gaussians (DoG) is studied by varying the number of supplementary intermediate scales when compared to the classical dyadic wavelet transform. The authors show that the scale progression does not have a significant influence on the retrieval performance, since the classical dyadic scheme allows the best performance. The optimal size of the visual vocabulary is also investigated. Results show that for classes with high intra–class variations such as healthy tissue, a high number of VWs is required. As soon as a sufficient number of VWs is reached, the performance remains stable with a slight decrease that can be explained as an effect of the curse of dimensionality.

An image retrieval system based both on visual and text attributes is proposed by Costa et al. in [15] to assist the diagnosis of hepatic lesions in CT. The system uses ROIs defined by the user as queries. The visual attributes consist of grey–level distributions and moments of the Hounsfield units in the ROIs as well as in the whole liver. Text attributes consist of 20 labels defined by clinicians to be relevant for the characterization of liver lesions in CT. The labels are attached to each lesion by two clinical experts. Intrinsic random forests [26] are used to assess the similarity between the lesions by counting the number of times that two instances appear in the same leaves. The proposed approach allows retrieval of similar lesions with high semantics and efficiency.

In [27], Safi et al. propose a computer-aided diagnosis system for the classification of pigmented skin dermoscopic images. In a first step, the global images are segmented into three regions indicating 1) healthy skin, 2) bright parts of the melanoma, and 3) dark parts of the melanoma using energy constraints in the CIE color space representation [28]. Then, visual features are extracted from the region corresponding to the dark parts of the melanoma. These consist of shape, color and texture properties as well as geometric attributes defined as important by the clinicians. A support vector machine (SVM) classifier with a polynomial kernel is used to classify skin regions in a compact feature representation obtained with the prevailing dimensions of principal component analysis (PCA). The methods are achieving an overall good performance on a dataset of 4240 benign and 232 malignant images.

2.3 Visual data management

An image retrieval system based both on visual information and text is used in [12] for the management of a large collection of histological images using the query by example paradigm. Visual content is represented using 500 VWs extracted from densely sampled patches expressed in terms of their discrete cosine transform of the three RGB channels. The occurrences of the VWs are expressed using a matrix $X_v \in \mathbb{R}^{n \times l}$, where n is the number of VWs and l is the number of instances. Text attributes are obtained from expert annotations and are represented as 46 binary attributes that form the matrix $X_t \in \mathbb{R}^{m \times l}$. Reductions of the dimensionality of the feature space are obtained by factorizing the matrices X_v and X_t. A distance based on the scalar product of the vectors with reduced dimensionality is used for retrieval. Several approaches are proposed and compared to obtain a multimodal representation of the images. The first one consists of creating a multimodal matrix from the concatenation of X_v and X_t. The second maps X_v on the factorized representation H of X_t using a linear transform such as $X_v = W_v H$. W_v i s obtained using multiplicative updating rules [29]. A third strategy consists of mapping X_v on X_t as $X_v = W X_t$ directly with $W \geq 0$. A baseline using the histogram intersection of the occurrence of VWs as similarity metric is shown to be outperformed by all multimodal retrieval approaches but the one based on the concatenation of the visual and textual attributes.

The combination of textual and visual attributes for image retrieval is also investigated by Rahman et al. in [19]. The ImageCLEF 2010 dataset of the medical image retrieval task [30] is used, in which each image has text attached in the form of an image caption and text from the scientific journals from which the image belongs. The text attached to images is represented using MeSH terms[3] resulting from a preprocessing step consisting of the removal of stop words. The

[3] `http://www.nlm.nih.gov/mesh/`

vector of text attributes is obtained with the vector space model [31] and the importance of each term is weighted using the $tf - idf$ scheme with local (i.e., at the document level) and global (i.e., document collection) weights. Visual attributes are obtained from an SVM–classification of image patches using color and texture information into a set of 30 visual concepts defined by annotations from experts. The similarity between images is computed as a linear combination of two cosine distance measures from each modality. The results returned for a given query are filtered by modality using a previously trained SVM classifier. The evaluation shows that the combination of text and visual features provides best results only when the modality filter is used.

The importance of the method for detecting salient points for further computation of VW is investigated by Haas et al. in [13]. Three detectors of interest points are compared: points with high response of DoG, a dense sampling using a regular Cartesian grid and centers of mass of ROIs segmented using the superpixel algorithm [24]. For each salient point, SIFT descriptors [22] are used to span the feature space in which the k–means algorithm identifies clusters and their centers as VWs. Inter–image distances computed as the χ^2 distance between the histograms of the occurrences of the various VWs over all salient points. The methods are evaluated on two datasets. The first evaluation uses the ImageCLEF 2009 medical image annotation data set and error evaluation[4] [32]. Error scores show the improved classification performance achieved by superpixel interest points when compared to dense sampling and DoG detectors. The methods are also evaluated on their ability to locate slices of lung CTs based on the position of the retrieved images. The position is determined by the median position value of the ten best retrieved images per query. Again, the superpixel approach outperforms dense sampling and DoG in terms of sum of squared errors of the vertical positions.

Cross–modality retrieval of ROIs consisting of bounding boxes of organs is addressed by Venkatraghavan et al. in [16]. First, a coarse localization of the regions in whole–body CT and magnetic resonance imaging (MRI) scans is obtained using sped–up robust features (SURF) [33] descriptors in Gabor–filtered images to determine the anatomical region (i.e., cranial, thoracic, abdominal, sacro–lumbar and the extremities). This initial localization is then refined using fuzzy uniformity index from texture descriptors obtained with 3D local binary patterns computed from a Gabor–filtered 3D volume. A sliding window search is used to exhaustively index volumetric datasets. Inverted files are used to index organs based on a semantic vocabulary containing 14 terms corresponding to organs (e.g., liver, left lung, right lung). The proposed methods are evaluated on a dataset of CT and MRI images from various body regions in terms of localization errors of the bounding boxes of the organs. A comparison with regression and decision forests shows that the proposed methods allow for a more precise localization of the organs. No quantitative evaluation of the retrieval performance is carried out.

[4] http://www.idiap.ch/clef2009/evaluation_tools/error_evaluation.pdf

Speed efficiency of CBIR is investigated by Donner et al. in [17] by using small versions of 2D and 3D images. 2D images from the ImageCLEF 2009 medical annotation data set [32] are downsampled to 32×32 thumbnails. 3D CT scans are downsampled to $16 \times 16 \times 16$ volumes. Various retrieval approaches are compared. First, PCA is applied to obtain a more compact representation, in which k–nearest–neighbors use an Euclidean distance and kD–trees to speed up the retrieval process. Second, correlations of rigidly aligned thumbnails are used as a distance metric between the images, which is only carried out in the axial plane for 3D volumes. Third, feature vectors obtained from distribution fields [34] are compared using a l_1–norm distance measure. At last, a histogram of oriented gradients (HOG) from SIFT descriptors extracted with either dense sampling or salient points detected with DoGs are used together with a χ^2 distance. The latter proved to outperform the others using the 2D ImageCLEF 2009 dataset. For 3D CT scans, the retrieval performance is evaluated as the distance between the center of the query image and the center of the most similar image. Again, the approach based on the histograms of gradients is providing best performance. No quantitative evaluation of the retrieval speed is carried out but the authors show that good retrieval performance can be achieved using only highly downsampled image thumbnails.

Clustering efficiency for the computation of VWs is addressed by Pauly et al. in [14] by using multiple random partitioning of the feature space based on extreme random subspace projection ferns. Local appearance with local distributions of color/intensity gradient directions are encoded using color/intensity values, LBP and HOG in 17×17 patches extracted at random locations. In this feature space, random ferns [35] perform a hierarchical partitioning with only one decision function per level. This allows to efficiently categorize an instance with a hierarchical set of binary tests. The authors modify the initial algorithms by using random splits at each level. Bags of VWs are obtained by a concatenation of the VWs from all partitions. The approach is evaluated on the modality classification task at ImageCLEF 2010 [30] and shows an improved accuracy and efficiency when compared to the k–means clustering algorithm.

Purely text–based image retrieval is proposed by Mata et al. in [36], where image retrieval tasks from the ImageCLEF 2009 and 2010 collections are used. The text queries are expanded using the MeSH controlled vocabulary of the National Library of Medicine. The queries are initially divided into N–grams. Then, various query expansion strategies are investigated. The first strategy uses cross–referencing functions provided by the MeSH vocabulary, such as *SeeRelatedDescriptor* and *ConsiderAlso*. The former associates the descriptor with other descriptors and the second returns terms having related linguistic roots. The second strategy uses directly the children terms provided by the hierarchical structure of the MeSH tree. The third strategy uses synonyms (so–called “entry terms” by MeSH) to expand the queries. The Lucene[5] text search engine is used to index and retrieve image documents. The results show that whereas the strategy based on the children terms for expansion performs best on the

[5] `http://lucene.apache.org/java/docs/index.html`

2009 dataset, no significant improvement is achieved in terms of mean average precision when compared with the baseline consisting of using the initial query without expansion.

3 Conclusions

The second workshop on content–based medical retrieval for clinical decision support showed a set of state-of-the-art works, from academic as well as industrial perspectives. We have found that the audience increased from previous meetings at the same venue. The domain is gaining interest in the MICCAI community.

One evident trend is the desire to show the contribution of the field to the clinical practice. As such, there is a desire to view and integrate content–based retrieval as one part of larger information access and management systems. Hospital information systems (HIS) and Picture Archival and Communication Systems (PACS) are the backbone of the hospital enterprise. The stored data are the asset for the hospitals and being able to use the knowledge stored in these archives is a key component to case–based reasoning and medical decision–making processes. Image analysis is a single component of this process. Text mining and general textual information retrieval are other components that need to be combined with visual retrieval for decision support in clinical settings. After the lively discussion at this workshop, we plan to continue holding the workshop at future MICCAI conferences, and we hope to see you all there.

4 Acknowledgments

We would like to thank the EU FP7 projects Khresmoi (257528), Promise (258191) and Chorus+ (249008) for their support as well as the Swiss national science foundation with the MANY project (number 205321–130046). A special thank you also belongs to IBM, which supported the workshop. We would also like to thank all reviewers as they helped assuring the high quality of the presented papers.

References

1. Müller, H., Michoux, N., Bandon, D., Geissbuhler, A.: A review of content–based image retrieval systems in medicine–clinical benefits and future directions. International Journal of Medical Informatics **73**(1) (2004) 1–23
2. Tagare, H.D., Jaffe, C., Duncan, J.: Medical image databases: A content–based retrieval approach. Journal of the American Medical Informatics Association **4**(3) (1997) 184–198
3. Hersh, W., Müller, H., Kalpathy-Cramer, J., Kim, E., Zhou, X.: The consolidated ImageCLEFmed medical image retrieval task test collection. Journal of Digital Imaging **22**(6) (2009) 648–655

4. Aisen, A.M., Broderick, L.S., Winer-Muram, H., Brodley, C.E., Kak, A.C., Pavlopoulou, C., Dy, J., Shyu, C.R., Marchiori, A.: Automated storage and retrieval of thin–section CT images to assist diagnosis: System description and preliminary assessment. Radiology **228**(1) (July 2003) 265–270
5. Caputo, B., Müller, H., Mahmood, T.S., Kalpathy-Cramer, J., Wang, F., Duncan, J.: Editorial of miccai workshop proceedings on medical content–based retrieval for clinical decision support. In: Proceedings on MICCAI Workshop on Medical Content–based Retrieval for Clinical Decision Support. Volume 5853 of Lecture Notes in Computer Science (LNCS)., Springer (2009)
6. 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)
7. 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: Working Notes of the 2007 CLEF Workshop, Budapest, Hungary (September 2007)
8. Kalpathy-Cramer, J., Müller, H., Bedrick, S., Eggel, I., de Herrera, A.S., Tsikrika, T.: The CLEF 2011 medical image retrieval and classification tasks. In: Working Notes of CLEF 2011 (Cross Language Evaluation Forum). (September 2011)
9. André, B., Vercauteren, T., Ayache, N.: Content–based retrieval in endomicroscopy: Toward an efficient smart atlas for clinical diagnosis. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
10. Burner, A., Donner, R., Mayerhoefer, M., Holzer, M., Kainberger, F., Langs, G.: Texture bags: Anomaly retrieval in medical images based on local 3D–texture similarity. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
11. Foncubierta-Rodríguez, A., Depeursinge, A., Müller, H.: Using multiscale visual words for lung texture classification and retrieval. In Greenspan, H., Müller, H., Syeda Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
12. Vanegas, J.A., Caicedo, J.C., González, F.A., Romero, E.: Histology image indexing using a non–negative semantic embedding. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. MCBR-CDS 2011, Lecture Notes in Computer Sciences (LNCS) (September 2011)
13. Haas, S., Donner, R., Burner, A., Holzer, M., Langs, G.: Superpixel–based interest points for effective bags of visual words medical image retrieval. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
14. Pauly, O., Mateus, D., Navab, N.: Building implicit dictionaries based on extreme random clustering for modality recognition. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)

15. Costa, M.J., Tsymbal, A., Hammon, M., Cavallaro, A., Sühling, M., Seifert, S., Comaniciu, D.: A discriminative distance learning–based CBIR framework for characterization of indeterminate liver lesions. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
16. Venkatraghavan, V., Ranjan, S.: Semantic analysis of 3D anatomical medical images for sub–image retrieval. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
17. Donner, R., Haas, S., Burner, A., Holzer, M., Bischof, H., Langs, G.: Evaluation of fast 2D and 3D medical image retrieval approaches based on image miniatures. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
18. Avni, U., Greenspan, H., Konen, E., Sharon, M., Goldberger, J.: X–ray categorization and retrieval on the organ and pathology level, using patch–based visual words. IEEE Transactions on Medical Imaging **30**(3) (2011) 733–746
19. Rahman, M.M., Antani, S., Demner-Fushman, D., Thoma, G.R.: Biomedical image retrieval using multimodal context and concept feature spaces. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
20. Vercauteren, T., Perchant, A., Malandain, G., Pennec, X., Ayache, N.: Robust mosaicing with correction of motion distortions and tissue deformations for in vivo fibered microscopy. Medical Image Analysis **10**(5) (2006) 673–692
21. Sivic, J., Zisserman, A.: Video Google: Efficient visual search of videos. In: Toward Category–Level Object Recognition 2006. (2006) 127–144
22. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2) (2004) 91–110
23. Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray–scale and rotation invariant texture classification with local binary patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence **24**(7) (July 2002) 971–987
24. Wildenauer, H., Mičušík, B., Vincze, M.: Efficient texture representation using multi–scale regions. In Yagi, Y., Kang, S., Kweon, I., Zha, H., eds.: Computer Vision – ACCV 2007. Volume 4843 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2007) 65–74
25. Ling, H., Okada, K.: Diffusion distance for histogram comparison. In: Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on. Volume 1. (june 2006) 246–253
26. Breiman, L.: Random forests. Machine Learning **45**(1) (2001) 5–32
27. Safi, A., Baust, M., Pauly, O., Castaneda, V., Lasser, T., Mateus, D., Navab, N., Hein, R., Ziai, M.: Computer–aided diagnosis of pigmented skin dermoscopic images. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)
28. Li, F., Shen, C., Li, C.: Multiphase soft segmentation with total variation and H^1 regularization. Journal of Mathematical Imaging and Vision **37**(2) (2010) 98–111

29. Cichocki, A., Zdunek, R., Amari, S.: New algorithms for non–negative matrix factorization in applications to blind source separation. In: Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on. Volume 5. (may 2006) 621–624
30. Müller, H., Kalpathy-Cramer, J., Eggel, I., Bedrick, S., Said, R., Bakke, B., Jr., C.E.K., Hersh, W.: Overview of the CLEF 2010 medical image retrieval track. In: Working Notes of CLEF 2010 (Cross Language Evaluation Forum). (September 2010)
31. Baeza Yates, R.A., Neto, B.R.: Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1999)
32. Tommasi, T., Caputo, B., Welter, P., Güld, M., Deserno, T.: Overview of the clef 2009 medical image annotation track. In Peters, C., Caputo, B., Gonzalo, J., Jones, G., Kalpathy-Cramer, J., Müller, H., Tsikrika, T., eds.: Multilingual Information Access Evaluation II. Multimedia Experiments. Volume 6242 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2010) 85–93
33. Bay, H., Ess, A., Tuytelaars, T., Gool, L.V.: Speeded–up robust features (surf). Computer Vision and Image Understanding **110**(3) (2008) 346–359
34. Sevilla, L., Learned-Miller, E.: Distribution fields. Technical Report UM–CS–2011–027, Dept. of Computer Science, University of Massachusetts Amherst (2011)
35. Özuysal, M., Calonder, M., Lepetit, V., Fua, P.: Fast keypoint recognition using random ferns. IEEE Transactions on Pattern Analysis and Machine Intelligence **32** (2010) 448–461
36. Mata, J., Crespo, M., Maña, M.J.: Using MeSH to expand queries in medical image retrieval. In Greenspan, H., Müller, H., Syeda-Mahmood, T., eds.: Medical Content-based Retrieval for Clinical Decision Support. Volume 7075 of MCBR-CDS 2011., Lecture Notes in Computer Sciences (LNCS) (September 2011)