Overview of the Third Workshop on Medical Content–Based Retrieval for Clinical Decision Support (MCBR–CDS 2012)

Henning Müller[1,2] and Hayit Greenspan[3]

[1]University of Applied Sciences Western Switzerland (HES–SO), Switzerland
[2]University Hospitals and University of Geneva, Switzerland
[3]University of Tel Aviv, Israel
henning.mueller@hevs.ch

Abstract. The third workshop on Medical Content–based Retrieval for Clinical Decision Support (MCBR–CDS 2012) took place in connection with the MICCAI conference (Medical Image Computing for Computer–Assisted Intervention) in Nice, France on October 1, 2012. This text gives an overview of the invited presentations and also the scientific papers presented at the workshop. In the description of the papers the comments and discussions at the workshop are taken into account, highlighting the current tendencies and scientific merits. The workshop finished with a panel that started with the need of clients of image retrieval software and went on across several other important areas such as the importance of high–quality annotated training and test data sets to advance current research. Such big data sets and a framework for researchers to work on them can have an important impact on the field of image–based decision support in the future.

Keywords: medical image analysis, medical information retrieval, clinical decision support, content–based medical image retrieval

1 Introduction

Visual image retrieval or content–based image retrieval (CBIR) has started in the early 1990s as an increasing amount of image data had become available in digital form and could thus be analyzed visually through computers [10]. In the coming decade much research work happened in the field and image retrieval advanced strongly in terms of architectures and also visual features with many application domains [19].

In the medical field, propositions for image retrieval and its usefulness were made quite early [12, 21]. Still, clinical applications were very rare despite a large amount of research in the field as stated a review article in 2004 [14]. Since then, medical image retrieval research has exploded as data sets have become available in benchmarks [13, 8]. A more recent review article highlights these developments and further current challenges in medical image retrieval [2].

Medical imaging is producing very large amounts of data and although many articles pretend to perform retrieval on an entire PACS (Picture Archival and Communication System) [5], no current system has indexed such large amounts. A report of the European Union estimates that 30% of world storage was occupied in 2010 by medical imaging and that mammographies in the USA alone accounted for 2.5 Petabytes in 2009 [1]. The influence of the analysis of big data can be major for medical image retrieval as the scalability to work with extremely large data sets could enable researchers to tackle rare diseases and really increase diagnosis performance based on learning from existing data.

The workshop Medical Content–Based Retrieval for Clinical Decision Support (MCBR–CDS) was held for the third time in connection with MICCAI in 2012. The workshop received 15 high–quality submissions and also asked two invited speakers to submit a paper to these proceedings. All papers were reviewed by at least three independent reviewers and in addition by one workshop organizer. Based on the review results ten papers were finally accepted for the workshop in addition to the two invited papers. These papers are published in the workshop proceedings of this volume of the Springer Lecture Notes in Computer Science. A panel finalized the workshop program. This panel led to vivid discussions on the role of medical image retrieval in clinical practice and also the sharing of annotated medical data to advance research towards large–scale or *big data*.

Author were able to modify their paper until two weeks after the workshop based on the comments received during the workshop and based on the discussions that took place during the entire day. The workshop presentations included several of the important current research areas including the increasing analysis of multidimensional data, the road towards the use of big data, but also the important topics of using high quality visual features and building real applications based on existing techniques including combinations of text and visual analysis. The workshop attendance reflected the fact that this is an important topic for the MICCAI community.

This paper starts with an overview of the papers presented at the workshop, starting with the invited talks in Section 2. Section 3 summarizes the discussions that took place at the panel session and throughout the workshop, and Section 4 closes the paper with conclusions.

2 Papers Presented at the Workshop

This Section describes the two invited talks of the workshop and also all scientific presentations.

2.1 Invited presentations

Gwenole Quellec was the first invited speaker starting the workshop with a talk on *heterogeneous information retrieval from medical databases*. His inspiring talk developed several ideas about medical image and video retrieval applications

from a theoretical but also from an application perspective. The importance of clinical data in connection with the visual features was highlighted for case–based retrieval, that can be considered much closer to clinical routine than image–based retrieval. Several applications for case–based retrieval were analyzed, notably the use of retinopathy images. A second part of the presentation analyzed challenges for retrieval of videos of medical interventions such as cataract operations as an example. The paper published in these proceedings [17] describes only part of the presentation, notably the video retrieval part. For the video analysis several videos of surgical acts of young and experienced surgeons are being analyzed. Specific phases of each operation were found and could then be detected automatically in the videos and separated. Inside each phase a real–time analysis is performed to be able to react quickly to deviations from an optimal operation. Operations considered good and operations considered poor are compared and thus for an on–going operation any deviation is detected in real time so the surgeon can be informed about a potential risk. A particular value of such a technique would be in the training of young surgeons but also giving constant feedback in real operations and warning from potential dangers..

Georg Langs presented his invited presentation on *VISCERAL: towards large data in medical imaging — challenges and directions* in the afternoon. This presentation explained the VISCERAL[1] project that aims at creating a benchmark for medical imaging on extremely large data sets using a cloud–based infrastructure that is shared by the participants [11]. The goals of the benchmark include two challenges for the research community on a data set of at least 10–20 TB of medical image data that are available to the project. The first challenge is the identification of organs or reference points in the human body, whether in full body scans or partial volumetric data. A focus of the project will thus clearly be on 3D data. A second comparison aims at retrieval using the radiology reports in potentially different languages and with various types of images to find similar cases. The discussion that followed the presentation showed the interest in the topic and highlighted that there are still many things that need to be defined by the project in collaboration with the research community. Work at the medical University of Vienna and inside the Khresmoi[2] project were then presented showing how important and efficient and effective data analysis is when going towards big data for a good retrieval quality. The detection of regions of interest or at least the identification of organs in the body can be an important first step to better analyze the visual medical data and extract semantic labels form the visual image information automatically.

2.2 3D Methods

3D data analysis has grown substantially in medical image retrieval over the past years and this was shown by half of the submissions to the workshop dealing with retrieval of tomographic data sets. It is the quickest growing medical data type.

[1] `http://visceral.eu/`, VISual Concept Extraction challenge in RAdioLogy
[2] `http://www.khresmoi.eu/`

In [4], Catalano at al. describe their article titled *exploiting 3D part–based analysis, description and indexing to support medical applications*. The text highlights the importance of the retrieval of medical 3D data, notably surface–based models for analyzing the different parts of objects and similarities between these parts. Modeling of medical 3D data is an important topic and then being able to use this information for similarity–based retrieval is equally important for many applications.

Another 3D retrieval application was presented by Indriyati Atmosukarto with *skull retrieval for craniosynostosis using sparse logistic regression models* [23]. This approach deals with malformations of the head bones regarding craniosynostosis. The 3D analysis of the skull can help to identify and more importantly quantify certain malformations. The analysis following operative interventions can help to track the evolution of the skull over time. For the interesting application and the solid theoretical and methodological quality this paper was awarded with the Khresmoi prize of the best workshop paper.

2.3 3D/4D Retrieval

Besides the general volume–based 3D analysis and applications, there are also several applications with clear medical application scenarios.

In *retrieval of 4D dual energy CT for pulmonary embolism diagnosis* Foncubierta et al. describe an application of using the 4D data of dual energy CT for the detection of pulmonary embolism in emergency radiology [6]. A difficulty is the extremely large amount of data that needs to be analyzed (11 times 400 slices) and also the difficulty to find in which energy bands the discriminative information is contained. As solid 4D texture is concerned it is also extremely difficult to visualize the data sets.

Simonyan et al. describe in *immediate ROI search for 3–D medical images* [18] a retrieval system in 3D databases that allows users to select 3D regions of interest and then search for visually similar volumes of interest in other image series. The paper gives examples for a theoretically sound framework using the ADNI (Alzheimer disease neuro imaging initiative) database of MRI images to demonstrate the system experimentally.

In *synergy of 3D SIFT and sparse codes for classification of viewpoints from echocardiogram videos* Qian et al. describe the analysis of echocardiogram 3D data [15]. 3D SIFT features are used on the noisy data sets of children ultrasound data of the heart. Modern ultrasound really allows to have high quality data where automatic analysis can become possible. Even 4D data sets, for example 3D data of the beating heart, have become available and could be intersting for similarity–based retrieval applications.

Quatrehomme et al. present in *assessing the classification of liver focal lesions by using multi–phase computer tomography scans* an interesting approach to the analysis of liver lesions [16]. The approach works on single slices but over time so analyzing the flow of a contrast agent. This time component adds the third dimension and shows to increase the performance of the classification

in an important way. The variety of applications show the large spectrum of applications in multidimensional data.

2.4 Visual Features

Visual features remain important, particularly transferring visual features from other domains to the use within medical imaging. In many benchmarks it was shown that the more visual features are used the better the results are [9]. Still, when moving towards big data it will simply become impracticable to work with too large a variety of features and thus optimized and compact visual features will become necessary.

In *customised frequency pre–filtering in a local binary pattern–based classification of gastrointestinal images*, Wimmer et al. present the use of Local Binary Patterns (LBP) for the analysis of gastrointestinal images [22]. LBPs have been used in various scenarios to represent texture information. In the case of gastrointestinal images (in this case for Celiac disease and polyps) a frequency–based pre–filtering of the images led to optimized results that outperformed LBP and several of its derivations.

Garcia Seco de Herrera et al. describe in *bag of colors for biomedical document image classification* the use of SIFT and bag of visual color features for document image classification in images types [7]. Using a standard data set of the ImageCLEF benchmark the two features combined showed to increase the classification performance more than any of the participants in the benchmark that only used the supplied training data. This shows that color carries a very important part of the document image information.

2.5 Multimodal retrieval

The last session of the day before the panel was on multimodal retrieval approaches, meaning in this case the combination of visual retrieval techniques and textual retrieval and not the combination of modalities such as PET and CT or MRI.

In *an SVD–bypass latent semantic analysis for image retrieval*, Stathopoulos et al. use latent semantic indexing on visual and textual information for image retrieval [20]. The results of mixing the modalities shows to perform well on the given ImageCLEF database. Another aspect of the search was the scalability, meaning to use simple visual features that could potentially scale to millions of images for the retrieval.

Castellanos et al. describe in *multimedia retrieval in a medical image collection: results using modality classes* an approach for using modality classes for retrieval from the ImageCLEF 2011 medical task [3]. Expanding queries with an automatically extracted image modality class overall increases the results. These gains strongly depend on the type of query and it can not be generalized to all types of queries. Modality class information can also be used to improve the figure captions or their context, so the corresponding figures can be retrieved easier in the future.

3 The Panel and Discussions

The presentations throughout the day have demonstrated the large variety in applications, in image types (CT, x–ray, MRI, dual energy CT, ultrasound, gastrointestinal videos, journal figures, ...) and also in techniques and visual features used. Medical image retrieval has found its way as part of several medical applications and will often do its work as part of a larger system, whether it is combined with text search, for getting decision support or in getting access to interesting cases for preparing courses.

The panel was organized as an open session around the topic *What is the CBIR role in Medical Decision Support?*. The penal started with an experience report of IBM on applications of medical image retrieval in working with clients on various integration projects, largely with the Kaiser Permanente group that organizes many hospitals in California. Several IBM projects were presented including the difficulty to actually have clinically relevant tools and systems that help physicians in doing their work better and quicker. Pure system performance in benchmarks does not correlate with user acceptance and many discussions and tests with the physicians were necessary to conserve only those parts of the systems that really have an added value for the physician. This means that the techniques need to be integrated in the work flow, they need to be fast and in many cases they will be invisible.

The open panel invited all participants of the workshop to join the discussion and propose ideas. One issue that was discussed is the need for high quality *annotated medical imaging data* in large quantities. Having access to such data would be a big advantage for the community, but those who prepare the data would need to get at least part of the benefits, which is currently not always the case. This can lead to data sets not being shared. Quality of the annotation and confidence in diagnosis were also mentioned as physicians often are confident of their own opinion and do not necessarily trust other physicians unless they have a sort of proof, for example in the form of biopsies. Inter rater disagreement are important to measure, also to have a baseline for computer–based decision support. The level of detail for annotation can vary strongly and this needs to be defined well to create useful data sets and not have essential information missing. Community efforts are expected to be needed where several partners create annotations together. It was also mentioned that there will always be new needs for data sets as benchmarking has also the risk to lead to standardization and new ideas.

The importance of *big data* was stressed as this could allow totally different approaches and maybe would lead with much simpler techniques to better results. Current computers have the possibility to deal with such extremely large amounts of data and much can still be discovered in this respect. Data protection of course needs to be respected and informed consent is in most cases necessary,even though data protection is different depending on the countries even inside Europe. Clear guidelines for the secondary use of medical imaging data in anonymized form need to be developed to ease research all while respecting privacy of the patients.

4 Conclusions

Medical image retrieval has remained a very dynamic research domain over the past ten years with many new directions and a strong evolution. Focus has come from theoretical models towards real applications and from small data sets to much larger data repositories. Whereas initial image retrieval applications focused on the image and human perception of visual similarity, modern applications are increasingly integrating the medical context into the retrieval process, such as mixing visual image data with clinical parameters, or several images for real case–based retrieval to help diagnosis and create links between a clinical cases and reports in the medical literature. Many 3D and even 4D retrieval applications have started and these applications will require even more than 2D retrieval the definitions of regions of interest to focus a more detailed analysis on the most important parts. Detecting small regions of potential abnormalities and then using these regions to find images or cases with similar lesions will remain an important direction for the coming years to improve current tools for medical decision support. Models for organs need to be built and links between images from the scientific literature in JPEG format and in the DICOM format in the patient record need to be available as well, to make sure that the various sources of information are in the end well connected.

Image retrieval based on visual and textual data might not be a visible part of all applications, but increasingly it is integrated into many tools, even though often in an invisible form. This indicates that some of the search and retrieval research should be conducted as part of specific applications or specific domains. Much can still be improved in terms of techniques and scalable approaches to deal with big data but it can be foreseen that several of the tools will be integrated into a variety of applications and systems, also with a commercial success.

5 Acknowledgments

This work was supported by the EU in the FP7 through the VISCERAL (318068), PROMISE (258191) and Khresmoi (257528) projects.

References

1. Riding the wave: How europe can gain from the rising tide of scientific data. Submission to the European Comission, available online at `http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf` (October 2010), `http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf`
2. Akgül, C., Rubin, D., Napel, S., Beaulieu, C., Greenspan, H., Acar, B.: Content–based image retrieval in radiology: Current status and future directions. Journal of Digital Imaging 24(2), 208–222 (2011)
3. Castellanos, A., García Serrano, A., Benavent, X., Benavent, J.: Multimedia retrieval in a medical image collection results using modality classes. In: Greenspan,

H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
4. Catalano, C.E., Robbiano, F., Parascandolo, P., Cesario, L., Vosilla, L., Barbieri, F., Spagnuolo, M., Viano, G., Cimmino, M.A.: Exploiting 3d part–based analysis, description and indexing to support medical applications. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
5. El-Kwae, E., Xu, H., Kabuka, M.R.: Content–based retrieval in picture archiving and communication systems. JDI 13(2), 70–81 (May 2000)
6. Foncubierta-Rodríguez, A., Poletti, P.A., Platon, A., Vargas, A., Müller, H., Depeursinge, A.: Retrieval of 4D dual energy CT for pulmonary embolism diagnosis. In: Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012 (October 2012)
7. García Seco de Herrera, A., Markonis, D., Müller, H.: Bag of colors for biomedical document image classification. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
8. 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), 648–655 (2009)
9. Kalpathy-Cramer, J., Müller, H., Bedrick, S., Eggel, I., García Seco de Herrera, A., Tsikrika, T.: The CLEF 2011 medical image retrieval and classification tasks. In: Working Notes of CLEF 2011 (Cross Language Evaluation Forum) (September 2011)
10. Kato, T.: Database architecture for content–based image retrieval. In: Jamberdino, A.A., Niblack, W. (eds.) Image Storage and Retrieval Systems. SPIEProc, vol. 1662, pp. 112–123. San Jose, California (February 1992)
11. Langs, G., Müller, H., Menze, B.H., Hanbury, A.: Visceral: Towards large data in medical imaging — challenges and directions. In: Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012 (October 2012)
12. Lowe, H.J., Antipov, I., Hersh, W., Smith, C.A.: Towards knowledge–based retrieval of medical images. The role of semantic indexing, image content representation and knowledge–based retrieval. In: Proceedings of the Annual Symposium of the American Society for Medical Informatics (AMIA). pp. 882–886. Nashville, TN, USA (October 1998)
13. Müller, H., Clough, P., Hersh, W., Deselaers, T., Lehmann, T., Geissbuhler, A.: Evaluation axes for medical image retrieval systems: the ImageCLEF experience. In: MULTIMEDIA '05: Proceedings of the 13th annual ACM international conference on Multimedia. pp. 1014–1022. ACM Press, New York, NY, USA (November 2005)
14. 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), 1–23 (2004)
15. Qian, Y., Wang, L., Wang, C., Gao, X.: Synergy of 3d sift and sparse codes for classification of viewpoints from echocardiogram videos. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
16. Quatrehomme, A., Millet, I., Hoa, D., Subsol, G., Puech, W.: Assessing the classification of liver focal lesions by using multi–phase computer tomography scans.

In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
17. Quellec, G., Lamard, M., Droueche, Z., Cochener, B., Roux, C., Cazuguel, G.: A polynomial model of surgical gestures for real-time retrieval of surgery videos. In: Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012 (October 2012)
18. Simonyan, K., Modat, M., Ourselin, S., Criminisi, A., Zisserman, A.: Immediate roi search for 3–d medical images. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
19. Smeulders, A.W.M., Worring, M., Santini, S., Gupta, A., Jain, R.: Content–based image retrieval at the end of the early years. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(12), 1349–1380 (December 2000)
20. Stathopoulos, S., Kalamboukis, T.: An svd–bypass latent semantic analysis for image retrieval. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
21. Tagare, H.D., Jaffe, C., Duncan, J.: Medical image databases: A content–based retrieval approach. Journal of the American Medical Informatics Association 4(3), 184–198 (1997)
22. Wimmer, G., Uhl, A.: Customised frequency pre–filtering in a local binary pattern–based classification of gastrointestinal images. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)
23. Yang, S., Shapiro, L., Cunningham, M., Speltz, M., Birgfeld, C., Atmosukarto, I., Lee, S.I.: Skull retrieval for craniosynostosis using sparse logistic regression models. In: Greenspan, H., Müller, H. (eds.) Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012, Lecture Notes in Computer Sciences (LNCS) (October 2013)