Overview of the First Workshop of Muldimodal Retrieval in the Medical Domain (MRMD 2015)

Henning Müller[12], Oscar Jimenez–del–Toro[1], Allan Hanbury[3], Georg Langs[4], and Antonio Foncubierta–Rodríguez[5]

[1]University of Applied Sciences Western Switzerland (HES–SO), Sierre, Switzerland
[2]University Hospitalsand University of Geneva, Switzerland
[3]Technical University of Vienna, Austria
[4]Medical University of Vienna, Austria
[5]ETHZ Z´urich, Switzerland

Abstract. The workshop Multimodal Retrieval in the Medical Domain (MRMD) took place in connection with the European Conference of Information Retrieval (ECIR) in Vienna, Austria on March 29, 2015. The workshop included two invited presentations and 7 accepted scientific papers. A session on the VISCERAL (VISual Concept Extraction in RAdioLogy) retrieval benchmark, an information retrieval benchmark combining semantics with 3D images, was included into the workshop and four groups presented their results on the benchmark using the same database and query topics to really lead to comparable results. The end of the workshop featured a discussion session with the invited speakers and workshop organizers regarding what is now needed to advance large–scale multimodal medical information retrieval.

1 Introduction

The workshop MRMD (Multimodal Retrieval in the Medical Domain) took place for the first time in March 2015. Its organization was supported by the VISCERAL[1] (Visual Concept Extraction Challenge in Radiology) project that organized several evaluation benchmarks in the medical domain, notably for image segmentation, lesion detection and information retrieval in a cloud–based infrastructure with the goal to move the algorithm to the data rather than the data to the algorihms [5].

The workshop received 11 scientific paper submissions of which 4 were accepted and 3 conditionally accepted with a second round review. All conditionally accepted papers became acceptable after the review comments were included into the resubmission and the papers were re–reviewed. In addition to the scientific presentations two invited speakers presented on closely related topics, one on semantic data representation and matching for image analysis and the other invited speaker on using very large amounts of imaging data and related clinical

[1] `http://visceral.eu/`

parameters to validate research algorithm in an infrastructure where only algorithms see the data and nothing needs to be seen by the researchers themselves (besides the structure of the data itself), thus preserving data privacy.

The VISCERAL project also invited the participants of the VISCERAL retrieval benchmark [15] to present their results in this workshop. Four groups presented techniques and approaches based on the same data and the same evaluation setting, thus stimulating discussions on techniques and optimizations. This allowed to compare techniques and identify also promising directions. Spanier et al. presented medical case–based retrieval of patient records using the RadLex hierarchical lexicon [13] concentrating thus on the semantic annotations of the images. Zhang et al. [16] then present an approach based rather on visual features extracted directly from the 3D data. Sungbin Choi et al. [3] combined visual and semantic features for retrieval. Finally, Jiménez del Toro et al. [7] also use a combination of visual features extracted form the organ marked in the query and combine it with the semantic features.

The topic of the workshop is very timely inside the information retrieval community with big data becoming a major topic in healthcare all around the world to improve care processes and potentially reduce costs. Retrieval of similar cases can in many areas thus be an important part to extract knowledge from patient data. Besides genonic data, imaging is clearly the biggest volume of medical data and a recent study of the EU estimates medical imaging to occupy around 30% of world storage [1]. Analysis of such approaches also to medical imaging has many challenges and it is important to get medical imaging researchers used to using large data sets despite the technical difficulties of it [9]. This can create real opportunities to improve care processes and generate new medical knowledge, moving towards digital medicine. Comparing several approaches on the same data has many advantages as results and optimizations can directly be discussed and compare by the participants without interpretation problems.

2 Invited Presentations

The first invited speaker in the morning just after the workshop opening was Camille Kurtz of the University Paris Descartes, France. He presented on *semantic features for the automated retrieval of radiological images in medical databases* [8]. He highlighted the importance of not only using low level visual features for medical image retrieval but adding semantic terms to this retrieval and particularly making the links between semantics and visual features. He explained a semantic distance that is used for the retrieval of liver images and that can create relations between annotated lesions and thus complete cases. Modeling of texture is matched with the semantic terms and the two together then led to optimal results for the retrieval process or for automatic annotation. Such tools can help physicians by highlighting and quantifying abnormalities.

The second invited talk was given by Eldad Elnekave, the Chief Medical Officer of Zebra Medical Vision[2] and titled *Clinical imaging: into the frontier of*

[2] `http://www.zebra-med.com/`

medical big data. Zebra Medical Vision is a startup company working in the space of big data in medicine and concentrating on the use of imaging data. Zebra has contracts with a large hospital group in Israel and access to 450 TB of imaging data and associated clinical data of 4 million patients over a period of 10 years. All data are stored locally in anonymized form and researchers can not directly access the data but only indirectly via executables that are run on the data and need to be prepared by knowling the structure of the original data. An objective is to create challenges for the scientific community in specific domains where decision support in medicine is needed. Scientists could thus work on very large data sets, test their algorithms of statistically interesting data sets and compare results to a strong baseline and base don a manually annotated ground truth. Zebra can then certify resulting software tools and hospitals can make sure to always have the highest quality decision support in a scenario like this, creating a win win situation for academia and the company. The software is currently in a beta version. Data preparations and annotations for first competitions are under way.

3 Presented Papers

There were seven papers orally presented at the workshop sessions, all related to multimodal data analysis but from a large number of different aspects and using many different data types.

Balikas et al. [2] present the BioASQ[3] challenge on question answering using medical data from the scientific literature (Mainly PubMed and PubMed central) and also the automatic annotation of medical text with semantic terms. As part of this EU project several scientific challenges were organized and are also planned for the future. This creates a testbed for algorithms and tools to be compared on the same data and evaluation scenario.

Similarly in the context of semantics and mainly text analysis of medical data, Lelong et al. [12] present an approach on information retrieval from a clinical medical record in French. The proposed system rewrites initial user queries and translates them into a semantic space that can as a consequence improve retrieval quality. The system was tested in a clinical scenario with real data and is planned to be used in the hospital environment directly where the group already maintains several tools.

3D imaging is the objective of [11], where Liu et al. present a content–based image retrieval system for brain MRI (Magnetic Resonance Imaging) images. Diffusion Tensor Imaging (DTI) is used in this case, also leading to a real 3D information retrieval system. 3D information retrieval is complex and needs to be focused on specific regions to separate information from background.

Similarly on 3D retrieval is the paper of Li et al. [10] but this time not on tomographic images but on echocardiology sequences, so videos. The system uses KAZE (Japanese for wind, highlighting the flow of non–linear information)

[3] `http://www.bioasq.eu/`

features to classify the videos. The system reports good performance on the retrieval of these compelx data.

In [14], Stathopoulos et al. present an approach to multi–modal indexing using the ImageCLEF database. The system uses latent semantic analysis (LSA) to combine the text and the visual features for retrieval. Several optimizations improve the performance compared to the best runs submitted at the competition.

Garcia Seco de Herrera at al. [6] then present an approach also used for multimodal retrieval, that includes besides text and visual features also the automatically determined image modality as a filter. For optimizing the automatic classification of image modality a training set expansion was done including a manual correction step. The resulting system leads to better classification results and slightly improved retrieval performance.

The last presented paper of Demner et al. [4] presents practical experiences with the automatic annotation of chest radiology reports for the further indexing and retrieval. The system is maintained by the NIH (National Insitutes of Health) and will make data available in important areas such as on chest xrays, where a very large number of exams is performance each year and where automatic analysia and reference cases can be of great help.

4 Discussions at the Workshop

At the end of the workshop a discussion session took place where the invited speakers and members of VISCERAL were at the front and all participants gave feedback on the question of what is currently needed to improve multimodal medical information retrieval. The preparation and annotation of large data sets available for research was one of the major points and the initiative of Zebra in this direction was highlighted. The model of making data only available for algorithms and not researchers could solve many of the current data availability problems as data could reside in medical insitutions. Challenges remain, as it would be best to get large data across many institutions, and having a distributed data analysis can complicate the structure. The manual annotation and curation of data are also seen as important and to share these efforts to maximize the use of existing data sets. Creating silver corpora seems like a necessity as manual annotations do not scale. It still needs to be determined how silver corpora can be validated, maybe only re–annotating regions or images where the automatic algorithms have a large disagreement. Like this effort can be spent in a way to achive the highest information gain. Most conferences in medical imaging now have room for challenges and also large funding organizations see benefits in creating challenges for problems of real importance to advance scientific results. Particularly long term funding for chalenges is necessary to be able to measure advances over time and give researchers a stable environment of data to work on. Longitudinal imaging data were regarded as extremely useful but often harder to obtain. Many disease patterns are best visible in evolution of imaging patterns

over time or for example response to therapy and this is rarely available in current imaging challenges.

A critical comment regarding the VISCERAL infrastructure was the difficulty of using virtual machines compared to running algorithms locally on an infrastructure of choice. Specific optimization, such as for GPUs are not trivial to duplicate in cloud environments, particularly if all participants are suppoed to work on precisely the same environment. This likely also limits participation in benchmarks run on cloud infrastructures. If a simple snapshot of a local machine could be created and transferred to the cloud as virtual machines, many problems could potentially be solved but VMs are currently not fully mobile. Maybe intermediate solutions such as docker[4] could help to limit the amount of additonal work and allow for simple transfer of tools for data analysis and retrieval to different locations. Maybe new developments in virtualization can also help in this respect and the virtualization approach was still regarded as most promising.

5 Conclusions

The first workshop of medical multimodal retrieval took place at ECIR 2015 in Vienna, Austria. The workshop had around 25 participants. The relatively small group of people was very active and the presentations, both invited and scientific, lead to lively discussions among the participants and many future ideas. The presented VISCERAL retrieval benchmark is the first benchmark using 3D imaging and semantic data for medical information retrieval. In this respect the multimodal nature and high complexity of the data may have been hard to work on for some research groups. The data will be kept available and we feel that there is still a high potential in using the data and improving current approaches based on real clinical data to find similar cases. Medical imaging is currently the biggest data producer in medicine and the automatic analysis of these large volumes is required to help clinicians. The importance of the data is also highlighted by companies such as Zebra Medical Vision that builds on the potential of using large data sets in a secure environment to improve clinical decision making and decision support.

Medical visual information retrieval has in this context the potential to bridge the gap between medical image processing and the information retrieval community. This is a highly interdisciplinary domain and thus sharing knowledge form each domain seems to be the best way to success.

Acknowledgments This work was partially supported by the EU in FP7 via the VISCERAL project (grant agreement 318068).

[4] `http://www.docker.com/`

References

1. Riding the wave: How europe can gain from the rising tide of dcientific 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. Balikas, G., Krithara, A., Partalas, I., Paliouras, G.: BioASQ: A challenge on large–scale biomedical semantic indexing and question answering. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
3. Choi, S.: Multimodal Medical Case–Based Retrieval on the Image and Report: SNUMedinfo at VISCERAL Benchmark. In: Multimodal Retrieval in the Medical Domain. Lecture Notes in Computer Science, vol. 9059. Springer (2015)
4. Demner-Fushman, D., Shooshan, S.E., Rodriguez, L., Antani, S., Thoma, G.R.: Annotation of chest radiology reports for indexing and retrieval. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
5. Hanbury, A., Müller, H., Langs, G., Weber, M.A., Menze, B.H., Fernandez, T.S.: Bringing the algorithms to the data: cloudbased benchmarking for medical image analysis. In: CLEF conference. Springer Lecture Notes in Computer Science (2012)
6. García Seco de Herrera, A., Markonis, D., Joyseeree, R., , Schaer, R., Foncubierta-Rodríguez, A., Müller, H.: Using semi–supervised learning for image modality classification. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
7. Jiménez-del-Toro, O.A., Cirujeda, P., Dicente Cid, Y., Müller, H.: RadLex Terms and Local Texture Features for Multimodal Medical Case Retrival. In: Multimodal Retrieval in the Medical Domain. Lecture Notes in Computer Science, vol. 9059. Springer (2015)
8. Kurtz, C., Idoux, P.A., Thangali, A., Cloppet, F., Beaulieu, C.F., Rubin, D.L.: Semantic retrieval of radiological images with relevance feedback. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
9. Langs, G., Müller, H., Menze, B.H., Hanbury, A.: Visceral: Towards large data in medical imaging – challenges and directions. Lecture Notes in Computer Science 7723, 92–98 (2013)
10. Li, W., Qian, Y., Loomes, M., Gao, X.: The application of KAZE features to the classification of echocardiogram videos. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
11. Liu, S., Hadi, N., Liu, S., Pujol, S., Kikinis, R., Zhang, F., Feng, D.D., Cai, W.: Content–based retrieval of brain diffusion magnetic resonance image. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
12. Soualmia, L.F., Lelong, R., Dahamna, B., Darmoni, S.J.: Rewriting natural language queries using patterns. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
13. Spanier, A.B., Joskowicz, L.: Medical Case–Based Retrieval of Patient Records Using the Radlex Hierarchical Lexicon. In: Multimodal Retrieval in the Medical Domain. Lecture Notes in Computer Science, vol. 9059. Springer (2015)

14. Stathopoulos, S., Kalamboukis, T.: Multi–modal indexing and retrieval using an LSA–based kernel. In: Multimodal Retrieval in the Medical Domain (MRMD) 2015. Lecture Notes in Computer Science, Springer (2015)
15. Jiménez-del Toro, O.A., Foncubierta-Rodríguez, A., Müller, H., Langs, G., Hanbury, A.: Overview of the VISCERAL Retrieval Benchmark 2015. In: Multimodal Retrieval in the Medical Domain. Lecture Notes in Computer Science, vol. 9059. Springer (2015)
16. Zhang, F., Song, Y., Cai, W., Depeursinge, A., Müller, H.: USYD/HES–SO in the VISCERAL Retrieval Benchmark. In: Multimodal Retrieval in the Medical Domain. Lecture Notes in Computer Science, vol. 9059. Springer (2015)