Workshop Multimodal Retrieval in the Medical Domain (MRMD) 2015

Henning Müller[1], Oscar Alfonso Jiménez del Toro[1], Allan Hanbury[2], Georg Langs[3] and Antonio Foncubierta–Rodríguez[4]

[1]University of Applied Sciences Western Switzerland (HES–SO), Switzerland
[2]Vienna University of Technology (TUW), Austria
[3]Medical University of Vienna (MUW), Austria
[4]Swiss Federal Institute of Technology (ETH) Zurich, Switzerland

Abstract. The workshop Multimodal Retrieval in the Medical Domain (MRMD) deals with various approaches of information retrieval in the medical domain including modalities such as text, structured data, semantic information, images, and videos. Goal is to bring together researchers of the various domains to combine approaches and compare experiences.
The workshop includes a special session on the VISCERAL benchmark that works on the retrieval of similar cases from a collection of 3D volumes of mainly CT and MRI data. Results of the participants will be compared at the benchmark and shoudl complement the genral topic of mutlimodal retrieval.

Keywords: Content–Based Image Retrieval, Multimodal retrieval, Informtion Retrieval infrastructures, VISCERAL

1 Introduction

Medical information is of interest to a wide variety of users, including patients and their families, researchers, general practitioners and clinicians, and practitioners with specific expertise such as radiologists [1]. There are several dedicated services that seek to make this information more easily accessible, such as Health on the Nets medical search systems for the general public and medical practitioners[1]. Despite the popularity of the medical domain for users of search engines, and current interest in this topic within the information retrieval research community, development of search and access technologies remains particularly challenging.

This workshop focuses on retrieval in the medical domain based on multimodal data. This can concern medical cases that refer to data about specific patients (used in an anonymised form), such as medical records, radiology images and radiology reports or cases described in the literature or teaching files. The workshop will consist of the following parts:

[1] http://www.hon.ch/

– Two invited talks on retrieval in the medical domain and infrastructures for evaluation on large–scale data.
– Presentations of submitted papers on the topic of the workshop, multimodal retrieval in the medical domain.
– Presentation of the results from the VISCERAL[2] Retrieval Benchmark, which benchmarks multimodal retrieval on large amounts of radiology image and text information [3].
– A discussion session on evaluation infrastructures for large–scale retrieval in the medical domain and potential new benchmarks [2].

2 Objectives and Outcomes

This workshop has the following objectives:

– Presentation of papers covering retrieval in the medical domain based if possible on large data sets and multimodal data.
– Presentation of the results from the VISCERAL Retrieval Benchmark.
– Discussion on evaluation infrastructures for large–scale retrieval in the medical domain and potential new benchmarks.

The target of the workshop is that the presentation of experiences and results from a pilot of a large–scale retrieval benchmark combined with presentations of work on related problems in the domain will lead to the proposal of new large–scale retrieval benchmarks in the medical or the information retrieval domain and innovative ways in which these benchmarks can be carried out. New retrieval techniques for text and multimodal data should presented at the workshop. The proceedings of the workshop will be published in the Springer LNCS series as post–proceedings. A report on the workshop will be produced, covering in particular the discussion and be submitted to the SIGIR Forum.

3 Structure of the Workshop and Paper Selection Process

One of the challenges of medical information retrieval is similar case retrieval in the medical domain based on multimodal data, where cases refer to data about specific patients (used in an anonymised form), such as medical records, radiology images and radiology reports or cases described in the literature or teaching files. The VISCERAL project aims at evaluating and promoting improvements of the state–of–the–art in this field, and is organizing the VISCERAL Retrieval Benchmark. The data set consists of 2311 volumes originated from three different modalities (CT,MR T1,MR T2). It serves the following scenario: a user is assessing a query case in a clinical setting, e.g., a CT volume, and is searching for cases that are relevant in this assessment for differential diagnosis. The algorithm has to find cases that are relevant in a large database of cases. For each topic (query case) there is:

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

– the patient 3D imaging data (CT, MRI);
– the 3D bounding box region of interest containing the radiological signs of the pathology;
– a binary mask of the main organ affected;
– the radiologic report extracted anatomy–pathology terms in form of csv files of RadLex terms.

The participants have to develop an algorithm that finds clinically–relevant (useful for differential diagnosis or same diagnosis) cases given a query case (imaging data only or imaging and text data), but not necessarily cases with the same final diagnosis.

Medical experts will perform relevance assessment of the top ranked cases by each approach, to judge the quality of retrieval. Experts will assess the relevance of the ranked cases. The evaluation measures used are the precision of the top–ranked X cases, where X is 10, 20, 30.

The benchmark is run on a cloud–based infrastructure that allows processing to be done on data stored on the cloud through Virtual Machines provided to the participants. Participants in this Benchmark will be encouraged to submit papers to the workshop that explore the data, identify approaches and understand how more data might be sourced. These will be presented, along with a discussion summarising all results of the benchmark.

In addition to papers related to this benchmark, further papers will be solicited describing approaches to other types of similar case retrieval in single or multiple modalities in the medical domain (e.g. similar patients based on medical records or medical records combined with laboratory values). Other retrieval approaches on medical data are euqally sollicited. Finally, bringing these researchers together will result in a guided discussion session. The discussion leads to new ideas for benchmarks on large–scale retrieval in the medical domain, and for infrastructures on which these benchmarks can be run and data can be shared.

All papers submitted to the workshop (VISCERAL Retrieval Benchmark and further papers) undergo a peer review by at least three members of the Programme Committee per paper. The acceptance decisions will be made by the organisers based on the recommendations of the reviewers targeting an acceptance rate of around 50%.

The workshop also features two invited speakers, one covering retrieval applications in the medical domain and the second related to evaluation infrastructures.

4 Intended Audience

This workshop is aimed at:

– Researchers working in (multimodal) information retrieval in the medical domain.

Table 1. Programme

9.00 – 10.45	Presentation of VISCERAL benchmark participants and global results
11.00 – 12.00	Invited talk on information retrieval evaluation infrastructures
12.00 – 13.00	Lunch
13.00 – 14.45	Presentation of the submitted papers
15.00 – 16.00	Invited talk on retrieval in the medical domain
16.00 – 17.30	Panel discussion on benchmarking infrastructures

- Researchers working on the creation of novel information retrieval evaluation approaches and evaluation infrastructures.
- Participants in the VISCERAL Retrieval Benchmark.

5 Conclusions

The MRMD benchmark aims to give a forum to researchers working on medical information retrieval in a variety of settings and using a vriety of techniques but favoring multimodal approaches. By combining the medical information retrieval techniques with a benchmark on large scale visual and semantic data we hope to create synergies and create new ideas in terms of research challenges, databases and evaluation approaches. The workshop also treats research infrastructures and this is also with the obecjtive to create discussions and discuss experiences of approaches to develop best practices inthis domain. We hope that these topics will attract a large audience nad particularly fruitful discussions.

6 Acknowledgments

This research was funded by the EU in the FP7 VISCERAL project (318068).

References

1. Hanbury, A., Boyer, C., Gschwandtner, M., Müller, H.: KHRESMOI: Towards a multi–lingual search and access system for biomedical information. In: Med–e–Tel, Luxembourg, 2011. pp. 412–416 (2011)
2. Hanbury, A., Müller, H., Langs, G., Menze, B.H.: Cloud–based evaluation framework for big data. In: Galis, A., Gavras, A. (eds.) Future Internet Assembly (FIA) book 2013, pp. 104–114. Springer LNCS, Springer Berlin Heidelberg (2014)
3. 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)