

Semantic Gastroenterological Images Annotation and Retrieval *Reasoning with a Polyp Ontology*

Yahia Chabane and Christophe Rey

Clermont Université, LIMOS, CNRS UMR 6158, F-63171 Aubiere, France

Keywords: Semantic Medical Image Retrieval, Polyp Ontology, Local Subsumption, Annotation.

Abstract: In gastroenterology, monitoring polyps is fundamental in order to detect a cancer. It may be difficult for surgeons to decide whether he should remove a polyp or not. A wrong decision may generate unjustified costs or be dangerous for the patient health. To help their diagnosis, physicians may need images of previously treated cases. For this purpose, we present in this paper a semantic image retrieval approach focused on endoscopic gastroenterological images. This approach is based on a slight extension of classical description logic reasonings, associated with a polyp ontology and a suited image annotation mechanism.

1 INTRODUCTION

In medicine, and for instance in gastroenterology, physicians and surgeons are used to basing some of their decisions on images they have taken during medical exams (e.g. whether to remove or not a polyp). But medical images can also be used during medical studies, throughout epidemiological research works, for medicolegal purposes, ... Thus they are a really important element in medical practises.

One step that is crucial in a medical images management platform is the research one. Of course, researching a medical image must be intuitive and images must be quickly obtained and be as relevant as possible. Since image retrieval is mainly based on a comparison between image annotations and image queries, annotating images must also be intuitive. In this paper, we adress these problems of annotating and retrieving medical images in the field of gastroenterology, focusing on endoscopic images of polyps. We follow a semantic image retrieval approach based on a new polyp ontology. This work is part of the Syseo project (Syseo, 2011), which aims at building an integrated gastroenterological DICOM images management platform.

In classical syntactic image retrieval approaches, images are described with a set of keywords (Datta et al., 2008), (Rui et al., 1999), (Bimbo, 1999). The quality of the image retrieval process comes from the relevance of these keywords, but also from the relevance of the keywords used to express the user query, and of course from the adequation between the anno-

tated keywords and the user ones. If the user does not use exactly the same keywords as those in the image description, then he may not retrieve this image even if it is relevant for him.

To add more flexibility, the semantic approach is based on a so-called "ontology" of the domain (Gruber, 2009), that is a dictionary where keywords are given a definition, expressed with other keywords having themselves a definition, and so on. We then talk about concepts more than keywords. So an ontology usually describes a whole domain with hundreds of concepts. As these concepts have a precise definition which links them to other concepts, the retrieval process can be based on the use of these links, which are called "semantic" since they are related to the concept definitions. Flexible retrieval techniques can then be designed as logical inference techniques called "reasonings". For instance, we can exploit the semantic link saying that a polyp is a kind of a lesion, to retrieve lesion images, as an extended research criteria (or if there is no polyp images). In this case the underlying reasoning is called "subsumption".

Current semantic image management systems generally make use of a knowledge representation formalism such as description logics (Baader et al., 2007) or semantic networks (Meghini et al., 2001) to be able to define ontologies and reasonings. In this paper we have chosen the description logics approach. Indeed, it allows us to reuse many theoretical results about languages expressivity and associated reasonings. Moreover the web ontology language OWL (OWL, 2007), which has been standardized by the

W3C for more than 10 years, is grounded on description logics. And last but not least, many reasoners implementing efficient reasonings are available (Hermit, 2013; RacerPro, 2013; Pellet, 2013; FaCT++, 2013).

In section 2, we explain what problems gastroenterologists face to manage their images, and informally expose how the semantic image retrieval mechanism we're proposing can help. Then, we formally present our proposal in section 3. Afterwards, in section 4, we quickly survey the image retrieval field with a focus on semantic medical image retrieval, before concluding.

2 PROBLEM AND SCENARIOS

To be precise, let's first define the suited vocabulary. A *concept definition* (or simply a definition) is an expression using concepts of the ontology that gives its meaning to another concept. This suggests the existence of *atomic concepts* which are only defined by their name. An *annotation* is the expression of what is observable in an image, expressed with concepts from the ontology. So an annotation is a kind of definition. A *property* is a part of a definition. For instance, a property of an annotation may be the part of the annotation which describes a specific observable aspect of the annotated image (color, texture,...). A *classification* is an existing thesaurus related to polyps that has been integrated into our ontology: each keyword of the thesaurus has been transformed into a concept of the ontology, and these have been given a definition according to gastroenterological concepts previously stored in the ontology.

The practical problem addressed in this paper is how to set up a semantic search mechanism suited for gastroenterologists in their everyday use. Such a mechanism should provide more flexibility and extendability wrt classical syntactic search. Concretely, this is especially interesting for a physician when he wants to match a set of information to a set of reference images (images for which it is known which kind of polyp and pathology they show) or reference concepts (classes from a polyp classification).

For instance, this matching can be useful in the following cases:

- as an help for annotation: during a surgery, when he takes a photo, a surgeon may want to describe it even basically; in this context he can give his observations (color, shapes, ...) to the system which will display corresponding reference images (those described by an annotation semantically linked with the same observations), help-

ing him annotating his image with the suited (and standard) concepts.

- as an help for a diagnosis: after having annotated the taken image, the surgeon may have to decide whether he removes or not the polyp. Then he can once again be helped by the system that would classify the image, according to its annotation, into various classifications (related to medical acts, pathologies, procedures,...). He can also be displayed images of the same kind of polyps that have been previously removed.
- as an help during medical education: learning a polyp classification can be eased by a system that is able to quickly relate a set of observations to the possible polyp classes it corresponds, and also to give the polyp classes that are not related to the observations by focusing on what is contradictory between the observation and the polyp description.

To make the previous cases a reality, we propose a three-parts system: (i) the ontology which contains the definitions of gastroenterological concepts, focused on polyps, and also containing all polyp classifications and their associated classes (see section 3.1), (ii) a set of reasonings (see section 3.3), and (iii) a set of reference images with their annotations which have been validated by experts. Now, we have isolated three main scenarios:

- S1 semantic images retrieval: given an image class (resp. an image annotation) the issue is to find the reference images belonging to this class (resp. being similar to this annotation). Here the semantic approach holds in the fact that classes of reference images are inferred from their annotation (and validated by experts).
- S2 exact classes retrieval: given an image annotation, and the name of a classification, the issue is to find the exact classes (in this classification) of the image, that is the classes which all definition properties can be inferred from the input image annotation.
- S3 approximated classes retrieval: given an image annotation, and the name of a classification, the issue is to find the approximated classes (in this classification) of the image, that is the classes from the definition of which we can infer all the properties of the input image annotation.

An interesting point with the previous scenarios, is that they can be combined. Combining S1 after S2 allows to retrieve all the reference image annotations that belong to the exact classes of an input image annotation. Similarly, combining S1 after S3 allows to

retrieve all the reference image annotations that belong to the approximated classes of an input image annotation. Moreover combining S3 after S1 allows to retrieve all the approximated classes of the reference images annotations belonging to an input class (or input image annotation). This can be helpful for the understanding of links between classes.

3 PROPOSED APPROACH

Setting up a semantic system, based upon description logics, implies to make a few critical choices: what language in the description logics family to express concepts definitions? What reasonings to compute inferences? What reasoner to make reasonings run? The difficult point resides in the well-known tradeoff between expressivity and computation performances. Indeed, the finer one can express knowledge, the harder (and thus slower) the reasonings are.

In our work, we are guided by our previous scenarios S1, S2 and S3: indeed, the first corresponds to the classical instance retrieval reasoning, while the other two are strongly based on subsumption. Since these reasonings are the fundamental ones for description logics, they have been exhaustively studied for a great panel of languages (Baader et al., 2007). And some of these languages have been implemented throughout efficient reasoners (HermiT, 2013; RacerPro, 2013; Pellet, 2013; FaCT++, 2013). So we decided to use a very expressive description logic, namely \mathcal{SHOIQ}^+ associated with its reasoner called HermiT (Motik et al., 2009).

Description logics (DL) are a knowledge representation and reasoning formalism (Baader et al., 2007) on which is based the OWL language (OWL, 2007), one of the main standards in semantic technologies. When modelling knowledge using DL, one define elements called individuals, sets of elements called concepts and relationships between these sets called roles. In our context, individuals correspond to images and concepts and roles to gastroenterological concepts and their relationships. We refer to (Baader et al., 2007) for more information on concept and role construction, language, reasoning and complexity in DL.

We now present the content of the polyp ontology, the annotation and query mechanism.

3.1 Ontology

The core hierarchy of the ontology is presented in figure 1 (where the black arrows associated to *Flat*,

Vienna, *PitPattern* and *Paris* means that there are other subconcepts unrevealed).

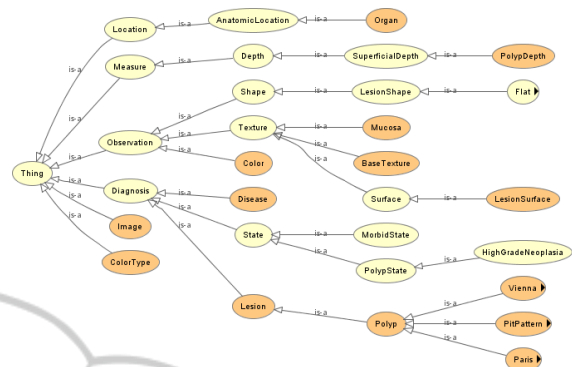


Figure 1: Core hierarchy of the polyp ontology. Snapshot taken in Protégé 4.2 (Protégé, 2013) with the OWL viz plugin.

The ontology is divided in three main parts. The first part is related to the observable properties of the image content (color, shape, texture, and anatomical properties like the considered organ and some measure about polyps). The second part is dedicated to medical comments (or even diagnosis) on the image. The last part is the definition of what is an image. This definition is tightly linked with the first two parts. It is described with the following concept description:

$$Image \equiv \exists \text{represents.} (\begin{aligned} &\exists \text{diagnosis.Diagnosis} \sqcup \\ &\exists \text{location.Location} \sqcup \\ &\exists \text{measure.Measure} \sqcup \\ &\exists \text{observation.Observation} \sqcup \\ &\exists \text{length.Float} \sqcup \\ &\exists \text{width.Float} \end{aligned})$$

This description means a polyp image may be associated to an information of type *Diagnosis* (e.g. the name of a class in a classification), to a location (an organ), to a measure (a polyp depth associated with a length and a width for the polyp dimensions), and to other observations (shape, texture, ...).

In the ontology, gastroenterological concepts come mainly from four standard classifications that have been integrated: the Paris classification (Paris, 2003) that describes polyp shapes, the Pit-Pattern classification (Kudo et al., 1994) that describes polyp surfaces, the Vienna classification (RJ et al., 2000) that describes pathological states associated to polyps and the MST3.0 (MST, 2009) that lists many gastroenterological concepts related to anatomy, observations and medical acts. Each concept coming from a classification and denoting a special set of polyps is called a class. Statistically speaking, the ontology gathers 58 defined concepts, 23 roles and 188 individuals, among them 100 images and 88 various indi-

viduals (names or gastroenterological organs, colors, textures). These data are linked through 58 subsumption relations, 40 disjoint concepts axioms and about 200 individuals assertions. It is contained into a 476 KBytes files. This is a quite small ontology since it is still in its infancy. In the future, we do not expect the TBox to contain thousands of concepts. On the contrary, the ABox size will grow linearly in the number of stored images.

The chosen language, \mathcal{SHOIQ}^+ , is a very expressive language since it allows the use of 12 concept and role constructors and 12 kinds of axioms. We refer to (Motik et al., 2009) for a precise description of these constructors and the associated semantics. It is also the language on which is built the Hermit reasoner which is one of the most effective reasoner up to our knowledge (Hermit, 2013).

3.2 Annotation and Query Building

The image annotation and queries are generated in the same manner. They are generated manually using an interactive interface. This interface allows navigation in the ontology. According to his observation (or need), the physician select most appropriate concepts and individuals for the representation of image. Formally, an image annotation (as a query) is a concept description A in conjunctive form defined as following:

$$A \equiv \exists \text{represents.} (\begin{aligned} & \prod_{i=0}^I (\exists \text{diagnosis}_i . \text{Diagnosis}_i) \sqcap \\ & \prod_{j=0}^J (\exists \text{location}_j . \text{Location}_j) \sqcap \\ & \prod_{k=0}^K (\exists \text{measure}_k . \text{Measure}_k) \sqcap \\ & \prod_{l=0}^L (\exists \text{observation}_l . \text{Observation}_l) \sqcap \\ & \exists \text{length. Float} \sqcap \\ & \exists \text{width. Float}) \end{aligned}$$

Such as :

$\text{Diagnosis}_i, \text{Location}_j, \text{Measure}_k, \text{Observation}_l$ are concepts descriptions and $\forall(i, j, k, l)$:

$$\text{Diagnosis}_i \sqsubseteq \text{Diagnosis}$$

$$\text{Location}_j \sqsubseteq \text{Location}$$

$$\text{Measure}_k \sqsubseteq \text{Measure}$$

$$\text{Observation}_l \sqsubseteq \text{Observation}$$

$\text{diagnosis}_i, \text{location}_j, \text{measure}_k, \text{observation}_l$ are roles names and $\forall(i, j, k, l)$:

$$\text{diagnosis}_i \sqsubseteq \text{diagnosis}$$

$$\text{location}_j \sqsubseteq \text{location}$$

$$\text{measure}_k \sqsubseteq \text{measure}$$

$$\text{observation}_l \sqsubseteq \text{observation}$$

One purpose of semantic annotation is to be carried out manually by experts.

The annotation(query) mechanism building is illustrated in figure 2. The user selects three concepts of ontology {stomach}, {orange} and {haemorrhagic}. The subsumers (belonging to the annotation concept definition) of these concepts will be determined in order to select the most appropriate roles for each concept. Thereafter, a concept description is built from these subsumers and roles. The result is the user annotation (or the user query).

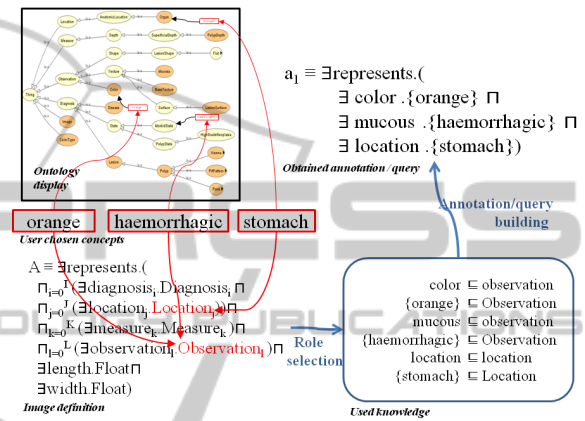


Figure 2: Annotation/query building.

Since the annotation / query building process depends on the A image concept description, then the same process (with another description) allows the generation of other types of annotations (queries), making the system easily extendable to other query needs.

3.3 Semantic Image Retrieval Process

The whole process of our semantic image retrieval approach is illustrated in figure 3.

First, DICOM images¹ are stored in a cloud database (1). The ontology (2) is linked to this database via a keyword database (3): in the keyword database are stored image identifiers linked with keywords which are concepts taken from the ontology. Moreover image identifiers are also stored in the ontology as individuals that are instance of their associated image annotations. Two modules (4) and (5) ensure the coherency between (1), (2) and (3).

Upon this knowledge infrastructure, the semantic image retrieval process runs as follows. First the system displays the subsumption hierarchy (the completed one from figure 1) computed from the ontology

¹Our prototype works on images stored into DICOM files. DICOM is a well-known file format that is widely used in medicine (DICOM, 1993).

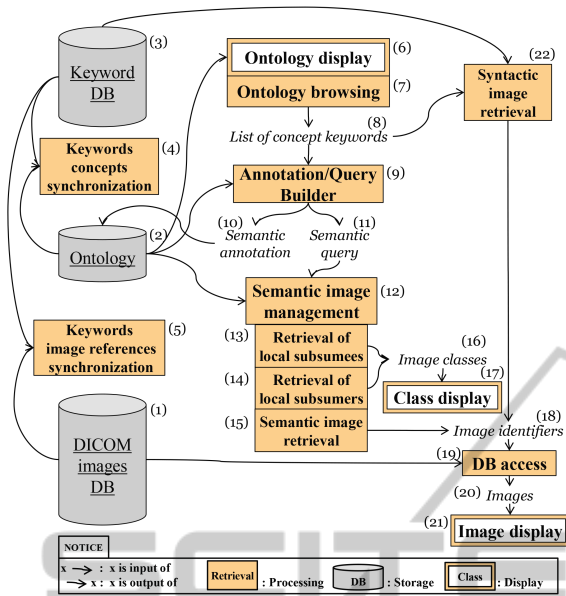


Figure 3: Semantic image retrieval approach.

(6). Then the user can browse it (7) and select a set of keywords which are concepts of the ontology (8). This set is then mapped to the previous definition of an image to obtain an image annotation. So, such an annotation is an instance of the generic definition of an image. Afterwards, this annotation (10) can either be stored in the ontology (this is the annotation scenario), or this annotation (11) can be viewed as a query (this is the semantic retrieval scenario).

In section 2, we defined three semantic retrieval cases (12): S1 (15), S2 (13) and S3 (14). We propose three DL reasonings to implement them (see tables 1 and 2). Reasoning R1 is well-known in the DL literature: we want to find all individuals (images identifiers) that belong to a given concept description (the annotation). Reasonings R2 and R3 are slight modifications of subsumption. In R2, we find all the subclasses of a given classification (e.g. Paris or Pit-Pattern) that are also subclasses of the query. So all properties defined in the query are true in these subclasses, knowing that these subclasses may have other properties not expressed in the query. In R3, we find all the superclasses of the query that are subclasses a given classification. So these superclasses share some of the properties defined in the query but not mandatorily all. That's why it is an approximation reasoning.

Once image identifiers have been obtained (18), it remains to look for (19) associated DICOM images (20). Then images can be displayed (21). Once image classes have been inferred (16), these can be displayed to the user (17). An very interesting feature in this process is that the semantic part can be inserted within a classical syntactic retrieval. Indeed, once the

Table 1: Scenarios and corrspeing reasonings.

#	Scenario	#	Reasoning	Fig. 3
S1	Semantic images retrieval	R1	Individual retrieval	(15)
S2	Exact classes retrieval	R2	Retrieval of local subsumees	(13)
S3	Approximated classes retrieval	R3	Retrieval of local subsumers	(14)

Table 2: Proposed reasonings definitions.

#	Input	Output
R1	TBox \mathcal{T} ABox \mathcal{A} Query C	All $i \in \mathcal{A}$ such that $\mathcal{T}, \mathcal{A} \models C(i)$
R2	TBox \mathcal{T} Query A Classification concept D	All concepts $C \in \mathcal{T}$ such that $C \sqsubseteq A \sqcap D$
R3	TBox \mathcal{T} Query A Classification concept D	All concepts $C \in \mathcal{T}$ such that $A \sqsubseteq C \sqsubseteq D$

list of keywords is known (8), a keyword-based search engine can be run (22) to retrieve images identifiers (18) from the keyword database (3).

3.4 Prototype and Experiments

The process described in the previous section was implemented using different technologies: the Struts 2 framework was used to implement the MVC pattern, Java Server Pages to create the interface, the OWL API for ontology management (the ontology being stored in an OWL file under an XML syntax), the HermiT reasoner (Motik et al., 2009) to be able to run classification and subsumption on the ontology, and Prefuse (Heer et al., 2005) to create the smooth interface used to display and browse the ontology.

Let's see now an execution example. By browsing the ontology as represented in figure 1, the user may select the concepts *Lesion* and *totallyFlat* (which in fact is an individual). By selecting these concepts like keywords, the user expresses that he's looking for images where man can see some lesion that is totally flat. From these keywords, the prototype may run a classical syntactic search. But, as seen in section 3.2, it may also built two semantic queries: the first will be built as the most general concept linking both *Lesion* and *totallyFlat*, and the second as the most general concept linking both *Lesion* and *totallyFlat* and being a subconcept of *Image*:

$$\text{Query} \equiv \text{Lesion} \sqcap \forall \text{observation. oneOf}\{\text{totallyFlat}\}$$

is the first query which can be associated

to a classification name (e.g. Vienna) so as the system can run reasoning R2 and R1.

$$\begin{aligned} \text{Query} \equiv & \exists \text{represents.} (\\ & \text{Lesion} \sqcup \\ & \forall \text{observation.oneOf}\{\text{totallyFlat}\} \end{aligned}$$

is the second query for which the system can retrieve all corresponding images (images that describe a totally flat lesion with other extra properties). On our current ontology, with a hundred of randomly generated images, the second query with reasoning R1 shows that 11 images correspond to these criteria, and the first query, with reasonings R2 and R3, shows that these criteria are present in the class Paris II B of the Paris classification, and in no other classification.

Anyhow, we are in the very early stages of experimentation. We plan to lead the testing of our approach in a twofold perspective. First we will qualitatively test the prototype on real images provided by gastroenterologists. The aim is to validate or invalidate parts of our ontology in order to improve it, and to test the accuracy and the usability of the proposed reasonings for practitioners. The second perspective is a quantitative study, focusing on reasoning performances. In our very first tests, we have randomly generated a hundred of images annotations and associated images identifiers. Since individuals retrieval could be achieved in an average of 30 seconds, we conclude that performances are not good. So we are now working to find the precise root cause of these bad performances.

4 RELATED WORKS

In this section, we survey the works addressing the semantic management of medical images with a focus on description logics based approaches.

4.1 Medical and Gastroenterological Ontologies

In medicine, building lexicons, taxonomies, thesaurii and classifications is not a new task. But what is new is adding them a semantics to get a real ontology, that is expressing concepts using a language that has a formal semantics, and defining concepts using complex constructions from other concepts. This is the knowledge representation principle description logics enable. The use of DL-based ontologies (especially OWL ontologies) is effective since about 15 years now (Staab and Studer, 2009; BioPortal, 2013). The first purpose of medical ontologies (OpenClinical, 2013) is to gather existing taxonomies so as to link together concepts having a same meaning but a

different name, which is frequent among taxonomies like (Galen, 2013; University, 2002; FMA, 2013; SnomedCT, 2007). We refer to (BioPortal, 2013) to a more complete panel of existing ontologies and taxonomies in medicine and biology. These are mainly used as knowledge references in information management systems which however are mainly used in an academic context (Horridge et al., 2011).

In the gastroenterological field, four classifications have been built, which are not ontologies. The MST classification (Minimal Standard Terminology for gastrointestinal endoscopy), release 3.0 (MST, 2009), contains a set of concepts related to gastroenterological anatomy, acts, observations and image capture hardware. The Paris classification (Paris, 2003) is focused on the description of polyp shapes. The Pit-Pattern classification (Kudo et al., 1994) describes polyp surface characteristics. The Vienna classification (RJ et al., 2000) contains terms related to pathological states that can be associated to polyps.

In the building of our polyp ontology, we have integrated these four gastroenterological classifications. We will integrate concepts from other medical ontologies in a second step, if it appears to improve the accuracy of the image retrieval process.

4.2 Image Annotation

A concrete usage of medical ontologies is image annotation, especially in the case of syntactic keyword-based image retrieval system. The Medico scenario in the Theseus project (Theseus, 2009) aims at setting up standards for the syntax and semantics in medical image annotation from ontologies. Our approach is quite similar in that we handle the annotation and retrieval problems using description logics. However, our domain is restricted to polyps, and moreover, our aim is less oriented towards diagnosis than towards giving physicians a semantic infrastructure to manage his/her medical images. The fact that it can be used as a diagnosis help is a consequence, but it is not the first objective.

The AIM project (AIM, 2010) aims at setting up an ontology-based standard for the annotation and the markup of medical images. Our approach differs in that we put the semantic capabilities at the heart of the system since we use a true ontology (not a lexicon) based on a DL and associated to precisely defined reasonings. The semantic features seem not to be a main objective in the AIM project.

Other works handle the issue of semantic image annotation (S. Dasmahapatra and Shadbolt, 2005; Rubin et al., 2008; Wennerberg et al., 2011). Our proposal is close to these works, differing in the used re-

trieval reasonings or in the fact that our ontology is not supposed to be distributed.

4.3 Semantic Image Retrieval Reasonings

DL image retrieval reasonings are usually quite standard, since they are basically grounded on classification (that is subsumption) and individual retrieval. What differs from one approach to another is the proximity notion that is used to qualify the good answer images wrt a query. We can find two classical approaches (Sciascio et al., 2000; Hu et al., 2003; Rubin et al., 2008; Opitz et al., 2009) which correspond to our R1, which is the classical individuals retrieval, and the composition of R2 followed by R1, which amounts to finding images associated to concepts that have the same properties as the query (and maybe others). Other approaches are based on non-standard DL reasonings (abduction and contraction) (Di Noia et al., 2005; Colucci et al., 2011), which imply, however, to use a less expressive DL. These reasonings enable a better ranking of answers than the previous ones. However we already handle the ranking issue by allowing the user to interact with the query interface.

5 CONCLUSIONS AND PERSPECTIVES

This work is, up to our knowledge, the first attempt to use a semantic approach to manage endoscopic images. If our tests bring validation, then it could be experimented also in other domains of medicine.

In summary, we have proposed a complete semantic image annotation and retrieval approach grounded on a new polyp ontology expressed in the \mathcal{SHOIQ}^+ description logic. It still need to be extensively tested since many issues are addressed: the reasonings that is at the core of the mechanism need to be validated and then optimized to get close to a real time processing, the polyp ontology content need to be validated and/or updated to make the reasonings more accurate, and the annotation interface need to be validated and criticized by end users (physicians and surgeons) so it can be improved to match their everyday use.

As our images are stored in the cloud, an interesting perspective of this work would be to benefit also from the computation capacities of the cloud by moving our reasoning computation there. Since a few years, some works adress this reasoning on the cloud problem (Mutharaju et al., 2010; Aslani and Haarslev,

2012).

Another perspective of this work, to improve performances of the individuals retrieval reasoning, would be to use an ontological query answering approach. In this approach, the aim is to translate the query and the ontology in a classical relational database context so as to benefit from the performances of the existing optimized (relational) database management systems (Calì et al., 2009; Lenzerini, 2011).

At last, another perspective of this work is related to the ranking of results. Now, this ranking is achieved by the interaction between the user and the browsing interface. It should be worthy to try to rewrite the polyp ontology into a less expressive DL so as to be able to apply more flexible reasonings like those proposed in (Di Noia et al., 2005; Colucci et al., 2011).

ACKNOWLEDGEMENTS

This work is supported in part by the Agence Nationale de la Recherche (under grant SYSEO ANR-10-TECSAN-005-01), the Conseil Régional d'Auvergne and SGAR.

REFERENCES

- AIM (2010). Annotation and image markup (aim) project. <https://cabig.nci.nih.gov/community/tools/AIM>.
- Aslani, M. and Haarslev, V. (2012). Concurrent classification of owl ontologies - an empirical evaluation. In *Description Logics*.
- Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., and Patel-Schneider, P. F., editors (2007). *The Description Logic Handbook: Theory, Implementation, and Applications (2nd Edition)*. Cambridge University Press.
- Bimbo, A. D. (1999). *Visual information retrieval*. Morgan Kaufmann.
- BioPortal (2013). Bioportal. <http://bioportal.bioontology.org/>.
- Calì, A., Gottlob, G., and Lukasiewicz, T. (2009). A general datalog-based framework for tractable query answering over ontologies. In *PODS*, pages 77–86.
- Colucci, S., Noia, T. D., Sciascio, E. D., Donini, F. M., and Mongiello, M. (2011). *Description Logic-Based Resource Retrieval*, pages 185–197. Encyclopedia of Knowledge Management.
- Datta, R., Joshi, D., Li, J., and Wang, J. Z. (2008). Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2).
- Di Noia, T., Di Sciascio, E., Donini, F. M., di Cugno, F., and Tinelli, E. (2005). Non-standard inferences for

- knowledge-based image retrieval. In *EWIMT 2005 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media Techniques*, IEE press, pages 191–197. IEE.
- DICOM (1993). Digital imaging and communications in medicine. <http://medical.nema.org/>.
- FACT++ (2013). The fact++ owl-dl and (partial) owl-dl reasoner. <http://code.google.com/p/factplusplus/>.
- FMA (2013). Foundational model anatomy. <http://sig.biostr.washington.edu/projects/fm/>.
- Galen (2013). Galen and the galen-core high-level ontology for medicine. <http://www.opengalen.org/>.
- Gruber, T. (2009). *Encyclopedia of Database Systems*, chapter Ontology. Springer-Verlag.
- Heer, J., Card, S. K., and Landay, J. A. (2005). prefuse: a toolkit for interactive information visualization. In *CHI*, pages 421–430.
- Hermit (2013). The hermit owl2 reasoner. <http://www.hermit-reasoner.com/>.
- Horridge, M., Parsia, B., and Sattler, U. (2011). The state of bio-medical ontologies. *Bio-Ontologies*. <http://bio-ontologies.knowledgeblog.org/135>.
- Hu, B., Dasmahapatra, S., Lewis, P. H., and Shadbolt, N. (2003). Ontology-based medical image annotation with description logics. In *ICTAI*, pages 77–.
- Kudo, S., Hirota, S., Nakajima, T., Hosobe, S., Kusaka, H., Kobayashi, T., Himori, M., and Yagyuu, A. (1994). Colorectal tumours and pit pattern. *J Clin Pathol*, 47.
- Lenzerini, M. (2011). Ontology-based data management. In *Proceedings of the 20th ACM international conference on Information and knowledge management, CIKM '11*.
- Meghini, C., Sebastiani, F., and Straccia, U. (2001). A model of multimedia information retrieval. *J. ACM*, 48(5):909–970.
- Motik, B., Shearer, R., and Horrocks, I. (2009). Hyper-tableau reasoning for description logics. *J. Artif. Intell. Res. (JAIR)*, 36:165–228.
- MST (2009). Mst - minimal standard terminology for gastrointestinal endoscopy. <http://www.worldendo.org/mst.html>.
- Mutharaju, R., Maier, F., and Hitzler, P. (2010). A mapreduce algorithm for el+. In *Description Logics*.
- OpenClinical (2013). Openclinical: knowledge management for medical care. <http://www.openclinical.org/ontologies.html>.
- Opitz, J., Parsia, B., and Sattler, U. (2009). Using ontologies for medical image retrieval - an experiment. In *OWLED*.
- OWL (2007). Owl, the web ontology language. <http://www.w3.org/2007/OWL>.
- Paris (2003). The paris endoscopic classification of superficial neoplastic lesions: esophagus, stomach, and colon: November 30 to december 1, 2002. *Gastrointestinal Endoscopy*, 58(6 Suppl):3–43.
- Pellet (2013). The pellet owl2 reasoner. <http://clarkparsia.com/pellet/>.
- Protégé (2013). The protégé open source ontology editor and knowledge-base framework. <http://protege.stanford.edu>.
- RacerPro (2013). The racerpro owl/rdf reasoner. <http://www.racer-systems.com/>.
- RJ, S., RH, R., Y, K., F, B., and al (2000). The vienna classification of gastrointestinal epithelial neoplasia. *Gut*, 47:251–255.
- Rubin, D. L., Mongkolwat, P., Kleper, V., Supekar, K., and Channin, D. S. (2008). Medical imaging on the semantic web: Annotation and image markup. In *AAAI Spring Symposium: Semantic Scientific Knowledge Integration*, pages 93–98. AAAI.
- Rui, Y., Huang, T. S., and Chang, S.-F. (1999). Image retrieval: current techniques, promising directions and open issues. *Journal of Visual Communication and Image Representation*, 10:39–62.
- S. Dasmahapatra, D. Dupplaw, B. H. P. L. and Shadbolt, N. (2005). Ontology-mediated distributed decision support for breast cancer. In *AIME 2005, LNAI 3581*, page 221?225. Springer-Verlag Berlin Heidelberg.
- Sciascio, E. D., Donini, F. M., and Mongiello, M. (2000). Semantic indexing in image retrieval using description logic. In *Proceedings of the 22nd International Conference on Information Technology Interfaces*.
- SnomedCT (2007). Systematized nomenclature of medicine - clinical terms. http://www.nlm.nih.gov/research/umls/Snomed/snomed_main.html.
- Staab, S. and Studer, R., editors (2009). *Handbook on Ontologies*, volume XIX of *International Handbooks on Information Systems*. 2nd ed. edition.
- Syseo (2011). Multimodal and multimedia image analysis and collaborative networking for digestive endoscopy. <http://www.syseo-anr.fr/Home->.
- Theseus (2009). Theseus project, medico scenario. <http://theseus.pt-dlr.de/en/920.php>.
- University, S. (2002). The institute for formal ontology and medical information science. <http://ifomis.org/>.
- Wennerberg, P., Schulz, K., and Buitelaar, P. (2011). Ontology modularization to improve semantic medical image annotation. *Journal of Biomedical Informatics*, 44(1):155–162.