

Living with the Semantic Gap: Experiences and Remedies in the Context of Medical Imaging

Yannis Kalfoglou, Srinandan Dasmahapatra, David Dupplaw, Bo Hu, Paul Lewis, and Nigel Shadbolt

Abstract—Semantic annotation of images is a key concern for the newly emerged applications of semantic multimedia. Machine processable descriptions of images make it possible to automate a variety of tasks from search and discovery to composition and collage of image data bases. However, the ever occurring problem of the semantic gap between the low level descriptors and the high level interpretation of an image poses new challenges and needs to be addressed before the full potential of semantic multimedia can be realised. We explore the possibilities and lessons learnt with applied semantic multimedia from our engagement with medical imaging where we deployed ontologies and a novel distributed architecture to provide semantic annotation, decision support and methods for tackling the semantic gap problem.

Index Terms—Semantics, Medical Imaging, Ontologies, Description Logics.

I. MEDICAL IMAGING DOMAIN

Advances in medical technology generate huge amounts of non-textual information, like images and other multimedia, along with more conventional media like text reports. Most of the existing systems focusing on extracting visual cues with the aid of image analysis algorithms may experience problems when a rather abstract and ambiguous query is asked [1]. Because low level descriptors cannot be uniquely associated with any other meaningful label unless explicitly declared or derived as the outcome of a classification procedure, retrieval based on knowledge level constructs is a non-trivial task to achieve in general. Our domain of exploration is medical imaging, in particular providing semantic support for the breast cancer screening processes. In the context of the MIAKT project (Medical Imaging and Advanced Knowledge Technologies)¹ we built the Breast Cancer Imaging Ontology (BCIO). It consists of several relatively independent modules at different levels of granularity with uniform interfaces to enable integration. Separations are defined vertically and horizontally. Because the patients are viewed through different apparatuses and instrumentation whose results are overlaid and compiled to give the whole picture, a natural vertical separation of the domain would be one module for each imaging method, X-ray, MRI, Ultra-Sound, etc. Each imaging module is composed of a set of image feature descriptors, a set of diagnosis descriptors capturing high-level abstract features

and a set of concepts for describing meta-image information. For instance, image analysts and/or radiographic technicians might focus more on low-level graphic features, e.g. shape, size and luminosity of an ROI (Region Of Interest) while a radiologist might step away from the fine details and concentrate on the interpretation of all ROIs within the context of the whole image. While looking into each individual imaging module, we define not only image descriptors but also image-capture related concepts. However, knowing what is on an image is sometimes not sufficient for domain experts to give a proper interpretation of the image, as knowledge on how the image is produced is equally critical. In practice, deciding whether such knowledge should be included or ignored is a trade-off between the complexity and the accuracy of the ontology. BCIO is trying to satisfy the requirements on both usability and extensibility. BCIO provides handles for the information pertaining to a particular case based on different aspects and different grain-sizes that are appropriate. This allows an expert to focus only on the facets relevant to her interest and/or expertise and makes available her interest BCIO is supported by a distributed architecture which enables a number of web-based services that provide discrete and disparate functionality to a generic application base. We provide annotation support where users can annotate an ROI of their choice with the aid of a graphical editor. We also support semantic querying on the BCIO concept and instance descriptions. BCIO is DL-based and we employ DL-based inference to provide automatic classification of query constructs using the underlying BCIO ontology.

II. DEALING WITH THE SEMANTIC GAP

When we consider the semantic gap as described by Hare and colleagues in [2]: "much of the interesting work which is attempting to bridge the gap automatically is tackling the gap between descriptors and the labels and not that between the labels and the full semantics"; the question is how this gap can be narrowed, with which technology and under which assumptions. We reformulate this question as a question of applying ontology mapping techniques to tackle it. As long as we have available a codified representation of the *full semantics* Hare and colleagues are referring to, then ontology mapping is a feasible approach with a lot of potential applications. We first apply a technique which allows us to wrap the different interpretations of an image into a single representation. But, there are situations where we need to preserve these different interpretations and align them for the sake of enabling interoperability. To do that, we employ semantic alignment. It is a subset of a bigger set of technologies that aim to find alignments of entire ontologies, like for instance ontology mapping and

Y.Kalfoglou and N.Shadbolt are with the Advanced Knowledge Technologies (AKT) group; email: {y.kalfoglou|nrs}@ecs.soton.ac.uk. Y.Kalfoglou is the corresponding author.

S.Dasmahapatra is with the Science and Engineering Natural Systems (SENSe) group. Email: sd@ecs.soton.ac.uk .

D.Dupplaw, B.Hu, and P.Lewis are with the Intelligence Agents Multimedia (IAM) group. Email: {d|p|b|h|p|h}@ecs.soton.ac.uk.

All authors are with the School of Electronics and Computer Science (ECS), University of Southampton, Southampton SO17 1BJ, UK.

¹ More on: www.aktors.org/miak2

alignment tools. To tackle the problem of finding alignments in order to narrow the semantic gap between semantics of the multimedia object and its label, we deploy the idea of semantic alignment which aims to use a subset of an ontology mapping system. In particular, we use semantic metrics [4] to discover alignments between ontological structures which could range from concept name alignments to simple string matching. We devised a modular architecture for deploying ontology mapping systems, most of which can provide us with semantic alignments of the structures we are interested in. The principle behind our modular architecture is that there is a variety of alignment systems out there and we are keen to use them conjunctively for the benefit of a better informed alignment. For example, in the context of mapping a specialists medical vocabulary, like the Foundational Model of Anatomy (FMA) large OWL ontology, to a generic medical model, like the OpenGALEN ontology, we found that we can improve the results of mapping by deploying different ontology mapping systems; each of which provides a cutting edge in different alignment algorithms [5]. The architecture is implemented as a multi-stage and multi-strategy system comprising of four modules, namely, *Feature Generation*, *Feature Selection and Processing*, *Aggregator* and *Evaluator*. In this system, different features of the input ontologies are generated and selected to fire off different kinds of feature matchers, which are an integral part of many ontology mapping systems. The resultant similarity values are compiled by multiple similarity aggregators running in parallel or consecutive order. The overall similarity is then evaluated to initiate iterations that backtrack to different stages.

III. CONSIDERATIONS

The MIAKT architecture was designed to provide general knowledge management in many domains, while allowing semantically marked-up services to be easily integrated into a knowledge management application. We applied three guidelines when designing the architecture: (a) it must allow institutions to retain control over their own data. In the medical imaging domain, this means that a hospital's radiography unit retains control of the images that they produce by having them stored on an institutional image server; (b) it must provide simple and fast extensibility. This means that new services can be imported into the architecture quickly and easily; (c) a major consideration which proved a challenge in its implementation, is that the architecture must provide enough flexibility to be able to be used in different application domains. Ensuring the distinction between domain and non-domain data, both ontological and otherwise, means that a great deal of care has to be taken in designing the interfaces between components. Utilising ontologies for the description of the system as well as the domain has helped us achieve this goal. Another consideration is that of annotation. Although annotation of images with the use of a rich semantic substratum (ideally codified in an ontology) has been explored in the literature (see, for example, [3]), we are looking forward to a more expressive annotation regime that goes beyond memorable object labels and attaches what Hare and colleagues refer to as "full semantics". There might be practical problems with the attachment and location and use of those codified semantics on the image in

question, but practitioners like to work with more and more expressive and detailed descriptions of an image (especially in specialist domains, like medical imaging). Finally, an issue which is related to that of annotation, is the use and analysis of user queries to assist with annotation and tackling the semantic gap. In their brief survey of user based queries analysis, Hare and colleagues point out a number of techniques that aim to classify user queries according to a given organisation. As most of that work is focused on inferring the context of the query, thus assisting in finding the most relevant image, we are interested to see whether such an analysis could be used alongside other contextual information to help narrow the semantic gap between the object labels (image in question) and the full semantics (image description).

Bridging the semantic gap in visual information retrieval is a key enabler for deploying semantic multimedia applications. Our experience with the MIAKT project and the breast cancer screening process support highlighted that this phenomenon exists and is fuelled by the inevitably different interpretations of images by different domain experts. We presented an approach to tackle it based on a simple wrapping technique, and we are considering deploying ontology mapping technology to cope with the more semantically rich heterogeneous descriptions of those images (if available). However, we also advocate that the use of a flexible and distributed architecture with reasoning support, such as the one we presented earlier, is an important infrastructure that needs to be in place to support semantic multimedia applications.

ACKNOWLEDGMENT

All the authors except Kalfoglou and Shadbolt are supported fully or partially by the OpenKnowledge and HealthAgents projects which are funded by the EU IST Framework 6 under Grant numbers IST-FP6-027253 and IST-FP6-027213. Kalfoglou and Shadbolt are fully supported by the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC) project which is sponsored by the UK EPSRC under Grant number GR/N15764/01. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the EU IST FP6 or the UK EPSRC or any other member of the UK AKT IRC.

REFERENCES

- [1] E. Di Sciascio, F.M. Donini and M. Mongelio. Structured Knowledge Representation for Image Retrieval. *Journal Of Artificial Intelligence Research*, 16:209-257, 2002.
- [2] J.Hare, P.Lewis, P.Enser, and C.Sandom. Mind the gap: Another look at the problem of the semantic gap in image retrieval. In *Proceedings of Multimedia Content Analysis, Management and Retrieval 2006 SPIE Vol.6073*, pp.607309-1, San Jose, California, USA., 2006
- [3] L.Hollink, G.Shreiber, B.Wielemaker, and B.Wielinga. Semantic annotations for image collections. In *Proceedings of the K-CAP'03 Workshop on Knowledge Markup and Semantic Annotation, Banff, Canada*, October 2004.
- [4] B.Hu, Y.Kalfoglou, H.Alani, D.Dupplaw, P.Lewis, and N.Shadbolt. Semantic Metrics. In *Proceedings of the 15th International Conference on Knowledge Engineering and Knowledge Management (EKAW'06)*, Podedbrady, Czech Republic, October 2006.
- [5] Y.Kalfoglou and B.Hu. CMS: CROSI Mapping System – Results of the 2005 Ontology Alignment Contest. In *Proceedings of the K-CAP'-5 Integrating Ontologies workshop, Banff, Canada, pages 77-84*, October 2005.