Ontology-based Medical Image Annotation with Description Logics

Bo Hu, Srinandan Dasmahapatra, Paul Lewis, and Nigel Shadbolt

Department of Electronics and Computer Science
University of Southampton, United Kingdom
e-mail: {bh, sd, phl, nrs}@ecs.soton.ac.uk

Abstract

The interpretation of medical evidence is normally presented in terms of a controlled, but diversely expressed specialist vocabulary and natural language phrases. Such informally expressed data require human intervention to ascertain its relevance in any specific case. In order to facilitate machine-based reasoning about the evidence gathered, additional interpretive semantics must be attached to the data; a shift from a merely data-intensive approach to a semantics-rich model of evidence. In this paper, we present a system to formally annotate medical images captured to aid the diagnosis and management of breast cancer, that enables a series of semantics-based operations to be performed. Our approach is grounded upon an imaging ontology specifying the domain knowledge and a Description Logic (DL) taxonomic inferential engine responsible for semantics-based reasoning and image retrieval.

1. Introduction

Advances in medical technology generate huge amounts of non-textual information (e.g. images) along with more familiar textual one. 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 [6]. They face difficulties in manipulating abstract interpretations of images if such semantics is not represented. For instance, it is difficult to answer whether "X-ray images with round calcifications" relate more closely to "images with calcifications" or "images with round shape abnormalities" due to that both "low-level" descriptors such as *shape* and "high-level" terms such *calcification* present in a single query. Text-based approaches resolve the ambiguity by hiring a keyword searching mechanism and presenting a range of potential results to the end-user who performs the necessary disambiguation. However, for large amounts of data, the search and potential-solution spaces make it increasingly problematic for medical staff to locate useful information. Although this "meaningless" mechanism benefits from ease of implementation, the potential loss of crucial materials which are relevant, but which may not have similar strings that are matched is a big price to pay [17].

Such problems can be tackled by representing the available information using languages such as DAML+OIL [15], OWL [16], which benefit from their underlying Description Logic (DL) foundation. They have a uniform syntax and unambiguous semantics defined for each and every conceptual constructor that enables the automated reasoning based on definitions of domain knowledge and relationships among them [2]. As regards semantics-based image annotation and retrieval, logic-based approaches are capable of formally describing facts that explicitly presented and deriving implicit knowledge from given premises [18]. The advantages of a DL-based approach are especially evident when considering medical images whose interpretation involves substantial domain knowledge.

In this paper we outline our experience of building a semantically rich system by accommodating image annotation and retrieval services around a rigidly defined ontology for medical images used in breast cancer treatment. Our approach is based on the assumptions that: given the breast cancer images, (i) expressing all the desired features using domain knowledge is feasible; (ii) manually marking up and annotating the regions of interest is practical; and (iii) representing and reasoning about the textual descriptions are performed with a reasonable complexity in a given querying context. This work is carried out within a project which seeks the applicability of the system developed to the

UK National Health Breast Screening Programme as a goal.

The presentation is structured as follows. Section 2 sketches the breast cancer screening process. Section 3 introduces the imaging ontology and the selected DL-based modelling language. Section 4 discusses the ontology-based image annotation and retrieval facilitated by DLs. Finally, Section 5 summarises our approach and suggests some further developments.

2. Breast Screening Process

Breast cancer is one of the leading causes of cancer death among women in the US [20] and is the most common cancer for women in the UK[1]. Diagnosis of breast cancer normally involves multi-disciplinary meetings with experts from different medical backgrounds, e.g. radiologists, surgeons, oncologists, histologists and other clinical staff. Appropriate medical images, with expert description and diagnosis is presented as evidence in these meetings. As with other disciplines, we expect considerable variability among experts or even between two interpretations carried out by the same expert. Hence, it provides the motivation to make uniform the vocabulary used in the breast cancer screening and diagnosis process.

A typical scenario of breast cancer assessment process starts with a report from routine X-ray mammographic examinations or self-reported abnormal symptoms followed by an X-ray mammographic examination. Therefore, a very good starting point of formally representing the breast screening process would be Xray mammography. Meanwhile, studies show that, when a definitely benign diagnosis cannot be made, breast magnetic resonance imaging (MRI) is normally treated as an expensive complementary method to X-rays to increase diagnostic confidence [23].

3. Breast Cancer Imaging Ontology

The aim of the Breast Cancer Imaging Ontology (BCIO) is to provide a commonly agreed vocabulary with formal definitions that can be used to represent breast X-ray and MRI images, abnormal findings and medical assessments. Such an ontology can not only facilitate the knowledge share, reuse and semantics-based reasoning over images but also help to standardise medical image interpretation and thus, to some extent, help improve breast cancer treatment.

1 Cancer Research UK, Retrieved on 10 June 2003 from World Wide Web: http://www.cancerresearchuk.org

3.1. Imaging Ontology

The American College of Radiology (ACR) has proposed a standard for X-ray mammography reports, called the Breast Imaging Reporting and Data System (BI-RADS) [1], which has been extensively evaluated using real-life cases [14, 13]. BI-RADS It provides a comprehensive lexicon for describing mammographic findings containing both X-ray-specific terms such as *image descriptors* (e.g. the shape of the lesion, the texture of the lesion), *lesion types* (e.g. calcification, mass), and terms related to the *breast cancer types* and *patient management*. The ontology for breast MRI is also based on existing studies under way for standardization published in ACR publications [5, 12, 19]. We considered this published vocabulary as a good starting point for building the ontology because of constraints on the availability of radiologists during this phase of the work. The applicability of these lexicons among UK practitioners will need further validation.

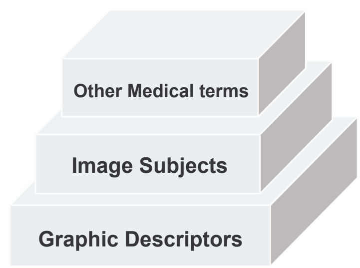

Figure 1. Multi-level model of BCIO

BCIO is a multi-level model (Figure 1) containing a set of graphic feature *descriptors*, a set of *image subjects* and a set of medical terms with references to the *image subjects*. The graphic *descriptors* are terms that describe regions of interest (ROI) on the image which can then be associated with instances of the *image subjects* in the ontology. For instance, as shown in Figure 2, the descriptive features of a delimited region corresponding to the concept Region-Of-Interest are represented using *descriptors* such as Texture, Shape, Margin, *etc.*. Subsequently, an instance of Region-Of-Interest can be used to construct other abstract objects, e.g. the instances of Medical-Image, by referring to the instances of Region-Of-Interest via the property contains.

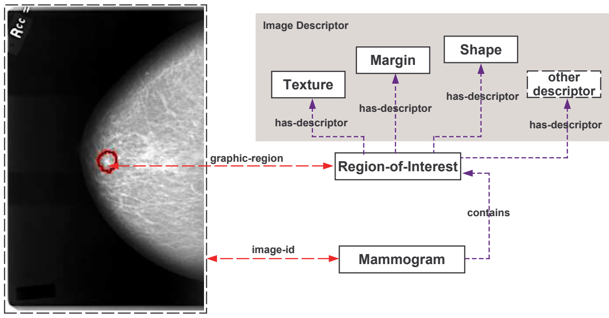

Figure 2. Region-Of-Interest and Mammogram

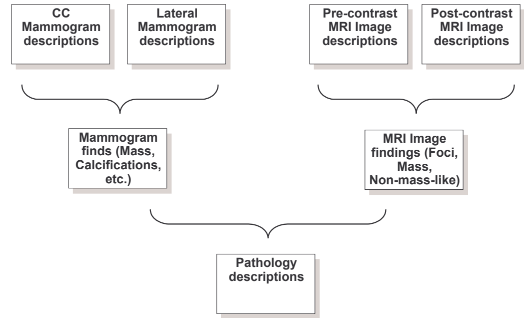

Figure 3. Integrating MRI and X-ray findings

3.2. Integration of MRI and Mammographic Ontologies

Although, the MRI and X-ray imaging parts of the ontology are kept independent to fulfil a modular design philosophy, in general, it is meaningful to fuse the data acquired from these two different imaging methods [3]. However, MRI and X-ray mammographic images are inherently disparate in several aspects that prevent a straightforward integration. These include: i) MRI presents volume-based dynamic information of the breast with the help of contrast agents while X-ray provides a high resolution, static 2D projection from lateral and/or cranial-caudal views; ii) patients take different positions: standing upright with the breast highly compressed for the X-ray mammographic imaging and lying facedown with the breasts uncompressed in the breast coils for the MRI; and iii) studies show that the radiologists tend to model the abnormalities presented on MRI and X-ray images in different ways [12]–Shape and Margin for describing X-ray abnormalities while combined shape-margin descriptors for MRI findings.

Such differences suggest the difficulty of mapping MRI and X-ray imaging techniques at the representation level. However, despite such disagreement, the two components can be associated at the pathology level as shown in Figure 3.

3.3. Description Logic-based Modelling

DLs are a family of knowledge representation and reasoning formalisms (KR&R) which have attracted substantial research recently, specially, after the DL ontology modelling languages (e.g. DAML+OIL [9]) are considered to be of crucial importance for the Semantic Web initiative [4]. DLs are based on the notions of concepts (i.e. unary predicates) and properties (i.e. binary relations). Using different constructors, complex concepts can be built up from simple ones. With its unambiguous semantics, DLs lend themselves to powerful reasoning algorithms which can automatically classify and memoise the domain knowledge to provide a cache for further inderences.

It is acknowledged that selecting a DL language with the right expressive power is crucial as a “trade-off” between the language expressiveness and the computational complexity is inevitable [2]. As regards our ontology, the $\mathcal{SHIQ}$ DL [2] is considered as the most suitable modelling language under both theoretical and practical considerations. On one hand, the characteristics of the imaging ontology suggest that the selected modelling language should be equipped with the expressive and deductive power for qualified role value restrictions (both existential and universal), role hierarchies, inverse roles and qualified role number restrictions, features of which are covered by the $\mathcal{SHIQ}$ language. On the other hand, several implemented DL-based inferential engines, such as FaCT [8] and RACER [7], support reasoning for $\mathcal{SHIQ}$ with acceptable complexity. Meanwhile, RACER also supports the reasoning with regard to $\mathcal{SHIQ}$ individuals—an important functionality for reasoning with BCIO. The syntax and semantics of $\mathcal{SHIQ}$ DL can be found in [2].

The worst-case computational complexity of $\mathcal{SHIQ}$ is ExpTime-Complete [10]. Nevertheless, with suitable optimisations, systems such as FaCT and RACER demonstrate very good practical performance. Moreover, when constructing BCIO, General Concept Introduction is ruled out—only simple concept names are allowed as the left hand operands in a concept introduction axiom—so as to reduce the computational complexity [2].

4. Processing Images with DLs

In our system, both image annotation and retrieval are supervised by a DL-based inferential engine as shown in Figure 4.

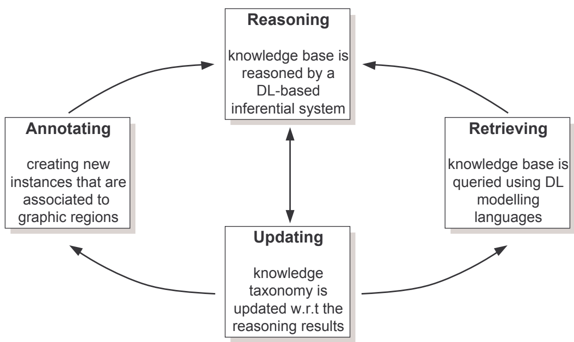

Figure 4. Control algorithm for image annotation and retrieval

4.1. Annotating Images with BCIO

The ontology can be loaded both from a DAML file on the disk or from a "live" database facilitated by MySQL[2] database server. In order to protect the original taxonomy, users are only granted limited power when using the graphical editor,and are only allowed to remove instances. and created When creating new instances, data-base style input forms guide the knowledge acquisition process (Figure 5).

When there are no images, users can populate the knowledge base with non-image-related instances, e.g. a particular Patient, a new Shape descriptor, *etc.*. Users can also create image-related instances in the image annotation interface (Figure 5). Currently, a ROI is manually delimited on a transparent marking panel covering the entire area of the image and assigned a unique string as the ID referred to by the property graphic-region. The marked up areas are represented using `Polygons` and encoded as (and decoded from) XML files (Figure 5) with the help of the Java XMLEncoder package (and XMLDecoder package, respectively). As a result, graphical features of the ROI can be stored and transferred persistently in XML format separated from the actual images and reloaded using a unique "marking panel" ID.

2 Available from http://www.mysql.com/.

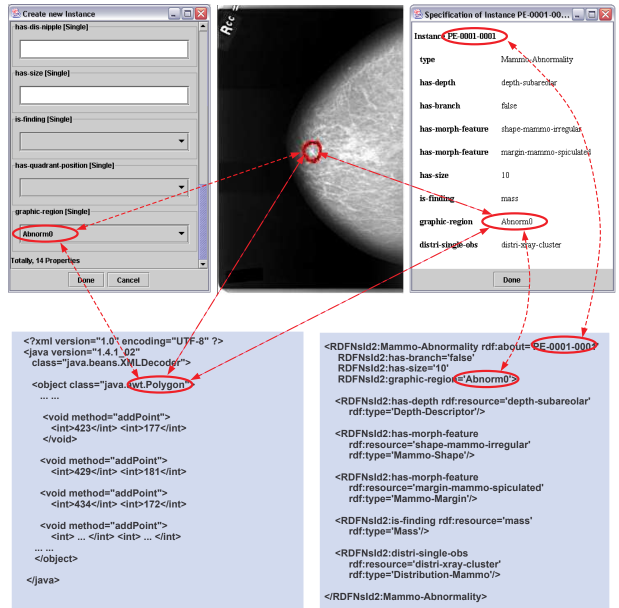

Figure 5. Creating a new instance of Mammogram Abnormality

4.2. Populating the Knowledge Base

As pointed out in previous section, extending BCIO axioms by end users after the creation of the ontology is not allowed, except for populating with new instances in order to prevent *ad hoc* extensions that may introduce inconsistencies. Although our system provides a certain degree of assistance for constructing ontologies, such assistance is limited and geared to be of help to those familiar with knowledge representation. Medical experts are likely to find constructing ontologies difficult, especially, the use of existential (**some**, $\exists$) and universal (**all**, $\forall$) quantifications [22]. As a result, the task of modifying the conceptual structure of ontologies is normally left to knowledge engineers. Moreover, we do not expect the medical staff to be keen and active in extending the BCIO axioms.

The knowledge acquisition with regard to concept instantiation is facilitated by a database-style input form. Because properties in BCIO are defined in a DL-based approach, it is not straightforward to present concepts in a database style. We solve the problem by normalising the property definitions prior to the input. The normalisation is carried out by associating each concept with a set of properties which contains all the properties defined explicitly for this concept and those defined explicitly for all its sub-concepts. Such a property normalisation procedure can be formalised as:

Let C and D be two DL concepts, Π_{C} and Π_{D} be the sets of properties defined explicitly for C and D respectively, Π_{C} is recursively normalised as the union of all normalised Π_{D} for every D subsumed (directly or indirectly) by C, i.e.

$$\mathtt{normalise}(\Pi_{\mathsf{C}}) = \bigcup_{\text{for all } \mathsf{D} \sqsubseteq \mathsf{C}} \mathtt{normalise}(\Pi_{\mathsf{D}})$$

For instance, given MRI-Image $\sqsubseteq$ Medical-Image defined as in Exp (1), if one wants to create an instance of Medical-Image, the set of properties must be normalised to include use-contrast-agent, although use-contrast-agent is explicitly defined only for the child MRI-Image rather than the parent concept.

$$\begin{aligned}\mathsf{MRI\text{-}Image} \doteq\ & \mathsf{Medical\text{-}Image} \sqcap \\ & (\mathbf{all}\ \mathsf{use\text{-}contrast\text{-}agent\ Contrast\text{-}Media}) \ldots\end{aligned} \quad (1)$$

When introducing new BCIO instances, it is not mandatory to instantiate the exact concept. A DL-based inferential engine can "fine tune" the hierarchial structure based on the concept and instance definitions. For instance, an individual image image-0001 can be specified initially to instantiate Medical-Image (Figure 6(a)). With proper values assign to its properties (e.g. (**all** literal-side Right), a DL-based inferential engine can re-classify and fine-tune image-0001 to the most suitable location in the conceptual hierarchies (e.g. image-0001∈RCC-Mamm-Image, Figure 6(b)).

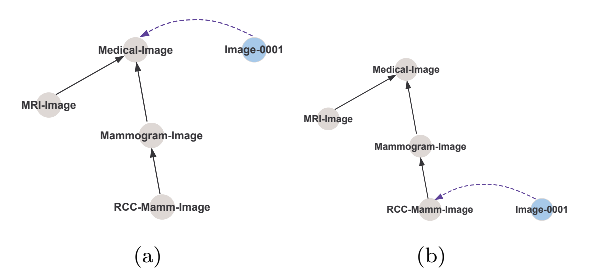

Figure 6. Introducing new image instance

4.3. Semantics-based Querying

Semantics-based querying can be performed with the help of BCIO-based concept and instance descriptions. DLs allow more freedom to carry out such enquiries. For instance, in addition to querying with the more precise description presented in Query (2) to retrieve instances, one can take a more fuzzy approach by asking for "all the round ROIs with clearly defined margin" or "all the image containing irregular ROIs" by sending queries Query (3) or Query (4) respectively.

$$\begin{aligned}\mathsf{Mamm\text{-}Abnor} \sqcap\ & (\mathbf{some}\ \mathsf{has\text{-}morph\text{-}feature\ shape\text{-}irregular}) \\ \sqcap\ & (\mathbf{some}\ \mathsf{has\text{-}morph\text{-}feature\ shape\text{-}irregular}) \\ \sqcap\ & (\mathbf{some}\ \mathsf{has\text{-}morph\text{-}feature\ margin\text{-}spiculated}) \\ \sqcap\ & (\mathbf{all}\ \mathsf{has\text{-}size}\ 10) \ldots\end{aligned} \quad (2)$$

$$\mathsf{Region\text{-}Of\text{-}Interest} \sqcap (\mathbf{some}\ \mathsf{has\text{-}morph\text{-}feature\ shape\text{-}round}) \quad (3)$$

$$\begin{aligned}\mathsf{Medical\text{-}Image} \sqcap\ & (\mathbf{some}\ \mathsf{contains} \\ & (\mathbf{some}\ \mathsf{has\text{-}morph\text{-}feature\ shape\text{-}irregular}))\end{aligned} \quad (4)$$

Such approximate approach is facilitated by the DL-based automated classification. When a query C_Q like Query (4) is sent to the knowledge base, a DL-based inferential engine treating C_Q as a normal concept definition classifies the knowledge base and locates the most general subsumee C_{sub} for C_Q with regard to the entire knowledge base, $\mathcal{T}$, i.e.

$$\begin{aligned}&\mathtt{mgs}(\mathsf{C}_Q) = \mathsf{C}_{\text{sub}} \models_{\mathcal{T}} \\ &\quad (\mathsf{C}_{\text{sub}} \sqsubseteq \mathsf{C}_Q) \wedge (\forall \mathsf{C}' \not\equiv \mathsf{C}_{\text{sub}} \in \mathcal{T}.(\mathsf{C}' \sqsubseteq \mathsf{C}_Q \Rightarrow \mathsf{C}' \sqsubseteq \mathsf{C}_{\text{sub}}))\end{aligned}$$

Hence, all the instances defined as direct or indirect instances of C_{sub} can be retrieved as the answer to the query. For instance, in the above example, Query (4) will first trigger the classification of concept (**some** has-morph-feature shape-irregular) which has all the irregular Region-Of-Interest as its children. Further reasoning can insert the concept equivalent to the query as a sub-concept of Medical-Image and a super-concept of all those Medical-Image containing irregular Region-Of-Interest. As a result, all the images that meet such restrictions are retrieved.

5. Conclusions

There is a growing interest in ontology-based image annotation and retrieval [11][21]. With carefully selected vocabulary, contents of an image or the meta-data of the image itself can be formally modelled with textual descriptions and thus processed by computers. The work presented in this paper is different from other apparently similar ones as we take full advantage of the DL-based reasoning when constructing the ontology and annotating and retrieving the images.

There are several reasons why a DL language is suitable for our ontology-based approach that motivated our choice. A DL language has uniform syntax, unambiguous semantics and clear separation between concepts (classes or templates) and instances (individuals or data). It is adopted as the underlying KR&R model of various web-ontology languages, such as DAML+OIL, OWL. Unfortunately, the "trade-off" between expressive and deductive powers and the abstract nature prevent DLs from directly supporting multimedia information, such as image, video, *etc.*. However, by annotating an image using properly defined descriptors, a DL-based inferential engine can answer semantics-based queries with regard to images and return appropriate pointers to the images.

In this paper, we also present an imaging ontology which is considered to be an initial effort towards a uniform and standard reporting system that seeks to, to some extent, reduce inter- and intra-expert variabilities.

The system discussed in this paper is not full-fledged. It is expected to adopt certain image analysis techniques to provide, as a complement to manual ROI delimitation, functionalities for automated ROI detection. It will also be interesting to extend the domain-specific BCIO to a more generic one so as to extend the capability of our system.

Acknowledgements

Research for this paper is funded by the British Engineering and Physical Sciences Research Council (EPSRC) under the MIAKT grant GR/R85150/01.

References

[1] American College of Radiology. *Breast Imaging Reporting and Data System: BI-RADS*, third edition.

[2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications.* Cambridge University Press, 2003.

[3] C. Behrenbruch, K. Marias, P. Armitage, M. Yam, N. Moore, R. English, J. Clarke, and M. Brady. Fusion of contrast-enhanced breast mr and mammographic imaging data. *Medical Image Analysis*, 2003. in press.

[4] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, pages 28–37, May 2001.

[5] J. Brown and *et.al.* Magnetic resonance imaging screening in women at genetic risk of breast cancer: imaging and analysis protocol for the UK multicentre study. *J. of Magnetic Resonance Imaging*, 18:765–776, 2000.

[6] E. Di Sciascio, F. Donini, and M. Mongiello. Structured Knowledge Representation for Image Regrieval. *Journal of Artificial Intelligence Research*, 16:209–257, 2002.

[7] V. Haarslev and R. Möller. RACER System Description. In *Proc. of the Intl. Joint Conf. on Automated Reasoning (2001)*, pages 701–706, 2001.

[8] I. Horrocks. *FaCT Reference manual Ver. 1.6.* Dept. of Computer Science, University of Manchester, UK, 1998.

[9] I. Horrocks. DAML+OIL: a reason-able web ontology language. In *EDBT 2002*, number 2287 in Lecture Notes in Computer Science, pages 2–13. Springer-Verlag, March 2002.

[10] I. Horrocks and U. Sattler. Optimised reasoning for $\mathcal{SHIQ}$. In *Proc. of the 14th European Conf. on Artificial Intelligence*, pages 277–281, July 2002.

[11] E. Hyvönen, A. Styrman, and S. Sarrela. Ontology-based image retrieval. In *Towards the semantic web and web services, Proceedings of XML Finland 2002 Conference*, pages 15–27, 2002.

[12] D. Ikeda, N. Hylton, K. Kinkel, M. Hochman, C. Kuhl, W. Kaiser, J. Weinreb, S. Smazal, H. Degani, P. Viehweg, J. Barclay, and M. Schnall. Development, standardization, and testing of a lexicon for reporting constrast-enhanced breast magnetic resonance imaging studies. *Journal of Magnetic Resonance Imaging*, 13:889–895, 2001.

[13] M. Lacquement, D. Mitchell, and A. Hollingsworth. Positive predictive value of the breast imaging reporting and data system. *Journal of the American College of Surgeons*, 189(1):34–40, July 1999.

[14] J. Lo, M. Markey, J. Baker, and J. C. Floyd. Cross-Institutional Evaluation of BI-RADS Predictive Model for Mammographic Diagnosis of Breast Cancer. *American Journal of Radiology*, 178:457–463, 2002.

[15] D. L. McGuinness, J. Hendler, and L. Stein. DAML+OIL: An Ontology Language for the Semantic Web. *IEEE Intelligent Systems*, 2002. 72–80.

[16] D. L. McGuinness and F. van Harmelen. *OWL Web Ontology Language Overview*, 20 March 2003.

[17] C. Meghini, F. Sebastiani, and U. Straccia. A model of multimedia information retrieval. *Journal of the ACM (JACM)*, 48(5):909–970, 2001.

[18] C. Meghini, F. Sebastiani, U. Straccia, and C. Thanos. A model of information retrieval based on a terminological logic. In R. Korfhage, E. Rasmussen, and P. Willett, editors, *Proc. of 16th ACM Intl. Conf. on Research and Development in Information Retrieval*, pages 298–307, 1993.

[19] E. Morris. Illustrated Breast MR Lexicon. *Seminars of Roentgenology*, XXXVI(3):238–249, 2001.

[20] M. Patlak, S. Nass, I. Henderson, and J. Lashof, editors. *Mammography and Beyond: Developing Technologies for the Early Detection of Breast Cancer.* The National Academies Press, 2001.

[21] A. Schreiber, B. Dubbeldam, J. Wielemaker, and B. Wielinga. Ontology-based photo annotation. *IEEE Intelligent Systems*, 16:66–74, 2001. May/June.

[22] R. Stevens, C. Wroe, S. Bechhofer, P. Lord, A. Rector, and C. Goble. Building ontologies in DAML+OIL. *Comparative and Functional Genomics*, 4(1), 2003.

[23] R. Warren and A. Crawley. Is Breast MRI Ever Useful in a Mammographic Screening Programme. *Clinical Radiology*, 57:1090–1097, 2002.