

Facilitating multi-disciplinary knowledge-based support for breast cancer screening

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

Intelligence, Agents, Multimedia Group (IAM),
School of Electronics and Computer Science,
University of Southampton,
Southampton SO17 1BJ, UK
{sd,dpd,bh,hgl,phl,nrs}@ecs.soton.ac.uk

Abstract In order to increase the accuracy of breast cancer screening, the diagnostics of imaging modalities (X-ray, ultrasound or magnetic resonance(MR)) is assessed alongside results of histopathological or cytopathological studies as well as clinical information about a patient. This procedure is called “triple assessment” in the UK. We have undertaken a project to develop a multiple-ontology based medical image-annotation and reasoning system to support this procedure. Our system integrates image-annotation tools for drawing, image analysis and feature extraction, with ontologies compliant with the proposed language standard, namely OWL. This integration is carried out by the mapping between images and instances in the ontologies. In this paper, we outline the ambit of the project in further detail and discuss some of the architecture and design issues that we have resolved in order to make our system modular and semantic web-enabled. The resulting knowledge base should not only support medical practitioners in the triple assessment process but also provide a resource from which personalised patient information could be delivered via the web direct to patient or carer. Finally, we describe the progress that we have made thus far on the implementation of this system.

Keywords: triple assessment, breast cancer, ontologies, image annotation
Biographies:

1 Introduction

In medicine, diagnostic practices depend crucially on the interpretation of evidence in the current literature. Enlarging this evidence-base can be facilitated if medical practitioners consent to a regimentation of nomenclature used to describe the details of the cases under investigation. This pooling together of distributed knowledge sources would lay the basis for greater uniformity in diagnostic intervention and treatment. Breast cancer is one such example of a disease which not only affects a large percentage of the population (one oft-cited

statistic is that one in eight women in the USA will contract the disease over the course of a lifetime of 80-plus years, with slightly lower numbers in Europe), but there is considerable variability in the diagnosis and management of the disease. Medical governing bodies in the EU and the USA have initiated measures to introduce greater uniformity across regions and hospitals in dealing with breast cancer, in both representation and intervention.

Information systems, and in particular Web technologies, are now ubiquitous in the medical workplace, and increasingly, distributed information sources are being accessed by both medical practitioners of every ilk and patients. A new direction in World Wide Web technologies is the Semantic Web [1], which promises to enable technologies that will make Web content that can be processed by machines. Such a provision will allow for intelligent Web-based services that execute protocols triggered by relevant content. In order to facilitate such a development, it is necessary to provide appropriate ontologies, which are descriptive regimentations of specific domains, for unambiguous recognition of content by machines. Ontologies express shared conceptualisations, so end-users must commit to an interpretation of the descriptive framework. This suggests that the methods and technologies developed in Semantic Web research may be fruitfully deployed to support medical practitioners in their efforts to pool together knowledge resources to tackle breast cancer. The resulting knowledge base should not only support taxonomic reasoning required in the diagnostic process but could also provide a resource from which personalised information can be delivered direct to patient or carer.

It is now mandatory for women over the age of fifty to undergo X-ray screening for breast cancer in the countries of the EU and the USA. In order to increase the reliability of interpretation of the X-ray images for diagnosis, potential abnormalities are assessed by means of pathology tests for cells and tissues sampled from areas of the breast considered suspicious by imaging studies. In the UK, the National Health Service has issued guidelines whereby radiologists with their X-ray images (ultrasound and magnetic resonance images as well, if available), histopathologists and/or cytologists with their diagnoses of cell and tissue cultures, and clinicians who know the medical history of the patient are required to perform an appraisal of every case in a consultative fashion. The procedure, called Triple Assessment, is what the project sets out to support with the help of an appropriate multiple ontology-based system. The paper describes the goals of aspects of the project in some detail, explaining how our use of Semantic Web technologies facilitates the development of a modular, open-ended system.

In section 2 we describe the relevant aspects of the medical decision-making problem in order to present the design requirements of our knowledge-based system. Much of the evidence presented at the triple assessment meeting is of a visual kind. Radiologists, cytologists and histologists present their assessment and display images of X-rays and sonograms and histopathological and cytological slides, pointing out the image features that are indicators of the disease. Therefore, our system must necessarily support indexing, retrieval and reason-

ing over image-data alongside the relevant domain-specific knowledge required to describe and reason over the medical concepts.

In section 3, we then describe the key components that make up the core of our proposed system, with the ontologies detailed in 4 and drawing and image-annotation tools in 5. We conclude with some discussion of the architecture in order to integrate the components that talk DAML and RDF in a distributed environment. This should allow for the integration of multiple ontologies; databases for holding image content and medical records and methods for image processing; retrieving information and reasoning over content in the context of medical guidelines and in terms of concepts deemed meaningful in the ontologies.

2 Triple Assessment: Procedures and System Requirements

Diagnostic measures based on X-ray images are not perfect, and analyses of extracted tissues may not reveal evidence of malignancy due to imprecise sampling. Such potential sources of misdiagnoses have led the National Health Service (NHS) in the UK to issue guidelines for a consultative process for screening and also for the subsequent management of malignant cases. For the treatment and management of patients, surgeons, clinicians, nurses, medical statisticians and oncologists are also present in a multi-disciplinary meeting. Our objective in this project is to support the initial *screening* phase of dealing with breast cancer. A hypothetical patient case illustrates the need for all these ingredients and how they might be called upon in the final system.

2.1 Knowledge modelling: Illustrative example

A typical scenario in this screening stage is as follows. A woman undergoes routine screening mammography and the X-rays suggest the presence of an abnormality called microcalcification by radiologists. The NHS guidelines for the screening process in such a case are shown in figure 2.1.

The presence of microcalcifications is often indicative of cancerous cells and some kind of biopsy is required to ascertain, for example, whether this is a case of malignant epithelial cells *within* the milk-ducts of the breast, a condition known as ductal carcinoma in situ (DCIS) or whether the malignancy is invasive and has spread beyond the ducts. Such histologic examination is undertaken by extracting a sample of tissue under imaging guidance, and often stereotactically guided, to ensure that the tissue being investigated is taken from the suspicious region of the breast. However, this is not always the case. To illustrate, if the microcalcifications have the following morphology: pleiomorphic (varying size and shape), fine linear, branching, or granular, then they are *candidates* for core-needle biopsy. However, if they are tiny, oval/round, uniform, and in a localised cluster these are *contraindications* for core-needle biopsy. The way that radiologists describe X-ray images are thus tied in to the mode in which subsequent action is taken by another specialist.

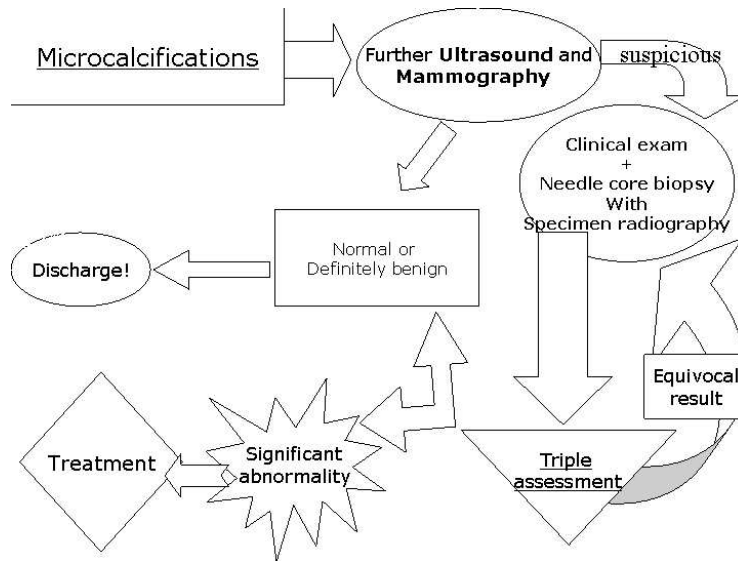


Figure 1. Guidelines for screening in case of microcalcifications

On the histologic end of the story, DCIS could be further divided into low, intermediate and high grades which have prognostic implications. Furthermore, they may be divided in five subtypes – comedo, papillary, micropapillary, cribriform, solid. These are also characterized by nuclear grade and necrosis. They are characterisations by the histopathologist which make sense within the community of histologists and do not necessarily make sense to radiologists or the consultant surgeon. The histologic slides themselves are described by a system of visual descriptors like rounded or eosinophilic which again are interesting only to those involved in the analysis of tissues. In addition, global markers, like the responsiveness of the tissue to hormonal treatment (estrogen and progesterone) needs to be noted as that gives indication of hormonal (tamoxifen) therapy. Finally, there is a clinical classification adopted, staging, that is recognized by specialists in all the disciplines involved. These range from Stage 0 to Stage IV and recurrent, and are used to plan treatment of the disease and patient management and care.

2.2 System Components

The support to be provided is in the domains of patient information; annotation and retrieval of images for various breast-imaging modalities like X-ray and MRI; and for histopathology, the annotation of slides with respect to the categories and concepts that make up the corresponding diagnostic reports. Most of the developments in this project have been in the radiology end of the process thus far, so we describe some of the imaging tools in the next section.

3 Image Types and Viewing Tools

Breast cancer screening utilises a number of different image types. Of primary importance is the X-ray. Different tissues in the body will attenuate X-rays passing through the body, and are useful for diagnosing tumours, especially micro-calcifications, and many other conditions. Digital or scanned-film X-ray images have high spatial resolution ($\sim 10^{-4}m$) and are therefore large ($\sim 10^6$ pixels). This has consequences for the kinds of storage the system will support.

Magnetic Resonance scans use radio waves and a strong magnetic field to provide clear and detailed images of internal organs and tissues. MR scans can be useful in the management of breast tumours because they provide detailed information that can allow the safe excision of lesions that formerly were felt to be inoperable, as well as for following the progression of tumours during treatment. A single MR image represents a 'slice across the patient and by moving the scanning equipment, acquiring several slices, a 3-dimensional image is created. MR images have a comparatively poorer spatial resolution ($\sim 10^{-3}m$), and this reduces their ability to detect small abnormalities such as micro-calcifications. A single MR slice of the human torso may contain $\sim 65,000$ pixels, but with around 25 slices in a time-series of 5-8 scans per session $\sim 10^7$ pixels are stored. For breast cancer screening a contrast-enhancing agent is injected into the patient and the MR radiographer will capture pre- and post-contrast scans. The rate and amount of uptake of the contrast-enhancing agent will depend on the number and size of blood vessels. As tumours are typically associated with a high concentration of blood vessels the MR scan can provide detailed information about tumour location, size, and, importantly, the nature and extent of their margins, the putative delineation of the lesion boundary.

3.1 Image Viewing

A key aim of the project is to provide a service to health care professionals that would allow them to view and annotate images of the breast. A number of image viewing and annotation tools are already in existence. In particular, there are several commercial, off-the-shelf products that provide these services. Further, there exist a small number of Java-based tools that can be run as applets within a grid-based image manipulation service (e.g. Inote by [17]) as envisaged in the project. However, the large image sizes associated with X-ray scans and the requirement to handle dynamic, three-dimensional MR scans provide a unique development challenge. In addition, there is a need to link with automatic medical image analysis algorithms and, possibly, computer aided diagnosis. Accordingly, a prototype for a suitable medical image viewing and annotation tool is being developed as part of the project.

The prototype software tool contains standard bitmap display routines for viewing X-ray images, and incorporates OpenGL to render the MR image file, adjusting for different scanning resolutions in the horizontal and vertical directions. The viewable dynamic range can be adjusted to highlight different structures in the breast, and the transparency of each cuboid can also be changed

to allow internal structures to be viewed. In addition, the viewing angle can be adjusted in real time using a mouse. The software is object-oriented such that its modular nature allows the straightforward addition of new code for viewing different types of imagery or for applying different image processing algorithms. Figure 2 shows a series of screen captures highlighting some of the facilities of the image viewing tool. In particular, Figure 2 (b) and (d) illustrate the use of a simple differencing operation to emphasize suspected tumours. In section 5, we use these facilities to associate regions in the images with appropriate concepts in the ontology. A retrieval system will then allow for the patient records or image annotations or other diagnostic details to be presented in response to a range of queries.




Figure 2. Screen captures from the prototype MR image viewing tool showing (a) a single slice, (b) a single slice showing differences between pre- and post-contrast scans in red, (c) a 3-D view showing all slices, and (d) a 3-D view showing differences between pre- and post-contrast scans in red.

In what follows we provide details of ontology development for the breast imaging modalities – X-ray and MRI. We are also in the process of designing an ontology for histopathology, which is not described below. We will also need to encode constraints into the ontology so that medical guidelines for reporting lesions are enforced. These are requirements for separation of categories which have different prognostic implications even though they are related by being part

of the same structural super-class. These details will be reported elsewhere. For related work, see [18].

4 Breast Cancer Imaging Ontology (BCIO)

The image types that we have incorporated into our prototypes are X-ray and MRI mammograms. There is likely to be considerable difference in the usage of terminology by expert practitioners in any discipline, which makes the task of creating consensual ontologies difficult. Access to clinical experts is constrained by available time. Our initial approach has been to build on existing lexicons and thesauri in the different domains under consideration.

4.1 BI-RADS Vocabulary

Based on the extensive field-work and substantial case studies, the American College of Radiology (ACR) has proposed a standard for breast mammography, the Breast Imaging Reporting and Data System (BI-RADS) [2]. BI-RADS provides a comprehensive lexicon for describing mammographic findings containing, among others:

- image descriptors, e.g., the shape of the lesion, the texture of the lesion,
- lesion types, e.g., calcification,
- location of lesion, e.g., outer upper quadrant,
- assessment of findings, e.g., suspicious.

We consider BI-RADS to be a good starting point for our Breast Cancer Imaging Ontology (BCIO). However, since our target users are affiliated with hospitals in the UK, the applicability of BI-RADS lexicon among these users requires their validation.

4.2 Modelling Language and Editor

Quite a few modelling languages are available now for the purpose of creating ontologies, e.g. RDF [8], RDF Schema (RDF(S)) [9], DAML+OIL [10], and OWL [11]. DAML+OIL is one of the ontology modelling languages supporting the migration from current data-intensive web-based resources to semantically described and interoperable web contents. Unlike RDF(S), DAML+OIL and OWL are languages which are defined on top of description logics (DL) and offer some of the advantages of having unambiguous semantics. We have developed our ontologies in DAML+OIL and in RDF(S), but we have ensured that we remain OWL-compatible, as OWL is the proposed standard for the ontology language for the Web. All these languages adhere to the syntactic constraints of the XML standards [15].

The underlying logical syntax of DAML+OIL and OWL, description logics, provide a means to formally define the semantics of the appropriate domain-specific concept (referred to as class hereinafter) and property constructors.

Therefore, DAML+OIL and OWL offer a variety of approaches to ontology construction ranging from a pure manual process to an automated classification based on a set of defined concepts and properties.

The freely distributed editor for DAML+OIL, OilEd, is used in the construction of BCIO. OilEd offers a frame-like user interface and limited reasoning services based on FaCT [5] and RACER [3] reasoning systems. The frame-like graphic user interface (GUI) helps reduce the necessary but tedious tasks needed to handle the XML-based syntax of DAML+OIL. Meanwhile, the DL-based reasoners can detect inconsistencies and exposes knowledge which is otherwise implicit.

4.3 Methodology

An ontology can be created in different ways by going through different stages [12]. There is no standard methodology for constructing an ontology, but only recommendations. In this project, we adopted an approach where the ontology was gradually refined from a rough hierarchical structure. The approach is, to a degree, modular in the sense that we maintain several separated tree structures, although under the same umbrella, the top-level class. The iterative refining procedures are as followings:

1. a hierarchical structure is established;
2. the incomplete ontology is reasoned by a DL-based reasoner to detect possible inconsistency and to derive subsumption relationships;
3. the original structure is populated with new classes, properties and instances.

The above process is repeated until a satisfactory ontology is obtained without inconsistency.

OilEd can perform the consistency check when building up the original BCIO. After the BCIO is created, currently, users may not change the hierarchical structures or add new sub-concepts to existing concepts. This prevents damage to the underlying structures of BCIO caused by *ad hoc* extensions to the original ontology.

4.4 The Breast Imaging Ontology

BCIO is composed of several modules. The modular design ensures that further extensions can be made to a particular module without disturbing the integrity and structure of others. The separation of modules is based on the characteristics and graininess of the knowledge.

The intended application of BCIO suggests that pure medical knowledge is not sufficient. There must be a bridging mechanism between the medical terms and the images. Such bridging is achieved by creating trees of simple classes (primitive concepts) representing the skeleton of the image “descriptors” (see Figure 3) and referring to the image oriented terms via instances of Region-of-Interest (DAML definition of Region-of-Interest is shown in Figure 4). Note

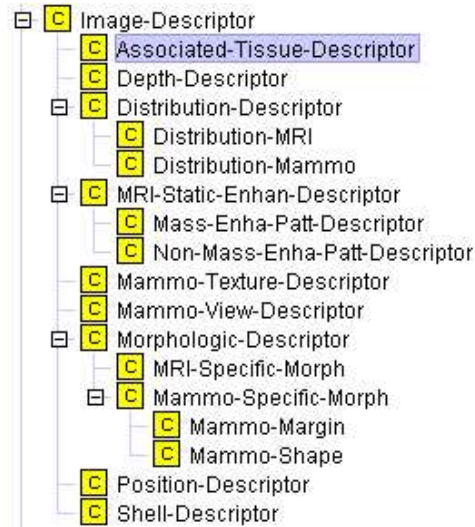


Figure 3. The Hierarchy of Image Descriptor

that under the umbrella of *Image-Descriptor* are not only the descriptors used to represent the physical existence, e.g. the morphological characteristics, but also those which are subject to interpretation and require medical experience and knowledge, e.g. the *Associated-Tissue-Descriptor*.

Currently, BCIO covers the vocabulary of both mammography (breast X-ray image) and breast magnetic resonance images (MRI for short, based on the draft version published by the MRI lexicon working group [7]). The significant differences between these two imaging methods [7] suggest that two separate sets of vocabulary should be maintained, although under the same parent class, i.e. *Image-Descriptor*.

Another part of BCIO is the set of terms related to medical tests, such as *Findings* (normal breast tissue and abnormal entities), *Medical-Exam* (the parent class of mammography and breast MRI), *Medical-Descriptor* (the collection of medical terms that cannot be defined necessarily and sufficiently). It also records image capture metadata – for instance, whether an X-ray is of the left breast and from the *medio-lateral-oblique* position, (LMLO) information of vital importance in locating lesions or tumours.

In addition, as part of the breast imaging and patient management process, inevitably, terms such as patients, clinicians and issues with regard to staging, etc., need to be included.

```

<daml:Class rdf:about="#Region_Of_Interest">
  <rdfs:label>Region_Of_Interest</rdfs:label>
  <rdfs:comment></rdfs:comment>
  <rdfs:subClassOf>
    <daml:Class rdf:about="#BC_Domain_Top"/>
  </rdfs:subClassOf>
  <rdfs:subClassOf>
    <daml:Restriction daml:cardinalityQ="1">
      <daml:onProperty rdf:resource="#graphic_region"/>
      <daml:hasClassQ rdf:resource="http://www.w3.org/2000/10/XMLSchema#string"/>
    </daml:Restriction>
  </rdfs:subClassOf>
  <rdfs:subClassOf>
    <daml:Restriction>
      <daml:onProperty rdf:resource="#graphic_region"/>
      <daml:toClass rdf:resource="http://www.w3.org/2000/10/XMLSchema#string"/>
    </daml:Restriction>
  </rdfs:subClassOf>
  <rdfs:subClassOf>
    <daml:Restriction>
      <daml:onProperty rdf:resource="#s_finding"/>
      <daml:toClass>
        <daml:Class rdf:about="#Findings"/>
      </daml:toClass>
    </daml:Restriction>
  </rdfs:subClassOf>
  <rdfs:subClassOf>
    <daml:Restriction daml:cardinalityQ="1">
      <daml:onProperty rdf:resource="#s_finding"/>
      <daml:hasClassQ>
        <daml:Class rdf:about="#Findings"/>
      </daml:hasClassQ>
    </daml:Restriction>
  </rdfs:subClassOf>
</daml:Class>

```

Figure 4. Definition of Region-of-Interest: property graphic-region is used to refer to delineated graphic regions on images

4.5 JENA-based Ontology Parser

Currently, extending BCIO axioms by end users (e.g. medical staff) after the creation of the ontology is not envisaged, except populating it with new instances. This approach is under both theoretical and practical considerations. On one hand, *ad hoc* extensions of the original ontology may disturb the original hierarchical structures and introduce inconsistencies. Although OilEd, the ontology editor, does provide a certain degree of assistance for creating ontologies, such assistance is limited and geared to be of help to those familiar with knowledge representation. Some medical experts are likely to find the effort required to create ontologies too time consuming[13]. As a result, the task of modifying ontologies is normally left to the knowledge engineers. On the other hand, if changing the taxonomy is allowed, a DL-based reasoner will have to be called to detect possible inconsistencies. Such extra facilities bring programming burdens, hence we do not expect the end users (medical staff) to be keen to extend the BCIO axioms.

However, populating the knowledge base with new instances (e.g. a particular patient) is to be expected. A JENA [4] based DAML parser is developed to facilitate a frame-based approach to populate the knowledge base. A frame provides attribute slots for any concept which can then be filled by corresponding values, a natural setting for a form-like interface. Historically, description logics arose in order to make the semantics of frames unambiguous, and here we have reversed this history to provide a workable user interface. We do this by associating to each class the properties that are defined for itself as well as those of all its subclasses. The resulting frame-based user interface is shown in Figure 5. As one may expect, some of the properties are not defined for the class **Region-of-Interest**, but a subclass of it (e.g. the property “has MRI enhancing pattern”, *has-enhn-pattern*).

4.6 Graphical Editing Tool

The image annotating tool benefits from a hyperbolic style ontology navigation interface (based on TouchGraph [14]). It allows users to navigate BCIO through the subsumption and instantiation relationships (i.e. the *is-a* and *instance-of* relationships).

As shown in Figure 6, populating the knowledge base is enabled by creating an instance (shown in light colour) of the selected concept node (shown in dark colour). The definition of the selected concept node is displayed simultaneously on the upper right window of the left half of the panel.

4.7 Ongoing work on the image ontologies

When populating the knowledge base, a RACER server is responsible for the inconsistency check of instances. RACER does not support reasoning with regard to *DatatypeProperty* [6] which makes the definition of concrete properties, e.g. age, name, staffid and size, very difficult. Such deficiency needs to be addressed

Definition of Concept **Region_Of_Interest**

has_quadrant_position [Multi]

quadrant_lower_inner
quadrant_lower_outer
quadrant_upper_outer
quadrant_upper_inner

graphic_region [Single]

is_finding [Single]

af_abnormal_signal_void
af_high_signal_in_duct
af_trabecular_thickening
af_axillary_adenopathy
af_abnormal_density
af_fluid_filled_dilated_ducts
af_architecture_distortion
af_cvsts

has_size [Single]

has_dis_nipple [Single]

Figure 5. Frame-based input form

if end users are allowed to add complex axioms or assertions into the knowledge base which may introduce inconsistencies.

5 Ontology-based Image Annotation Tools

In section 3 the tools developed to render the relevant images were presented. With the ontology in place, we now describe how they may be used to establish relations between (regions in) the images and appropriate concepts in the ontology.

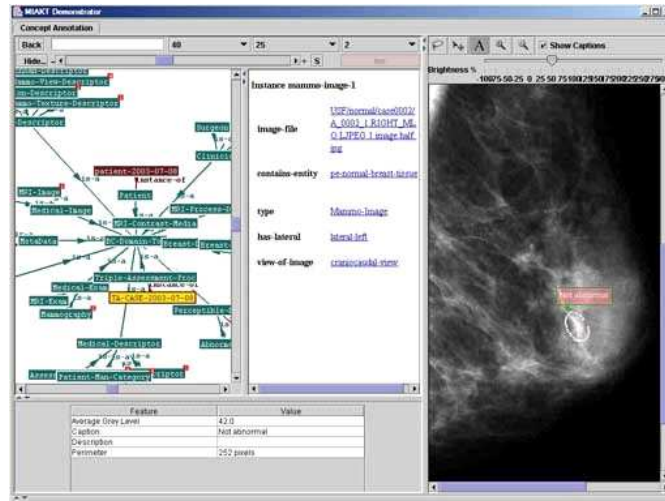


Figure 6. Graphic User Interface

5.1 Image Annotation

The prototype image-viewing tool provides the facility to annotate the X-ray and MR images. The user is able to highlight a region of interest (ROI) using either a mouse click-and-drag approach to draw a box (as illustrated in Figure 7), or a mouse hold-and-draw approach to map the margins of a suspected lesion more precisely (as shown in Figure 8). Several options are available to support these processes, such as full image magnification, localised magnification, and localised contrast enhancement.

The annotation of the ROI is achieved in a number of ways. Firstly, the user can select a class and an associated parameter value from the appropriate image ontology (see section 4.4) and add this to an annotation window containing the description of the ROI. Ontologies can be imported into the viewing tool as RDF or plain text. This option ensures that text-based annotations remains in a standard form even when different personnel use the software. Secondly, the user may decide to add shape and textural descriptors that are computed using machine vision algorithms applied to the boundaries and interiors of ROIs. It is anticipated that as the project matures more descriptors from the medical image analysis domain will be added to those currently implemented (e.g., as used by [16]). This option provides a list of continuous values for descriptors including area, perimeter, complexity, and eccentricity in addition to first and second order statistics of brightness. Figure 8 shows the prototype software tool being used to view and annotate a left, medio-lateral oblique (LMLO) X-ray. Three ROIs have been drawn and the user is annotating one of these using the X-ray image ontology and a selection of machine vision descriptors.

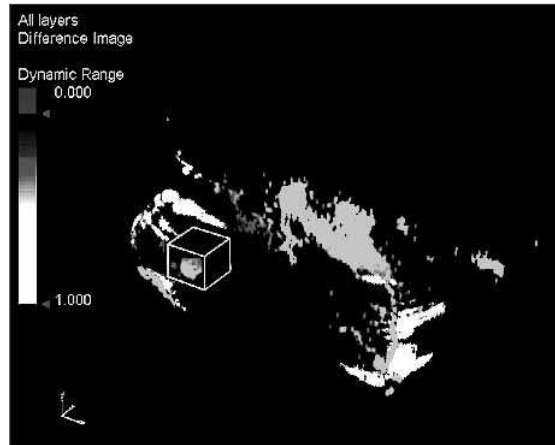


Figure 7. Screen capture from the prototype MR image annotation tool showing the three-dimensional difference view from Figure 2(d) (here in black and white) and a user-added region of interest (ROI) box around a suspected tumour.

Once the images have been annotated the ROIs and their descriptions are stored as RDF or plain text. The ROIs are given a unique label that can be derived from the patients name, the date of the X-ray scan and the viewing geometry (e.g., LMLO). For X-ray imagery the location of each ROI boundary pixel is stored using polar coordinates referenced to the ROIs centre. The centre is determined simply by calculating the mean Cartesian coordinates of the boundary of the ROI in the image. For MR imagery the ROI boundary and its centre of gravity can be specified using coordinates in three dimensions.

6 Conclusion

In this paper we have described work towards the realisation of a multidisciplinary knowledge based support system for the triple assessment procedure in the breast cancer screening programme.

Some of the difficulties of knowledge modelling in the domain have been described together with work on the development of suitable ontologies for the domains involved. We have introduced the development of tools for annotating images both manually and automatically to link knowledge implicit in the various image modalities with the ontological information in order to provide a more powerful platform for assisting directly with the decision making processes.

Future extensions to the prototype software that we have experimented with include the capability to provide computer aided diagnosis using a classification system, such as a Bayesian classifier, using the annotations provided by the user as inputs (e.g., as in [19]). In addition, a facility that allows the user

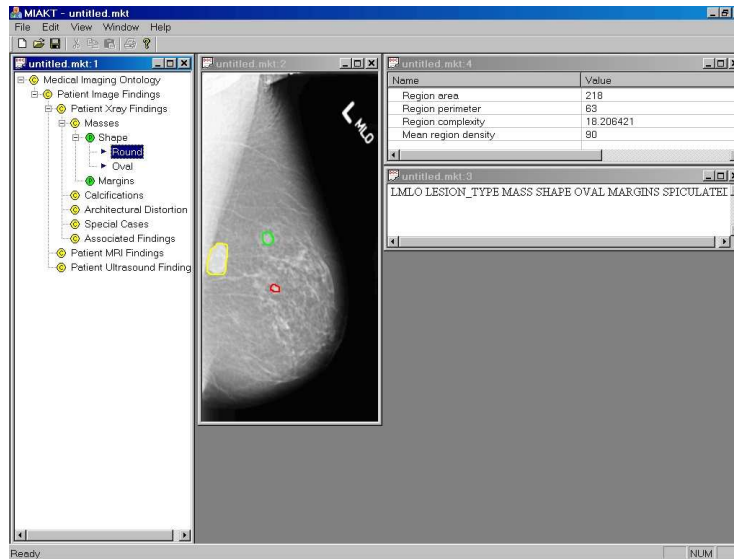


Figure 8. The prototype image viewing and annotation tool showing part of the X-ray image ontology (left window), an X-ray view with precisely-drawn ROIs in distinct colours (middle window), a user-added description for one ROI (bottom right window) and shape descriptors provided by machine vision algorithms (top right window).

to ask the software tool to suggest the most likely annotation for the ROI is under development. In this case the software tool will use a database of exemplar images, previously annotated with both the ontological-based and machine vision descriptors, and employ a nearest neighbour algorithm to determine the appropriate annotation to apply to the ROI.

In the longer term the services will be provided via a distributed grid-enabled infrastructure, in order to facilitate the rapid registration of the images. Although the annotation tools described here are currently implemented discretely we are putting together a multimedia database environment which is Semantic Web-enabled. This will support content and concept based retrieval and navigation and will provide facilities for handling combinations of textual and visual information handling through close integration between the ontologies for the application domain and image based information within the system. It is planned that this system will support not only the medical professionals involved in the manner indicated, but also the patients who need to be informed and aware of their condition. A key component that is also being developed using this framework as a conceptual backbone will provide personalised patient information using the factual and knowledge components captured in the system.

Acknowledgements

The authors are grateful to their collaborators on the project – Kalina Bontcheva, Michael Brady, Liliana Cabral, Fabio Ciravegna, John Domingue, David Hawkes, Enrico Motta, Christine Tanner, Jonathan Whiteley, Yorick Wilks and Yalin Zheng – for many valuable discussions. This project (MIAKT) is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) for support under grant number GR/R85150/01. MIAKT is part of the e-Science initiative and is a collaboration between two IRCs supported by the EPSRC – AKT (GR/N15764/01) and MIAS (GR/N14248/01).

References

1. T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, may 2001.
2. American College of Radiology. *Breast Imaging Reporting and Data System: BI-RADS*, third edition.
3. V. Haarslev and R. Möller. *RACER User's Guide and Reference Manual Version 1.6*. University of Hamburg, Computer Science Department, July 2001. Technical Report.
4. Hewlett-Packard Company. *Jena Semantic Web Toolkit*. Copyright 1994-2003, Available at <http://www.hpl.hp.com/semweb/jena.htm>.
5. I. Horrocks. *FaCT Reference manual Ver. 1.6*. Dept. of Computer Science, University of Manchester, UK, 1998. Available from <http://www.cs.man.ac.uk/~horrocks/FaCT>
6. I. Horrocks, F. van Harmelen, P. Patel-Schneider, T. Berners-Lee, D. Brickley, D. Connolly, M. Dean, S. Decker, D. Fensel, R. Fikes, P. Hayes, J. Heflin, J. Hendler, O. Lassila, D. McGuinness, and L. Andrea Stein. *DAML+OIL (March 2001)*. World Wide Web Consortium. Available at <http://www.daml.org/2001/03/daml+oil-index.html>
7. D.M. Ikeda, N.M. Hylton, K. Kinkel, M.G. Hochman, C.K. Kuhl, W.A. Kaiser, J.C. Weinreb, S.F. Smazal, H. Degani, P. Viehweg, J. Barclay, and M.D. Schnall. Development, standardization, and testing of a lexicon for reporting contrast-enhanced breast magnetic resonance imaging studies. *Journal of Magnetic Resonance Imaging*, 13:889–895, 2001.
8. O. Lassila and R.R. Swick. *Resource Description Framework (RDF) Model and Syntax Specification*. W3C, 22 February 1999. Available at <http://www.w3.org/TR/REC-rdf-syntax/>.
9. B. McBride. *RDF Vocabulary Description Language 1.0: RDF Schema*. W3C Working Draft, 23 January 2003. Available at <http://www.w3.org/TR/rdf-schema/>.
10. D. L. McGuinness, J. Hendler, and L.A. Stein. DAML+OIL: An ontology language for the semantic web. *IEEE Intelligent Systems*, 72-80, 2002.
11. D. L. McGuinness and F. van Harmelen. *OWL Web Ontology Language Overview*, 20 March 2003.
12. N. F. Noy and D. McGuinness. *Ontology development 101: A guide to creating your first ontology*. Technical Report SMI-2001-0880, School of Medical Informatics, Stanford University, USA, 2001.

13. 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.
14. TouchGraph LLC. *TouchGraph*. Available at <http://www.touchgraph.com/>.
15. World Wide Web Consortium. *Extensible Markup Language (XML)*. Available at <http://www.w3.org/XML/>.
16. Chan, H-P., Sahiner, B., Lam, K.L., Petrick, N., Helvie, M.A., Goodsitt, M.M., and Adler, D.D. Computerized analysis of mammographic microcalcifications in morphological and texture feature spaces, *Med. Phys.*, 25 (10), 1998.
17. Robert Bingler. Inote: An Image Annotation Tool in Java, Available at <http://www.iath.virginia.edu/inote/>
18. R. D. Shankar, S. W. Tu, M. A. Musen. Use of Protg-2000 to Encode Clinical Guidelines. 2002. Stanford Medical Informatics Report, SMI-2002-0944 Available at http://www-smi.stanford.edu/pubs/SMI_Reports/SMI-2001-0944.pdf
19. Elizabeth Burnside, Daniel Rubin, Ross Shachter. A Bayesian Network for Mammography. Stanford Medical Informatics Report, SMI-2001-0867 Available at http://www-smi.stanford.edu/pubs/SMI_Reports/SMI-2001-0867.pdf