

Multimedia Distributed Knowledge Management in MIAKT

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

IAM Group, University of Southampton, Southampton, SO17 1BJ, UK
[dpd|sd|bh|phl|nrs]@ecs.soton.ac.uk,
WWW home page: <http://www.aktors.org/miakt>

Abstract. Digital media facilitates tight integration of multi-modal information and networking allows this richly textured knowledge to be shared. We present the system we have developed in the MIAKT (Medical Imaging with Advanced Knowledge Technologies) project that provides knowledge management, and facilities for semantic annotations on mammographic images in the context of clinical and histopathological information.

This paper also describes the novel generic architecture we have built on semantic web technologies to facilitate the annotation of images with ontological concepts, and storage thereof, in any domain. Functionality of a specific domain application is provided through web-resources, which are called through a task invocation system which abstracts the actual service implementation from the client application implementation.

1 Introduction

The drive towards semantic web [4] technologies has provided a research area that brings together semantic annotation and image feature extraction. Automating semantic annotation of images is a difficult process in most domains, the annotation requiring some level of intervention from users. In the medical domain, images are rarely clearly defined, and often regions of interest are difficult to spot by a trained expert. By combining a number of different technologies, including the semantic web technologies, into a generic system, we can begin to provide some support for both the manual and automatic annotation of these images, as well as providing a means for retrieval and reuse of the data.

Breast cancer screening is now mandatory for women over the age of 50. This process consists of the capturing of an x-ray mammogram and a radiologist examining it for any areas considered abnormal. They are then assessed, if necessary, by means of pathology tests (biopsies) by a histopathologist. Data from the radiologist, the histopathologist, and the clinician (who has knowledge of the history of the patient) are brought together to make a consultative appraisal of each particular case in a Multi-Disciplinary Meeting (MDM). This process is known as the Triple Assessment Procedure and the work presented here, as part of the MIAKT (Medical Imaging with Advanced Knowledge Technologies) project, is intended to support this collaborative meeting and manage the knowledge that goes with it, using the Semantic Web technologies.

To achieve this we have developed a novel architecture for delivery of applications to users based on ontological application descriptions. The application's data sources and computation sources are distributed which provides access to the application from any available application server via a roving client application. An important and convincing argument for the use of such a distributed framework is that all parties involved in an application's data or functionality pool retain control of their respective property whilst still being able to access the relevant parts from remote application clients. It eases both the integration issues as well as the intellectual and ethical issues for institutions to retain rights to their data, property or system, and provide a service to which interested parties can connect.

Image annotation is conducted locally on a user's machine to ensure adequate user feedback. However, the images are retrieved from remote servers, image features are generated using remote analysis services, and the results of the annotations are stored in the ontological database that contains the patient record of the patient concerned.

In the following section we briefly review related work. Section 2 presents an overview of the generic knowledge management framework used to provide the middleware to the multimedia and knowledge management process, and in sections 3 and 4 we describe how we use this in the MIAKT application. In the MIAKT scenario we describe the image distribution system and the image annotation tools. Section 5 gives conclusions and a brief mention of future work.

2 The Generic Framework

The novel, distributed architecture, that the MIAKT application is built upon, uses web-based services to provide discrete and disparate functionality to a generic application base shown in figure 1. The architecture is deliberately abstracted from any particular application domain (and its description) providing a generic structure for rapidly prototyping new knowledge-based applications that require media annotation in new domains. Abstracting the architecture from the application domain provides a considerable challenge in the designing of an API that ensures components are still interoperable in disparate domains.

A user transparently interacts with the architecture through an application client that is also built around a generic architecture that can be rapidly implemented into a specific application by mediation through an 'application ontology'. This application ontology is distinct from the domain ontology and provides application settings for a specified domain such as which media viewers are used to display and annotate images.

The core of the framework is based on the invocation of web-services through a task invocation sub-system that provides configurable functionality for the target application. The services available to a specific application are described in that application's ontology. The methods made available by these services are automatically discovered through description mining, or by server interrogation depending on the implementation of the target service and the nature of the 'han-

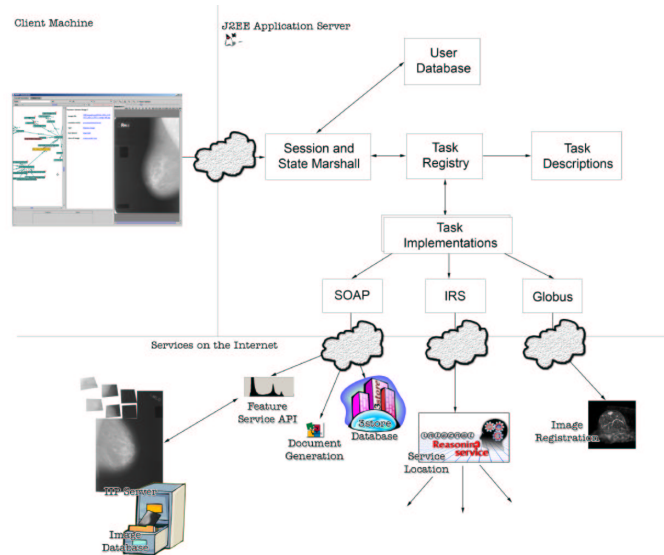


Fig. 1. The MIAKT framework

plers' or description provided. The methods are associated through a mapped repository with task names which are called from a client unconcerned with the task's implementation. Currently the architecture supports both SOAP [11] (web-services) and the Internet Reasoning Service (IRS) [8] task implementations and they are imported into a task registry using WSDL [10] mining for the SOAP tasks, and server interrogation for the IRS. The architecture makes it simple to add new service providers, and it is possible that invocations of services on Globus servers will be supported in the future.

It is important over such a communal service architecture to have interoperable function calls, and all data given to, and returned by, services are to be of primitive types : strings, integers, etc. Complex data is marshalled to and from XML by domain-specific handlers.

To store the domain data in instantiated ontologies, our database service is based on an RDF-triple database called 3Store [3]. This database is accessed over a SOAP webservice, and to maximise interoperability, results are returned as XML and parsed in the client to extract pertinent information or display results.

The on-demand delivery of the application description to the generic client provides a means to customise the application to a given domain, while the distributed nature of the framework provides potentially unlimited interoperability, giving access to any web-service based from any application domain. In the next section we describe how this generic base is put to use for medical image and knowledge management in the project in which it was developed.

3 The MIAKT Application Architecture

For the medical knowledge management domain, a number of data sources which are used to provide the underlying data for this domain are required. For the support of the multi-disciplinary meeting we require at least the following:

- Patient records including information about the patient, what examinations they have undergone, and what results were concluded from those examinations.
- Multimedia data such as X-Rays and MRI mammograms that are taken during examinations and are required for marking up suspicious areas and then relating those to the patient’s medical data, including biopsies.

Using the relevant media viewers and analysis services, the client application can automatically provide a method for associating annotations made on the multimedia data to semantic concepts in the domain data, which is the patient information in this application. The following sections describe these data sources and their usage.

3.1 Patient Records

The Breast Cancer Imaging Ontology (BCIO), developed in the MIAKT project, is designed in a modular manner representing different levels of resolution of the application domain. Highly abstract terms, such as “Medical Image” or “Image Descriptor”, are on one end of the descriptive grain-size while concrete descriptors, like “Spiculated Margin” describing the shape of a region of interest, are on the opposite end. Between them are several levels of interim concepts constructing a referencing bridge. Such an approach makes it possible to replace a particular part of the ontology to adapt to minor, or fundamental, changes of the application domain.

The BCIO ontology is based on a standardised lexicon called BI-RADS (Breast Imaging Reporting and Data System) developed by the American College of Radiographers (ACR) [7]. We have utilised recommended guidelines by the ACR and the National Health Service in the UK to extend this lexicon and develop the ontologies. The ontologies are compliant[2] with the Web Ontology Language (OWL)[1] standard.

We currently source our data from anonymised, legacy data [12], but in practise this would be entered by radiographers as new patients arrive. The patient records would, in practise, be stored on 3store databases located at the institution controlling the data, and accessed by webservices with appropriate safeguards to ensure privacy and compliance with ethical procedures.

3.2 Images and Multimedia

The way in which images and other multimedia data are integrated into the framework can have an important effect on the flexibility of the image annotation

systems, and the system as a whole. Therefore, the framework does not stipulate any particular conceptual position in the architecture for storage, or analysis of multimedia data. It is possible to have the data stored separately from both the application client, where the user is viewing it, and the analysis algorithms which are calculating and storing feature vectors for features in the data. Indeed, multimedia data, like the various modalities of images that are produced in the medical domain, can be stored on institutionalised servers which are able to deliver the data to the client on demand by image servers. This provides the potential to integrate with current hospital image repositories such as PACS (Picture Archiving and Communications System).

Details of the multimedia data are entered as instance information along with the patient record, thereby linking the remote data sources with the patient examination record which facilitates immediate access to the relevant multimedia data at the client.

The digital x-ray mammogram images, fundamental to this domain, have very large dimensions (on average about 2500x4000 pixels, or about 4Mb when compressed). The images are transferred using the Internet Imaging Protocol (IIP) over a standard servlet interface, which gives the ability to view and manipulate them over relatively low bandwidth connections despite this size. The IIP servlets deliver image tiles, on-demand, from various precalculated resolutions of the image, to the client's IIP image viewer. MRI images are delivered slice-by-slice from an MRI image server to the client's MRI image viewer.

The framework is not confined to using any particular protocol for serving images, and indeed it would be undesirable to limit flexibility in this way. Access to image servers is initiated from only those processes that understand the respective image modality, such as an image viewer in the application client or a feature vector generator on a feature service. This relieves the application server from the transfer of any potentially large images, because image transfer is conducted directly between the image client and the image server.

```
<source-information>
  <source type="IIP Image">
    <image-server>http://imgsrvr/fcgi-bin/iipsrc.fcgi</image-server>
    <image-filename>/images/case0042/LEFT_MLO.LJPEG.tif</image-filename>
  </source>
  <source-region type="PointSet2D">
    <boundary>(100:100),(200:200),(200:100)</boundary>
  </source-region>
</source-information>
```

Fig. 2. Complex object marshaling using primitive strings containing XML, showing the marshaled result for a simplified example of a region of interest on an image.

To allow the comparison and classification of images based on their content, image descriptors are generated. Currently some generic shape and colour descriptors are integrated into the system, and our collaborators at Oxford are developing X-ray-specific feature modules that deliver descriptors more relevant to the domain, providing a better means for classification. We have developed an API that provides general functionality for generating and comparing feature vectors on images and publishing these analysis algorithms as a webservice. It provides a defined interface for client processes to interact with any feature modules that are offered as a service on the server. Feature vectors are created from annotations made in the relevant viewers. When an annotation is created the feature module is automatically called to create the relevant vectors. When calling the feature service, the inputs and outputs, such as the definition of a source image or image region, are marshaled into an XML object such as the simplified example shown in Figure 2. The flexibility of this feature service architecture means it would be a simple integration process to replace the default relational database with specialised feature-based indexing databases.

Using these frameworks, the MIAKT application supports specific medical image analysis algorithms, one of which is the registration (alignment) of images [6] from different time frames, or even different modalities such as histopathology slides or MRI slices, which provides a good method for abnormality detection in MRI images (using subtraction of registered images). These types of registration process are highly computationally intensive and have been implemented using Grid technologies, which are currently accessed using a standard web-service invocation mechanism. In the MIAKT application, there are also services for the generation of descriptors for masses in both MRI images and X-ray images. The descriptors these services calculate are based on the domain-ontology. In the ontology the shape of a mass can be described as ‘irregular’, ‘round’, etc., and data including this information is generated from an X-ray image analysis service. Subsequently, classification services using Bayesian networks, attempt to classify an abnormality as malignant or benign based on its descriptors. A similar service is available for MRI images that returns a final finding ‘malignant’ or ‘benign’ based on a set of low-level, image-content features taken from an annotation roughly delineated by a user.

4 The Client Application

A generic application client provides the knowledge management user interface, including the image annotation tools. Images are stored as instances in the domain ontology, and the application client gives the user access to these instances via an interactively navigable ontology visualisation tool, as described below. The domain ontology is retrieved from a location on the network stipulated by an application ontology instance.

We provide two major ways to navigate the ontology. Firstly there is the hierarchical view, which shows the concepts in the ontology based on the typical subsumption relationship allowing quick navigation to concepts. A Touch-

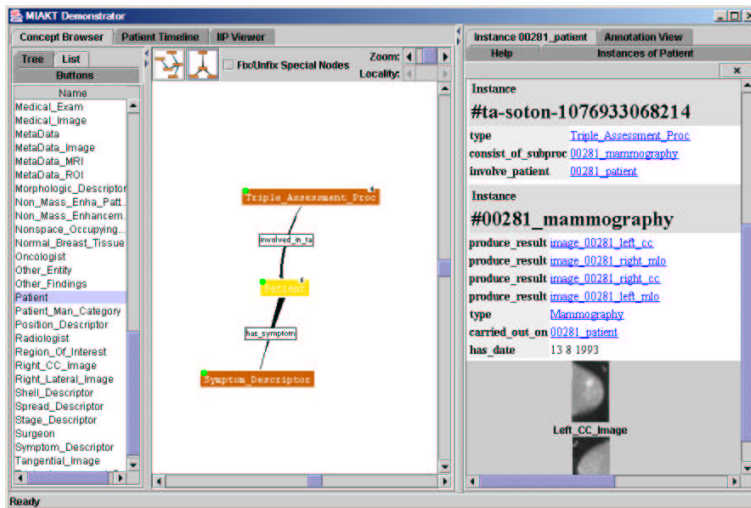


Fig. 3. The generic client application has hierarchical, list and TouchGraph [9] views of the ontology. Here a patient instance is being viewed using an application dependant handler for the mammography and the generic handler for other instances.

Graph [9] view of the ontology allows the full network structure to be viewed and manipulated in real-time (Figure 3). Instances of particular concepts can also be retrieved from this view and information about a specific instance can be recalled by clicking on the relevant instance identifier. Instances are displayed as a list of slots and values. To allow customisation of the client application, the instance handlers that display these lists can be extended for particular concept types, thereby allowing the client application to provide context-based, and domain-based instance visualisation. Media viewers are then dynamically loaded into the application client and instantiated based on the instance type, media type, and/or the delivery mechanism of the media.

The media viewers are implemented using a defined interface which allows them to be invoked to produce annotations of regions of interest in a given medium. For example, the IIP image viewer allows users (in the case of MIAKT, radiologists) to draw around regions of interest in the image, as shown in Figure 4, and these form the basis of annotations. An annotation observer process receives annotations from the media viewers and automatically invokes feature modules, both local and remote, that are able to take the given region of interest as input and produce feature vectors. Domain-dependent feature analysis modules may output concepts relevant to the domain-ontology, thereby allowing direct (but manually verified) insertion into the instances of concepts from the domain. For example, the margin of a mass may be classified using shape features (irregular, round, etc.). Non-domain feature vectors may be inserted into

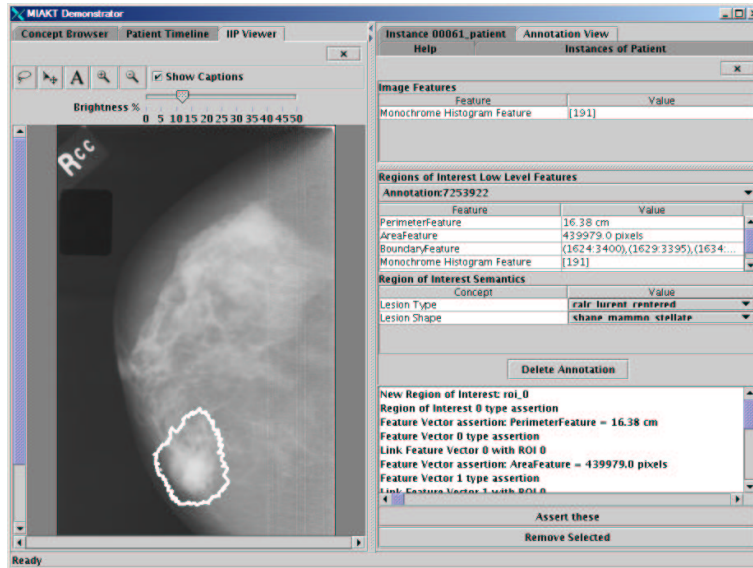


Fig. 4. The IIP image viewer allows image annotations to be generated using drawing tools. When feature vectors are available, they are displayed alongside the image.

the domain instances where appropriate, or under a generic ‘Image Descriptor’ banner.

The architecture provides automatic activation of modules that perform media processing using a regimented API between the application and the observers that receive annotations from viewers. The framework’s indifference to local and remote activation of media modules facilitates sites with large computational power, or storage capacity, to be used to generate descriptive vectors from media which is remote to both the feature module and the client. For example, the MIAKT project uses image analysis modules running remotely in our partner sites.

Once features have been generated by feature modules, they are automatically mapped to domain concepts to be associated with the ontology as instances in the database. This is currently achieved by feature and domain-specific classification code, although we are investigating using generic classification techniques for this step, that classify feature vectors into a set of controlled classes specified by the ontology. This classification provides default values for the semantic descriptions of the relevant annotation, which the user can validate. The insertion of instances into the database is done manually by the user, thereby allowing the user to disregard features which are giving no added information, or incorrect information. To aid the speed of this process, and to allow the user to make alterations to the instances, assertions into the database are pooled prior to a batch assertion. We are investigating ways to make this insertion an easier process; us-

ing form-based input is a well-understood method of knowledge storage for the medics we have contacted, but providing a method for the generation of general forms provides some challenge. It is possible that domain-based context-based form generators would be necessary, but at the expense of generic flexibility.

4.1 Other Services

To enrich the value of the MIAKT application, other services have been included which are available through the server-side task invocation sub-system.

Consultation during a multi-disciplinary meeting, on the best course of action for the patient, relies on the outcomes of the examinations that the various medical staff performed on the patient. These outcomes are based mainly on the doctor's experience, but also rely on their full attention and concentration. It is possible for human errors to be made. For this reason we have developed naive Bayes and MLP-based classification algorithms that, based on patient records for previous patient cases, attempt to classify the type of lesion from its ontological description. In the near future, we hope to extend this classification to image-content-based features. On our current data sets they are giving correct results around 75% of the time, although currently, this accuracy does seem to be limited by our datasets.

Using technologies such as GATE [5], collaborators at Sheffield have developed a technology which allows natural language documents to be generated based on the ontological instance data. By applying this technology to a patient's case notes, the effort of writing up routine patient reports is reduced for the busy medical staff who currently have to do this by hand.

The UMLS (Unified Medical Language System) is a repository of thousands of medical terms along with their description, mediated through a meta-thesaurus. We have made this service available through a web-service to the client application to allow descriptions to be sought for relevant medical terms.

5 Conclusions and Future Work

This paper has described how the web-based application architecture that has been developed in the MIAKT project can be used to provide knowledge management for applications where semantic image annotation is necessary, and in particular, how it has been used to provide multi-media knowledge management in the medical domain. As well as image annotation, the system provides for multi-platform service invocation based on the instances of an application ontology, which we believe is a generic and flexible protocol for multi-application deployment.

In future developments the application ontology will be formalised into an abstract process model based description of the application which will provide a mechanism for generating unique application clients that are suited to users of different applications. Also the MIAKT application will be built-upon to provide greater support to the application of the multi-disciplinary meetings with

scheduling of hospital resources, further image analysis modules and classification services.

Our generic architecture lends itself to many different domains and we are looking forward to using the system to prototype different applications in different domains to prove its genericity.

6 Acknowledgements

The authors are grateful to their collaborators on the project - Kalina Bontcheva, Michael Brady, Liliana Cabral, Fabio Ciravegna, John Domingue, David Hawkes, Hugh Lewis, Enrico Motta, Maud Poissonier, Christine Tanner, Yorick Wilks and Yalin Zheng - for many valuable discussions. This project (MIAKT) is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) 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. McGuinness, D., van Harmelen, F.: Ontology Web Language (OWL) Overview. Available at <http://www.w3.org/TR/owl-features/> (August 2003).
2. McGuinness, D., Hendler, J., Stein, L.A.: DAML+OIL: An Ontology Language for the Semantic Web. *IEEE Intelligent Systems*, (2002) 72–80.
3. Harris, S., Gibbins, N.: 3store: Efficient Bulk RDF Storage. In *Proceedings 1st International Workshop on Practical and Scalable Semantic Web Systems*, Sanibel Island, Florida, USA, (2003) 1–15.
4. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. *Scientific American*, May 2001.
5. Bontcheva, K.: Reuse and Problems in the Evaluation of NLG systems. In *Proceedings of EACL.03 Workshop on Evaluation Initiatives*, April 2003.
6. Zheng, Y., Tanner, C., Hill, D.L.G., Hawkes, D., White, M., Khazen, M., Leach, M.: Alignment of Dynamic Contrast Enhanced MR Volumes of the Breast for a Multicenter Trial: An Exemplar GRID Application. In *Medical Imaging 2004*, San Diego, CA, USA. SPIE, February 2004 (to appear).
7. American College of Radiology. Breast Imaging Reporting and Data System: BI-RADS. Available at http://www.acr.org/departments/stand_accrad/birads/.
8. Motta, E., Domingue, J., Cabral, L., Gaspari, M.: IRS-II: A Framework and Infrastructure for Semantic Web Services. *2nd International Semantic Web Conference (ISWC2003)*, Florida, USA, (20th–23rd October 2003) 306–318.
9. TouchGraph LLC. TouchGraph. Available at <http://www.touchgraph.com/>.
10. The WSDL Specification, WWW Consortium, March 2001. Available at <http://www.w3.org/TR/wsdl>.
11. The SOAP Specifications, WWW Consortium. Available at <http://www.w3.org/TR/soap/>.
12. Heath, M., Bowyer, K.W, Kopans, D. et al.: Current Status of the Digital Database for Screening Mammography, in *Digital Mammography*, Kluwer Academic Publishers (1998), 457–460.