6-Hunter.qxp 15/12/2005 15:19 Page 264

A framework to enable the semantic inferencing and querying of multimedia content

Jane Hunter*

Distributed Systems Technology Centre, University of Queensland,
St Lucia, Queensland 4072, Australia
E-mail: jane@dstc.edu.au
*Corresponding author

Suzanne Little

ITEE Department, University of Queensland, St Lucia,
Queensland 4072, Australia
E-mail: slittle@dstc.edu.au

Abstract: Cultural institutions, broadcasting companies, academic, scientific and defence organisations are producing vast quantities of digital multimedia content. With this growth in audiovisual material comes the need for standardised representations encapsulating the rich semantic meaning required to enable the automatic filtering, machine processing, interpretation and assimilation of multimedia resources. Additionally generating high-level descriptions is difficult and manual creation is expensive although significant progress has been made in recent years on automatic segmentation and low-level feature recognition for multimedia. Within this paper we describe the application of semantic web technologies to enable the generation of high-level, domain-specific, semantic descriptions of multimedia content from low-level, automatically-extracted features. By applying the knowledge reasoning capabilities provided by ontologies and inferencing rules to large, multimedia data sets generated by scientific research communities, we hope to expedite solutions to the complex scientific problems they face.

Keywords: images; inferencing; multimedia; scientific; semantic.

Reference to this paper should be made as follows: Hunter, J. and Little, S. (2005) 'A framework to enable the semantic inferencing and querying of multimedia content', *Int. J. Web Engineering and Technology*, Vol. 2, Nos. 2/3, pp.264–286.

Biographical notes: Jane Hunter is a Distinguished Research Fellow at the Distributed Systems Technology Cooperative Research Centre. For the past ten years she has been developing data models, ontologies, metadata standards (Dublin Core, MPEG-7, and MPEG-21), schemas, software tools, and query languages to enable the indexing, archival, analysis, integration and management of large multimedia collections. She is also the liaison between MPEG (Moving Pictures Experts Group) and the W3C.

Suzanne Little is a PhD student at the School of Information Technology and Electrical Engineering at the University of Queensland. Her current research focus is the application of semantic web technologies to multimedia indexing and querying. She has an honours degree in information technology from the University of Queensland.

Copyright © 2005 Inderscience Enterprises Ltd.

1 Introduction and objectives

The semantic web (W3C, 2001a) is an activity of the W3C which aims to extend the current web by providing tools that enable resources on the web to be defined and semantically linked in a way that facilitates automated discovery, aggregation and re-use across various applications. Within the semantic web activity, a web ontology language (OWL) (W3C, 2003b,c) has been developed for defining structured, web-based ontologies that will provide richer integration and interoperability of data among descriptive communities. Adoption of standardised metadata vocabularies and ontologies, expressed in standardised machine-processable languages such as OWL are enabling people, machines, services, agents to find, consume, process and produce data, information and knowledge.

In parallel with advancements in the development of the semantic web, is a rapid increase in the size and range of multimedia resources being added to the web. Cultural, educational, government, scientific and research organisations are making enormous contributions to the amount of heterogenous multimedia information on the internet through the digitisation and online publication of image, audio, video and data sets. To enable multimedia content to be included in the semantic web, it needs to be described semantically. Generating descriptions of multimedia content is inherently problematic because of the volume and complexity of the data, its multidimensional nature and the potentially high subjectivity of human-generated descriptions. Significant progress has been made in recent years on automatic segmentation or structuring of multimedia content and the recognition of low-level features within such content. However, comparatively little progress has been made on machine-generation of semantic descriptions of audiovisual information. Within this paper we describe our recent and ongoing research efforts within the FUSION (Fuel cell Understanding through Semantic Inferencing, Ontologies and Nanotechnology) project (FUSION Project, 2003) at DSTC which attempts to combine standards for multimedia content description with recent advances in semantic web technologies, to develop systems that maximise the potential knowledge which can be mined from large heterogeneous multimedia information sets on the internet. More specifically we describe our research into the inferencing of high-level, domain-specific, semantic descriptions of multimedia content from low-level, automatically-extracted (MPEG-7 (ISO/IEC 15938-5 FDIS Information Technology, 2001)) features, using ontologies and pre-defined inferencing rules. Such semantic descriptions enable sophisticated semantic querying of multimedia resources in terms familiar to the user's domain and ensure that the information and knowledge within the multimedia content has a much greater chance of being discovered and exploited by services, agents and applications on the web.

Previous research which has attempted to apply semantic web technologies across a broad range of disciplines, domains and databases on the web (e.g. the Open Archives Data Providers (Open Archives Initiative, 2003)), has met with limited success. Trying to define and apply rules to automatically infer relationships between separate atomic multimedia resources from across disciplines and databases distributed across the internet, proved extremely difficult (Little et al., 2002). With little or no filtering or quality control of the metadata and minimal agreement on vocabularies, there is insufficient interoperability of the metadata at both the syntactic and semantic level. Hence we decided to develop and evaluate our semantic inferencing framework within a specific deep and

6-Hunter.qxp 14/12/2005 13:11 Page 266

narrow scientific domain in which quite well-established rules and vocabularies are already accepted. In order to test our hypothesis and evaluate our tools, we chose a particular eScience materials engineering application and corresponding community – the optimisation of fuel cells. This application requires the analysis and assimilation of microstructural images, performance data and manufacturing parameters which we aim to support through sophisticated semantic querying capabilities.

Although one of the strengths of the proposed approach is the interoperability provided by the use of (XML Schema-defined) metadata standards and OWL ontologies, this approach also introduced an additional problem – understanding and closing the gaps within the semantic web's layer architecture (Miller, 2001). In particular we had to develop mechanisms for linking the XML instances and XML Schema (W3C, 2001c,d,e) layer with the resource description framework (RDF) (Lassila and Swick, 1999) and ontology and inferencing layers – interfaces which are still ambiguous or not precisely defined within the semantic web community.

The remainder of the paper is structured as follows:

- Section 2 describes previous research which has focussed on the generation of semantic image or multimedia descriptions from low-level features. As far as we are aware, no one has attempted to do this using ontologies and inferencing rules.
- Section 3 describes the overall system architecture.
- Section 4 describes the ontologies which we developed using OWL:
    - The MPEG-7 ontology which defines the semantics of domain-independent multimedia content descriptions.
    - The FUSION ontology which defines the semantics of terms used by the Fuel Cell community.
    - The OME (Open Microscopy Environment (Open Microscopy Environment, 2003)) ontology which defines the semantics of the OME metadata vocabulary for describing microscopy data.
    - The ABC Top Level Ontology (Lagoze and Hunter, 2001) which we used to link and harmonise the MPEG-7, OME and FUSION ontologies.
- Section 5 describes the automatic image analysis programs which we invoked to extract low-level MPEG-7 features and to generate MPEG-7 descriptions.
- Section 6 describes how we defined and applied inferencing rules to infer high-level semantic descriptions from low-level MPEG-7 features.
- Section 7 describes the more sophisticated kinds of queries which could be applied, given the semantic descriptions, to improve the assimilation and presentation of large-scale, complex multimedia datasets and expedite domain-experts' understanding or knowledge.
- Section 8 provides our conclusions and describes future directions and plans for this research.

6-Hunter.qxp 14/12/2005 13:11 Page 267

2 Previous work

Most recent research in image or video indexing and retrieval has focused on query-by-example (QBE) (Deng and Manjunath, 1997; Marchand-Maillet, 2000; Zhang et al., 1997). However, semantic querying or query-by-keyword (QBK) has recently motivated research into semantic indexing of images and video content. A number of research efforts have investigated the application of automatic recognition techniques to extract a number of different low-level visual or audio features which together can be used to generate improved semantic descriptions of multimedia content. Recent attempts include work by Naphade et al. (2000) which propose a statistical factor graph framework to bridge the gap between low-level features and semantic concepts. Chang et al. (1998) use a library of examples approach, which they call semantic visual templates. Adams et al. (2003) manually annotate atomic audio and video features in a set of training videos and from these develop explicit statistical models to automatically label the video with high-level semantic concepts. Zhao and Grosky (2002) employ a latent semantic indexing technique which integrates a variety of automatically extracted visual features (global and sub-image colour histograms and anglograms for shape-based and colour-based representations) to enable semantic indexing and retrieval of images.

In their recent paper, Marques and Barman (2003) used machine-learning techniques to semantically annotate images with semantic descriptions defined within ontologies. Our approach is to replace the machine-learning component with domain-specific inferencing rules defined by domain-experts through an intuitive user-friendly interface. Hatala and Richards (2003) applied a similar approach to generating semantic descriptions for learning objects – relevant values for a particular metadata field are suggested by applying ontologies and rules – but they have not applied this approach to image content. As far as we are aware, the work described here is the first attempt at applying semantic inferencing rules to the semantic indexing of image data.

Although many of these earlier approaches are similar to ours, in that they are using computers to automatically extract low-level features and then deduce high-level descriptions from specific combinations of low-level features, they do not employ the formal languages, standardised metadata schemas, ontologies and inferencing rules which we are using to implement this approach. We believe that the combined use of semantic web technologies – metadata standards such as MPEG-7 and OME defined using XML Schema, ontologies expressed in OWL and rules expressed in RuleML (The Rule Markup Initiative, 2003), will enhance the potential interoperability of the multimedia content being described, and further enhance its ability to be discovered, processed, filtered, exchanged and combined by people, machines, services and agents on the web.

3 The system architecture

Figure 1 illustrates the overall system components and architecture for the knowledge management system which we have developed for the FUSION project.

Figure 1 Overall system architecture and components

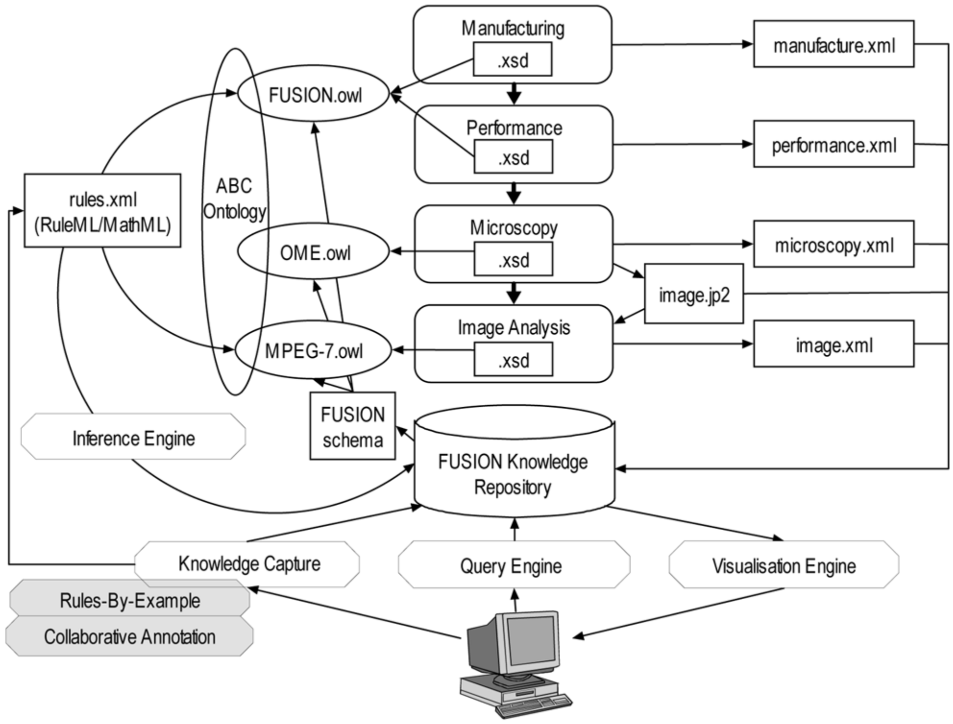

XML Schemas are used to define the manufacturing data, performance data, microscopy details, and image features associated with cross-sectional images (JPEG, 2000) of fuel cells. The metadata schemas combine the MPEG-7 (Multimedia Content Description) and OME (Open Microscopy Environment) standards with a metadata schema developed specifically to satisfy the requirements of fuel cell analysts (FUSION). In addition, ontologies were developed to define the semantic relationships between the terms used in each of these schemas and RuleML was used to define the relationships between low-level automatically-extracted MPEG-7 features and high-level FUSION concepts or terms familiar to the fuel cell experts. All of the generated and validated metadata is stored in a central knowledge repository. A query engine, visualisation engine and knowledge capture and extraction tools sit on top of the knowledge repository and enable users to access, interpret, assimilate and mine the stored data.

Each of these components is described in detail in the following sections.

4 Metadata schemas and ontologies

Metadata is the value-added information which documents the administrative, descriptive, preservation, technical and usage history and characteristics associated with resources. It provides the underlying foundation upon which search engines and digital asset management systems rely to provide fast, precise access to relevant resources across networks and between organisations. The metadata associated with multimedia objects is

6-Hunter.qxp 14/12/2005 13:11 Page 269

infinitely more complex than simple metadata for resource discovery of simple atomic textual documents and the problems and costs associated with generating such metadata are correspondingly magnified.

Metadata standards enable interoperability between systems and organisations so that information can be exchanged and shared. In 2001, the Moving Pictures Expert Group (MPEG), a working group of ISO/IEC, published MPEG-7 (Technology, 2001), the 'Multimedia Content Description Interface', a standard for describing multimedia content. The goal of this standard is to provide a rich set of standardised tools to enable both humans and machines to generate and understand audiovisual descriptions which can be used to enable fast efficient retrieval from digital archives (pull applications) as well as filtering of streamed audiovisual broadcasts on the internet (push applications). MPEG-7 is intended to describe audiovisual information regardless of storage, coding, display, transmission, medium or technology. It addresses a wide variety of media types including: still pictures, graphics, 3D models, audio, speech, video and combinations of these (e.g. multimedia presentations). Within the specification, MPEG-7 definitions (description schemes and descriptors) are expressed in XML Schema (W3C, 2001c,d,e).

Various communities and initiatives are developing domain-specific metadata standards e.g. CIDOC/CRM (ICOM-CIDOC Data Model Working Group, 1995) for museum objects, FGDC (Federal Geographic Data Committee, 2003) for geospatial data, NewsML (International Press Telecommunications Council, 2003) for news objects and IEEE LOM (IEEE Learning Technology Standards Committee, 2002) for educational resources. Within the scope of the FUSION application we are interested in generating XML descriptions of microscopy images of cross-sections of fuel cells. We do this by designing an application profile which draws on and combines two existing metadata standards with a third metadata schema developed specifically for this domain:

- MPEG-7 (Multimedia Content Description Interface) – for describing the low level features and encoding format of images.
- OME (Open Microscopy Environment) – a standard for describing the source of microscopy data (this overlaps with mpeg7:CreationInfoDS).
- FUSION – which provides the metadata description of a fuel cell as required by fuel cell scientists or analysts. This includes descriptions of the thickness, surface area, density, porosity and conductivity of the anode, cathode and electrolyte which compose the fuel cell.

Appendices A, B and C provide examples of three simple XML descriptions of a JPEG-2000 image which has been taken from a cross-section of a fuel cell. These descriptions are compliant with the MPEG-7, OME and FUSION XML Schemas respectively.

The objectives of the work described below are to enable the MPEG-7 description to be generated automatically using image recognition programs. The OME metadata will also be generated automatically by manipulating the output from the software provided with the microscope. The FUSION description will be generated by defining and applying inferencing rules which will infer the anode, cathode and electrolyte components and their attributes from the automatically-extracted low-level MPEG-7 features (regions, colour, shape, texture).

4.1 Harmonisation through the ABC ontology

In order to enable semantic interoperability between the MPEG-7, OME and FUSION metadata vocabularies, and to define the semantic and inferencing relationships between the terms in these different schemas, we needed to develop ontologies for each of these vocabularies and merge or harmonise them. We used the top-level ABC ontology (Lagoze and Hunter, 2001) developed within the Harmony project to do this. The ABC ontology provides a global and extensible model that expresses the basic concepts that are common across a variety of domains and provides the basis for specialisation into domain-specific concepts and vocabularies. Other possible upper-level ontologies which could be used for this purpose include SUMO (Niles and Pease, 2001) and DOLCE (Gangemi et al., 2002). We harmonised the ontologies manually through a process of analysing and comparing the semantic relationships between the classes in each ontology (Doerr et al., 2003; Hunter, 2003). Figure 2 illustrates the harmonisation process at a very high level.

Figure 2 Ontology harmonisation using ABC

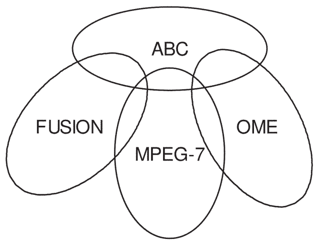

4.2 The MPEG-7 ontology

Figure 3 illustrates graphically a subset of the MPEG-7 OWL ontology (Hunter, 2001) which we have developed to define the semantics of terms used within the MPEG-7 standard for describing multimedia content. Figure 3 shows only those classes for which 'colour' is a visual descriptor and the five sub-properties of colour. A complete description of the MPEG-7 ontology is available at http://metadata.net/harmony/MPEG7/mpeg7.owl.

Figure 3 The MPEG-7 visual descriptor 'colour'

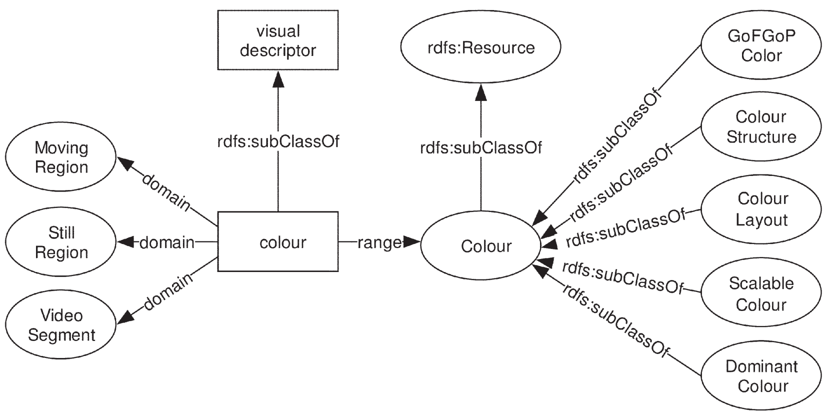

6-Hunter.qxp 14/12/2005 13:11 Page 271

4.3 The FUSION ontology

Figure 4 illustrates the key classes or concepts and their relationships, as defined in the ontology that we developed for the fuel cell community. This was developed in collaboration with domain experts using the approach recommended in (Noy and McGuinness, 2001).

Figure 4 The fuel cell ontology

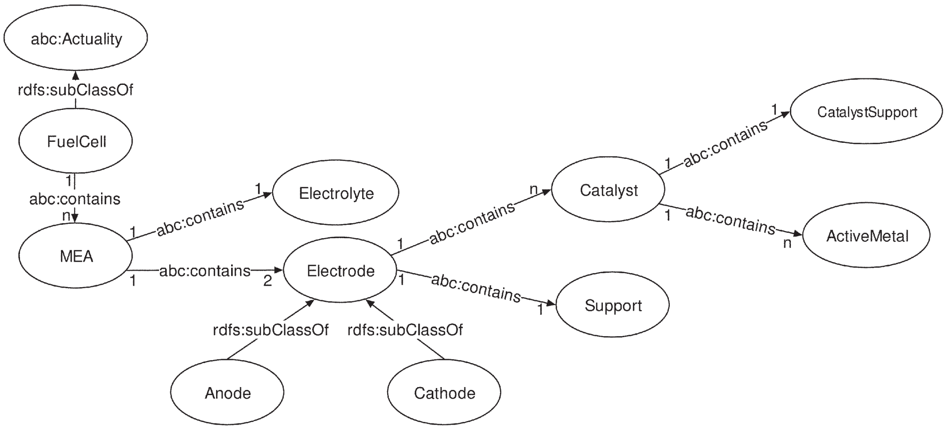

4.4 The OME ontology

Figure 5 illustrates a subset of the classes, relationships and properties that we defined when developing an OWL ontology for describing microscopy data from the OME CoreSemanticTypes (Open Microscopy Environment).

Figure 5 A subset of the OME ontology showing an instrument and its components

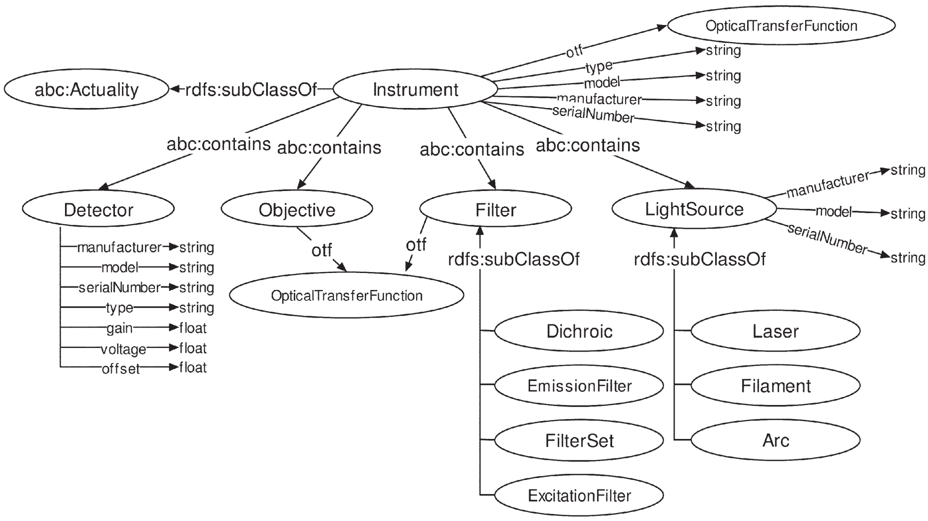

4.5 Linking ontologies to XML data

Once we have defined XML Schemas for the metadata descriptions and the corresponding ontologies in OWL, we need a way of linking the semantic definitions in the ontologies to the corresponding elements in the schemas or instances. In an earlier paper (Hunter and Lagoze, 2001) we described how the 'xx:semantics' attribute in XML Schema documents can be used to link element definitions to an ontology term thereby linking the semantic definitions to the preferred encoding or XML representation. The advantages of this approach are the separation of encodings from semantics and the consequent freedom to structure and name elements in XML to suit the application or user community without altering the meaning of the term. Appendix D contains the XML Schema for a fuel cell description and illustrates the use of the 'xx:semantics' attribute to link this schema to its corresponding OWL ontology. The code in Figure 6 is a brief excerpt from this.

Figure 6 Fragment of sample XML Schema for FUSION data

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmnls="http://www.w3.org/2001/XMLSchema"
xmnls:xs="http://www.w3.org/2001/XMLSchema"
xmlns:xx="http://www.example.org/XMLRDFSchemaBridge"
targetNamespace="http://metadata.net/fusion/FUSION">

  <xs:annotation>
    <xs:documentation>Draft XML Schema for FUSION</xs:documentation>
  </xs:annotation>

  <xs:element name="fuelcell" type="fuelcell"/>
  <xs:complexType name="fuelcell"
    xx:semantics="http://metadata.net/fusion/FUSION#fuelcell">
    <xs:sequence>
      <xs:element name="anode" type="anodeType"/>
      <xs:element name="electrolyte" type="electrolyteType"/>
      <xs:element name="cathode" ref="cathodeType"/>
    </xs:sequence>
    <xs:attribute name="id" use="required"/>
  </xs:complexType>
 ...
</xs:schema>
```

5 Automatic feature extraction

Automatic feature extraction for images is a significant research area and a wide variety of tools and mechanisms exist for analysing scientific and other types of images (e.g. see Internet Analysis Tools Registry http://www.cma.mgh.harvard.edu/iatr/display.php?search_type=all)). Within the scope of this project, we are not particularly interested in the underlying method by which these programs analyse images but we are interested in the outputs or features that such programs produce and whether they can be used by our domain experts or inferencing engines to infer the occurrence of higher-level semantic concepts e.g. regions or objects of a particular colour, texture and shape. Although we have designed our system so that it can invoke automatic image analysis tools which have been set up as web services, we have chosen initially to use MATLAB (MathWorks, 2003) because it is a popular and powerful tool which is currently being widely used by

6-Hunter.qxp 14/12/2005 13:11 Page 273

microscopists and scientists for analysing scientific images. In an ideal system, SOAP (W3C, 2003d) and WSDL (Christensen et al., 2001) would be employed to make the automatic feature extraction tools available as web services. This would allow greater choice and flexibility as more advanced multimedia extraction and analysis tools become available.

MATLAB is capable of producing a large amount of low-level data about the features and objects within an image e.g. area, mean density, standard deviation density, perimeter/length, X/Y centre, integrated density, etc. The data produced by MATLAB needs to be mapped to the image analysis (MPEG-7) description and related (using the inferencing rules) to the high-level semantic terms in the FUSION description. In particular we want to map regions extracted from the images to the fuel cell components such as the anode, cathode, electrolyte, catalyst and substrate. This process is described in detail in the next section.

6 Semantic inferencing rules

6.1 Rule description techniques

The most popular markup language for rules and the current most likely candidate for adoption by the semantic web community is RuleML (Boley et al., 2001, The Rule Markup Initiative, 2003), the Rule Markup Language. Developed by a consortium of industry and academic partners called the Rule Markup Initiative, RuleML provides a markup language which allows for descriptions of rules of the form if $x(a,b)$ [and|or] $y(b,c)$ then $z(b,d)$ and facts of the form $x(a,b)$ in XML format. Some translation and tool support is available for RuleML. A Java reasoning engine called Mandarax (Dietrich, 2003) is available, which allows querying and application of RuleML rules and facts. Other possible rule standards are the eXtensible Rule Markup Language (XRML[1]) and the Description Logic Markup Language (DLML[2]), neither of which are as well developed or widely supported as RuleML. Figure 7 provides an example of a simple rule expressed using RuleML.

In addition to rules defining relationships based on purely semantic lines, our particular application also requires rules to describe mathematical relationships within and between ontology terms and data elements. This is a common requirement for multimedia data associated with scientific applications. For example, an electrode contains catalyst particles and data on their average size is required – this information can be calculated from the size recorded for each of the catalyst particles using a mathematical formula. MathML (W3C, 2003a), the W3C's math markup language, provides a comprehensive method for describing both presentation and content of mathematical formulae. While MathML has also set theory and Boolean logic operators and could, therefore, also be used to define logic rules, it lacks the reasoning engine support of RuleML and can only (natively) define relationships such as equals, greater than etc., and cannot describe higher level semantic relationships such as *depicts*. Figure 8 illustrates a example of MathML describing a simple formula.

Figure 7 RuleML representation of the simple example

```
<?xml version="1.0" encoding="UTF-8"?>
<rulebase xmlns="http://www.ruleml.org/"
          xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xmlns:mpeg7="urn:mpeg:mpeg7:schema:2001"
          xmlns:mathml="http://www.w3.org/1998/Math/MathML"
          xmlns:xml="http://www.w3.org/XML/1998/namespace"
          xmlns:fusion="http://metadata.net/FUSION">
  <imp>
<!-- A StillRegion depicts Catalyst if the colour value is greater than
(112, 112, 112) and the shape is logically equivalent to a circle, the
texture is . and the magnification at which the image was taken is greater
than or equal to 40 -->
    <_head>
      <atom>
        <_opr> <rel>mpeg7:Depicts</rel> </_opr>
        <var>Region</var>
        <var>fusion:Catalyst</var>
      </atom>
    </_head>
    <_body>
      <and>
        <atom>
          <_opr> <rel>mathml:gt</rel> </_opr>
          <var>Region</var>
          <ind>mpeg7:DominantColor</ind>
          <ind>(113,113,113)</ind>
        </atom>
        <atom>
          <_opr> <rel>mpeg7:Shape</rel> </_opr>
          <var>Region</var>
          <ind>circle</ind>
        </atom>
        <atom>
          <_opr> <rel>mathml:geq</rel> </_opr>
          <ind>ome:Magnification</ind>
          <ind>10000</ind>
        </atom>
        <atom>
          <_opr> <rel>mpeg7:StillRegion</rel> </_opr>
          <var>Region</var>
        </atom>
      </and>
    </_body>
  </imp>
</rulebase>
```

Figure 8 Example of a formula/rule in MathML

```
<!-- The mean width of an anode is the mean of the difference between y1
and y2 coords of the object found in the media. -->
<mathml:reln><mathml:eq/>
  <mathml:ci>/anode/width/mean</mathml:ci>
  <mathml:ci>
    <mathml:mean>
      <mathml:reln><mathml:minus/>
        <mathml:ci>//mpeg7/object/region/coords/y2</mathml:ci>
        <mathml:ci>//mpeg7/object/region/coords/y1</mathml:ci>
      </mathml:reln>
    </mathml:mean>
  </mathml:ci>
</mathml:reln>
```

CellML (Cuellar et al., 2002) (a markup language for describing cell biology) uses MathML to describe mathematical relationships within marked-up data. CellML uses a subset of MathML to embed mathematical formulae into the documents which utilise variables identified by a *variable* tag elsewhere in the markup. The use of variable identifiers within the MathML content indicator field (*mathml:ci*) allows complex relationships to be captured and data to be recalculated more effectively, reducing user effort in building complete, accurate data collections.

Unlike CellML, the FUSION project accesses data from a broad range of sources (manufacturing, performance, microscopy and image analysis data) that are related through merged ontologies. Hence FUSION requires both semantic (ABC, MPEG-7, FUSION) and mathematical relationships to be defined. Therefore both RuleML (semantics) and MathML (mathematics) are required to achieve integration of comprehensive, complete datasets with a minimum of user effort. In addition the variables used in the rules and formulae may be in different documents and can be either abstract (ontological terms) or specific (xml elements). We use XPath (W3C, 1999) to identify and retrieve the specific variable values or elements from the XML instances and the xx:semantics attribute to link this back to the ontological definitions.

To apply the defined rules set and infer high-level semantic annotations, we can use either pre-existing engines (e.g. Mandarax for RuleML with a wrapper for formatting data input) or develop our own application. At this stage we have decided to use Mandarax, but because of the tenuous interface between OWL ontologies and RuleML it may be necessary to develop extensions or wrappers to Mandarax to incorporate knowledge implicitly defined with the existing OWL ontologies (e.g. subsumption relationships or transitive, symmetric or inverse relationships).

6.2 A simple semantic inferencing example

Suppose the user is interested in verifying the theory that small catalyst particle size affects the performance of fuel cells. The system does not explicitly record such information. However, by using rules to infer semantic descriptions from the features extracted from the microscopy images plus the implicit relationships between terms defined in the harmonised ontologies, this information can be derived.

Appendix A contains an MPEG-7 XML description generated from applying MATLAB image analysis to a microscopy image and mapping the output to MPEG-7. Using RuleML to express a simple rule such as the one described in Figure 9, it can be inferred that the feature in the image being referred to, http://metadata.net/FUSION/media/image.jp2, depicts a *Catalyst* as defined in the FUSION fuel cell ontology.

Figure 9 Simplified rule example

```
mpeg7:depicts (mpeg7:StillRegion, fusion:Catalyst)
      IF mpeg7:DominantColor > (112,112,112)
      AND mpeg7:Shape is circle
      AND ome:magnification >= 10000.
```

Figure 7 illustrates the complete RuleML representation of this example. In this example, the relationship being described is mpeg7:Depicts and the variables are identified as mpeg7:Color, mpeg7:Texture and mpeg7:Shape. The mpeg7 ontology provides the additional knowledge that mpeg7:DominantColor, mpeg7:ScalableColor, mpeg7:ColorLayout etc., are all rdfs:subClassOf mpeg7:Color. To find elements within the data that contain information about colour, the xx:semantics attribute in the XML Schema is used to link to the ontology and discover those tags that are relevant to mpeg7:Color and to extract their values in order to apply the rule.

Once the features which depict catalysts have been identified, the second step is to determine whether or not they qualify as 'small'. The user defines *small* to be catalysts of size less than 50 microns^2. The FUSION ontology defines *area* to be a sub-class of *size*. By comparing the *area* value for the catalyst features with the required value, fuel cells with 'small' catalyst particles can be discovered. The fuel cell ID associated with the microscopy information can be used to retrieve the performance data for each of these fuel cells. By appropriately presenting the images and performance data in parallel, the user will be able to determine whether their hypothesis appears reasonable and worthy of further exploration.

The fuzzy nature of image indexing and retrieval means that users are frequently not interested in exact matches between colour, texture or shape but only require similarity-based search and retrieval. In this first prototype, we have implemented this by including a default threshold of 5% (which can be modified) in matching instances against the values specified in the rules. However, we are in the process of developing a more intuitive graphical rule editing tool (Rules-By-Example) which enables domain experts to graphically define their semantic inferencing rules using visual examples drawn from palettes of colour, texture and shape e.g. a StillRegion of an image is a 'cathode' if its colour is like *this*, its texture is like *this* and its shape is like *this*. This interface will essentially provide a mechanism for capturing tacit domain-expert knowledge about image interpretation.

This simple example shows how, by using inferencing rules to assign semantic meaning to image features and ontologies to define further semantic relationships, information not explicitly recorded by the system can be mined and used to enable more sophisticated search and retrieval. Such inferencing may be done on-the-fly or in advance and stored in a knowledge repository (e.g. a relational or XML database) to enable more efficient querying and retrieval. In our system we are carrying out the inferencing in advance and then storing and indexing the inferred information in a database. At this stage we are using MySQL but we will shortly be evaluating the Tamino XML database (SoftwareAg, 2003).

Because the initial semantic inferencing rules may be speculative or even incorrect, it is necessary to include mechanisms for providing feedback and enabling the rules to be rejected or adapted based on this feedback. Feedback will be provided through the integration of sample database set that has been manually annotated by domain experts and that can be used to test and improve the initial rule definitions. Users can see the effect of applying a new rule immediately, by viewing the results retrieved from the sample database based on a corresponding query. This process is described in the next section.

7 Semantic querying and presentation

The high-level semantic descriptions which have been inferred together with the knowledge which is both explicitly and implicitly defined within harmonised ontologies (ABC, MPEG-7, FUSION, OME), enable more sophisticated querying to be performed in

6-Hunter.qxp 14/12/2005 13:11 Page 277

the ‘natural language’ of the community that is attempting to mine knowledge from the multimedia content.

For example, consider the following which is typical of the kind of queries carried out by fuel cell analysts: ‘How does the mean catalyst size in electrodes of width <20 microns effect electrode conductivity?’

In order to answer this query we need to:

1 find all electrodes of width <20 microns

2 find performance/conductivity data for those fuel cells

3 find the mean catalyst size for those electrodes.

These three steps are described in detail below.

Step 1. Find all electrodes with width <20 microns.

- The image regions which represent electrodes are determined by applying inferencing rules to the fuel cell images.
- The width is calculated from the *feature.length* property value for image regions that are electrodes – this may be difficult due to different magnifications of the microscopy images, and the need for conversions between units e.g. 16,500 nm = 16.5 microns which is <20 microns and fits the requirements.
- The image.xml files are retrieved for images which contain electrodes of width <20 microns.
- The image.xml files are linked to the microscopy.xml files which contain the Fuel Cell ID.

Step 2. Find performance/conductivity data for those fuel cells.

- The Fuel Cell ID is used to access the related performance.xml data which records the AC impedance test results which describe electrode conductivity.

Step 3. Find the mean catalyst size for those electrodes.

- Using the Fuel Cell ID, all related images can be retrieved.
- Using the rule which describes which features are ‘catalysts’, their area (which is a subclass of *size*, this is known from the ontology) can be derived.
- Using the MathML encoded rule for calculating *mean*, the mean value for all of the returned catalyst sizes for each electrode can be calculated.

The final step is to present all of the relevant related retrieved data in an integrated, synchronised and coherent presentation which will enable scientists to detect trends or patterns or visualise the results of hypotheses that would not be possible through traditional search interfaces. An earlier research prototype (Little et al., 2002), developed through a collaboration between DSTC and CWI[3], will be refined specifically to provide the presentation engine for this project. The system dynamically generates multimedia presentations in the SMIL 2.0 (W3C, 2001b) (Synchronised Multimedia Integration

Language) format by aggregating retrieved results sets based on the semantic relationships between them (which have been inferred by applying pre-defined inferencing rules to the associated metadata). CWI's Cuypers system (van Ossenbruggen et al., 2001) is used to apply constraints such as page size, band width, user preferences etc., to adapt the presentation components to the user's device capabilities and platform. In the earlier OAI-based prototype, we hardwired the mappings from semantic relationships between retrieved digital objects to spatio-temporal relationships in the presentations. For the FUSION project, we plan to develop a graphical user-interface which allows the fuel-cell experts to interactively map semantic relationships to their preferred spatio-temporal presentation modes. For example, sequences of images, in increasing order of magnification, may be displayed either as a slide show or tiled horizontally or vertically and either to the left or right of corresponding performance data. The domain experts can interactively define and modify their preferred modes of layout and presentation and their preferences will be reflected in the results of their next search (Figure 10).

Figure 10 The automatic SMIL presentation generator

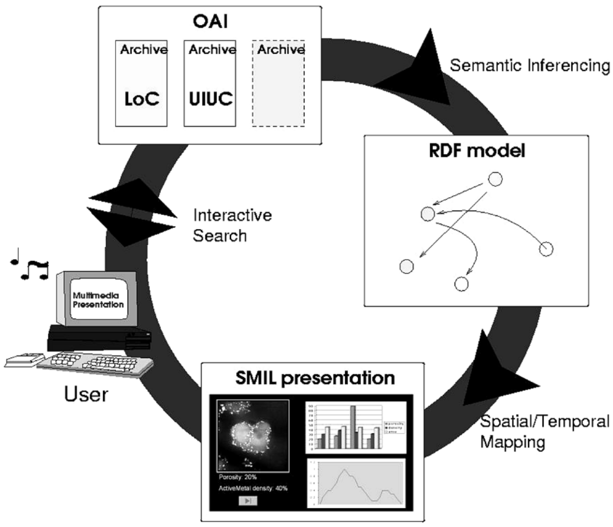

For example, the SMIL visualisation tool will be useful in enabling a better understanding of how fuel cells degrade over time. Sequences of micrographic images retrieved from identically-manufactured fuel cells after increasing duration of usage, could be displayed as animations, synchronised in parallel with the display of corresponding performance data. When SMIL presentations reveal new information or previously unrecognised patterns or trends, users will be able to save and annotate them, for later retrieval, as evidence to support new hypotheses or theories.

8 Conclusions and future work

Although the work which we have described in this paper is at a relatively preliminary stage, we have defined and developed an overall framework, based on semantic web technologies, that will enable the semantic indexing and retrieval of multimedia content, and more specifically scientific image content. We believe that the proposed application of semantic web technologies will enhance prior research which has already demonstrated that semantic indexing of multimedia is improved by combining a variety of low-level audiovisual cues automatically extracted from the content. Employing formal languages (XML Schema, OWL), standardised metadata schemas (MPEG-7, OME), ontologies and inferencing rules to deduce high-level semantic descriptions from low-level features, will expedite the semantic indexing process for multimedia content, enhance its potential interoperability, and hence enhance its ability to be discovered, processed, filtered, exchanged and combined by people, machines, services and agents on the web. Furthermore, the proposed architecture decouples the automatic feature extraction tools from the domain-specific inferencing rules, ontologies and semantic querying and presentation engines, thus maximising the system's flexibility, scalability and extensibility.

In the future we plan to:

- Develop a graphical rule editing tool (Rules-By-Example) which supports the definition of rules (in RuleML) using Query-By-Example (QBE)-type tools e.g. a StillRegion of an image is a 'cathode' if its colour is like this, its texture is like this and its shape is like this.
- Investigate the integration of sample database set manually annotated by domain experts which can be used to improve the initial rule definitions i.e. implement a rule-based system which learns and adapts based on users' manual annotations and feedback.
- Investigate and compare alternative database or repository options for storing the data, metadata and images e.g. D-Space, Tamino or a traditional relational database such as Oracle or MySQL.
- Further develop the system as a web services architecture by enabling both the automatic feature extraction tools and the sets of inferencing rules and Mandarax reasoning engine to be accessed or discovered as web services and made available to all members of a particular community. Such an architecture enables collaborative development and sharing of rules for the benefit of the whole community but still allows the researchers to retain their (copyright) images and data.
- Develop a user interface to the SMIL Presentation/Visualisation Engine which enables the users to define their preferred mappings from semantic relationships between atomic multimedia resources to spatio-temporal presentation modes.
- Continue the implementation of a prototype knowledge management system for the fuel cell community and carry out detailed evaluation and usability studies and system refinements.

Acknowledgements

The work reported in this paper has been funded in part by the Cooperative Research Centre for Enterprise Distributed Systems Technology (DSTC) through the Australian Federal Government's CRC Programme (Department of Education, Science and Training). Thanks also to staff from the Centre for Microscopy and Microanalysis at the University of Queensland for their feedback and the images used in this project.

6-Hunter.qxp 14/12/2005 13:11 Page 280

References

Adams, B., Iyengar, G. et al. (2003) 'Semantic indexing of multimedia content using visual, audio and text cues', *EURASIP Journal on Applied Signal Processing*.

Boley, H., Tabet, S. and Wagner, G. (2001) 'Design rationale of RuleML: a markup language for semantic web rules', *Semantic Web Working Symposium (SWWS)*.

Chang, S.F., Chen, W. and Sundaram, H. (1998) 'Semantic visual templates: linking visual features to semantics', *IEEE International Conference on Image Processing (ICIP '98)*, Chicago, IL.

Christensen, E., Curbera, F. et al. (2001) 'Web Services Description Language (WSDL) 1.1', *W3C Note*, available from: www.w3.org/TR/wsdl.

Cuellar, A., Nielsen, P. et al. (2002) 'CellML specification 1.1 draft', *Draft Specification, 6 November 2002*, available from: www.cellml.org/public/specification/20021106/cellml_specification.html.

Deng, Y. and Manjunath, B.S. (1997) 'Content-based search of video using colour, texture and motion', *International Conference on Image Processing (ICIP '97)*, Washington, DC, USA, October 26–29.

Dietrich, J. (2003) *The Mandarax Project*, available from: www.mandarax.org/.

Doerr, M., Hunter, J. and Lagoze, C. (2003) 'Towards a core ontology for information integration', *Journal of Digital Information*, Vol. 4.

Federal Geographic Data Committee (FDGC) (2003) available from: www.fgdc.gov/.

FUSION Project (2003) available from: http://metadata.net/sunago/fusion.htm.

Gangemi, A., Guarino, N. et al. (2002) 'Sweetening ontologies with DOLCE', *EKAW 2002*, Siguenza, Spain.

Hatala, M. and Richards, G. (2003) 'Value-added metatagging: ontology and rule based methods for smarter metadata', *Rules and Rule Markup Languages for the Semantic Web (RuleML2003)*.

Hunter, J. (2001) 'Adding multimedia to the semantic web – building an MPEG-7 Ontology', *International Semantic Web Working Symposium (SWWS)*, Stanford, July 30–August 1.

Hunter, J. (2003) 'Enhancing the semantic interoperability of multimedia through a core ontology', *IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Conceptual and Dynamical Aspects of Multimedia Content Description*.

Hunter, J. and Lagoze, C. (2001) 'Combining RDF and XML Schemas to enhance interoperability between metadata application profiles', *The 10th World Wide Web Conference (WWW10)*, Hong Kong.

ICOM-CIDOC Data Model Working Group (1995) *Conceptual Reference Model*, CIDOC/CRM, available from: http://cidoc.ics.forth.gr/.

IEEE Learning Technology Standards Committee (2002) *IEEE LOM Working Draft 6.1* (IEEE 1484.12.1–2002).

International Press Telecommunications Council NewsML (2003) available from: http://www.newsml.org/.

ISO/IEC 15938-5 FDIS Information Technology (2001) *Multimedia Content Description Interface – Part 5: Multimedia Description Schemes*, July.

JPEG (2000) JPEG-2000. available from: http://www.jpeg.org/JPEG2000.html.

Lagoze, C. and Hunter, J. (2001) 'The ABC ontology and model (v3.0)', *Journal of Digital Information*, Vol. 2.

Lassila, O. and Swick, R. (1999) 'Resource description framework (RDF) model and syntax specification', *W3C Recommendation*, 22nd February, available from: www.w3.org/TR/REC-rdfsyntax.

Little, S., Geurts, J. and Hunter, J. (2002) 'The dynamic generation of intelligent multimedia presentations through semantic inferencing', *The 6th European Conference on Research and Advanced Technology for Digital Libraries, (ECDL 2002)*, Rome, Italy, September.

Marchand-Maillet, S. (2000) 'Content-based video retrieval: an overview', *Technical Report*.

Marques, O. and Barman, N. (2003) 'Semi-automatic semantic annotation of images using machine learning techniques', *International Semantic Web Conference*, Florida, October 2003.

6-Hunter.qxp 14/12/2005 13:11 Page 281

MathWorks MATLAB (2003) available from: www.mathworks.com/products/imageprocessing.

Miller, E. (2001) *The W3C Semantic Web Activity*, July 30–31.

Naphade, M.R., Kozintsev, I. et al. (2000) 'A factor graph framework for semantic indexing and retrieval in video', *IEEE Workshop on Content-based Access of Image and Video Libraries (CBAIVL'00)*, Hilton Head, SC, June 16.

Niles, I. and Pease, A. (2001) 'Towards a standard upper ontology', *The 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, Ogunquit, Maine, October 17–19.

Noy, N. and McGuinness, D. (2001) 'Ontology development 101: A guide to creating your first ontology', *Technical Report*.

Open Archives Initiative OAI (2003) available from: www.openarchives.org.

Open Microscopy Environment OMECoreSemanticTypes.ome, 4 Sept 2003, available from: http://openmicroscopy.org/XMLschemas/STD/RC2/OMECoreTypes.ome.

SoftwareAg Tamino XML Server (2003) available from: www.softwareag.com/tamino/.

Technology, I. I-F. I. (2001) *Multimedia Content Description Interface – Part 5: Multimedia Description Schemes*, July.

The Rule Markup Initiative RuleML (2003) available from: www.dfki.uni-kl.de/ruleml.

van Ossenbruggen, J., Geurts, J. et al. (2001) 'Towards second and third generation web-based multimedia', *The 10th International World Wide Web Conference*, Hong Kong, May 1–5.

W3C (1999) *XML Path Language (XPath), W3C Recommendation*, J. Clark and S. DeRose (Eds) 16 November 1999, available from: www.w3.org/TR/xpath.

W3C (2001a) *Semantic Web Activity Group*, available from: www.w3.org/2001/sw.

W3C (2001b) *Synchronized Multimedia Integration Language (SMIL 2.0) Specification, W3C Recommendation*, A.Cohen (Ed.) August 7, 2001, available from: www.w3.org/TR/smil20.

W3C (2001c) *XML Schema Part 0: Primer, W3C Recommendation*, D.C. Fallside (Ed.) 2 May, 2001, available from: www.w3.org/TR/xmlschema-0/.

W3C (2001d) *XML Schema Part 1: Structures, W3C Recommendation*, H.S. Thompson, D. Beech, M. Maloney and N. Mendelsohn (Eds) 2 May 2001, available from: www.w3.org/TR/xmlschema-1.

W3C (2001e) *XML Schema Part 2: Datatypes, W3C Recommendation*, P.V. Biron and A. Malhotra (Eds) 2 May 2001, available from: www.w3.org/TR/xmlschema-2/.

W3C (2003a) *Mathematical Markup Language (MathML), W3C Proposed Edited Recommendation*, D. Carlisle, P. Ion, R. Miner and N. Poppelie (Eds) 04 August 2003, available from: www.w3.org/TR/2003/PER-MathML2-2003080.

W3C (2003b) *OWL Web Ontology Language Overview, Candidate Recommendation*, D.L. McGuiness and F. van Harmelen (Eds) 18 August 2003, available from: www.w3.org/TR/owl-features.

W3C (2003c) *OWL Web Ontology Language Reference, W3C Candidate Recommendation*, M. Dean and G. Schreiber (Eds) 18 Aug 2003, available from: www.w3.org/TR/owl-ref/.

W3C (2003d) *SOAP Version 1.2 Part 0: Primer, W3C Recommendation*, N. Mitra (Ed.) 24 June 2003, available from: www.w3.org/TR/2003/REC-soap12-part0-20030624/.

Zhang, H., Wang, A. and Altunbasak, Y. (1997) 'Content-based video retrieval and compression: a unified solution', *IEEE International Conference on Image Processing*, Santa Barbara, CA, October 1997.

Zhao, R. and Grosky, W.I. (2002) 'Negotiating the semantic gap: from feature maps to semantic landscapes', *Pattern Recognition*, Vol. 35, pp.51–58.

Notes

1 eXtensible Rule Markup Language (http://xrml.kaist.ac.kr).

2 Description Logic Markup Language (http://co4.inrialpes.fr/xml.dlml.

3 Centrum voor Wiskunde en Informatica, Amsterdam (http://www.cwi.nl).

6-Hunter.qxp 14/12/2005 13:11 Page 282

Appendix A MPEG-7 description of an image of a fuel cell


```xml
<?xml version="1.0" encoding="iso-8859-1"?>
<Mpeg7 xmlns="urn:mpeg:mpeg7:schema:2001"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xmlns:mpeg7="urn:mpeg:mpeg7:schema:2001"
     xmlns:xml="http://www.w3.org/XML/1998/namespace"
     xsi:schemaLocation="urn:mpeg:mpeg7:schema:2001 .\Mpeg7-2001.xsd">
   <Description xsi:type="ContentEntityType">
     <MultimediaContent xsi:type="ImageType">
        <Image>
           <MediaLocator>
              <MediaUri>http://metadata.net/FUSION/media/image012.jp2
              </MediaUri>
           </MediaLocator>
           <CreationInformation>
              <CreationUri>http://metadata.net/FUSION/microscopy012.xml
              </CreationUri>
           </CreationInformation>
           <MediaInformation>
              <MediaProfile>
                <MediaFormat>
                   <Content href="urn:mpeg:mpeg7:cs:ContentCS:2001:2">
                      <Name>visual</Name>
                   </Content>
                   <FileFormat href="urn:mpeg:mpeg7:cs:FileFormatCS:2001:3">
                      <Name xml:lang="en">JPEG2000</Name>
                   </FileFormat>
                   <FileSize>10483</FileSize>
                   <VisualCoding>
                      <Format
href="urn:mpeg:mpeg7:cs:VisualCodingFormatCS:2001:1" colorDomain="binary">
                         <Name xml:lang="en">JPEG2000</Name>
                      </Format>
                      <Pixel aspectRatio="0.75" bitsPer="8"/>
                      <Frame height="288" width="352" rate="25"/>
                   </VisualCoding>
                </MediaFormat>
              </MediaProfile>
           </MediaInformation>
           <Object id="object1">
              <Mask xsi:type="SpatialMaskType">
                <SubRegion>
                   <Polygon>
                      <Coords mpeg7:dim="2 5"> 5 25 10 20 15 15 10 10 5
15</Coords>
                   </Polygon>
                </SubRegion>
              </Mask>
              <VisualDescriptor xsi:type="ScalableColorType" numOfCoeff="16"
numOfBitplanesDiscarded="0">
                <Coeff> 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 </Coeff>
              </VisualDescriptor>
              <VisualDescriptor xsi:type="RegionShapeType">
                <MagnitudeOfART>0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 </MagnitudeOfART>
              </VisualDescriptor>
              <VisualDescriptor xsi:type="HomogenousTextureType">
                ...
              </VisualDescriptor>
           </Object>
        </Image>
     </MultimediaContent>
   </Description>
</Mpeg7>
```

Appendix B OME description of the microscope used to capture the fuel cell image


```
<?xml version="1.0" encoding="UTF-8"?>
<OME xmlns="http://www.openmicroscopy.org/XMLschemas/OME/FC/ome.xsd"
xmlns:Bin="http://.../XMLschemas/BinaryFile/RC1/BinaryFile.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://...ome.../XMLschemas/OME/FC/ome.xsd
                    http://...ome.../XMLschemas/OME/FC/ome.xsd
                    http://...ome.../XMLschemas/STD/RC2/STD.xsd">
  <Instrument ID="urn...Instrument:123456">
    <Microscope Manufacturer="Zeiss" Model="foo" SerialNumber="bar"
Type="Upright"/>
    <LightSource ID="urn...LightSource:123456" Manufacturer="Olympus"
Model="WMD Laser" SerialNumber="1...4">
      <Laser Type="Semiconductor" Medium="GaAs">
        <Pump ID="urn...LightSource:123789"/>
      </Laser>
    </LightSource>
    <LightSource ID="urn...LightSource:123123" Manufacturer="Olympus"
Model="Realy Bright Lite" SerialNumber="1jhf16">
      <Arc Type="Hg"/>
    </LightSource>
    <Detector ID="urn...Detector:123456" Type="CCD" Manufacturer="Kodak"
Model="Instamatic" SerialNumber="frf8u198"/>
    <Objective ID="urn...Objective:123456" Manufacturer="Olympus"
Model="SPlanL" SerialNumber="456anxcoas123">
      <LensNA>2.4</LensNA>
      <Magnification>10000</Magnification>
    </Objective>
    <Filter ID="urn...Filter:123456">
      <FilterSet Manufacturer="Omega" Model="SuperGFP"
LotNumber="123LJKHG123"/>
    </Filter>
    <OTF ID="urn...OTF:123456" PixelType="int8" OpticalAxisAvrg="true"
SizeX="512" SizeY="512">
      <ObjectiveRef ID="urn...Objective:123456"/>
      <FilterRef ID="urn...Filter:123456"/>
      <Bin:External Compression="bzip2" SHA1="012...456"
href="OTF123.otf"/>
    </OTF>
  </Instrument>
...
</OME>
```

6-Hunter.qxp 14/12/2005 13:11 Page 284

Appendix C FUSION description of a fuel cell image

```
<?xml version="1.0" encoding="UTF-8"?>
<fuelcell xmlns="http://metadata.net/fusion/FUSION">
  <anode>
    <density context="fusion:catalyst" unit="per micron squared">42
    </density>
    <porosity unit="microns">42</porosity>
    <surfaceArea context="fusion:fuelcell" unit="microns squared">42
    </surfaceArea>
    <surfaceArea context="fusion:electrolyte" unit="microns squared">42
    </surfaceArea>
    <width>
      <mean unit="microns">42</mean>
      <sdev unit="microns">8</sdev>
    </width>
    <catalyst id="1">
      <density context="fusion:activeMetal" unit="per micron squared">42
      </density>
      <nearestNeighbour degree="1" unit="">42</nearestNeighbour>
      <shape>circle</shape>
      <size unit="microns">42</size>
      <surfaceArea unit="microns squared">42</surfaceArea>
      <size context="fusion:activeMetal">
        <mean unit="microns">42</mean>
        <sdev unit="microns">8</sdev>
      </size>
      <activeMetal id="1">
        <nearestNeighbour degree="1" unit="microns">42</nearestNeighbour>
        <shape>circle</shape>
        <size unit="microns">42</size>
        <surfaceArea unit="microns squared">42</surfaceArea>
      </activeMetal>
      <catalystSupport id="1">
        <composition>carbon</composition>
        <porosity unit="percentage">42</porosity>
        <size unit="microns">42</size>
      </catalystSupport>
    </catalyst>
  </anode>
  <electrolyte>
    <conductivity unit="">42</conductivity>
    <composition>nafion</composition>
    <porosity unit="percentage">42</porosity>
    <width>
      <mean unit="microns">42</mean>
      <sdev unit="microns">8</sdev>
    </width>
  </electrolyte>
  <cathode>same as anode</cathode>
</fuelcell>
```

6-Hunter.qxp 14/12/2005 13:11 Page 285

Appendix D XML Schema for FUSION

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmnls="http://www.w3.org/2001/XMLSchema"
xmnls:xs="http://www.w3.org/2001/XMLSchema"
xmlns:xx="http://www.example.org/XMLRDFSchemaBridge"
targetNamespace="http://metadata.net/fusion/FUSION">
   <xs:annotation>
     <xs:documentation>Draft XML Schema for FUSION</xs:documentation>
   </xs:annotation>
   <xs:element name="fuelcell" type="fuelcell"/>
   <xs:complexType name="fuelcell"
xx:semantics="http://metadata.net/fusion/FUSION#fuelcell">
     <xs:sequence>
        <xs:element name="anode" type="anodeType"/>
        <xs:element name="electrolyte" type="electrolyteType"/>
        <xs:element name="cathode" ref="cathodeType"/>
     </xs:sequence>
     <xs:attribute name="id" use="required"/>
   </xs:complexType>
   <xs:complexType name="anodeType"
xx:semantics="http://metadata.net/fusion/FUSION#anode">
     <xs:sequence>
        <xs:element name="density" type="densityType"/>
        <xs:element name="porosity" type="porosityType"/>
        <xs:element name="surfaceArea" type="surfaceAreaType"/>
        <xs:element name="width" type="widthType"/>
        <xs:element name="catalyst" type="catalystType"/>
     </xs:sequence>
   </xs:complexType>
   <xs:complexType name="catalystType"
xx:semantics="http://metadata.net/fusion/FUSION#catalyst">
     <xs:sequence>
        <xs:element name="density" type="densityType"/>
        <xs:element name="nearestNeighbour" type="nearestNeighbourType"/>
        <xs:element name="shape" type="shapeType"/>
        <xs:element name="size" type="sizeType"/>
        <xs:element name="surfaceArea" type="surfaceAreaType"/>
        <xs:element name="activeMetal" type="activeMetalType"/>
        <xs:element name="catalystSupport" type="catalystSupportType"/>
     </xs:sequence>
     <xs:attribute name="id" use="required"/>
   </xs:complexType>
   <xs:complexType name="activeMetalType">
     <xs:sequence>
        <xs:element name="nearestNeighbour" type="nearestNeighbourType"/>
        <xs:element name="shape" type="shapeType"/>
        <xs:element name="size" type="sizeType"/>
        <xs:element name="surfaceArea" type="surfaceAreaType"/>
     </xs:sequence>
     <xs:attribute name="id" use="required"/>
   </xs:complexType>
   <xs:complexType name="catalystSupportType">
     <xs:sequence>
        <xs:element name="composition" type="compositionType"/>
        <xs:element name="porosity" type="porosityType"/>
        <xs:element name="size" type="sizeType"/>
     </xs:sequence>
   </xs:complexType>
   <xs:complexType name="electrolyteType"
```

6-Hunter.qxp 14/12/2005 13:11 Page 286


```
xx:semantics="http://metadata.net/fusion/FUSION#electrolyte">
    <xs:sequence>
      <xs:element name="conductivity" type="conductivityType"/>
      <xs:element name="composition" type="compositionType"/>
      <xs:element name="porosity" type="porosityType"/>
      <xs:element name="width" type="widthType"/>
    </xs:sequence>
  </xs:complexType>
     <xs:simpleType name="compositionType"
xx  xx:semantics="http://metadata.net/fusion/FUSION#composition">
       <restriction base="string"/>
     </xs:simpleType>
     <xs:complexType name="conductivityType"
   xx:semantics="http://metadata.net/fusion/FUSION#conductivity">
       <xs:attribute name="unit" type="xs:string"/>
     </xs:complexType>
     <xs:complexType name="densityType"
   xx:semantics="http://metadata.net/fusion/FUSION#density">
    <xs:attribute name="context" type="xs:anyURI" use="required"/>
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:complexType name="nearestNeighbourType"
xx:semantics="http://metadata.net/fusion/FUSION#nearestNeighbour">
    <xs:attribute name="degree" type="xs:integer" use="required"/>
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:complexType name="porosityType"
xx:semantics="http://metadata.net/fusion/FUSION#porosity">
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:simpleType name="shapeType"
xx:semantics="http://metadata.net/fusion/FUSION#shape">
    <restriction base="string"/>
  </xs:simpleType>
  <xs:complexType name="sizeType"
xx:semantics="http://metadata.net/fusion/FUSION#size">
    <xs:sequence>
      <xs:element name="mean" type="meanType"/>
      <xs:element name="sdev" type="sdevType"/>
    </xs:sequence>
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:complexType name="surfaceAreaType"
xx:semantics="http://metadata.net/fusion/FUSION#surfaceArea">
    <xs:attribute name="context" type="xs:anyURI" use="required"/>
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:complexType name="widthType"
xx:semantics="http://metadata.net/fusion/FUSION#width">
    <xs:sequence>
      <xs:element name="mean" type="meanType"/>
      <xs:element name="sdev" type="sdevType"/>
    </xs:sequence>
  </xs:complexType>
  <xs:complexType name="meanType"
xx:semantics="http://metadata.net/fusion/FUSION#mean">
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
  <xs:complexType name="sdevType"
xx:semantics="http://metadata.net/fusion/FUSION#sdev">
    <xs:attribute name="unit" type="xs:string"/>
  </xs:complexType>
</xs:schema>
```