Realizing the Hydrogen Economy through Semantic Web Technologies

Jane Hunter, *Distributed Systems Technology Centre*
John Drennan, *Centre for Microscopy and Microanalysis, University of Queensland*
Suzanne Little, *University of Queensland*

*The FUSION (**F**uel Cell **U**nderstanding through **S**emantic **I**nferencing, **O**ntologies and **N**anotechnology) project applies, extends, and combines Semantic Web technologies and image analysis techniques to develop a knowledge management system to optimize the design of fuel cells.*

Diminishing reserves of cheap, recoverable crude oil, together with increasing tensions between the Middle East and the West, will likely threaten our access to affordable oil in the future. Alternative fossil fuels such as coal, tar sand, and heavy oil will only worsen global warming.[1] Hydrogen, however, is plentiful and clean, stores energy more effectively than batteries, burns twice as efficiently in a fuel cell as gasoline does in an internal-combustion engine, leaves only water behind, and can power cars. So, many observers view as inevitable the transition from an economy powered by fossil fuels to one based on hydrogen.

But many challenges remain before fuel cells will replace oil and other fossil fuels. Materials scientists must determine how to increase the cells' efficiency, reduce their production costs, determine how they degrade over time, extend their life, and recycle their components. Overcoming these challenges requires a deep understanding of the components' microscopic structure, detailed analytical information about how the various components interact, and the ability to predict how the components will perform under load conditions.

To solve these problems, materials scientists are turning to knowledge management techniques, particularly Semantic Web technologies, to make sense of and assimilate the vast amounts of microstructural, performance, and manufacturing data that they acquire during their research.

At the University of Queensland, the Distributed Systems Technology Centre and the Centre for Microscopy and Microanalysis are collaborating on a project that's applying Semantic Web technologies to optimize fuel cell design. (For the DSTC URL and other URLs pertinent to this article, see the related sidebar.) We call it FUSION (*F*uel Cell *U*nderstanding through *S*emantic *I*nferencing, *O*ntologies and *N*anotechnology).

Applying Semantic Web technologies to fuel cell understanding

Because fuel cell efficiency depends on the fuel cell layers' internal structure and the interfaces between them, an analysis of electron-microscopy images of cross-sectional samples through fuel cells can reveal valuable information. (For more on how fuel cells work, see the "Hydrogen Power" sidebar.) Simple macrolevel information such as the thickness, surface area, roughness, and densities of the cell layers can help us determine gas permeation of the electrode materials. Nanolevel information about the electrode's internal interface structure provides data on exchange reaction efficiency. Figure 1 illustrates the image data obtainable at different magnifications, which must be analyzed and semantically indexed to fully mine the potential knowledge in the images.

Besides the complex multilayered microstructural information that the images reveal, another consideration is the manufacturing conditions and processing parameters used to produce the cell configurations. Fuel cell production involves a highly complex set of procedures that introduce a further set of variables (temperature, time, pressure, and so on) that influence the cell's performance.

Additionally, for each cell configuration, performance data is available in the form of complex graphs that plot voltage against current density (see Figure 2).

The problem's crux is to enable the assimilation of the three highly complex data sets—microstructural,
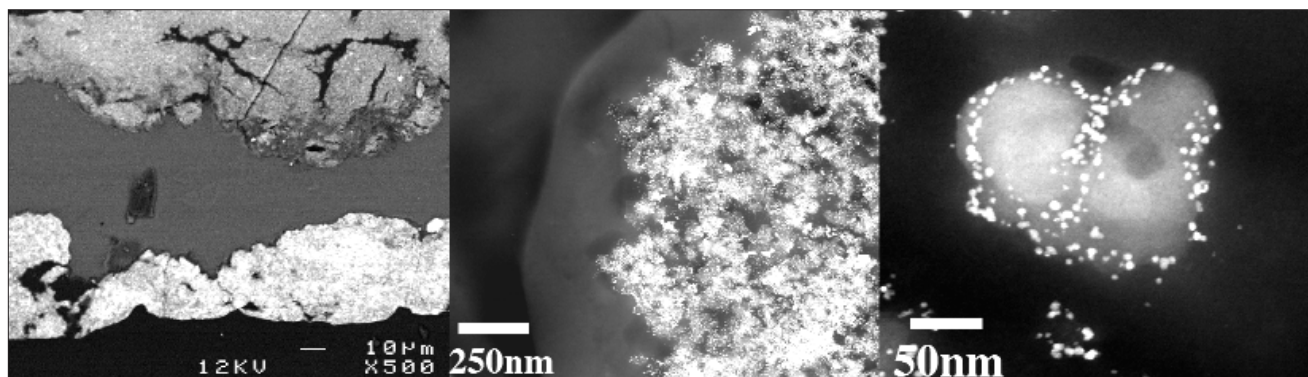
Figure 1. Microscopic images of a fuel cell at three different magnifications.

processing, and cell performance data—so that we can interrogate the information to reveal trends that could help improve fuel cell design and efficiency. However, the amount of information is such that human processing is impossible; the task requires more sophisticated means of data mining. Making sense of the vast amounts of heterogeneous, complex, and multidimensional information that fuel cells present to materials scientists provides a challenging and ideal application for testing Semantic Web technologies' capabilities and the Semantic Web community's claims.[2]

The Fusion project has two main goals. First, we aim to apply and evaluate existing Semantic Web technologies. Second, we aim
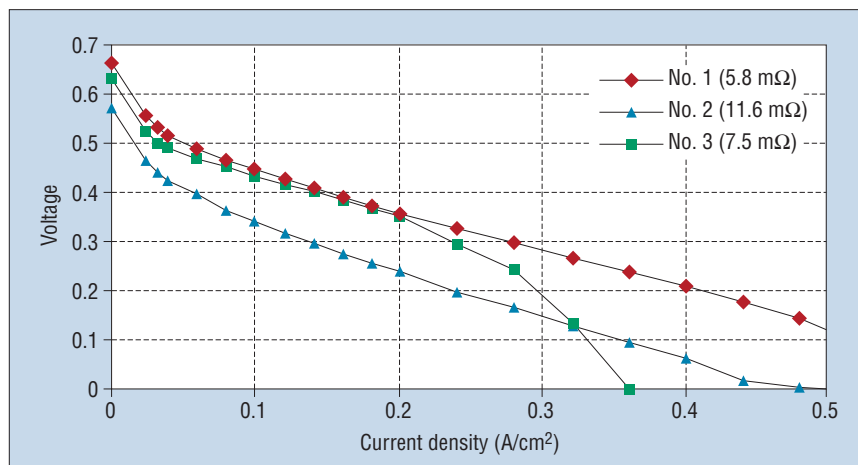

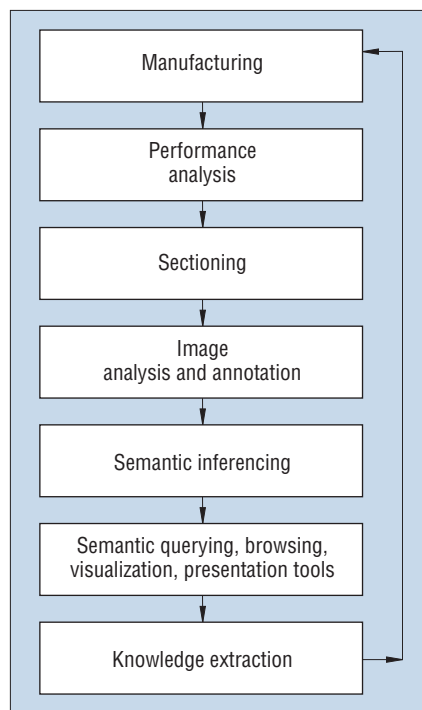**Figure 2. Typical performance data for a fuel cell.**

Figure 3 Fuel cell workflow.

to extend and refine these technologies to develop a knowledge management system for the fuel cell research and industry communities. The project has eight key phases:

1. Analyze fuel cells' workflow or life cycle from the manufacturing stage to the sectioning, analysis, and knowledge extraction stages (see Figure 3).
2. Determine the metadata capture requirements at each step, relevant metadata standards, and the availability of automatic metadata extraction tools.
3. Develop metadata schemas and controlled vocabularies to describe and record the fuel cell data (manufacturing parameters, performance results, and cross-sectional images at different levels of magnification) to document data sources and control metadata quality.
4. Develop metadata input and capture tools based on the manufacturing and analysis workflow (see Figure 3). Then, integrate these tools with automatic image-processing technologies and populate the underlying, networked databases.
5. Develop a fuel cell ontology. Then, merge it with ontologies for multimedia description and microscopy, using a common top-level ontology.
6. Enable the definition and application of inferencing rules that generate high-level semantic descriptions of the fuel cells (defined in the fuel cell ontology)

from automatically extracted low-level features.
7. Develop a semantic querying, browsing, and presentation and visualization interface based on the fuel cell ontology. The interface will use an interactive presentation generator system[3] to assimilate semantically related images and data into coherent multimedia presentations.
8. Integrate collaborative image and video annotation tools. These tools will let domain experts label a small domain-specific image database to improve image querying and to attach new knowledge generated from using the system.

We describe some of these phases in more detail throughout this article.

Applying ontologies and semantic inferencing

Figure 4 illustrates the system components and architecture for the knowledge management system we've developed to manage the manufacturing, performance, and image data captured from fuel cell components.

We use a subset of the ISO/IEC MPEG-7 (Multimedia Content Description)[4] standard to describe the low-level image features. We use the OME (Open Microscopy Environment) standard to capture associated microscope details and settings—essential source information within any scientific context. We developed FUSION metadata schemas to satisfy fuel cell analysts' descriptive requirements. In addition to the XML Schemas we use to define and validate the metadata descriptions, we developed ontologies to define the semantics and relationships between the concepts used in each of these schemas. We describe these ontologies in detail later.

RuleML rules, which domain experts define through a GUI, are used to relate combinations of low-level automatically extracted MPEG-7 features to high-level FUSION concepts or terms. When the system applies these rules to images in the database, it can generate high-level semantic descriptions. The system stores all the generated and validated metadata in a central knowledge repository. A query engine, visualization engine, and knowledge capture (annotation) tools sit on top of the knowledge repository. Together, these components let fuel cell analysts access, interpret, assimilate, and mine the stored data.
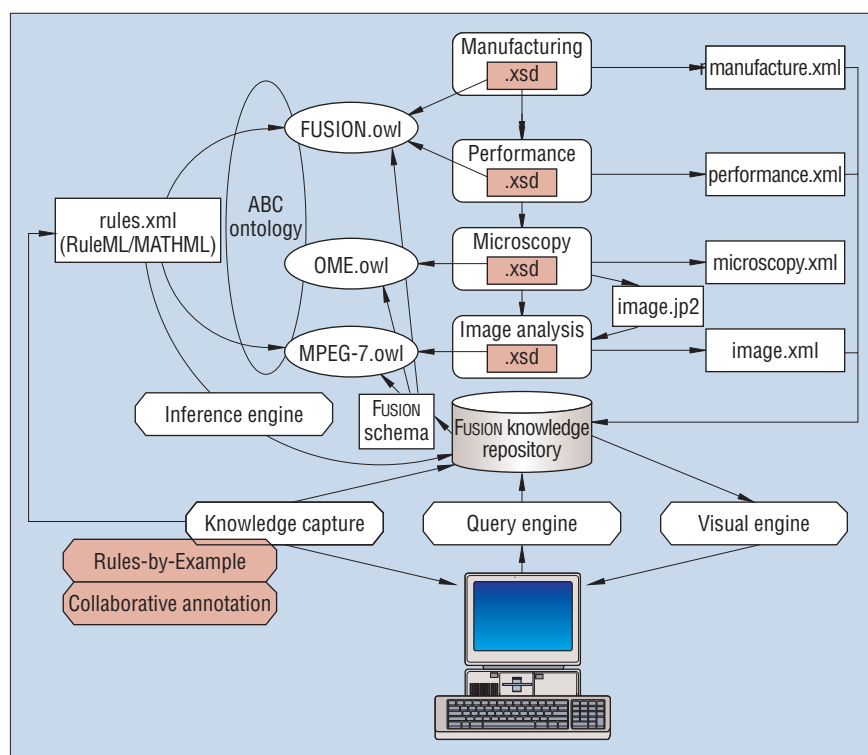

**Figure 4. The FUSION system architecture.**

Hydrogen-powered fuel cells could provide us with a flexible, pollution-free method of producing electricity. Figure A illustrates how hydrogen produces electricity in a fuel cell. Hydrogen passes over an active electrode, which catalyzes the production of hydrogen ions. These ions travel through a membrane to an electrode on the system's air side, reacting with the oxygen to produce water. The process involves the transport of electrons that are effectively harvested to produce energy. These systems work at elevated temperatures and have a predicted efficiency of over 60 percent.[1]

Reference

1. J. Larminie and A. Dicks, *Fuel Cell Systems Explained*, 2nd ed., John Wiley & Sons, 2003.

Figure A. How Proton Exchange Membrane (PEM) fuel cells operate. (This figure is based on an illustration at the Smithsonian Institution Web page http://fuelcells.si.edu/basics.htm.)
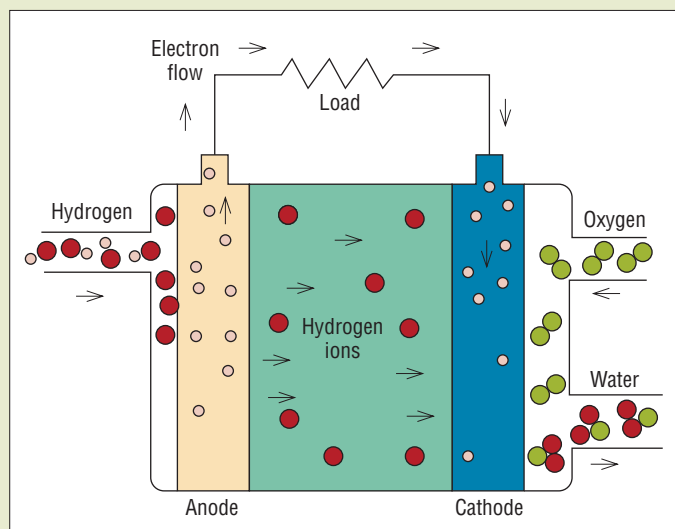
The MPEG-7 ontology

Figure 5 illustrates a subset of the MPEG-7 OWL ontology.[5] We developed this subset to define the semantics of terms used in the MPEG-7 standard for describing low-level visual features. Figure 5 shows only those classes for which "color" is a visual descriptor and color's five subproperties. A complete description of the MPEG-7 ontology is available at http://metadata.net/mpeg7.

The OME ontology

The OME[6] is an open-source collaborative software project that aims to develop a common environment for the analysis, management, and exchange of biological microscopic images. A key component of the OME is an XML-encoded file standard for the storage, analysis, and exchange of image data output from microscopy procedures. A common standard for recording microscope-related metadata is essential for capturing microscopy images' source. This standard includes information such as the microscope's manufacturer, serial number and model, instrument settings, detectors, filters, and light sources. We developed an OME ontology from the OME CoreSemanticTypes. It defines the semantics associated with microscope output data and settings and lets us relate them to the image descriptors defined by the MPEG-7 and FUSION ontologies. Figure 6 illustrates a subset of the OME OWL ontology.
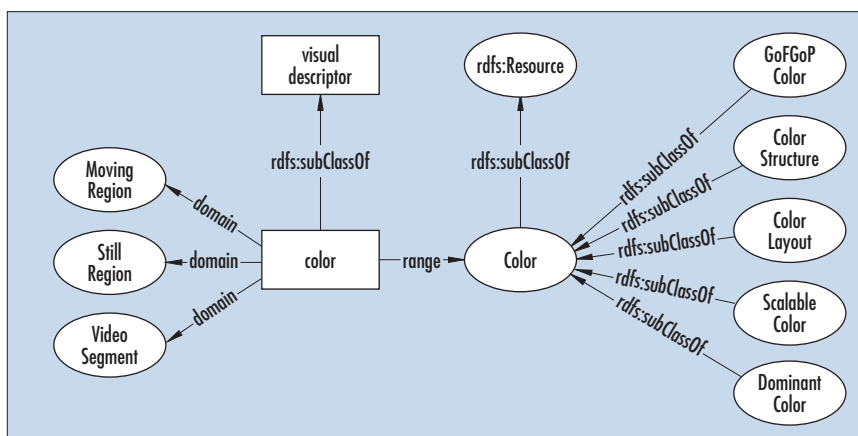

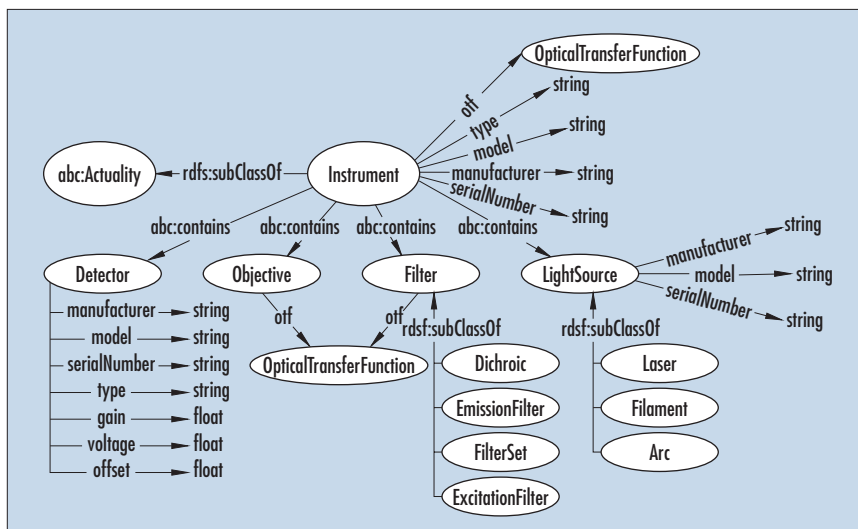**Figure 5. The MPEG-7 visual descriptor "color."**

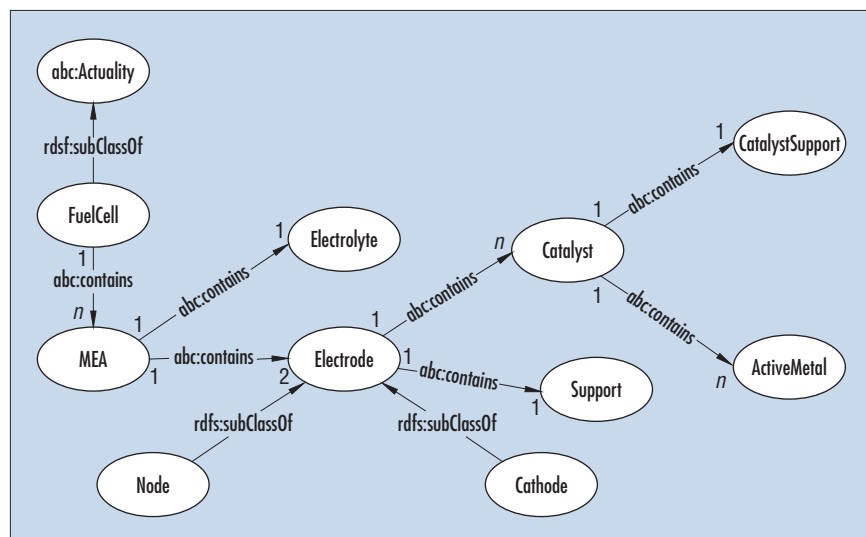**Figure 6. A subset of the OME ontology showing an instrument and its components.**

Figure 7. The key classes and their relationships in the Fusion fuel cell ontology.
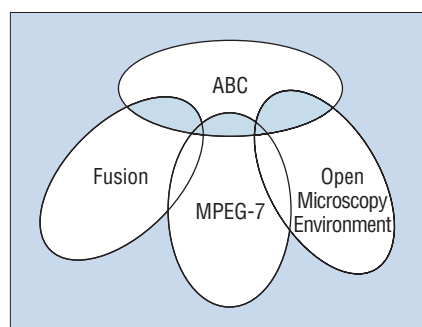

Figure 8. Using ABC to merge ontologies.

The Fusion ontology

Using the approach recommended by Natalya Noy and Deborah McGuiness,[7] we developed the Fusion fuel cell ontology in collaboration with domain experts. We use it to generate the high-level semantic descriptions that these experts use when searching, retrieving, and annotating images. Figure 7 illustrates the key classes or concepts and their relationships, as defined in Fusion.

To enable semantic interoperability between the MPEG-7, OME, and Fusion metadata vocabularies and to define the semantic and inferencing relationships between the terms across these different ontologies, we harmonize or relate them using the top-level or core ABC (A Boring Core) ontology.[8] This ontology, developed within the Harmony project, provides a global, extensible model that expresses the basic concepts that are common across a variety of domains. It also provides the basis for specialization into domain-specific concepts and vocabularies. Other possible upper-level ontologies that we could use for this purpose include SUMO[9] (Suggested Upper Merged Ontology) and Dolce[10] (Descriptive Ontology for Linguistic and Cognitive Engineering). Figure 8 illustrates the harmonization process.

Inferring semantic descriptions of images

In recent years, significant progress has occurred in the automatic recognition of low-level features in images. However, comparatively little progress has occurred in machine generation of high-level semantic descriptions of images. In the Fusion project, we've developed a unique, user-assisted approach to generating ontology-based semantic descriptions of images from low-level automatically extracted features. Our approach lets domain experts define rules (specific to their domain) that map particular combinations of low-level visual features (color, texture, shape, and size) to high-level semantic terms defined in their domain ontology.

Such semantic descriptions enable more sophisticated semantic querying of the images in terms familiar to the user's domain (as Figure 9 shows). These descriptions also ensure that services, agents, and applications on the Web have a greater chance of discovering and exploiting the information and knowledge in the images. The use of ontologies reduces the semantic descriptions' potential subjectivity, and the inference rules' documentation provides a mechanism for capturing and sharing human domain knowledge, in the form of the domain experts' analysis methods.

The Rules-by-Example interface

To define semantic-inferencing rules for images, users must be able to specify values for low-level visual features. Consider, for example, a fuel cell expert labeling electron-microscopy images, to enable the search and retrieval of particular types of fuel cell components:

IF [(*color* is like *this*) AND (*texture* is like *this*) AND (*shape* is like *this*)] THEN (the *object* is a *platinum substrate*)
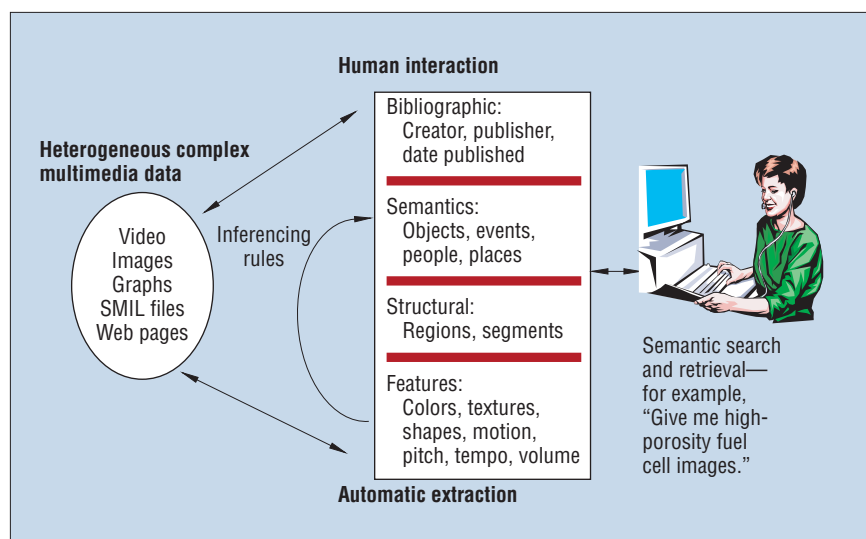

Figure 9. Inferring semantic descriptions from combinations of low-level features.

The simplest, most intuitive way to specify such rules is to provide a *query-by-example* (QBE)-type interface by which users can specify color, texture, shape, and size through visual examples, drawing from images in their collection of interest. So, we developed our *Rules-by-Example* interface—a GUI that provides users with color palettes, texture palettes, predefined shapes, or drawing tools, by which they can define their own semantic-inferencing rules. The RBE is generated from and dependent on the chosen back-end (OWL) ontology (specified during system configuration), which defines and constrains the semantic descriptions that can be generated. We can migrate the system to a different domain simply by choosing a different back-end ontology.

Figure 10 illustrates the initial prototype interface we developed to investigate the RBE's user requirements and usability. The menu options are populated dynamically from the back-end ontologies. For example, a property's possible operators and values are generated from its range definition; for instance, if the type is a string, then relationships such as "less than" are irrelevant. Such dynamic population of the GUI provides user support while maintaining the domain independence that's a strength of this system. Standard logical operators (AND/OR) enable the specification of more advanced logic-based rules.

To specify the values required for atoms in the rule body, the user can select an existing image (or set of images) and select image regions or segments as examples. For instance, in the screen shot in Figure 10, the color specification in this rule states that for the rule to be true (that is, for StillRegion to depict ActiveMetal), the color must be one of the set on the palette. This palette displays all the colors used in the selected segment, and the user can adjust this as he or she sees fit. An additional complexity when searching for visual features is the fuzzy-matching problem, which usually requires similarity rather than exact matches. To support this, the interface provides the ability to set a threshold value for individual image attributes or features.

The RBE interface lets the user quickly develop, apply, and refine highly complex rules without understanding complex low-level MPEG-7 terms or values. Manual creation of such rules would be extremely difficult and time consuming. In addition, the interface lets the users focus the system on the recognition of objects, regions, or features of highest priority or interest. Initial
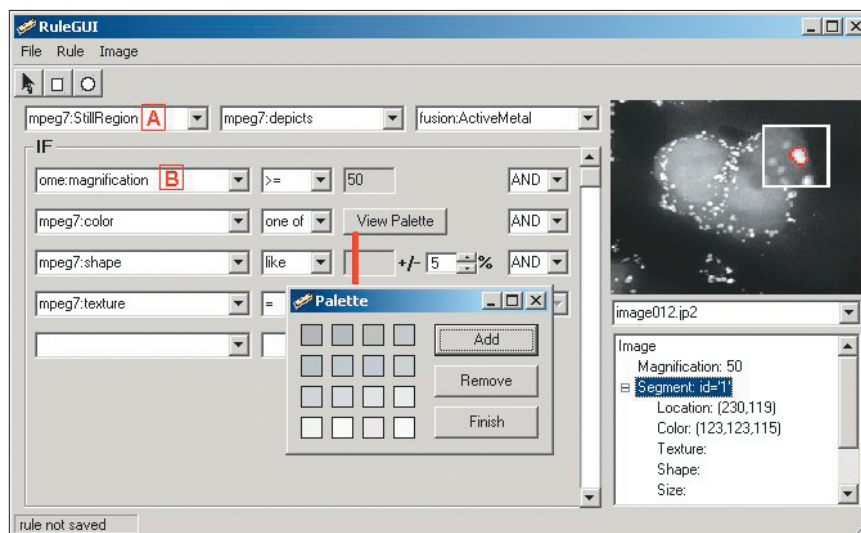

Figure 10. The Rules-by-Example interface prototype.

user feedback has been promising, and we're making further refinements before implementing the final interface design in Java.

Once the user is satisfied with a rule definition, he or she can save it in the RuleML format (possibly augmented with MathML) for processing by an inferencing engine. Possible candidates for the inferencing-engine core include JESS (Java Expert System Shell) and Mandarax, a Java RuleML engine. (Such engines will require extensions if the rule definitions include MathML expressions.) Input to the inferencing engine consists of a set of semantic-inferencing rules, a collection of images, and their automatically extracted low-level MPEG-7 features.

To automatically analyze the fuel cell images, we use MATLAB, a popular, powerful tool widely used by scientists and microscopists. MATLAB can produce a large amount of low-level data about the features and objects in an image—for example, color, texture, shape, area, mean density, standard deviation density, perimeter and length, center coordinates, and integrated density. Although we're currently calling MATLAB analysis methods directly, we plan to employ SOAP, WSDL, and OWL-S to make MATLAB's automatic feature extraction tools available as Web Services. This would allow greater choice and the flexibility to plug in more advanced multimedia extraction and analysis tools and services as they become available.

The data that's output from MATLAB analysis methods is transformed to MPEG-7 descriptions using Perl scripts that we've developed. Given the automatically extracted

MPEG-7 descriptions of images together with the user-specified rules, the inference engine can generate semantic descriptions, determine new semantic relationships (not explicitly defined), and populate and expand the knowledge repository. We can evaluate the results by comparing the automatically generated semantic descriptions with manually attached annotations for the same images.

Given the semantic descriptions generated by applying the rules defined in the RBE GUI, domain experts can perform much more complex and sophisticated queries over the image data. As we mentioned before, they can use terms that are useful and familiar to their domain—for example, "give me all the images depicting high-porosity platinum substrates."

Querying, presentation, and visualization tools

Existing Web-based search and retrieval interfaces, which present the results of a topic or keyword search as a list of URIs, are totally inadequate for e-science applications. New search, browse, and navigation interfaces are needed that can present complex multidimensional, multimedia data in novel ways that let scientists detect trends or patterns that wouldn't be revealed through traditional search interfaces.

An earlier research prototype, developed through a collaboration between the DSTC and the CWI (the Netherlands' National Research Institute for Mathematics and Computer Science),[11] is undergoing refinement for this project. The system dynamically
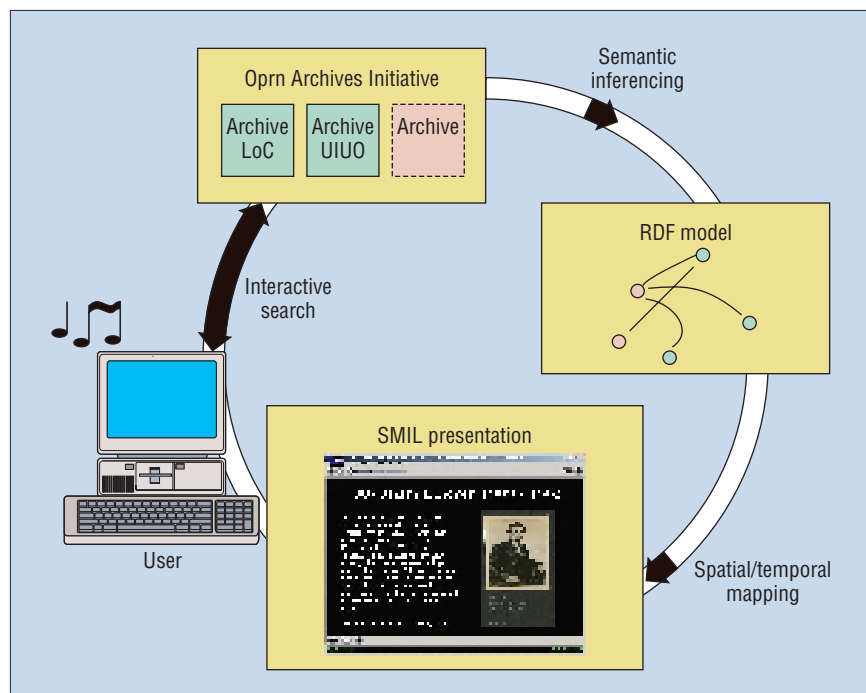
Figure 11. The automatic Synchronized Multimedia Integration Language presentation generator.

generates multimedia presentations in the SMIL 2.0 (Synchronized Multimedia Integration Language) format by aggregating retrieved results sets on the basis of the semantic relationships between them. The system infers the semantic relationships by applying predefined inferencing rules to the associated metadata. CWI's Cuypers system applies constraints such as page size, bandwidth, and user preferences to adapt the presentation components to the user's device

capabilities and platform.[3] In the earlier Open Archives Initiative-based prototype[11] (see Figure 11), we hardwired the mappings from semantic relationships between retrieved digital objects to spatiotemporal relationships in the presentations.

In the FUSION project, we're developing a GUI that lets fuel cell experts interactively map semantic relationships to their preferred spatiotemporal presentation modes. For example, the GUI can display sequences of

images, in increasing order of magnification, as a slide show. Or, it can tile them horizontally or vertically and to the left or right of the corresponding performance data. The domain experts can interactively define and modify their preferred modes of layout and presentation, and their next search's results will reflect their preferences.

For instance, the SMIL visualization tool (the OAI-based prototype we mentioned earlier) will help enable a better understanding of how fuel cells degrade over time. Sequences of micrographic images retrieved from identically manufactured fuel cells after increasing duration of use could be displayed as animations, synchronized in parallel with the display of corresponding performance data.

When SMIL presentations reveal new information or previously unrecognized patterns or trends, users will be able to save and annotate them, for later retrieval, as evidence to support new hypotheses or theories.

Knowledge capture and annotation tools

Assuming that domain experts (in this case, materials scientists) elicit new knowledge by using our system, we'll need tools that can record the newly extracted information or ideas and the evidence or contextual information to support them.

In the DSTC's FilmEd project (http://metadata.net/filmed), we've been developing the prototype Vannotea system (see Figure 12).[12] Using the GrangeNet broadband research network and access grid nodes that support large-scale group-to-group collaboration and high-quality audio-video, Vannotea users can open an MPEG-2 video file or JPEG 2000 image and share tools to collaboratively segment, browse, describe, annotate, and discuss the particular video or image of interest.

Vannotea is an extension of the W3C's Annotea prototype.[13] Whereas Annotea enables the stand-alone annotation of Web pages, Vannotea supports the annotation of audiovisual documents (or segments thereof) in a high-bandwidth collaborative environment. We've extended Annotea's RDF-based annotation schema and Xpointer to support the annotation of audiovisual documents. The user specifies an annotation's context through spatiotemporal extensions to XPointer that enable the location of specific segments, keyframes, or regions in keyframes. This approach also lets us use existing annotation server implementations such
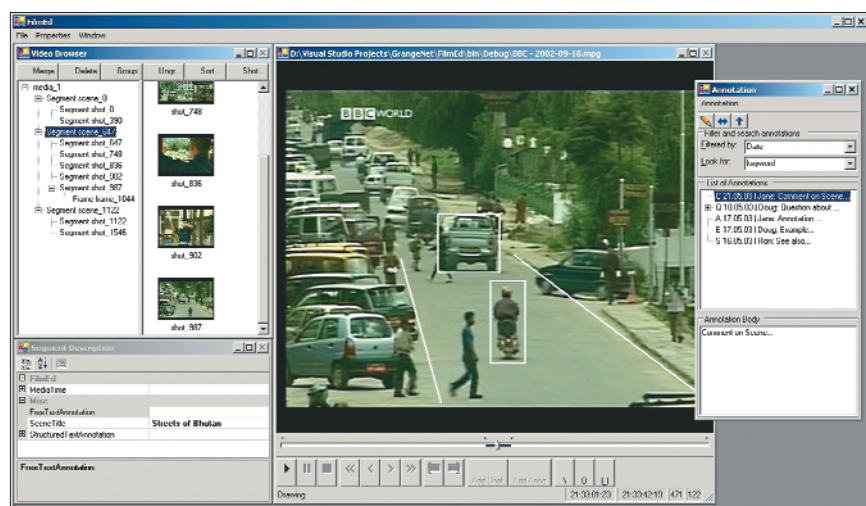

Figure 12. The Vannotea video annotation tool.

as Zope or the W3C Perllib server. These are RDF databases that sit on top of MySQL and provide their own query language, Algae.

We initially designed Vannotea for annotating high-quality MPEG-2 video and JPEG 2000 images. However, enabling the collaborative annotation and discussion of documents of other media types—for example, text, Web pages, audio, and SMIL files—will be relatively trivial. We envision a demand for tools such as Vannotea in both FUSION and other e-science projects.[14] Such tools will enable the collaborative annotation of telemicroscopic video and images as well as the storage and annotation of SMIL presentations generated dynamically as a result of semantic searches.

This article describes how we've been extending, refining, and applying several key Semantic Web technologies to a real-world scientific problem—fuel cell optimization. Although this research is still in progress, we believe that this unique application of Semantic Web technologies to knowledge mining in the materials-engineering domain will prove extremely valuable. Initially, it will help improve the design and performance of solid polymer electrolyte fuel cells. In the longer term, it will help set new paradigms for knowledge management and sharing across a range of scientific and microstructural-engineering applications. ◨
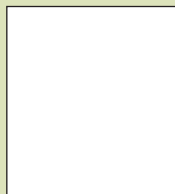
The Authors

Jane Hunter is a Distinguished Research Fellow at the Distributed Systems Technology Cooperative Research Centre. For the past 10 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's also the liaison between MPEG (Moving Pictures Experts Group) and the W3C. Contact her at DSTC, Level 7, GP South, Univ. of Queensland, St Lucia, Queensland 4072, Australia; jane@dstc.edu.au.

John Drennan is the director of the University of Queensland's Centre for Microscopy and Microanalysis. An expert on the microstructural/property relationships in ion-conducting solids, he has published over 90 refereed papers in the field and has made a significant impact on the microstructural design of solid electrolyte materials for fuel cell applications. Contact him at the Centre for Microscopy and Microanalysis, Univ. of Queensland, St. Lucia, Queensland 4072, Australia; j.drennan@mailbox.uq.edu.au.

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. Contact her at the School of Information Technology and Electrical Eng., Univ. of Queensland, St Lucia, Queensland 4072, Australia; s.little@dstc.edu.au.

References

1. J. Rivkin, *The Hydrogen Economy: The Creation of the World Wide Energy Web and the Redistribution of Power of Earth*, Tarcher Putnam, 2002.

2. J. Hendler, "Science and the Semantic Web," *Science*, vol. 299, no. 5606, 24 Jan. 2003, pp. 520–521.

3. J. van Ossenbruggen et al., "Towards Second and Third Generation Web-Based Multimedia," *Proc. 10th Int'l World Wide Web Conf.*, 2001.

4. *ISO/IEC 15938-5 FDIS Information Technology, Multimedia Content Description Interface—Part 5: Multimedia Description Schemes*, Fraunhofer Inst. für Integrierte Publikations und Informationssysteme, July 2001.

5. J. Hunter, "Adding Multimedia to the Semantic Web—Building an MPEG-7 Ontology," *Proc. Int'l Semantic Web Working Symp.* (SWWS), 2001.

6. J.R. Swedlow et al., "Informatics and Quantitative Analysis in Biological Imaging," *Science*, vol. 300, no. 5616, 4 Apr. 2003, pp. 100–102.

7. N. Noy and D. McGuinness, *Ontology Development 101: A Guide to Creating Your First Ontology*, tech. report SMI-2001-0880, Stanford Medical Informatics, Stanford Univ., 2001.

8. C. Lagoze and J. Hunter, "The ABC Ontology and Model (v3.0)," *J. Digital Information*, vol. 2, no. 2, 2001.

9. I. Niles and A. Pease, "Towards a Standard Upper Ontology," *Proc. 2nd Int'l Conf. Formal Ontology in Information Systems* (FOIS 2001), 2001.

10. A. Gangemi et al., "Sweetening Ontologies with DOLCE," *Proc. 13th Int'l Conf. Knowledge Eng. and Knowledge Management* (EKAW 2002), LNCS 2473, Springer-Verlag, 2002, pp. 166–181.

11. S. Little, J. Geurts, and J. Hunter, "The Dynamic Generation of Intelligent Multimedia Presentations through Semantic Inferencing," *Proc. 6th European Conf. Digital Libraries* (ECDL 2002), LNCS 2458, Springer-Verlag, 2002, pp. 158–175.

12. R. Schroeter, J. Hunter, and D. Kosovic, "Vannotea—A Collaborative Video Indexing, Annotation and Discussion System for Broadband Networks," to be published in *Proc. 10th Int'l Multimedia Modelling Conf.* (MMM 2004), IEEE CS Press, 2004.

13. J. Kahan et al. "Annotea: An Open RDF Infrastructure for Shared Web Annotations," *Proc. 10th Int'l World Wide Web Conf.* (WWW10), Int'l World Wide Web Conf. Committee, 2001; www10.org/cdrom/papers/frame.html.

14. *Technology Theme: Collaborative Visualisation*, Grid Outreach Programme, 2003, www.gridoutreach.org.uk/themes/cv_intro.htm.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.