IML: An Image Markup Language

Lober WB, Trigg LJ, Bliss D, Brinkley JM

Division of Biomedical and Health Informatics, Department of Medical Education
University of Washington School of Medicine, Seattle WA 98105

Image Markup Language is an extensible markup language (XML) schema used to describe both image metadata and annotations. It describes both data pertaining to an entire image, and data that are tied to specific regions or features of the image. Developed for a specific domain in Medical Education, this paper describes extensions to take advantage of the Dublin Core metadata standard, and of an XML schema for vector graphics representation. We have developed a prototype system of open source tools implementing an authoring system, a client system, and an image annotation database which can be queried though the Web.

IMAGES AND IMAGE ANNOTATION

Images offer a valuable means of conveying information, and are a vital component of both medical education and clinical care. Annotations added to these images, whether as grease pencil notations or as structured metadata, greatly aid in their interpretation and in sharing information about interpretations. Traditional methods of commenting on a medical image have included both labeling the image directly and writing or dictating an impression or report about the image.

In the past, these annotation methods have been both effective and expedient; however, they have limitations similar to those of paper medical records. Briefly, traditional methods of image annotation do not facilitate dynamic or customized display of annotations, electronic distribution and duplication of reports, or retrieval based on attributes of the annotation..

We feel that the ability to create highly structured, feature specific, and searchable annotations will facilitate even greater utility of medical images both in teaching and clinical practice.

GLOBAL ANNOTATIONS

Global metadata is data about the entire image, or data about the creation or history of the image annotation file itself. This includes information such as the title of the image, its source, the author of the annotations, the date the annotations were done, the version of the image annotation format used, etc.

Many previous efforts to incorporate metadata into images have focused on global metadata. These include applications of Dublin core metadata[1] to image libraries[2,3] and the DICOM[4] Structured Reporting Standard (Supplement 23). An emerging standards effort called MPEG 7[5] ("Multimedia Content Description Interface") promises to offer a general scheme for annotating images, video, audio, and text, but it is still in the early stages, with only 11 image demonstration projects as of the July 2000 meeting.

REGIONAL ANNOTATIONS

Regional metadata is data that describes, illustrates, or in some way refers to an individual feature or specific region within an image.

There are several examples of image annotation systems that allow annotation of images using regional metadata. These examples exist both inside and outside of the medical domain, but there is no widely used standard which allows structured encoding of both global metadata pertaining to the entire image, and of regional or feature specific metadata.

The Digital Anatomist (DA) atlases comprise a web-based annotated image collection[6] depicting anatomic structures. The atlases are used in undergraduate medical education courses at the University of Washington, and are accessed by other users throughout the world. The annotation model is based on a “gross anatomy” metaphor with regional features surrounded by outlines, and a “pin” stuck in the image connected by a string to a label. Images may be linked together through these labels for navigation.

I ^{2}Cnet[7] offers a Web-based autonomous repository of medical images and related information which permits medical specialists to interact with images in the repository and with each others’ annotations by adding comments, annotations, illustrations, etc. Annotations in that system are described as graphical objects specifying a region of interest in an image, which is linked to notation text. These notation texts may be grouped as overlays, which include an annotation message, and pointers that in turn are associated with an object containing authoring and similar metadata. The application supports authoring, storage, searching, and communicating about the images, and the annotations are stored as ASCII text.

Inote[8] is an image annotation tool developed by the Institute for Advanced Technology in the Humanities at the University of Virginia. The authoring tool is a standalone application written in Java, which allows the user to attach textual annotations to various regions in an image to and to store those annotations in a separate XML file. It can create “details” of several types such as rectangle, polygon, circle, and point. One or more annotations may be attached to each of these details. The details are organized into overlays.

PAIS[9] is an image annotation system developed by the first author that consists of two components called Author and Client. As with Inote, the Author component is a Java application, and the Client component is a Java applet. To author an image collection, one copies a series of GIF or JPEG images into a directory that is part of a web site. Using the authoring tool, annotations and links are interactively and dynamically “drawn” on the image. These annotations are stored as XML files which are associated with the image. To view the annotated images, the client software accesses the web server and downloads both an annotation and its associated image. The annotated image is rendered dynamically on the client. Navigation is through linked labels, as with the Digital Anatomist model.

BlueNOTES[10], VisualMed[11], and OpenLab Core[12] are examples of commercial image annotation products. They use proprietary formats to store their data. As far as we were able to ascertain, the annotation data remains editable, as it might with PowerPoint [13]), but it is not independently searchable or retrievable.

When medical educators and practitioners look at images, they see both global and local characteristics. Both education and clinical practice would be enhanced by a standardized way of commenting on these image features.

GOALS

An image annotation system useful across domains should meet the following goals: 1) annotation metadata should be structured, and searchable; 2) the representation of annotation data should be concise and flexible; 3) the scheme must encompass several divergent models of annotation currently in use, such as drawn outlines, overlaid layers, and superimposed measurements; 4) both global and feature specific metadata should be represented; 5) existing metadata standards should be used when possible; 6) existing software standards and tools should be used when possible.

In this paper, we present both a generalized scheme for image markup based on the extensible markup language (XML)[14], and several open source tools we have built to aid in creating and displaying annotated images. These include authoring tools, a client, and an image annotation database. In addition, we have used XML schema standards to document IML in a way that allows for both unambiguous interpretation of the format and for automated validation of a specific set of image markup metadata against an IML 2 template.

ANNOTATION SYSTEM OVERVIEW

Borrowing from the design of the Digital Anatomist, IML is based on the principle that image data and annotation data should be stored separately. An IML-compliant authoring tool captures the annotation information and stores the data in a separate XML file, rather than writing annotation graphics and text directly “on top” of an image (as one might see in a system using layers of the type created by a graphics program such as Adobe Photoshop).

When an annotated image is requested by client software, the annotation file is downloaded and parsed. The image referred to by the annotation file is retrieved, and then the annotation data are dynamically “drawn” on top of the image to fulfill the client’s request. This rendering process may take place either on the server or the client side of the transaction. A significant advantage of this approach over that of layered images is that both server or client may exercise a great deal of discretion over which annotations are supplied, and how they are rendered, based on what it knows about the end-user and their intent in using the image. Annotations customized to the level of study and interest of the audience can be “automagically” displayed. For example, different sets of annotations may be appropriate for students of anatomy in 7th grade, in medical school, and during a cardiothoracic surgery fellowship.

EVOLUTION OF IML

IML began as a replacement for an existing image markup language format called Frame Format. This format, used by the Digital Anatomist system, is a LISP-like, text-based format for expressing regional and limited global annotation data. In order to implement functionality similar to that of the Digital anatomist software, the initial implemented version of IML (IML 1.1) followed this format closely. We developed IML 1.2 in response to data representation issues uncovered as we began to work with relational database storage of IML. We have subsequently extended the format in the IML 2 specification to take advantage of additional features offered by XML, and to more closely comply with existing and emerging standards such as Dublin Core Metadata[15] and Scalable Vector Graphics[16].

CONCEPTUAL MODEL OF ANNOTATION

The present IML schema is detailed in Figure 1. The root, or highest-level element in the schema is called AnnotatedImage. This element must have a reference

(URL) to an image associated with it, as well as a version attribute that identifies the specific IML version. There is also a required ID element associated with the root element; however, the authoring tool may generate this automatically. There are a number of optional elements under this AnnotatedImage element to specify characteristics such as Dublin Core metadata, an edit history for the annotation file, or regional, layered, or measurement-based annotations.

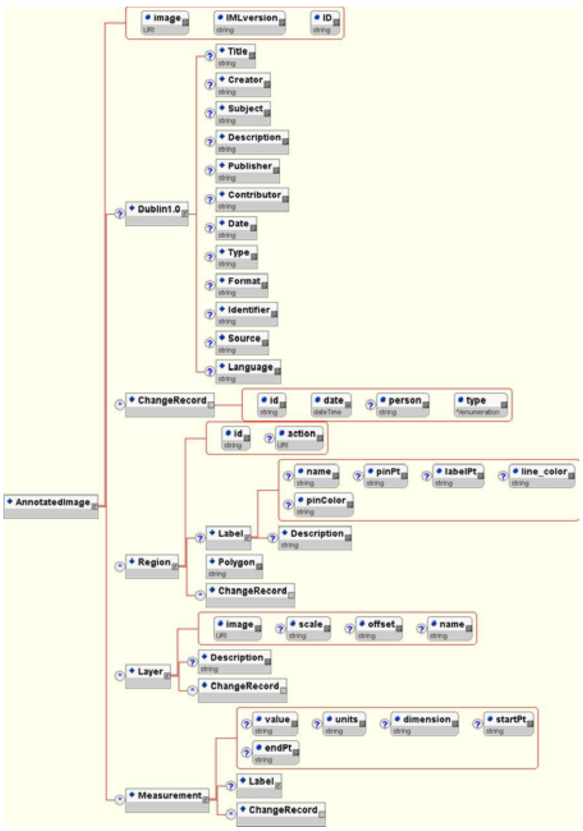

Figure 1. IML 2.1 Schema

IML 2 implements the ten text elements of the Dublin core 1.0 metadata standard[17]. The Dublin element is optional, as are all of its constituent elements.

The edit history is implemented through an element called ChangeRecord, which contains required attributes for date, type of change, and an internally maintained ID. It also contains an optional attribute to identify the person making the change, though an individual authoring tool may opt to require this attribute.

The AnnotatedImage element may contain any number of annotations. There are three types of annotations: outlines drawn around features of the image, other images ("layers") overlaid on part or all of the image, and one or two-dimensional measurements made between features of the image.

Region annotations contain an outline defining an area of interest in the image. That outline is automatically assigned a required ID, and may have an action associated with it (in the form of a URL which client software might access if a user selects the outlined region). In addition, an outline may have a label, which is a construct that may include a name, a description, a "pin" point that can have a location and a color, and a colored line connecting the pin to the label for the region. A region may also have one or more ChangeRecords associated with it to provide a revision history specific to that element.

Layer annotations contain a reference to a required image with optional scaling, offset, and name attributes. Additionally, a layer may have an optional description and series of ChangeRecords.

Measurement annotations are patterned after ultrasound annotations, with the ability to identify points or lines and record an associated measurement value and units. In addition, measurements may contain a standard Label and a series of ChangeRecords.

EXISTING SOFTWARE TOOLS

Three software tools have been developed to date for the IML image annotation system. These tools are all still in the active development stage, and are available under the open source GPL[18] license. Presently, the software tools use the IML 1.2 format, and are designed to demonstrate the interaction of authoring, client, and database applications for image annotation.

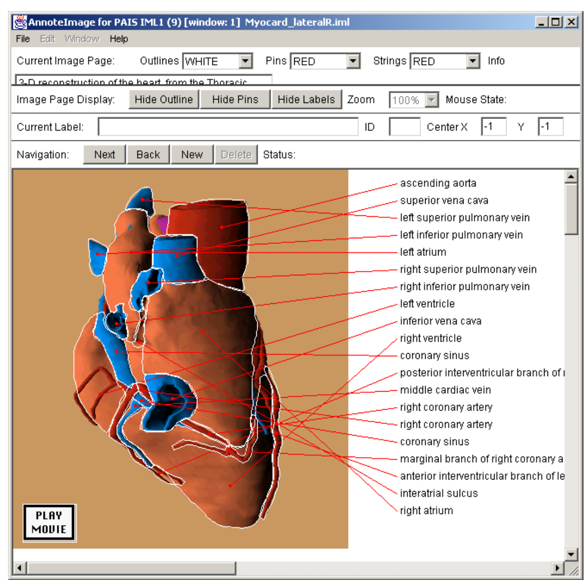

Figure 2. Authoring Tool

The **authoring tool** (Figure 2) allows graphic annotations to be drawn on top of an image, capturing both the points outlining the region of interest and the text metadata entered by the author. It is written as a Java[19] 1.1 application and must be installed locally on a user's system.

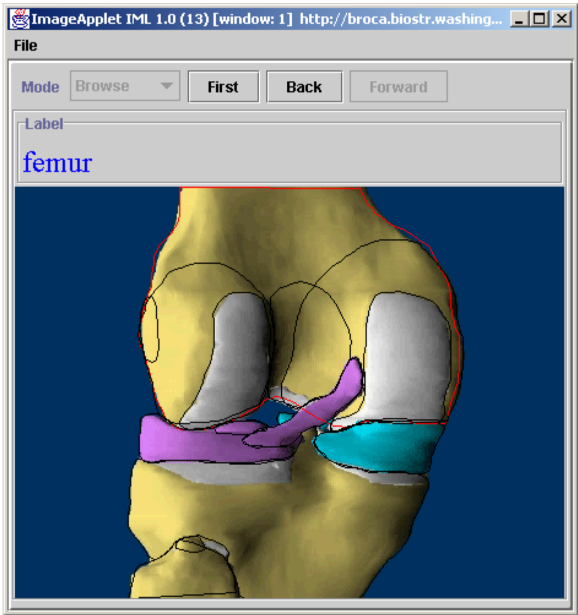

Figure 3. Client Tool

The **client tool** (Figure 3) is a Java 2 applet, which is automatically downloaded by a user's browser when the annotated image collection is first accessed. The image is displayed on the client's machine, and the annotations are drawn on top of that image by the applet. This allows for rollover hot spots in the image and other client-side user interface features. Annotation text is displayed as the mouse moves over a feature.

The **image annotation database** comprises a repository of image annotations. Images are not stored in the database; however, each annotation contains a Web link to its related image. The client tool accesses the database to retrieve new annotations as it follows navigation links, but the real benefit of the image annotation database is the ability to search the database fields to find annotations which meet certain criteria. These queries might include, for example, annotated images containing at least one regional annotation labeled "Left main coronary artery" or "Reed-Sternberg cell." We have tested database searching but have not yet implemented it in the present client.

We implemented the database using MS SQL Server 2000[20] It may be queried through the use of standard Web server protocols, and can return its results either as XML or HTML documents.

The database presently contains image annotations from two complete Digital Anatomist atlases, the Thorax and the Knee. There are 442 annotated images, with a total of 7351 regional annotations. A sample query done though a direct SQL interface, which returned all regions containing the word "coronary", yielded 54 regional annotations from 24 distinct annotated images in 160 ms. The same query, framed as an XML query through a web browser, and returning a series of IML 1.2 documents and ready to be processed by an IML client, took 3.0 seconds. A sample IML 1.2 document is shown in Figure 4.

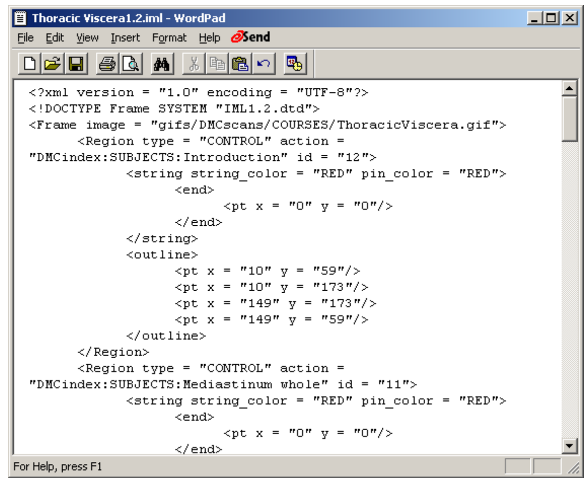


```
<?xml version = "1.0" encoding = "UTF-8"?>
<!DOCTYPE Frame SYSTEM "IML1.2.dtd">
<Frame image = "gifs/DMCscans/COURSES/ThoracicViscera.gif">
      <Region type = "CONTROL" action =
"DMCindex:SUBJECTS:Introduction" id = "12">
            <string string_color = "RED" pin_color = "RED">
                  <end>
                        <pt x = "0" y = "0"/>
                  </end>
            </string>
            <outline>
                  <pt x = "10" y = "59"/>
                  <pt x = "10" y = "173"/>
                  <pt x = "149" y = "173"/>
                  <pt x = "149" y = "59"/>
            </outline>
      </Region>
      <Region type = "CONTROL" action =
"DMCindex:SUBJECTS:Mediastinum whole" id = "11">
            <string string_color = "RED" pin_color = "RED">
                  <end>
                        <pt x = "0" y = "0"/>
                  </end>
```

Figure 4. Example IML 1.2 document

DISCUSSION

Images are a rich source of educational and clinical information. IML offers a way to extend the accessibility of this source of data and to structure it in a searchable format. IML has now undergone two substantial revisions. It began as a format that duplicated the existing Digital Anatomist image annotation data representation scheme and integrated emerging XML standards and development tools in creating annotation software[21].

It has evolved to include a significantly expanded conceptual model for outline style annotations, other annotation types such as layers and measurements, a standardized model of metadata based on the Dublin Core standard, features compatible with an XML-capable database, and new XML features such as scalable vector graphics (SVG)[22].

We continue to be interested in the expanding scope of XML, and are particularly interested in incorporating more elements of SVG to allow richer specification of graphic element drawing styles.

There are some limitations in the present representation scheme. We have taken an annotation scheme designed for images from a specific domain and generalized it. We will surely encounter many authors who have used a different style of annotation, and are dependent on a particular font, a search through data tables in a proprietary format, or some other feature that our system lacks.

In general, we have tried to represent content relationships rather than display format, though we have included what we felt were essential display elements. One solution may be to separate these elements, in much the same fashion that the XSL (extensible stylesheet language[23]) has developed into separate transformation (XSLT) and display (XSL-FO) languages. We anticipate that closer adherence to SVG may facilitate a similar strategy for IML.

We have adopted the form of the Dublin Core and SVG standards without implementing all of the details. While our Dublin Core data representation is consistent with the standard, we intend to take advantage of that group's publication of XML representations, using RDF (Resource Definition Format) of their data. Using RDF for within IML will result in easier external exposure of the Dublin Metadata to metadata search engines.

Finally, performance needs to be improved for XPATH queries (which are queries conveniently represented as XML element specifications in the query URL) both in terms of function and speed. The current version of XML query support (Microsoft XML for SQL Server) does not support sub-string queries. This is a serious problem that we believe will be rectified in a future release. We have identified a workaround that does allow XML to be returned from substring queries, but turns those queries into a several step process. XML queries presently run about 20 times slower than their native SQL counterparts. We have found XPATH performance acceptable, but are investigating ways to improve it. As we discussed earlier, the performance of URL-based native SQL queries that return XML documents is excellent.

Conclusions

Our experience with the prototype system thus far demonstrates its feasibility. The system allows images to be annotated, and allows the annotations to be stored in a database. There they can be queried by the image client or by another program. The latter approach is more interesting, as it permits queries of the annotation database by an intelligent agent capable of highly tailored retrieval of image information based on metadata content searches.

Our prototype system demonstrates that structured image annotation is a viable and practical means of storing image metadata with currently available technologies. Our code is released as open source, and we hope it will prove useful to other investigators.

Further information:
Bill Lober, MD lober@u.washington.edu
http://depts.washington.edu/cirg/iml

The work supported under NLM extramural fellowship F38 LM 00086.

---

[1]Dublin Core Metadata Element Set, Version 1.0: Reference Description, Sept, 1998. (http://dublincore.org/documents/1998/09/dces/)
[2]CONTENT database, Center for Information Systems Optimization, University of Washington (http://content.engr.washington.edu/)
[3]University of Washington Digital Collections, (http://content.lib.washington.edu/index.html)
[4]Supplement 23, Digital Imaging and Communications in Medicine (DICOM), October, 1999. (http://medical.nema.org/Dicom/supps/sup23_lb.pdf)
[5](http://www.darmstadt.gmd.de/mobile/MPEG7)
[6]J. F. Brinkley, et. al, "The Digital Anatomist information system and its use in the generation and delivery of Web-based anatomy atlases," *Comp and Biomed Rsrch*, vol. 30, pp. 472-503, 1997.
[7]Chronaki,C.E.;Zabulis,X.;Orphanoudakis,C. I[2]Cnet medical image annotation service. Med Inform(Lond).1997Oct-Dec;22(4):337-47.
[8]Institute for Advanced Technologies in the Humanities, http://jefferson.village.virginia.edu/inote/
[9]PAIS ref from CV.
[10] http://www.blueshoe.com/bluenotes/
[11] http://www.vmedsys.com/english/html/clinical/frame.html
[12] http://www.meyerinst.com/html/openlab/z_core.html
[13] http://microsoft.com
[14] Extensible Markup Language (XML) 1.0. *W3C Recommendation* 10-February-1998. (http://www.w3.org/TR/1998/REC-xml-19980210.html)
[15] Dublin
[16] Ferraiolo, J (Ed.), Scalable Vector Graphics (SVG) 1.0 Specification, *W3C Candidate Recommendation (http://www.w3.org/TR/SVG/),* 02 November 2000.
[17] Dublin
[18] GNU General Public License (http://fsf.org/copyleft/gpl.html)
[19] http://java.sun.com
[20] http://microsoft.com/products
[21] Lober WB, Brinkley, JF, A Portable Image Annotation Tool for Web-based Anatomy Atlases, *Proc AMIA Symp 1999*;:1108, November 1999
[22] (SVG) 1.0 Specification
[23] http://www.w3.org/Style/XSL/