The 9th International Symposium on Virtual Reality, Archaeology and Cultural Heritage
VAST (2008)
M. Ashley, S. Hermon, A. Proenca, and K. Rodriguez-Echavarria (Editors)

The Arrigo Showcase Reloaded – towards a sustainable link between 3D and semantics

S. Havemann[1], V. Settgast[1], R. Berndt[1], O. Eide[2], D.W. Fellner[1,3]

[1]Institute of Computer Graphics and Knowledge Visualization (CGV), TU Graz, Austria, www.cgv.tugraz.at
[2]Unit for Digital Documentation, University of Oslo, Norway, www.edd.uio.no
[3]TU Darmstadt & Fraunhofer IGD, Germany

Abstract

It is still a big technical problem to establish a relation between a shape and its meaning in a sustainable way. We present a solution with a markup method that allows to label parts of a 3D object in a similar way to labeling parts of a hypertext. A 3D-markup can serve both as hyperlink and as link anchor, which is the key to bi-directional linking between 3D objects and web documents.
Our focus is on a sustainable 3D software infrastructure for application scenarios ranging from e-mail and internet over authoring and browsing semantic networks to interactive museum presentations. We demonstrate the workflow and the effectiveness of our tools by re-doing the Arrigo 3D showcase. We are working towards a "best practice" example for information modeling in cultural heritage.

Categories and Subject Descriptors (according to ACM CCS): I.3.4 [Graphics Utilities]: Graphics editors I.3.6 [Methodology and Techniques]: Languages,Standards

1. Introduction

Cultural Heritage is inherently 3-dimensional. On the one hand the remains from our past form a VAST collection of 3D objects. On the other hand, every item of cultural value is embedded in time-varying social, political, geographical, traditional, and personal contexts. Each artefact has its own history, but also belongs to an object class, which develops over time as well. In fact, each cultural artefact belongs to many classes: Shape, material, manufacturing technology, possessor, all provide equally valid views on an object.

This shows that when attempting to embed 3D objects into their semantic context, things quickly become extremely complicated. As a consequence, Cultural Heritage is a very rich and interesting, but also very demanding field from the information modeling point of view.

1.1. The Arrigo Showcase

The motivation for the work described in this paper comes from a beautiful museum exhibit, the *Arrigo Showcase*. It was created by the group of Paolo Cignoni and Roberto Scopigno of the Visual Computing group of CNR Pisa, Italy, and it represents one of the rare cases where *interactive* 3D technology was successfully used in a museum exhibition. Using a mouse, visitors could interactively explore detailed 3D models of all 12 statues of the Arrigo ensemble. The *VirtualInspector* technology developed in Pisa [CPCS08, CPCS06] maintains interactivity even with close-up views of massive 3D datasets with millions of vertices. It was deployed for the public display of the *David* statue scanned in the *Digital Michelangelo* project by Marc Levoy from Stanford. The display was located right next to the real statue, which is 5 meters tall. Visitors could discover details in 3D that were barely visible on the real statue. This way the 3D visualization does not replace, but enhances the appreciation and understanding of the real artefact.

The interesting part of the **Arrigo Showcase** is that it encourages to explore not only the 3D dataset, but also its semantic embedding. Some parts of the statues, like the mutilated hands or the hollow back, lead to obvious questions. With information icons on the 3D model visitors could access an information pane with a short textual explanation

Figure 1: *The Arrigo 3D presentation. Beautifully made, but a closed solution. All texts shown are just bitmap images that are inconvenient to change.*

(Fig. 1). In that way semantic information is made accessible via 3D.

1.2. Drawbacks of the Arrigo Approach

In 3D software there is always a tradeoff between **fine-tuning** and **generality**. Optimal results are possible only by adjusting all components: input devices, display hardware, rendering algorithms, graphical user interface, content quality, graphic design, interaction metaphors, and look and feel.

As beautiful as it may be, the Arrigo Showcase suffers from fundamental limitations with respect to the following requirements:

- **Sustainability**: Computer hardware can become outdated, as well as display algorithms and file formats. But cultural objects keep their semantics and should be available independently of the technical development status.
- **Extensibility**: Adding new information icons, including additional views on the subject, alternative interpretations. Linking into the showcase, and links from the showcase to information outside. Adding such links much later.
- **Changeability**: To replace a system component is possible only with clear, well documented specifications of the interfaces between the modules. Modules should comply to standards rather than to ad-hoc design decisions.
- **Versatility**: Ideally, data and systems cover many application scenarios, e.g., using the same data in a museum exhibition and for scientific research. Data should be internet accessible, and should be reliably citeable in a scientific publication or in an e-mail.
- **Openness**: Every piece of information leads to other pieces; information can be hierarchically grouped or arranged in a graph (*hypertext*); each of the information icons is in fact part of a larger *semantic network*.

1.3. Requirements Analyis through Use Cases

It is not an easy task to overcome the limitations of the Arrigo system. An analysis was carried out to identify a number of use cases that a more general system would have to cover.

- **Scenario: Web browser with coupled 3D browser**
  Following a hyperlink that points to a 3D object should bring up a 3D browser displaying this object. 3D objects may contain embedded hyperlinks (information marks) leading to other 3D objects, or to web documents. Such software shall be easy to use and ideally be free.
- **Scenario: Museum Presentation**
  The 3D browser and a web browser are seamlessly coupled together in a fullscreen application. The 3D layout and the web page are beautifully tuned to match. The web browser can show multimedia content (images, video, Flash) triggered by links on the 3D model. The 3D view (camera, shading) can be modified using hyperlinks.
- **Scenario: Scholarly Research**
  The 3D browser and the web browser are connected to a semantic data base via a combined frontend. It gives access to a semantic network of entities (generalized "multimedia" documents). Relations between entities can be browsed, and new relations can be established.
- **Scenario: Authoring semantic networks**
  There must be an easy way to create 3D attachments. An authoring application allows importing raw datasets, objects can be spatially arranged via drag-and-drop, and relations between 3D objects can be defined. Ideally, scenes can also be provided with behaviours (*e.g. drag-and-drop behaviour*).
- **Scenario: Development of 3D-Applications for CH**
  CH requires a plethora of specialized applications to cover all aspects from excavation management over GIS databases to artefact collections and museum archival software. They would all greatly benefit from functionality like in the Arrigo showcase.

1.4. Information Model: Design Decisions

The first step is to define the information infrastructure. It must be simple, reliable, versatile, extensible, and standard compliant. Our basic decisions and definitions are:

- **Integration into semantic networks**
  Cultural information is modeled today as a network of en-

OpenSG/XML/GML	OpenSG engine with GML scripting, maintaining XML attachments to 3D scene graph nodes
ActiveEpoch	*OpenSG/XML/GML* as ActiveX control that consists only of an OpenGL window
EpochViewer	Thin layer to wrap ActiveEpoch into an application that resolves links to 3D object parts
EBHO	Internet browser extension, redirect .dae to *EpochViewer*, including relative links (#Head)
AuthoringTool	Markup application combining *ActiveEpoch* control with an *InternetExplorer* control
PresentationTool	Integrated 3D/web browser with simplified interaction for museum exhibitions

Figure 2: *The components of our software infrastructure. The ActiveEpoch control encapsulates the 3D functionality.*

tities connected by semantic relations. 3D should be part of them.

- **XML for structure and semantics**
  These semantic networks are encoded in XML. Many XML technologies exist to manipulate XML encoded knowledge.
- **All references are URLs**
  An entity is nothing but a *uniform resource locator*, a URL. A relationship (e.g., a RDF-triplet) links two entities (two URLs) together by a predicate (another URL).
- **Bi-directional linking HTML ↔ 3D**
  A 3D object is uniquely identified by a URL pointing to a model file with the appropriate MIME type (.dae):
  http://www.CH-models.org/statues/arrigo5.dae
- **3D-Links on sub-object level (part level)**
  A part of a 3D object is uniquely identified by a relative link, similar to linking to an anchor in an HTML document:
  http://www.CH-models.org/statues/arrigo5.dae#Head
- **3D annotations in XML: Collada**
  .dae-links refer to Collada XML documents which have a *library part* and a *scene part*. 3D objects in the library may have information attached. Our attachments consist of a *region in space* (link anchor) and a *URL (link target)*.
- **3D objects in binary formats**
  The actual 3D datasets are not encoded in XML but in native 3D formats (obj, ply etc). The 3D objects in the library part of a Collada file only refer to them via URL.
- **3D scenes in XML: Collada**
  The scene part of a Collada file defines the spatial arrangement of the objects defined in the scene part. It consists of hierarchical transformations, a *scene graph*.
- **Multiple interpretations simultaneously**
  Different .dae-Files can provide the same 3D model with different sets of annotations (interpretations). One model part can have multiple hyperlinks attached, different views or aspects of that part.
- **Mechanisms and behaviour are possible**
  The 3D scripting language GML [Hav05] has an XML-based encoding (.xgml). GML code can in fact be inserted directly into Collada to provide objects or whole sub-scenes with behaviour.

The next step is identifying a set of functional units that can be combined in different ways to realize the various application scenarios. The resulting software modules are shown in Fig. 2, they will be further explained shortly.

This paper shall demonstrate how this infrastructure works by re-doing the Arrigo Showcase. The purpose of this exercise was to assess the usefulness and the effectiveness of our tools. Furthermore, we consider it one step towards providing best-practice examples of sustainable information modeling in Cultural Heritage – which is our main goal.

2. Previous and Related Work

The conceptual basis for our work are *The London Charter* (TLC) [Lon06] and CIDOC-CRM [CDG*05]. The TLC introduces the notion of *intellectual transparency*, which reflects the necessity of maintaining the distinction between measurement (wall with measured height) and interpretation (wall inferred from foundation wall). As a consequence, it prescribes to collect *paradata*, provenance and processing history, throughout the whole work flow.

The *Conceptual Reference Model* CIDOC-CRM from the "International Committee for Documentation of the International Council of Museums" (ICOM-CIDOC) is the standard for relational semantic networks in Cultural Heritage. It provides 84 *entity classes* (actor, place, time-span, man-made object etc). It also defines 141 relations, e.g. participated in, performed, at some time within, took place at, is referred to by and has created.

Niccolucci et al. have demonstrated the integration of 3D into semantic databases based on the X3D format using the MAD/SAD framework [ND06]. Our current work aims at an extension of their ideas, but is based on Collada instead of X3D. Collada is an exchange standard from the *digital content creation* (DCC) industry originally defined by SONY to streamline content exchange for the Playstation. [AB06]. It is now an open standard hosted by the Khronos group.

Our work can also be seen in the context of the *Seven open problems in 3D documents* [HF07], where maintaining a consistent relation between a shape and its meaning was pointed out as one major problem that is unsolved with current 3D technology. One step in that direction was the integration of XML into the scene graph [HSKF06] and the availability of a sustainable presentation framework for museums (aka *3D-Powerpoint*) [HSLF07]. The current work can be understood as a generalization of these approaches.

Concerning the processing of textual data into semantically enriched representations we have seen some projects connecting the Text Encoding Initiative (TEI) to other information resources using conceptual modelling. The Henry III Fine Rolls Project used RDF/OWL [CSVP07] and New Zealand Digital Library used Topic Maps [Tuo06]. And both used CIDOC-CRM in their ontology building. But the integration has been to objects such as historical persons and places and not to 3D representation, and the linking has been on the level of single objects. The distinction between an object and a part of an object (a statue and the arm of the statue) has not been taken into consideration.

3. Text Processing

The Text Encoding Initiative (TEI) is a consortium of institutions and individuals from all over the world. The TEI is also a set of guidelines for the encoding of textual material, and it is a set of computer readable files [tei08]. The guidelines and the computer readable files specify a set of rules, documents have to adhere to, in order to be accepted as TEI documents. These rules are expressed as XML schema.

Although the main area of usage for TEI has traditionally been in the field of edition philology, it has been used in the museum and cultural heritage sector as well [EL08]. Since 2004, the work of connecting TEI encoded texts to ontological information has been coordinated by the TEI consortium through the Ontologies Special Interest Group (http://www.tei-c.org/Activities/SIG/Ontologies/).

3.1. The TEI encoded document

Text encoding is a process of making implicit information explicit. A text printed in a book contains implicit information helping the human reader understanding it: Paragraphs, sentences, punctuation, etc. In order to assist machine processing of documents, some of this information has to be made explicit.

The TEI document is encoded by the use of XML elements. Each element in a TEI tag set has a description and usage examples as well as a name. This means that the semantic status of each element is fixed. The tag set also defines the relation between the different elements in the document's abstract structure so that only certain elements are legal at each place in the structure. The abstract structure of the document can be expressed as a tree structure. In addition to this tree structure, the document includes connections between individual nodes (id-idref pairs).

TEI documents consists of two main parts:

- The TEI header, which is a bibliographical description of the document,as well as possible sources, e.g. printed texts or hand-written manuscripts. The header may also contain thesaurus-like information and world knowledge, such as lists of classes of linguistic features or lists of persons with names used to refer to them as well as other information.
- The TEI body, in which the main text is stored.

Texts also contain references to objects in the real world or in a possible world, such as a name referring to a physical person or strings referring to 3D models of statues and its parts.

In order to connect positions in the text in which such references are made to other resources, they have to be explicitly marked up. This is what we do in the the semantic mark-up of documents. We include XML elements in the text so that these elements can be addressed. An example is the string “The hands must have held the imperial insignia”. The two substrings of special interest are “The hands” and “the imperial insignia”, and these two are tagged and have *xml:id* values so that they can be referred to.

```
<div type="section" n="1">
    <head>
        <rs type="artefact" xml:id="a1" key="Statue7">Arrigo VII enthroned</rs>
    </head>
    <div>
        <head>
            <rs xml:id="a2" type="artefactpart" key="Statue7#hands">Hands</rs>
        </head>
        <p>The <rs xml:id="a3" type="artefactpart" key="Statue7#arms">arms of the
            statue</rs> have been reduced to stumps. Originally, they must have been
            placed in such a way as to soften the diagonal folds of the dress on the
            chest. <rs xml:id="a4" type="artefactpart" key="Statue7#hands">The
            hands</rs> must have held <rs xml:id="a5" type="artefactpart"
                key="Statue7#insignia">the imperial insignia</rs>. <rs xml:id="e1"
                type="event" key="Statue7#first-mutilation">The first mutilation of the
                statue</rs> may have occurred when <rs xml:id="e3" type="event"
                key="Statue7#first-mutilation">somebody tore them of</rs>: <rs
                xml:id="a9" type="artefactpart" key="Statue7#insignia">they</rs> were
            probably made of <rs type="material" xml:id="m1">precious metal</rs>. The
            sceptre and the globe, together with the crown, symbolized power. They are
            therefore characteristic features of the iconographic representation of
            kings and emperors.</p>
    </div>
    <div>
        <head>
            <rs xml:id="a10" type="artefactpart" key="Statue7#head">Head</rs>
        </head>
        <p><rs xml:id="a11" type="artefactpart" key="Statue7#head">The head of the
                emperor</rs> has broken off from <rs xml:id="a12" type="artefactpart"
```

Figure 3: *Text processing output with semantic markup.*

The information extraction needed for this semantic markup is a human process, even though parts of it may be automated. It is based on a certain reading of a text. In our example, the various elements of the statues were of great importance, and were marked up in a very detailed way, whereas other parts of the text were less important. The result of this process is an XML document in which specific parts of the information are made explicit. Even if it will be possible to automate larger parts of this process, the choice of which information categories should be marked up will still have to be made manually.

For internet use, the TEI Consortium has provided a set of XSLT style sheets to convert TEI documents to HTML.

3.2. CIDOC-CRM modelling

The TEI document described above contains information about the objects discussed in the text, but no formalized relation between them. There are different strategies for including such information. To ensure that the best tool is used for each part of the work, the semantic model was stored in an CIDOC-CRM compliant XML document [CDG*05]. The tools existing for conceptual modelling in TEI are not suitable for complex cultural heritage systems such as the Arrigo showcase.

```
<crm:P62.Depicts>
  <!-- Shortcut -->
  <crm:E21.Person rdf:about="Arrigo VII">
    <crm:P1.is_identified_by>
      <crm:E82.Actor_Appellation rdf:about="Arrigo VII">Arrigo VII
      </crm:E82.Actor_Appellation>
    </crm:P1.is_identified_by>
    <crm:P67.is_referred_to_by>
      <crm:E73.Information_Object rdf:ID="p1"/>
    </crm:P67.is_referred_to_by>
    <crm:P67.is_referred_to_by>
      <crm:E73.Information_Object rdf:ID="p2"/>
    </crm:P67.is_referred_to_by>
  </crm:E21.Person>
</crm:P62.Depicts>
<crm:P46.is_composed_of>
  <crm:E22.Man-Made_Object rdf:about="Hands">
    <crm:P47.is_identified_by>
      <crm:E42.Object_Identifier
        rdf:ID="http://havemann.cgv.tugraz.at/fedora/get/arrigo/Statue7#hands"/>
    </crm:P47.is_identified_by>
    <crm:P102.has_title>
      <crm:E35.Title rdf:about="Hands">Hands</crm:E35.Title>
    </crm:P102.has_title>
    <crm:P67.is_referred_to_by>
      <crm:E73.Information_Object rdf:ID="a2"/>
    </crm:P67.is_referred_to_by>
```

Figure 4: *Processed text output as RDF triplets.*

The actual conceptual modelling is based on a human understanding of the meaning of the text. Each place the text contains a string such as “The hands”, a CIDOC-CRM *E22 Man-Made Object* is created in the model with the *xml:id* value of the element in the TEI document connected through an CIDOC-CRM *P67 is referred to by* property. Further, each set of hands is connected to the statue they form part of through other CIDOC-CRM properties. For each statement, the *xml:id* value connecting it to the correct place in the TEI document is included.

We must stress that in such a system, it is necessary to connect each statement in the CIDOC-CRM model to the place in the document from which it is taken. The formalized, machine readable information is stored in CIDOC-CRM. A requirement of this linking is to show a human reader the place from which the information is fetched. This is vital, not just to make a nice interconnected information system, but to ensure the scholarly quality of the work. The critical user of systems built on these principles will easily be able to check the source of the statements in the model, and maybe criticise the published semantic model for what he believes to be misinterpretation of the text. Thus, the function of the links is important at the same level as the footnotes in a traditional scholarly paper.

In the showcase described in this paper, information from only one document is included. In a production system this will rarely be the case. If artefacts are important enough to create respective 3D models, one would suppose more than one source of information about them is available.

3.3. Document production

In the showcase, a pre-existing text was TEI encoded and the CIDOC-CRM model was developed manually. In a production system, it would be better to include the encoding and the conceptual modelling into the text production process. To be able to create such complex documents, the author would need access to a text production system in which such information could be easily included.

For the TEI encoding, good XML tools do exist and are used by many projects, e.g. Oxygen from SyncRO Soft. Tools for conceptual modelling also exist, but they are not integrated into XML tools used by the TEI community. In our opinion, development of user friendly tools in this area is necessary to persuade scholars to create texts in which explicit semantic models are included. The only alternative, post-processing of texts, will be more expensive and produces less quality. It has to be done for legacy texts, but hopefully not for texts created in the future.

3.4. Information integration

The use of the tools created by text encoding and conceptual modelling is based on the links created between parts of the 3D models and fragments of TEI texts that goes through entities and properties of the CIDOC-CRM models. Such links can be used as seen in Figure 7. But the information of different types that have been created are also connected to other resources, e.g. in the following ways:

- A part of the statue is connected to other parts of the same statue and to other statues by geometrical relations. This includes connections based on speculations about historical arrangements.
- A text fragment is related to other text fragments in the same document. It is also related to texts in other documents through implicit as well as explicit intertextuality.
- The statements in the CIDOC-CRM model are related to other statements in the same model, as well as to similar statements and co-referring statements both locally and in other models.

In this way, the work with 3D models of culture historical artefacts is located where it belongs: In the network of interconnected knowledge sources, both classical and modern, expressed in all sorts of media.

4. Semantic 3D Markup Technology

After the production of a 3D dataset, Collada files for annotation can be generated with the AuthoringTool. The software is capable of retrieving 3D objects from a local file system or remotely from a URL. This also works for follow up files like material definitions and textures which may not be included in the model file.

The 3D data is embedded into a scene to add the necessary additional data. The conceptual components of the scene are:

- **Hierarchical structure**: The scene graph structure is used to arrange one or multiple 3D models.
- **Reference to 3D data**: The 3D data, potentially massive datasets, is referenced using a URL.
- **Transformation**: Spatial location and orientation of the imported 3D model or sub-scene.

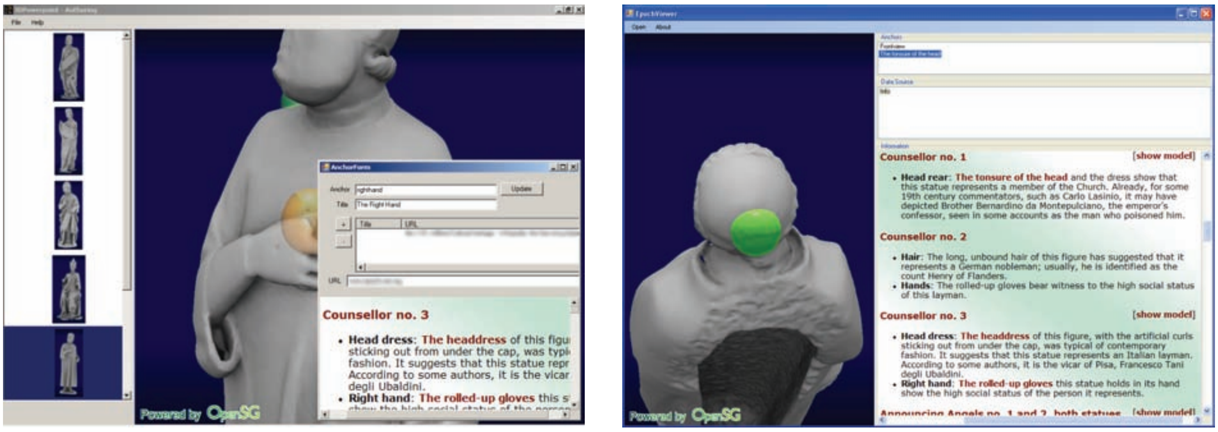

Figure 5: *The AuthoringTool (left) and one of the presentation tools with integrated web browser (right).*

- **Annotation data**: All of the neccessary data to define a markup is directly attached to each 3D model.

In the first step of the authoring process the model can be arranged properly. It may be neccessary to re-align the data to match a desired up-axis for consistency, or to align with other objects. Also scaling to the original size may be neccessary. Depending on the used shape aquisition method the original data may or may not be in a unified metric coordinate system. All transformations are stored only in the scene file, the original 3D data are not modified.

The scene is written to a "lightweight" version of the Collada format (just scene, no shape data) so that referenced shape data remain in a binary format. This way, optimized binary formats beyond the Collada standard can be supported. All those formats have to be well defined and documented for long term archival. This was achieved by using the open source scene graph engine OpenSG as a unified basis for the whole process of 3D aquisition. The supplemental data containing the markup definition and associated URLs is stored in a so called extra element provided by the Collada scheme. Extras containing custom data can be attached to every scene graph object. Note that Collada compliant applications are not allowed to discard unknown extra data.

Having set up the Collada scene for the 3D model it is now possible to add markup areas with the authoring software. One of the easiest ways to define an area in 3D is using a sphere. AuthoringTool allows to create spheres by clicking on the model's surface and dragging the mouse to modify the radius of the sphere. Other primitives like boxes may also be considered as markup geometry. Additional geometry definitions for annotation areas can be added in the future with little expense.

The user will most likely navigate to a good point of view to define the area. This behavior helps to store a suitable camera definition for the markup area. There are some algorithms available to find such view points automatically, but in general it is best to let the knowledged user choose it manually in the authoring process. The camera is described quite similar to the common Collada specification. The annotation geometry and the camera are defined in the object space of the model. Independent from the transformation of the scene it is possible later on to reuse the annotation definition for other scenes.

Now the user has to set an identifier for a later reference of the annotation. This identifier is similar to an anchor point of a web page. It has to be unique within the scene. One or more URLs can then be attached to an anchor leading to the additional data. This step depends a lot on the purpose of the scene. For a museum application the URLs lead to appealing presentations of details on the model. For scientific markup, a data base entry with collected facts may be more appropriate. The authoring software supports the user in finding the associated data with the integrated web browser. Drag and Drop functionality helps to copy URLs easily from one place to another.

The information is stored directly with the scene graph object as XML code. Later modifications of the scene will not destroy the defined markup and the connection to the model. An example of the extra element is listed in Figure 6.

5. Applications: Authoring, Presentation, Browsing

The heart of the EpochViewer framework is an ActiveX control, which provides the 3D widget area. It combines the scene graph system OpenSG with the scripting language GML [GBHF05]. Since the first version the operator set of GML has been extended with support for reading the lightweight Collada scene description and basic XML manipulation. The control can be used in custom VisualBasic, C# applications or embedded directly into an HTML page. It is also available in all ActiveX-compliant software (like Mi-

```
...
<annotation id="head"
            title="Head␣of␣Arrigo">
  <annotation_geometry>
    <annotation_sphere>
      <center>13.1 31.96 2.32</center>
      <radius>3.73</radius>
    </annotation_sphere>
  </annotation_geometry>
  <annotation_camera name="eyepoint">
    ... similar to Collada
  </annotation_camera>
  <annotation_url title="Some␣Text">
   http://www.txt.com/statue/
  </annotation_url>
  <annotation_url title="A␣Photo">
   http://photos.org/statue/
  </annotation_url>
</annotation
...
```

Figure 6: *Example ⟨annotation⟩ within the Collada scene "Statue07.dae": The annotation "head" with one sphere as the annotation's markup geometry, a camera definition and links to two external information sources.*

crosoft Office, Internet Explorer, .NET Framework), which makes it the ideal platform for generating custom tailored user interfaces.

We would love to use other operating systems and/or component technologies, but currently there are no reasonable alternatives available in terms of easy deployment and extensibility. However it would not be difficult to fit our approach into alternative component frameworks.

Inspired by the Arrigo Showcase [CPCS06] the following sample applications (see Figure 5) were presented at the Epoch final review meeting in Rome:

- **AuthoringTool**
  The AuthoringTool is an exemplary authoring tool for creating semantically enriched 3D objects. This application focuses on people with special knowledge of an object (e.g. a historian or an archaeologist), so the functionality was stripped down to the absolutely necessary.
- **PresentationTool**
  The PresentationTool combines the ActiveEpoch control with a browser control to display the URL targets and a list of anchors. Annotations can either be examined by clicking on the markup geometry in the 3D scene or by selecting an anchor from the list. By choosing an annotation, the point of view is automatically moved to the predefined position.
- **Browsing**
  In order to integrate our framework into the web world it is important to stay consistent and intuitive in terms of URL scheme and default browser behaviour.
  The URL scheme described in RFC 1738 [BLMM94] uses the character "#" to delimit a URL from its fragment/anchor. This mechanism is used to address anchors within an HTML document. Using that scheme we can easily reference the anchor "head" within the "Statue7.dae" Collada file (see Figure 6):
  <a href="http://www.cgv.tugraz.at/Statue07.dae#head">
  Unfortunately it is up to the browser implementation how to react to a fragment/anchor in a web address. The browser default behaviour on documents it cannot handle itself is to download the content to a temporary file and to start an application which is registered to handle that type of content, but without transmitting the anchor information.
  For the Internet Explorer in Microsoft Windows this can be achieved by a Browser Helper Object (BHO), which is a plugin to customize or add functionality, e.g. the Adobe Acrobat plugin uses a BHO in order to correctly process the fragments for PDF files (some.pdf#page=15).
  Other browsers allow to implement the same behaviour, e.g. with the use of an extension for Mozilla-based browsers.

6. Conclusion and Future Work

The concepts presented in this paper are easy to understand and the implementation of tools is not too much of a challenge. So we consider it to be a simple solution to a complex problem: the semantic annotation of 3D data on a sub-object level. Parts of a dataset, e.g. the left arm of a statue, can be linked bi-directionally to other data sources. The additional

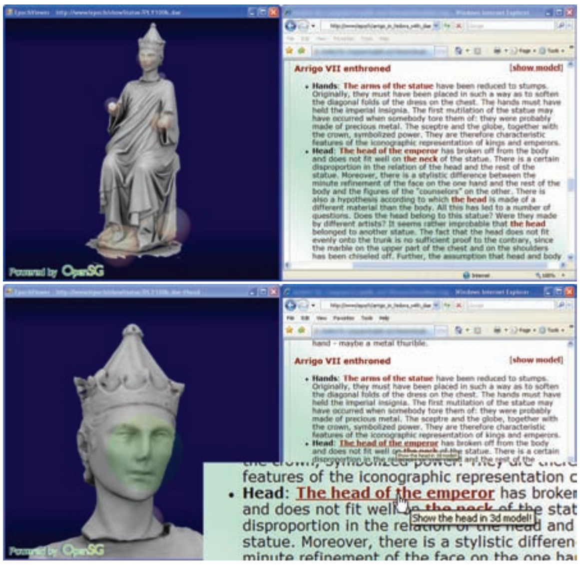

Figure 7: *A typical viewing session: After the user has clicked on the head's icon it is highlighted in the viewer and the camera moves to the predefined setup. Clicking on the arms in the 3D view would trigger the browser to jump to the associated text.*

information can be of any kind like texts, images, videos, etc. as long as a standard internet browser can display it.

Information in any form is accessible independently of the application scenario and can be referenced in a unified way. The creation of annotations as well as referencing parts of the model is easy to use. And the integration into standard internet browsers makes working with digital 3D objects as comfortable as working with text files. Exchanging annotated geometry with collegues is trivial via e-mail.

Using this solution, 3D models based on culture historical artefacts will be integrated into the wider context of knowledge sources that will become the **semantic web of cultural heritage**. This will push the process of changing such 3D models from "technical toys for the few" to elements in the toolbox used by researchers, curators and the general public as part of their daily work with cultural information.

Next steps include the integration of data base systems with large amounts of 3D data and annotations. It has to be possible to collect existing information using data base queries to create Collada files with annotations automatically. In fact, the queries can be directly attached to the 3D models. Adding annotations back to the database will allow a collaborative information exchange between researchers.

The presented software, both the ActiveEpoch control and the sample applications for authoring and browsing, is available from www.cgv.tugraz.at/Epochviewer.

Acknowledgements

The authors want to thank Roberto Scopigno, Paulo Cignoni and the ISTI-CNR for the Arrigo models. And the authors wish to thank the European Commission, in particular the EPOCH network of excellence in cultural heritage (www.epoch.eu) for their generous support.

References

[AB06] ARNAUD R., BARNES M. C.: *Collada: Sailing the Gulf of 3D Digital Content Creation*. AK Peters Ltd, 2006.

[BLMM94] BERNERS-LEE T., MASINTER L., MCCAHILL M.: *Uniform Resource Locators (URL)*. RFC 1738, Internet Engineering Task Force, Dec. 1994. ds.internic.net/rfc/rfc1738.txt.

[CDG*05] CROFTS N., DOERR M., GILL T., STEAD S., STIFF M.: *Definition of the CIDOC Conceptual Reference Model*, version 4.2 ed. CIDOC Documentation Standards Working Group, June 2005. Also ISO/PRF 21127, available from cidoc.ics.forth.gr.

[CPCS06] CALLIERI M., PONCHIO F., CIGNONI P., SCOPIGNO R.: Easy Access to Huge 3D Models of Works of Art. In *Fourth Eurographics Italian Chapter 2006* (Graz, Austria, 2006), Fellner D., (Ed.), Graz University of Technology, Eurographics Association, pp. 29–36. Catania, 22-24 February 2006.

[CPCS08] CALLIERI M., PONCHIO F., CIGNONI P., SCOPIGNO R.: Virtual Inspector: a flexible visualizer for dense 3D scanned models. *IEEE Computer Graphics and Applications 28*, 1 (Jan.-Febr. 2008), 44–55.

[CSVP07] CIUALA A., SPENCE P., VIEIRA J. M., POUPEAU G.: Expressing complex associations in medieval historical documents: The Henry III fine rolls project. In *Digital Humanities 2007. Conference abstracts* (2007).

[EL08] EIDE O., LIGHT R.: The text encoding initiative: what's in it for museums? In *CIDOC 2007 conference papers* (2008), CIDOC.

[GBHF05] GERTH B., BERNDT R., HAVEMANN S., FELLNER D. W.: 3D Modeling for Non-Expert Users with the Castle Construction Kit v0.5. In *Proc. VAST 2005* (Nov 2005), Mudge, Ryan, Scopigno, (Eds.), Int. Symp. on VR, Arch. and Intelligent Cult. Heritage, Eurographics / ACM Siggraph, Eurographics Press, pp. 49–57.

[Hav05] HAVEMANN S.: *Generative Mesh Modeling*. PhD thesis, Institute of Computer Graphics, Faculty of Computer Science, Technical University Braunschweig, Germany, November 2005.

[HF07] HAVEMANN S., FELLNER D. W.: Seven research challenges of generalized 3D documents. *IEEE Computer Graphics and Applications 27*, 3 (2007), 70–76. (special issue on 3D documents).

[HSKF06] HAVEMANN S., SETTGAST V., KROTTMAIER H., FELLNER D.: On the integration of 3D models into digital cultural heritage libraries. In *VAST 2006 Project Papers* (Nicosia, Cyprus, Nov 2006), Ioannides, Arnold, Niccolucci, Mania, (Eds.), Int. Symp. on VR, Archae. and Cultural Heritage, Epoch publication, pp. 161–169.

[HSLF07] HAVEMANN S., SETTGAST V., LANCELLE M., FELLNER D.: 3d-powerpoint - towards a design tool for digital exhibitions of cultural artifacts. In *VAST 2007* (Brighton, UK, Nov 2007), Arnold, Niccolucci, Chalmers, (Eds.), 8th Int. Symp. on VR, Archaeology and Cultural Heritage, Eurographics Association, pp. 39–46.

[Lon06] LONDON CHARTER INITIATIVE (HUGH DENARD): The london charter, June 2006. www.londoncharter.org.

[ND06] NICCOLUCCI F., D'ANDREA A.: An Ontology for 3D Cultural Objects. In *Proc. VAST 2006* (2006), pp. 203–210.

[tei08] *TEI P5: Guidelines for Electronic Text Encoding and Interchange. Version 1.1.0. Last updated on 4th July 2008*. TEI Consortium, 2008.

[Tuo06] TUOHY C.: Topic maps and tei — using topic maps as a tool for presenting tei documents. In *TEI Day in Kyoto 2006* (2006), pp. 85–99.