

3D-centered media linking and semantic enrichment through integrated searching, browsing, viewing and annotating

S. Pena Serna¹, R. Scopigno², M. Doerr³, M. Theodoridou³, C. Georgis³, F. Ponchio² and A. Stork^{3,4}

¹Fraunhofer IGD, Germany

²Visual Computing Laboratory, ISTI-CNR, Italy

³Institute of Computer Science, FORTH-ICS, Crete, Greece

⁴TU Darmstadt, Germany

Abstract

The digitally documented world heritage is archived in databases or repositories, where collections of metadata, images, multimedia objects or nowadays even digital 3D artifacts can be stored and queried. Modeling and linking all this information is complex and involves refined categorizations and relations, which are usually accessed by either simplistic or overwhelmingly complex interfaces. Hence, finding the right level of abstraction for a general interface is very challenging. This becomes even more demanding, if in addition to collection exploration, semantic enrichment is required. This work focuses on the design and implementation of an integrated interface, in which four dedicated activities are combined and provided: i) flexible definition of relational queries, ii) browsing of multiple results and query refinement, iii) inspection and analysis of multimedia objects, and iv) building and enrichment of semantic relationships (data paths) between objects. This integrated interface can handle different kinds of multimedia objects, allowing for querying and annotating text, 2D images or 3D artifacts. We present the software design of this interface and the corresponding underlying model in the semantic network. This work is a general step toward interfacing to 3D Linked Open Data.

Categories and Subject Descriptors (according to ACM CCS) : H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Hypertext navigation and maps

1. Introduction

The impressive evolution of digital sampling technologies (digital cameras, 3D scanning systems, modeling tools) allows for producing excellent digital objects of works of art and Cultural Heritage (CH) objects in general. We envision a new approach to the study, conservation and restoration of our Cultural Heritage based on such digital assets. One of the major challenges is to allow anyone to access those digital representations from distance, enabling him/her to study in detail any CH object from any remote location, thanks to the high quality and accuracy of the digital models which may be available. We do not propose to replace the on-site analysis of the single CH object, but to couple it with new possibilities for searching over remote repositories, retrieval, visual analysis and, moreover, for making the visual (and numeric) comparison of artworks possible, even if these are located in different spaces or are too fragile for frequent manual handling. In order to support such a scholarly work in

the CH domain, we need innovative repositories capable to store high-quality multimedia (MM) representations of objects, their cultural and historical context, and the interconnection between all of them [DTT*10, PBH*10]. Modeling and linking all this information is complex and involves refined categorizations and relations, which are usually accessed by either simplistic or overwhelmingly complex interfaces. Hence, finding the right level of abstraction for a general interface is very challenging.

The paper presents the architecture and implementation details of what we refer to as the Integrated Viewer/Browser (IVB). The IVB is devoted to searching, querying, browsing, viewing, segmenting and annotating multimedia objects stored in the semantic repository (RI), which is composed of a content management system (OR), an RDF triple store for the metadata (MR), several kinds of content retrieval indices, and a common access and query manager. The paper is organized as follows: Section 2 reviews the scientific

work related to our contributions; Section 3 describes the software design principles of the IVB, in order to integrate the back end with the RI, and the front end with the Graphical User Interface (GUI) components; Section 4 presents the back end, including the semantic repository and the corresponding semantic models; and Section 5 deals with the GUI components of the main actions: searching, browsing, viewing, segmenting, and annotating.

2. Related Work

As mentioned before, in this paper we address searching, browsing, viewing, and annotating. Since the spectrum is too wide to describe all related work, we focus mainly on related work concerning the innovative core of our contribution: 3D media annotation and semantic enrichment; and shortly on segmentation.

Annotating 3D objects is commonly understood as marking a point or outlining a part of an object and associating it with a text string. Many CAD systems and viewers designed for engineering and architecture support such simple forms of annotation [JIDG^{*}99, JGD02]. Software vendors such as Siemens PLM Solution [Sie11] and Adobe [ADO09] also provide annotations to 3D CAD models. An example of work concerning annotation over digital CH artifacts was proposed by Hunter and Gerber [HG10], where authors have developed an annotation tool for 3D museum artifacts, attaching tags or comments to either 3D points or views of the whole object. Rodriguez-Echavarría et al. [REMA09] developed an approach for annotating 3D content representing sculptures and monuments, using COLLADA and linking parts of the 3D geometry to CIDOC-CRM [CDG^{*}09] URIs. Pittarello and Gatto [PG11] recently presented a new architecture for managing combined top-down (ontology driven, typical for the semantic web) and bottom-up (user driven, typical for the social web) annotations, including an integrated interface for annotating, browsing and searching 3D virtual worlds.

In order to attach the semantic annotations to a specific area of the 3D object, techniques such as geometric segmentation [DFPC08, ARSF07] are employed. In this context, ShapeAnnotator [ARSF07] focuses on the segmentation of 3D objects, providing a set of algorithms the user can choose from, and labeling these segments based on an ontology (or taxonomy) of the object type being segmented. Shamir [Sha08] gives a good overview about mainly automatic mesh segmentation approaches and he shows that this is an active research field without a single solution that serves all purposes. One of the first large scale test studies in the field of segmentation [CGF09] was based on a user study, where several people were selected and asked to segment different objects (by means of a crowd-sourcing approach). Segmentation in the CH domain is a complex task, since CH objects can be arbitrary complex and segments are not always defined by geometric features.

One work closely related to the one proposed in this paper, is the approach proposed for redesigning the Arrigo information system [HSB^{*}09]. The Arrigo system was originally designed as interactive virtual kiosk for a museum exposition [BBC^{*}04], where the visitor can explore 3D objects and the annotations provide him with detailed information about the artifacts or parts of them. In the following redesign of the system, the 3D information and annotations have been encoded and transmitted using a COLLADA container and the annotations are expressed in a CIDOC CRM conform model. In the case of the Arrigo system, the goal was the presentation of 3D objects and of associated static and pre-defined information, whereas the work presented here addresses the full pipeline of creating annotations, linking media objects, searching and enriching them.

In summary, we can state that there is currently no comparable approach which covers searching (incl. query development), browsing, viewing, segmenting and annotating in a similar way as we do, by addressing the needs of CH professionals for documentation and research purposes as well as providing a basis and (re)source for virtual museum presentations.

3. IVB Design Principles

The IVB deals with a very ambitious purpose, enabling an integrated interface, in which four dedicated activities are combined and provided: i) flexible definition of relational queries, ii) browsing of multiple results and query refinement, iii) inspection and analysis of multimedia objects, and iv) building and enrichment of semantic relationships. Such an ambitious purpose requires robust design principles ranging from software architecture to user perspective, e.g.: a) platform independence, b) separation of back and front end code, c) clear communication between GUI components, d) fast integration of new components or functionalities, and e) a simple and flexible User Interface (UI).

3.1. Plug-in System

A plug-in system is necessary given the large number of multimedia supported formats and the flexibility it introduces in the development cycle with multiple developers. It also contributes to the platform independence, the fast integration of new components, and the separation of legacy code from the UI. The plug-in system is designed to avoid any code modification or recompilation inside the application. It consists of a *Plugin SDK* (Figure 1), which contains all necessary components for:

1. Loading a specific plug-in (based on UUID or Filename)
2. Searching for loaded plug-ins (with a UUID or Filename)
3. Automatically loading all plug-ins and create one instance from each of them
4. Integrating a window or widget of a plug-in into the application (for rapid prototyping)

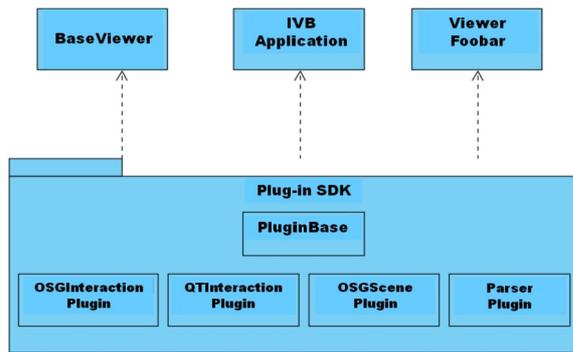


Figure 1: Plug-in System overview with classes *PluginLoader*, *Kernel*, and *IPlugin*.

5. Integrating windows / widgets of other applications (re-parenting)
6. Integrating web-based applications (over QWebview)

The three main classes of the Plug-in System provide the needed flexibility, in order to comply with the required design principles. These classes are:

- **PluginLoader:** it is responsible for finding, loading and optionally creating objects from plug-ins. It also knows the specific file extensions or naming of shared libraries on the operating system;
- **Kernel:** this interface is the minimal interface that each application needs to implement. It provides the application methods that are callable by external plug-ins;
- **IPlugin:** this interface is the minimal interface that needs to be implemented for each plug-in. It does not have any additional external library dependencies, and it is meant for simple plug-ins.

3.2. Architecture

The back end of the IVB is the repository communication library, which manages user authentication, processes queries, and retrieval of results and multimedia objects, including the engine for formulating SPARQL queries and for parsing results. The front end is the GUI layer, which assists the user in the query formulation process, in navigating the query results, in segmenting and visualizing multimedia objects, and in creating annotations. The IVB is written in C++, it is based on the Qt framework [QT11], and its visualization is powered by OpenSG [OSG11]. The communication between the RI and the IVB is depicted in Figure 2.

Since the purpose of the IVB is to support CH scholars, curators or practitioners in their study activity, we aim to go beyond search and visualization functionalities. We enable the semantic enrichment of the multimedia collection stored in the repository, by providing instruments to annotate or link any MM document or CH object. To better explain our vision, we introduce the following terms:

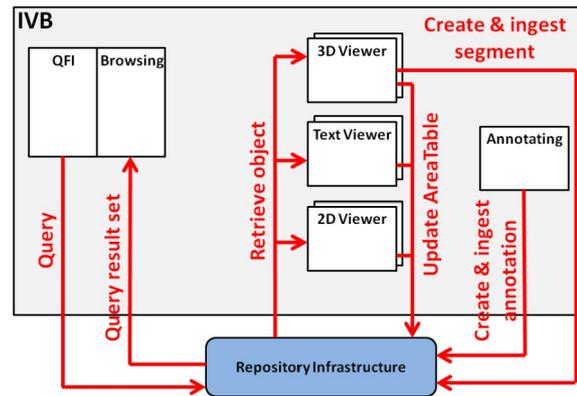


Figure 2: Communication between the RI and the IVB.

- **Segment:** a segment is a subset of a 3D object. Segments are generated out of manual (or automatic) structuring processes and become objects on their own. The input 3D object to the segmentation process and the result exist as entities in our repository and are linked to each other by a segmentation event. A segment can become an *area*;
- **Area:** an area is an abstraction of positions/regions defined in the virtual space of different MM objects, i.e. texts, images and 3D objects. The geometric definition of an area exists in addition to the geometric elements of the object, e.g. it is not referred to as a subset of the vertices of the 3D object itself;
- **Annotation:** an annotation is a set of commented and classified relations between areas and other contextual elements, i.e. annotations are independent from the media object underneath and its dimensionality (1D, 2D, 3D).

3.3. Interface Layout

The IVB provides the four main actions: searching, browsing, viewing, and annotating. Combining these actions in a logical and neat interface requires a clear layout. Figure 3 illustrates the interaction between the IVB components. *Searching* sends the user-defined query to the RI, the RI gives back the results to *browsing*, this transmits the user-selected UUID of the multimedia object to *viewing* and then *viewing* retrieves the multimedia object from the RI. *Viewing* creates the areas in the RI and communicates the UUIDs of the areas to *annotating*, where a user-defined annotation is generated and ingested into the RI. Furthermore, the results in *browsing* can be used by the user for query refinement.

4. IVB Back End

The back end of the IVB is mainly composed of: a) the Repository Infrastructure, b) the Metadata Schema, and c) the Annotation Model.

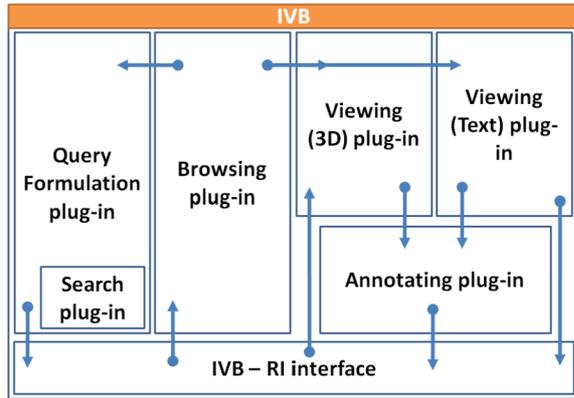


Figure 3: Conceptual interface layout of the IVB components.

4.1. Repository Infrastructure

The Repository Infrastructure (RI) was designed and is being implemented as a sustainable, distributed repository for a massive quantity of large digital objects and their metadata, capturing scientific descriptions and monitoring all processes of data generation and processing. The RI is designed to allow different kinds of users (researchers, academics, web users, etc) to store, manage, query and retrieve whatever temporary and permanent digital objects, which are created in the course of digitization and post-processing.

The essential internal components of the RI are shown in Figure 4: (a) the Object Repository (OR), a distributed mass data storage layer, (b) the Metadata Repository (MR), an integrated semantic network layer, (c) the Content Retrieval Indices (CRI) for different search modalities by several content modules, and (d) the Query Manager (QM), a single homogeneous access point to querying the three components.

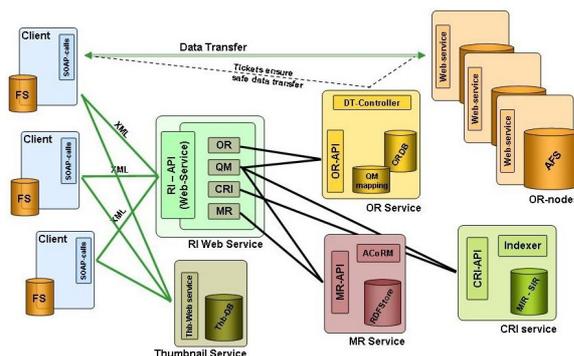


Figure 4: Overview of the Repository Infrastructure architecture.

The clients communicate with the RI via SOAP through a central webservice and the RI dispatches the requests to its

components (the MR, OR, and CRI Services). Data transfer is performed directly between clients and OR nodes, as initiated by the OR Service which controls the distributed OR nodes (DT-Controller). The RI offers a central thumbnail database and provides a webservice for handling http thumbnail management requests. The Object Repository (OR) connects to a (potentially large) number of distributed OR nodes for the physical data storage. The central OR Service provides access security, data integrity and risk-of-loss control, including metadata backup. The Metadata Repository (MR) is an RDF triple store that aims at providing a common place to reason on, query, manipulate and export provenance metadata concerning any temporary or permanent object stored in the OR and related metadata about the modeled reality. The MR is based on a homogeneous global schema, an extension of the CIDOC CRM (ISO 21127:2006) [CDG09] that models not only human material history and cultural objects, but also provenance metadata (CRM_{dig}) [TTD 10].

4.2. Metadata Schema

Cultural and scientific data cannot be understood without knowledge about the meaning of the MM objects and the ways and circumstances of their creation. Generally, we use metadata to assess meaning, relevance, and quality of the MM assets, giving us also the capabilities to improve and re-process those digital models [DT11]. This knowledge comprises the *provenance* of the data. Provenance metadata are event centric and are best described in a historical order, in order to ensure that there are no references to non-existent (non-recorded) events or objects. Metadata representing the different types of events that incur in the lifetime of a MM asset can be structured by answering the following four main questions about:

- Who: the people (organizations) playing role in the event
- Where: the place the event occurred
- When: the time the event occurred
- What: the things involved in the event

These events form an open-ended directed graph of things used or produced by events, rather than linear chain or inverse tree from sources to final result. This graph should represent also the physical acquisition context and processes of the MM assets and the description of the CH objects acquired [DTT* 10]. We chose CIDOC CRM as the core conceptual schema and its further development CRM_{dig} .

4.3. Annotation Model

The annotation interface of the IVB enables the formulation of individual or aggregated annotations over any of the supported MM types (text, 2D images, 3D objects, etc.), based on a common and uniform approach. A powerful and flexible underlying annotation model based on CIDOC CRM has been added to the IVB: it describes the metadata of the

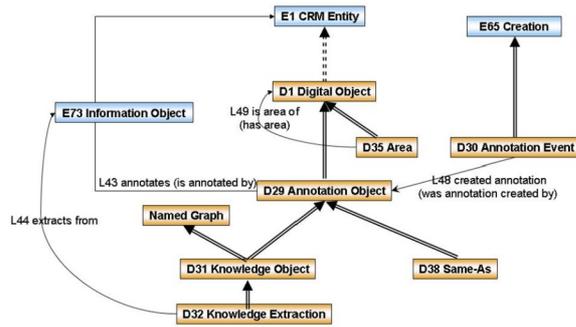


Figure 5: Annotation Model within the RI.

event of annotating as a named graph, containing the annotation statements as part of the same semantic CIDOC graph for later searching and reasoning. In order to associate RDF triples with parts of 3D objects or other media, we refer to the generalized notion of *area*.

The Annotation Model (Figure 5) has two basic entities: the Annotation Event and the Annotation Object. The Annotation Event is the parent event that creates the Annotation Object. The Annotation Object is the entity describing the association between the annotated objects. We define two sub-classes of the Annotation Object: the Knowledge Object modeled as a Named Graph and the Same-As, which is used to declare co-reference links. The Knowledge Extraction is a specialization of Knowledge Object that is used to model information that will be (semi)automatically extracted from legacy data. Our annotation model can be seen as generalization of the Open Annotation Model [OAM11], because it uses Named Graphs, which describe any form of linking, and as specialization of it, since `oac:hasBody` can point to a Named Graph.

5. IVB Front End

The front end of the IVB has 4 components: a) Query Formulation Interface (searching), b) browsing, c) viewing, and d) annotating. Several instances of the first three components are available, but only one instance of the annotating component exists. This is needed to keep the synchronization between the areas generated by the different viewers for the corresponding annotations.

5.1. Query Formulation Interface – QFI

The QFI is designed as an easy to use and flexible interface to the repository, able to support both the query formulation process for multiple retrieval modes and a complex semantic metadata network, as well as the presentation of results, which may include many different types of multimedia content. The user is guided from the selection of the first query

parameters on through the process of formulating and refining retrieval or associative query expressions. In order to embed reasoning capabilities into the RI, we use the "Onto-text OWLIM" Reasoner to materialize transitive properties that cannot be queried by SPARQL, to relate inverse links and to materialize deductions expensive at run-time.

We enable the formulation of effective queries, by means of adequately describing our metadata and supporting the necessary reasoning without requesting users to learn the semantic model. We proposed and are implementing a new query paradigm for semantic networks. A set of generic, intuitive "fundamental" relationships are presented to the user through the QFI interface. With this approach, we are able to achieve high recall with a compact set of very simple, comprehensive questions. We firstly divide the entities of our application into a set of relevant fundamental categories (FC). These are presented to the user as choices for querying and serve as domains and ranges of fundamental relationships (FR) that can be queried as described below. We try to cover the domain with as few FRs as possible, but still to be able to make some powerful distinctions such as discerning places from people with the same name.

We have selected the following FCs: **Thing, Actor, Event, Place, Time, and Concept**. Given those FCs, we have declared the following FRs, whose interpretation depends on the combination of FCs they are applied to:

1. *has type*
2. *is part of*
3. *is similar or the same with*
4. *has met*
5. *from, has founder or has parent*
6. *at*
7. *refers to*
8. *is referred by / is referred to.*

The QFI is composed of the following fields:

- a free text search box where keywords can be entered;
- a set of checkboxes corresponding to the Fundamental Categories;
- a list of Fundamental Relationships specified through a menu and a text field.

A user formulating a query in our system may first type in a keyword (Figure 6). To refine the query, a user can select one or more FC that qualify the types of entities he is looking for. Then, the user can further refine the query by composing a sort of Where-Clause, selecting FRs and query parameters for their ranges, which must be URIs of nodes in the semantic network. There are two possible ways to specify the URI: the user can search by keyword directly in the parameter / term field and select the desired result (similarly to many auto-suggesting interfaces), or he can use the result of the previous query by either using drag-and-drop or copy-paste actions on the presented results (in browsing). This mechanism relieves the user completely from understanding the

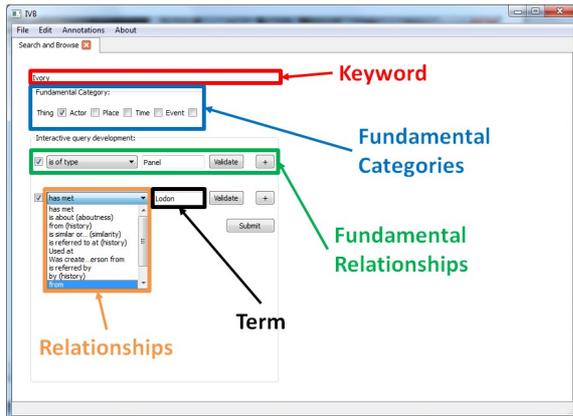


Figure 6: Query Formulation Interface.

schema of the semantic network, and yet allows for formulating powerful queries. Some intelligence is built in the system where: the list of available relationships depends on the selected FCs; and the term suggestions filters the results, presenting only a valid combination of FCs and FRs. Multiple FCs can be specified and they will be combined using the AND operator, while each FR accepts multiple parameters and joins them using the OR operator.

5.2. Browsing

The result of a query is a ranked list of CRM entities (URIs), which are complemented with additional metadata information for the user. Each entity is primarily identified by a label and it is visually represented by a thumbnail. Depending on the Fundamental Category and existing properties of the entity, other important information is also displayed. In order to facilitate faster and deeper inspection of the semantic network, each entity is enriched with a set of predefined queries. The browsing component in Figure 7, is based on an HTML template engine and allows for easy and flexible customization of the result presentation pages and a seamless integration with the web. The solution adopted enables the specification of layouts for each class of object in a JavaScript-based template language. These templates are filled in with query results and rendered using Qt's implementation of the WebKit rendering engine. This mechanism provides a simple way to rapidly update and iterate the interface and exposed metadata, according to the user requirements.

5.3. Viewing

The IVB application supports the native visualization of many different multimedia objects: images, 3D objects, text documents, audio and video. Different objects can also be visualized using external applications while the IVB takes care of retrieving the files from the repository. Beside the simple visualization feature over all types of media documents,

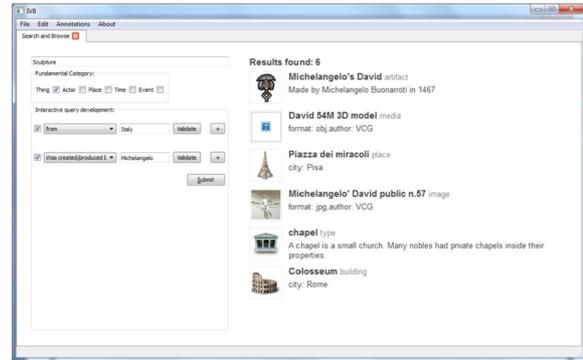


Figure 7: Searching and Browsing interface of the IVB.

the IVB supports the visualization and creation of *areas* (see Figure 8). An effective visualization of the defined areas is implemented by rendering these in overlay and linking each one to the semantic information about the area (labels, creator, annotations) on a side box. Each viewer is responsible for the rendering of the areas while common functionalities are exposed by the IVB application for displaying general semantic information, creating new areas, and maintaining consistency.

The RI will usually hold large 3D objects resulting from 3D scanning, which can be problematic to be explored remotely due to the bandwidth limitations. Therefore, supporting streaming and efficient rendering of very large mesh- or point-based 3D datasets is a mandatory goal for the IVB. Many different multiresolution approaches have been proposed in the literature ([KDG⁺08]), we just remind the recent trend that focuses on clustered multiresolution structures [YSGM04, CGG⁺04, CGG⁺05, Pon08]. We have adopted the Nexus library [Pon08], in order to fulfill this requirement, which implements in a common framework a multi-resolution approach, progressive streaming and compressed encoding of the models, all of them coupled with an efficient memory management caching system.

5.4. Annotating

Annotating different multimedia objects in a consistent manner requires a generalized concept. We have proposed a unique and uniform way to define areas on the variety of multimedia objects stored in the RI. We suggest the use of the generic and extensible METS schema [MET11], by means of defining the term "area" to describe a part of interest in basically any media object stored in the RI. However, since the METS schema does not provide `<mets:area>` attributes for 3D models and HTML documents, we have defined extensions to METS. We keep the new schema as close to the original METS schema as possible. We managed this by extending `<mets:area>` element, by introducing two new attributes (`AreaID` and `ExtMeshAreaID`) and by adding new

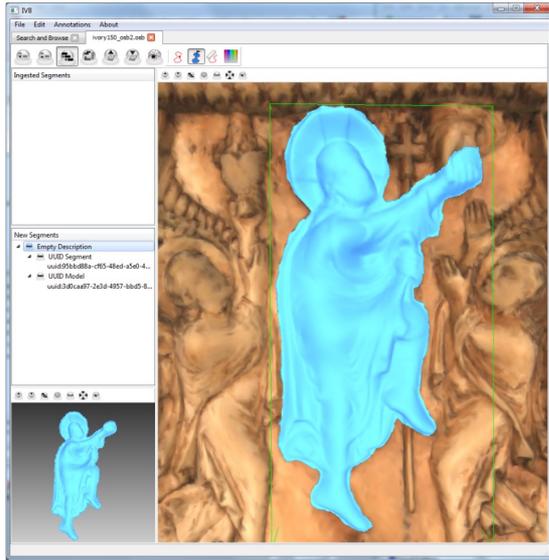


Figure 8: Segment on the digitized Ivory Panel from the V&A Museum in London, which is used as area for annotating.

values for SHAPE (values: COLLADA, RELRECT, RELCIRCLE, RELPOLY and EXTMESH) and BETYPE (value: DOM-RANGE). In the *METS 3D Extension*, possible areas can be specified on 3D objects as:

- Sphere, specified as radius and translation;
- Axis aligned box, specified as dimensions and translation;
- Mesh subsection, specified as a set of vertices and triangles.

In all those cases, the 3D area should refer to the object coordinate system and not directly to the 3D object files. We implemented this by adding a new SHAPE attribute value (COLLADA) to the < mets:area > element. The COLLADA node must be inside a < collada > element declaring TYPE attribute, the type of the node (e.g. BOX, SHPERE etc). This extension is supported by the viewers, which deal with the definition and creation of the areas. The UUIDs of the areas are sent to the annotating component, where four different modules enable the creation of annotations:

1. Annotation Module: it generates the UI of the annotating component, it receives areas from the viewers within the IVB and it allows the user to create new annotations. The user can select from different annotation templates based on the CIDOC definitions;
2. Data Collector: the data which was entered by the user during the annotation process is collected by this module. All data will be extracted from the input fields and the data collector converts it in the corresponding format;
3. TriG-File Generator: the RI requires the annotations to

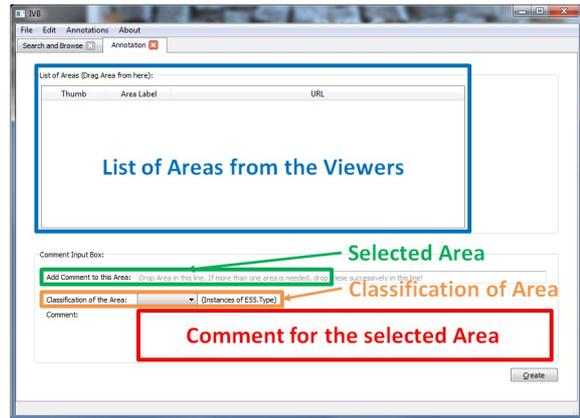


Figure 9: Annotating component interface for commenting.

- be described in the TriG format; therefore this module converts the user annotation input into the TriG format;
4. RI Ingesting Module: this module ingests the TriG files into the RI, in order to enrich the semantic network and to be retrieved by other clients or components (e.g. viewing).

Figure 9 shows the interface of the annotating component. The interface collects the areas generated by the different viewers and the user decides, which area(s) will be used for the desired annotation. Figure 9 illustrates the interface for commenting an area, in this case the area can be classified (e.g. damage) and the corresponding comment can be entered. The annotating component also provides a similar interface for building relationships between multimedia objects, where the user selects, which are the *source* and *destination* areas for a defined relationship.

6. Conclusions and Future Work

The exploration and enrichment of digital 3D Cultural Heritage collections is an active field of research in the scientific community. We proposed the interface and the underlying semantic models for exploring and enriching such collections, by means of integrating searching, browsing, viewing and annotating. Our contribution is a breakthrough for the current scholarly work in the CH domain, it will enable the study, conservation and restoration of the world's digitized Cultural Heritage based on the analysis and understanding of 3D objects. Although the development of this work has been assisted by a multidisciplinary team of CH professionals, metadata experts and computer graphics scientists, the IVB has not been tested in a professional environment. Hence, we will start a testing phase with domain experts toward its maturing and further development. Additionally, we will incorporate more functionalities, such as shape- and material-based search, image browsing, 3D profiling and measuring, and additional annotation models.

Acknowledgements

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 231809 (IST IP "3D-COFORM").

References

- [ADO09] ADOBE: Adobe Developer Connection - Acrobat 3D Developer Center. <http://www.adobe.com/devnet/acrobat3d.html>, 2009. 2
- [ARSF07] ATTENE M., ROBBIANO F., SPAGNUOLO M., FALCIDIENO B.: Semantic annotation of 3D surface meshes based on feature characterization. In *Proceedings of the semantic and digital media technologies 2nd international conference on Semantic Multimedia* (Berlin, Heidelberg, 2007), SAMT'07, Springer-Verlag, pp. 126–139. 2
- [BBC*04] BARACCHINI C., BROGI A., CALLIERI M., CAPITANI L., CIGNONI P., FASANO A., MONTANI C., NENCI C., NOVELLO R. P., PINGI P., PONCHIO F., SCOPIGNO R.: Digital reconstruction of the Arrigo VII funerary complex. In *VAST 2004* (Bruxelles, BE, Dec. 7-10 2004), Y. Chrysanthou K. Cain N. S., Niccolucci F., (Eds.), Eurographics, pp. 145–154. 2
- [CDG*09] CROFTS N., DOERR M., GILL T., STEAD S., STIFF M.: *Definition of the CIDOC Conceptual Reference Model*. Tech. rep., Mar 2009. http://www.cidoc-crm.org/docs/cidoc-crm_version_5.0.1_Mar09.pdf
- [CGF09] CHEN X., GOLOVINSKIY A., FUNKHOUSER T.: A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug. 2009). 2
- [CGG*04] CIGNONI P., GANOVELLI F., GOBBETTI E., MARTON F., PONCHIO F., SCOPIGNO R.: Adaptive tetrapuzzles: efficient out-of-core construction and visualization of gigantic multiresolution polygonal models. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 3 (Aug. 2004), 796–803. 6
- [CGG*05] CIGNONI P., GANOVELLI F., GOBBETTI E., MARTON F., PONCHIO F., SCOPIGNO R.: Batched multi triangulation. In *Proceedings IEEE Visualization* (Conference held in Minneapolis, MI, USA, October 2005), IEEE Computer Society Press, pp. 207–214. 6
- [DFPC08] DE FLORIANI L., PAPALEO L., CARISSIMI N.: A java3D framework for inspecting and segmenting 3D models. In *Proceedings of the 13th international symposium on 3D web technology* (New York, NY, USA, 2008), Web3D '08, ACM, pp. 67–74. 2
- [DT11] DOERR M., THEODORIDOU M.: CRMdig: A generic digital provenance model for scientific observation. *TaPP 2011: 3rd USENIX Workshop on the Theory and Practice of Provenance* (Jun 2011). 4
- [DTT*10] DOERR M., TZOMPANAKI K., THEODORIDOU M., GEORGIS C., AXARIDOU A., HAVEMANN S.: A repository for 3D model production and interpretation in culture and beyond. In *VAST10: The 11th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (Paris, France, 09/2010 2010), Eurographics Association, Eurographics Association, pp. 97–104. 1, 4
- [HG10] HUNTER J., GERBER A.: Harvesting community annotations on 3D models of museum artifacts to enhance knowledge, discovery and re-use. *Journal of Cultural Heritage* 11, 1 (2010), 81–90. 2
- [HSB*09] HAVEMANN S., SETTGAST V., BERNDT R., EIDE O., FELLNER D. W.: The arrigo showcase reloaded: towards a sustainable link between 3D and semantics. *ACM J. Comput. Cult. Herit.* 2 (July 2009), 4:1, 4:13. 2
- [JGD02] JUNG T., GROSS M. D., DO E. Y.-L.: Annotating and sketching on 3D web models. In *Proceedings of the 7th international conference on Intelligent user interfaces* (New York, NY, USA, 2002), IUI '02, ACM, pp. 95–102. 2
- [JIDG*99] JUNG T., LUEN DO E. Y., GROSS M. D., AUGENBROE G., EASTMAN C., JUNG T., LUEN DO E. Y., GROSS M. D.: Immersive redlining and annotation of 3D design models on the web. In *8th International Conference on Computer Aided Architectural Design Futures* (1999). 2
- [KDG*08] KASIK D., DIETRICH A., GOBBETTI E., MARTON F., MANOCHA D., SLUSALLEK P., STEPHENS A., YOON S.-E.: Massive model visualization techniques: course notes. In *ACM SIGGRAPH 2008 classes* (New York, NY, USA, 2008), ACM, pp. 1–188. 6
- [MET11] METS Schema, Metadata Encoding and Transmission Standard. <http://www.loc.gov/standards/mets/>, 2011. 6
- [OAM11] Open Annotation Model. <http://www.openannotation.org>, 2011. 5
- [OSG11] OpenSG: Portable scenegraph system to create realtime graphics programs. <http://www.opensg.org/>, 2011. 3
- [PBH*10] PAN X., BECKMANN P., HAVEMANN S., TZOMPANAKI K., DOERR M., FELLNER D. W.: A distributed object repository for cultural heritage. In *VAST10: The 11th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (Paris, France, 09/2010 2010), Eurographics Association, Eurographics Association, pp. 105–114. 1
- [PG11] PITTARELLO F., GATTO I.: 'toboa-3d: an architecture for managing top-down and bottom-up annotated 3d objects and spaces on the web. In *Proceedings of the 16th International Conference on 3D Web Technology* (New York, NY, USA, 2011), Web3D '11, ACM, pp. 57–65. 2
- [Pon08] PONCHIO F.: *Multiresolution structures for interactive visualization of very large 3D datasets*. PhD thesis, Clausthal University of Technology, Dec. 2008. 6
- [QT11] QT: Cross-platform application and UI framework. <http://qt.nokia.com/>, 2011. 3
- [REMA09] RODRIGUEZ-ECHAVARRIA K., MORRIS D., ARNOLD D.: Web based presentation of semantically tagged 3d content for public sculptures and monuments in the uk. In *Proceedings of the 14th International Conference on 3D Web Technology* (New York, NY, USA, 2009), Web3D '09, ACM, pp. 119–126. 2
- [Sha08] SHAMIR A.: A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6 (2008), 1539–1556. 2
- [Sie11] SIEMENS: PLM - Product Lifecycle Management. <http://www.plm.automation.siemens.com/>, 2011. 2
- [TTD*10] THEODORIDOU M., TZITZIKAS Y., DOERR M., MARKETAKIS Y., MELESSANAKIS V.: Modeling and Querying Provenance by Extending CIDOC CRM. *Distributed and Parallel Databases* (2010). 4
- [YSGM04] YOON S.-E., S ALOMON B., GAYLE R., MANOCHA D.: Quick-vdr: interactive view-dependent rendering of massive models. In *ACM SIGGRAPH 2004 Sketches* (New York, NY, USA, 2004), SIGGRAPH '04, ACM, pp. 22–. 6