SCULPTEUR: Multimedia Retrieval for Museums

Simon Goodall[1], Paul H. Lewis[1], Kirk Martinez[1], Patrick A.S. Sinclair[1], Fabrizio Giorgini[2], Matthew J. Addis[3], Mike J. Boniface[3], Christian Lahanier[4], and James Stevenson[5]

[1] Electronics and Computer Science, University of Southampton, SO171QP, UK. {sg,phl,km,pass}@ecs.soton.ac.uk
[2] Giunti Interactive Labs, Sestri Levante, 16039, Italy. f.giorgini@giuntilabs.it
[3] IT Innovation Centre, Southampton, SO167NP, UK. {mja,mjb}@it-innovation.soton.ac.uk
[4] Centre de Recherche et de Restauration des Musées de France, Paris, France. Christian.Lahanier@culture.gouv.fr
[5] Victoria and Albert Museum, South Kensington, London. J.Stevenson@vam.ac.uk

Abstract. The paper describes the prototype design and development of a multimedia system for museums and galleries. Key elements in the system are the introduction of 3-D models of museum artefacts together with 3-D as well as 2-D content based retrieval and navigation facilities and the development of a semantic layer, centred on an ontology for museums, which aims to expose the richness of knowledge associated with the museum collections and facilitate concept based retrieval and navigation integrated with that based on content and metadata. Interoperability protocols are designed to allow external applications to access the collection and an example is given of an e-Learning facility which uses models extracted to a virtual museum.

1 Introduction

Museums, galleries and other cultural heritage institutions are finding it beneficial, to maintain and exploit multimedia representations of their collections using database indexing and retrieval technology. Their requirements present researchers with a growing series of challenges, particularly as the range of representations is expanding to include, for example, 3-D models and digital videos and the technologies available are evolving rapidly, notably in the area of web services and the semantic web.

This paper provides a progress report on a project which is designed to meet some of these new challenges, especially in the area of retrieval and navigation of multimedia information in the context of emerging semantic web technology. The project, SCULPTEUR[2,1], involves five major European galleries, the Uffizi, the National Gallery and the Victoria and Albert Museum in London, the Musée de Cherbourg and the Centre de Recherche et de Restauration des Musées de France (C2RMF) which is the Louvre related art restoration centre. Each of

P. Enser et al. (Eds.): CIVR 2004, LNCS 3115, pp. 638–646, 2004.

these maintains a substantial digital archive of its collections. Other technical partners include Centrica in Italy and GET-ENST in Paris. The project builds on the work of an earlier museum database project, ARTISTE[11,3].

One of the goals of SCULPTEUR is to extend the retrieval and navigation facilities of the digital museum archives to 3-D multimedia objects. Increasingly museums are recognising the value of 3-D visualisation of their artefacts not only for researchers, curators and historians but also potentially as an information source for a wider public and as a basis for e-learning and commercial activity. A second new goal of the project is concerned with the perceived benefit of structuring and integrating the knowledge associated with the museum artefacts in an ontology. Metadata associated with the artefacts is being mapped to the ontology to form an integrated knowledge base and graphical tools are being developed to provide browsing of the concepts, relationships and instances within the collections. Integrated concept, metadata and content based retrieval and navigation facilities are being implemented to explore the knowledge base. Other goals include the development of a web agent to locate missing metadata and also exploitation of the system by development of an e-learning product to make use of it.

In section 2 related work is presented and in section 3 we describe some of the specific needs of our users, identified at the start. Sections 4 and 5 discuss issues relating to building the ontology based architecture and multimodal, multimedia retrieval respectively. Section 6 presents an example of interoperability using an e-Learning example and finally, in section 7, some conclusions and challenges are described.

2 Related Work

The SCULPTEUR project is related to and builds on previous work in several fields. A major source of inspiration comes from previous work on the ARTISTE and MAVIS projects[10,11] and a large body of other published work on content based image retrieval systems[12]. 3-D model capture is an important element and one of the partners, GET-ENST, has developed techniques for accurately generating 3-D models from multiple views[19]. Various authors have published algorithms for 3-D model matching using a variety of feature vectors extracted from mesh based 3-D representations based on for example, 3-D Hough transforms[18], 3-D moments[17] and surface features such as chord distributions[14] and radial axis distributions[15]. The work from Princeton has been particularly influential in this area and a recent paper on 3-D benchmarking compares a range of matching algorithms[16] in terms of retrieval performance using the Princeton Benchmark data set. The idea for the semantic layer and knowledge base draws on previous work at Southampton[20] and semantic web technology[6]. Various other groups have reported the use of ontologies for image annotation.

3 Identifying User Needs

The ARTISTE project combined metadata based retrieval with both general purpose content based retrieval tools such as colour, spatial colour and texture matching with specialist content based facilities to meet some specific needs of the participating museums. These included, for example, sub-image location[21], low quality query image handling[22] and canvas crack analysis and retrieval[9]. The ARTISTE system was the starting point for our work in SCULPTEUR and the new needs of the museums were identified, particularly in relation to their increasing use of 3-D digital representations of their artefacts. These included the ability to compare and retrieve 3-D objects on the basis of size, colour, texture and 3D shape. Examples of more specialised requirements included the ability to retrieve objects by sub-parts and by detail measurements, the ability to classify objects using these features and the ability to correlate for example figurines and moulds in which they may have been made by the external and internal 3-D profiles respectively. All these requirements and others relating to the forms of retrieval and navigation required and the ability to integrate with external applications such as e-Learning packages, impose substantial demands on the architecture and functionality required in the new system.

4 Building the Knowledge Base

The current SCULPTEUR architecture is shown in figure 1. The system is implemented as a web based client-server application and the server-side repositories hold the raw multimedia objects, feature vectors for content-based retrieval, textual metadata and the ontologies which are implemented using Protégé[8] and held as RDF[13]. Tools have been developed for importing new or legacy objects or mapping directly to existing collections in situ. Above the repositories, components of the architecture provide semantic integration between concepts and relations in the ontology and metadata and media objects in the repositories plus search, browsing and retrieval services for both SCULPTEUR driven activity and eventually for externally invoked activity via the interoperability protocols, SRW[5] and OAI[4]. The users' desktop is essentially a standard browser augmented with 2-D and 3-D model viewers, facilities for query formulation such as colour pickers, and uploading facilities for query models.

The starting point for the development of the ontology was the conceptual reference model (CRM)[7] developed by the documentation standards committee of the International Council of Museums (CIDOC). The aim of the CRM is to support the exchange of relevant cultural heritage information across museums, based on shared semantics and common vocabularies. Working closely with the museums is necessary to extend the core CRM to enhance the particular areas relevant to each museum and to develop mappings between museum metadata values and concepts in the ontology. The problems and difficulties of developing the semantic layer in this way should not be underestimated. The importing and mapping of legacy museum data, both in terms of concepts and instances

in the ontology is a complex manual process involving collaboration between technologists, domain experts and CRM experts. Problems occur at various levels from establishing coherent semantics at the highest level to interpreting or unravelling obscure coding and formats at the lowest. Tools are being developed to use the mappings to automatically build associations between media objects and concepts, making them instances of the concepts in the ontology.

Using the ontology to develop a semantic layer in this way, provides a bridge across the semantic gap, facilitating search for media objects via the concepts and relationships in the ontology as well as via more usual content and metadata based searching. In addition, the semantic layer aims to expose the richness of information surrounding the media objects themselves, allowing the system to move from one in which search is solely focussed on the media objects to one in which any of the entities in the semantic layer can become the focus of the investigation. Thus, the system will not only allow queries of the type "Find me all 3-D objects in the collection with a shape similar to this query object" but also queries focussed on other entities such as "Find me all countries that produced artists working in the 17th Century."

The architecture also includes a system ontology which captures the concepts and relationships associated with the system itself. The system consults the ontology to determine which tools and components to use for a particular task and will eventually be able to expose its facilities more elegantly to external agents wishing to use them.

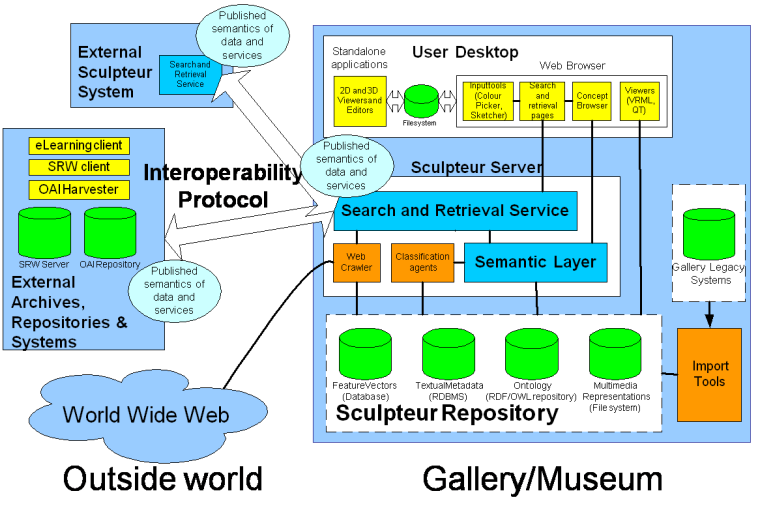

Fig. 1. The SCULPTEUR Architecture

5 Multimodal Multimedia Retrieval

In addition to the 2-D content-based image retrieval (CBR) tools of ARTISTE, the prototype SCULPTEUR system is designed to support content-based retrieval of 3-D object models by model matching. In one sense 3-D content-based retrieval is more straight forward than 2-D as the objects are explicitly represented rather than embedded in a pixel (or voxel) matrix. However, mesh based representations are not immediately amenable to comparison for CBR applications and 3-D CBR techniques compare the similarity of the shape of the object as represented by some feature vector extracted from the mesh.

So far in SCULPTEUR several published 3-D matching algorithms have been implemented and integrated with retrieval facilities in the first prototype to provide 3-D CBR. These include the D2 shape distribution descriptors from the Princeton Shape Retrieval and Analysis Group[14], the histogram descriptors from Paquet and Rioux[15] and the Area to Volume Ratio descriptor[23] which is a single statistic giving the ratio of the surface area of the model to its enclosed volume. The D2 descriptor records the distribution of distances between random points on the surface of the model and is rotation and translation invariant and robust to changes in mesh resolution. In our implementation, a 64 bin histogram was used for D2. There are three versions of the histogram descriptors of Paquet and Rioux . They define a cord as the vector between the centre of mass of an object and a point on its surface. Their first histogram records the distribution of the cord lengths for all points within the mesh. The other two variations record the distribution of angles between cords and the first and second principal axis respectively. Each of these was implemented as a 16 bin histogram. Before including them in the prototype, an evaluation was made of the five 3-D algorithms for CBR. Using the Princeton Benchmark[16] base dataset, (training group) consisting of 907 models representing about 90 object classes, precision-recall graphs of the five algorithms were created. They are shown in figure 2 together with the precision-recall graph for random retrieval. It can be seen that the D2 descriptor gives the best retrieval results in terms of precision-recall and as expected, the more basic Area to Volume Ratio descriptor gives the poorest retrieval results. The three Paquet and Rioux histograms are in between but the histogram of the angle between the cord and principal axis version gives the best retrieval results of these three.

An example of retrieval results in SCULPTEUR is shown in figure 4. The top ten best matches are shown for a 3-D retrieval using a vase as the query and the D2 descriptor for matching. The first match is the query object as expected. The test dataset used here consists of around 300 models from both the museum partners and from other collections.

The concept browser, shown in figure 3, provides graphical navigation of the semantic layer. Concepts in the ontology are represented by the nodes in the graph and the relationships between concepts are represented by the graph edges. In the current prototype, the feature vectors are only associated indirectly with concepts via the media objects, but the interface does allow the combination of concept based retrieval with the other retrieval modes.

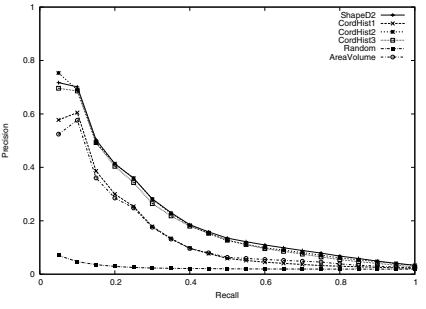

Fig. 2. Precision/Recall for shape descriptors

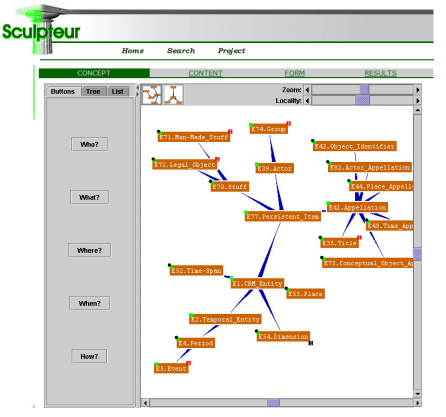

Fig. 3. Concept Browser

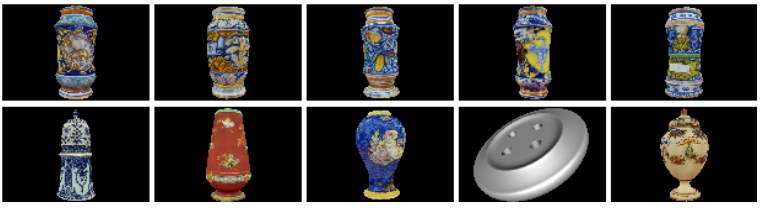

Fig. 4. Top ten results for a query vase (first image is also the query)

Certain metadata fields contain complex or non-atomic values that express relationships between records, and the concept browser is able to display these fields with the graphical interface in a new window.

6 Interoperability: An E-learning Application

One of the goals of the SCULPTEUR system is to provide an easy interface to external applications which may wish to utilise the media objects or other knowledge in the system. As an example of such an application, Giunti, Interactive Labs, the project coordinators, are integrating their e-Learning package using the interoperability protocol, SRW, implemented in the system. The integration is motivated by the recent increased interest by cultural institutions in reusable multimedia components for learning (called Cultural Learning Objects, CLO) and new technologies, capable of online learning contents delivery and management. The result is a content authoring tool able to create and manage 3D virtual learning environments of Cultural Learning Objects.

The user interface for the Giunti system, shown in figure 5, assists curators, instructional designers and educators to build virtual exhibitions of 3-dimensional Cultural Learning Objects, define learning paths and package the 3D virtual learning environment according to the new e-learning specifications defined by IMS (Instructional Management System[24] which eases the exchange of cultural contents and educational material from one museum to another.

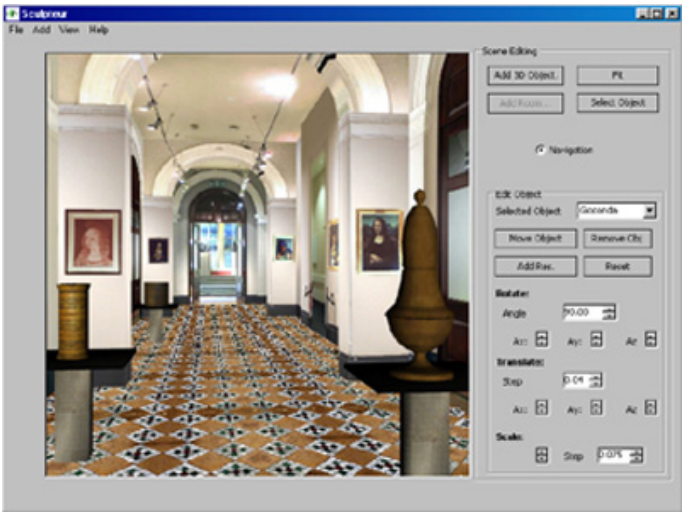

Fig. 5. User interface for the Giunti e-Learning System

According to the defined learning path, the tool automatically generates the SCORM run-time environment APIs[25] for the tracking of users' actions and the communication with a Learning Management System and embeds these calls in the generated VRML. This allows a Learning Management System to control the learning paths on-the-fly in accordance with the experiences made by the user while he/she navigates the 3D virtual environment. In order to establishing an appropriate virtual environment the e-Learning package will query the SCULPTEUR system for appropriate artefacts to place in the virtual museum and these will be delivered through the SRW interface.

7 Conclusions and Some Outstanding Challenges

The paper has presented a prototype multimedia system which we are in the process of developing. The project aims to capitalise on emerging semantic web technologies and novel 3-D retrieval to provide museums and galleries with more versatile facilities for exploring and exploiting their digital collections. The project is on-going and several major challenges remain.

The ontology, at the heart of the semantic layer, serves a number of purposes. Notably, it provides a basis for interoperability between digital libraries and, for example, with e-learning facilities. Through the concept browser, it also aims to provide a high level navigation and retrieval interface for the collections. However, to be a realistic shared conceptualisation of the museum domain, the ontology is of necessity, a large and complex representation. One of the significant challenges is to make a more intuitive, easy to use interface for exploration and navigation which still exposes the richness of the collections. The ontology also needs closer integration with the content and metadata based retrieval to provide enhanced retrieval capabilities, for example by automatic query expansion through the semantic layer. We have also developed a prototype web crawler which seeks for missing information from the ontology on the web. In the absence of widespread uptake of semantic web technology, the information extraction process is mainly via natural language processing and many problems and opportunities for improvement in functionality remain here.

The first prototype of SCULPTEUR is currently under evaluation by the five galleries involved in the consortium. Feedback from this process will enable us to continue to evolve the system towards a more useful exploration and retrieval facility for museum collection managment.

Acknowledgements. The authors wish to thank: the European Commission for support through the SCULPTEUR project under grant IST-2001-35372. Thanks also to our collaborators including F. Schmitt and T. Tung of GET-ENST, R. Coates of the V&A museum, J. Padfield of the National Gallery, R. Rimaboschi of the Uffizi, J. Dufresne of the Musée de Cherbourg and M. Cappellini of Centrica for many useful discussions, use of data and valuable help and advice; Patrick Le Boeuf of the Bibliothèque Nationale de France for assistance with mapping to the CRM; TouchGraph (www.touchgraph.com) for software used in the concept browser; Hewlett Packard's Art & Science programme for the donation of server equipment, the ARCO Consortium[26] for the VRML model of the VAM Art Decò corridor and other models.

References

1. Addis, M., Boniface, M., Goodall, S., Grimwood, P., Kim, S., Lewis, P., Martinez, K., Stevenson, A.: SCULPTEUR: Towards a new paradigm for multimedia museum information handling. 2nd International Semantic Web Conference, pp 582–596, October 2003.
2. SCULPTEUR IST-2001-35372 http://www.sculpteurweb.org
3. ARTISTE IST-1999-11978 http://www.artisteweb.org/
4. Open Archives Initiative http://www.openarchives.org/
5. ZING Search and Retrieve Web service http://www.loc.gov/z3950/agency/zing/srw/
6. The Semantic Web http://www.semanticweb.org/
7. Crofts, N., Dionissiadou, I., Doerr, M., Stiff, M.: Definition of the CIDOC Object-Oriented Conceptual Reference Model, V.3.1. July 2001

8. Gennari, J., Musen, M., Fergerson, R., Grosso, W., Crubzy, M., Eriksson, H., Noy, N., Tu, S.: The Evolution of Protégé: 2002. Technical Report, SMI-2002-0943
9. Abas, F., and K. Martinez: Craquelure Analysis for Content-Based Retrieval. IEEE DSP, pp 111–114, 2002 conference. July 2002.
10. Dobie, M. R., Tansley, R. H., Joyce, D. W., Weal, M. J., Lewis, P. H., Hall W.: A Flexible Architecture for Content and Concept Based Multimedia Information Exploration. Proc. The Challenge of Image Retrieval, Newcastle, 1999, pp 1–12.
11. Addis, M., Boniface, M., Goodall, S., Grimwood, P., Kim S., Lewis, P., Martinez, K., Stevenson, A.: Integrated image content and metadata search and retrieval across multiple databases, Proc. of Second Int. Conf. on Image and Video Retrieval 2003, pp. 91-100, Illinois, USA.
12. Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., Jain. R.: Content-Based Image Retrieval at the End of the Early Years, PAMI volume 22 of 12, pages 1349-1380 2000.
13. RDF Resource Description Framework http://www.w3c.org/RDF/
14. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Matching 3D models with shape distributions, in Proc. Shape Modelling International, IEEE Press, Genova, Italy, 2001, pp. 154–166.
15. Paquet, E., Rioux, M.: Nefertiti: a Query by Content Software for Three-Dimensional Models Databases Management, in Proc.Int Conf. on Recent Advances in 3-D Digital Imaging and Modeling, May 12-15, 1997, pp. 345–352.
16. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The Princeton Shape Benchmark, to appear in Shape Modeling International, Genova, Italy, June 2004
17. Elad, M., Tal, A., Ar., S.: Content Based Retrieval of VRML Objects - An Iterative and Interactive Approach. Proceedings of the sixth Eurographics workshop on Multimedia 2001. Manchester, UK, pp. 107–118
18. Zaharia, T., F.Preteux, F.: Hough transform-based 3D mesh retrieval Proceedings SPIE Conference 4476 on Vision Geometry X, San Diego, CA, August 2001, pp. 175–185.
19. Esteban, C. E., Schmitt,F.:Silhouette and Stereo Fusion for 3D Object Modeling. Fourth International Conference on 3-D Digital Imaging and Modeling , pp. 46–54 October 06-10, 2003, Banff, Alberta, Canada.
20. Lewis, P. H., Davis, H. C., Dobie, M. R. and Hall, W. : Towards Multimedia Thesaurus Support for media-based Navigation. In Proceedings of Image Databases and Multi-Media Search, Series on Software Engineering and Knowledge Engineering 8, pp 111–118. 1997
21. Chan, S., Martinez, K., Lewis, P., Lahanier, C. and Stevenson, J. : Handling Sub-Image Queries in Content-Based Retrieval of High Resolution Art Images,. Proceedings of International Cultural Heritage Informatics Meeting 2, pp. 157-163, 2001.
22. Fauzi, M. F. A. and Lewis, P. H. : Query by Fax for Content-Based Image Retrieval. Proceedings of International Conference on the Challenge of Image and Video Retrieval, LNCS vol. 2383,pp. 91-99, London, United Kingdom 2002.
23. C. Zhang and T. Chen, Efficient Feature Extraction for 2D/3D Objects: in Mesh Representation, ICIP 2001, Thessaloniki, Greece, pp 935–938, 2001.
24. IMS Global Learning Consortium, Content Package Specifications, http://www.imsglobal.org/content/packaging/index.cfm
25. Advanced Distributed Learning, Sharable Content Object Reference Model, http://www.adlnet.org/index.cfm?fuseaction=AboutSCORM
26. ARCO project, http://www.arco-web.org