Content-Based Indexing and Retrieval of Cultural Heritage Data: An Integrated Approach to Documentation with Application to the EROS

Paquet, Eric; Lahanier, C.; Pitzalis, D.; Peters, S.; Aitken, G.; Viktor, H.L.

Publisher's version / Version de l'éditeur:

The 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2006) [Proceedings], 2006

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

NRC · CNRC

*Content-Based Indexing and Retrieval of Cultural Heritage Data: An Integrated Approach to Documentation with Application to the EROS **

Paquet, E., Lahanier, Cl, Pitzalis, D., Peters, S., Aitken, G., Viktor, H.L.
November 2006

* published at The 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2006). Nicosia, Cyprus. October 30 – November 4, 2006. NRC 48792.

Canada

The 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2006)
M. Ioannides, D. Arnold, F. Niccolucci, K. Mania (Editors)

Content-based Indexing and Retrieval of Cultural Heritage Data: an Integrated Approach to Documentation with Application to the EROS Database

E. Paquet[1,3] C. Lahanier[2†] D. Pitzalis[2] G. Aitken[2] S. Peters[1] and H. L. Viktor[3]

[1]Visual Information Technology, National Research Council, Ottawa, Canada
[2]Centre de Recherche et de Restauration des Musées de France, Paris, France
[3]School of Engineering and Information Technology, University of Ottawa, Ottawa, Canada

Abstract

Over the last few decades, Cultural Heritage documentation has been characterized by the massive use of digital media. Recently, the use of three-dimensional scanner technologies has provided us with the opportunity to obtain an unambiguous body of information characterizing the three-dimensional shapes of the artefacts. These vast repositories are now structured in databases, for easy access. Such databases contain not only the artefacts, but also relevant information such as restoration reports, data regarding quantitative analysis, chemical formulae, etc. It follows that storing such information is not enough. Rather, it should be indexed in order to be searched and retrieved easily and rapidly. In addition, the data should be preserved as technologies evolve over time, in order to ensure long-term preservation and access. This paper presents a framework for indexing and retrieval of 2D and 3D Cultural Heritage data. In our approach, novel archiving and indexing techniques, developed by the National Research Council of Canada, are employed. We present the results as applied to the EROS (European Research Open System) Database of the C2RMF. This database consists of an impressive collection of scientific and technical data about paintings and artefacts found in all the museums of France. Our results indicate that our content-based approaches are able to accurately index and retrieve diverse images and 3D objects, based on the artefacts as well as their fragments. That is, using for example a fragment of a picture, we are able to retrieve the correct image even in conditions where lighting, orientation and the surroundings of the reference are different. The content-based retrieval system is also able to retrieve different views of the same object, e.g. of a Chalcidian amphora. In addition, our approach is able to find groups of similar images or objects, such as white figurines from the same period or 3D scans of an Anadyomene Venus.

Categories and Subject Descriptors (according to ACM CCS): H.3.1 [Information storage and retrieval]: Indexing methods

1. Introduction

Cultural Heritage applications are now characterized by their massive utilisation of digital media [LAC*04]. This has been employed to document sites, artefacts and restorations. Up to recently, such documentation was mostly based on pictures, reports and physical and chemical analysis. In recent years it has been realized that to describe a work of art only with pictures is not enough: an unambiguous body of information characterizing the three-dimensional shape of the artefacts is also needed, for example, to evaluate the deterioration of the shape over time. That is why 3D scanning has become a common practice in Cultural Heritage [LDP02, LFJ*05].

With the improvement of these acquisition techniques [GBT*02] and technology, the amount of information, both in terms of required storage space and number of items has become enormous. For that reason, it has become necessary to structure this large amount of information in databases. To structure data is not enough, we need

† R. Pillay[2]

to index them in order to search and retrieve easily and rapidly [FP02].

Text-based indexing has been making tremendous progress over the last few years and we refer the reader to the literature for a review on the subject [AML*05, PPLA05]. In the present paper, we would like to concentrate on content-based indexing and retrieval of images and 3D models. Our motivation is twofold. Firstly, as many of our examples will show, words are not enough to formulate many queries which, if formulated in terms of pictures and 3D models, are self-evident. Secondly, even if an adequate text-based description is available e.g. historical information, images and models constitute a valuable complement. In addition, if no textual information is available, content-based indexing can rapidly be created since, as opposed to its textual counterpart that need human intervention and judgement, the creation of the indexes is entirely automatic.

Our paper is organized as followed. After some general considerations on images, an algorithm for content-based indexing and retrieval is presented. We will then address the issue of content-based indexing and retrieval to 3D artefacts based on three-dimensional shape. Such algorithms are applied to the EROS Database [ALPP05, ALPP06, PPLA05] of the C2RMF (Centre of Research and Restoration of the France Museums). Some particularly interesting and relevant results are later discussed. The combination of all the above constitute an integrated approach to Cultural Heritage documentation.

2. General Considerations on Images

Images and models are of the outmost importance in virtual collections. They are (and will remain in the foreseeable future) the easiest, fastest and most economical mean for creating virtual collections. Furthermore, most three-dimensional models are covered with textures. The texture constitutes an important visual descriptor for the model under consideration and convey essential historical, artistic and archaeological information.

Images are difficult to describe [TS06]. They convey a large amount of complex and ambiguous information. The ambiguity is due to the fact that an image is a two-dimensional projection of the three-dimensional world and due the fact that the illumination of this world is arbitrary and cannot be controlled. Because of this ambiguity and complexity, it is difficult to segment images and to understand them. For the above-mentioned reasons, we propose a statistical approach in which the overall composition of the image is described in an abstract manner.

2.1. Indexing and Retrieval of Images

We now depict the algorithms developed by the National Research Council of Canada. The colour distribution of each image is described in terms of hue and saturation. This colour space (HSV) imitates many characteristics of the human visual system: the hue corresponds to our intuition of colour (for example red, green or blue), while saturation corresponds to the colour strength (for example light red or deep red).

Next, a set of points is sampled from the image. A quasi-random sequence generates the points. In the present implementation, the Sobol sequence has been selected. Each point of this sequence becomes the centre of a structuring element. For each centre position, the pixels inside the corresponding structuring element are extracted and the associated hue and saturation images are calculated. The statistical distribution of the colours within the window is characterized by a bi-dimensional histogram. The first dimension of this histogram [FTT05] corresponds to the hue or the saturation quantified on a discrete and finite number of channels. The second dimension corresponds to the relative proportion of each channel within the window. This two-dimensional histogram is computed and accumulated for each point of the sequence, i.e. the current histogram is the sum of the histograms at the current and at the previous position. From this process, a compact descriptor or index is obtained.

This index provides an abstract description of the composition of the image i.e. of the local distribution of colours throughout the image. This is very important. This index does not represent a global description of the image nor is it based on a particular segmentation scheme. Instead, it characterizes the statistics of colour distribution within a small region that is moved randomly over the image. Consequently, there are no formal relations in between the different regions, which means that the different components of a scene can be combined in various ways while still being identified as the same scene. That is why that algorithm is robust against occlusion, composition, partial view and viewpoint. Nevertheless, this approach provides a good level of discrimination.

As we know, an image is worth a thousand words, which means that it is difficult to describe an image based solely on words. For that reason, our retrieval approach is based on the so-called “query by example” or “query by prototype” paradigm. To this end, we created a search engine that can handle such queries. In order to initiate a query, the user provides an image or prototype to the search engine. This prototype is described or indexed and the later is compared with a metric to a database of pre-calculated indexes, which correspond to the images of the virtual collection. The search engine finds the most similar images with respect to the prototype and displays them to the user. The user then acts as an expert: he chooses the most meaningful image from the results provided by a search engine and reiterates the query process from the chosen image. The process is repeated until convergence is achieved.

2.2. Indexing and Retrieval of 3D Objects

The indexation of three-dimensional artefacts differs fundamentally from the indexation of images [IJL*05, TV04]. If the three-dimensional information has been acquired accurately at a sufficiently high resolution, the three-dimensional geometry constitutes an unambiguous body of information in the sense that there is a one-to-one correspondence between the virtualized geometry and the physical geometry of the artefacts. As explained in the previous section, the situation is entirely different for images. Shape also constitutes a language of its own right. In addition to verbal language, humanity has developed a common shape language. This is particularly evident in fields like art and architecture. For that reason, the "query by prototype" approach is a powerful paradigm for the retrieval of similar artefacts. As far as the overall structural design is involved, the three-dimensional artefacts retrieval system is very similar to its image counterpart: the artefacts of the collection are indexed offline and a database of indexes is created. In order to interrogate this database, the query is initiated with a prototype artefacts. From the proto-artefacts, an index is calculated and compared with the help of a metric to the indexes of the collection in order to retrieve the most similar objects in terms of three-dimensional shape. As stated before, the user can act as an expert in order to reiterate the process until convergence.

Consequently, the main differentiation between the two systems (image versus 3D) is the index. We now describe our algorithm for three-dimensional artefacts description. We assume that each artefacts has been modelled with a mesh. This is a non-restrictive representation for virtualized artefacts since most acquisition systems generate such a representation. In the present case, a triangular mesh representation is assumed. If the mesh is not triangular, the mesh is tessellated accordingly. Our objective is to define an index that describes an artefacts from a three-dimensional shape point of view and that is translation, scale and rotation invariant. The later invariants are essential because the artefact can have an arbitrary location and pose into space.

The algorithm can be described as follows. The centre of mass of the artefact is calculated and the coordinates of its vertexes are normalized relatively to the position of its centre of mass. Then the tensor of inertia of the artefact is calculated. This tensor is a 3 x 3 matrix. In order to take into account the tessellation in the computation of these quantities, we do not use the vertexes per se but the centres of mass of the corresponding triangles; the so-called tri-centres. In all subsequent calculations, the coordinates of each tri-centre are weighted with the area of their corresponding triangle. The later is being normalized by the total area of the artefact, i.e. with the sum of the area of all triangles. In this way, the calculation can be made robust against tessellation, which means that the index is not dependent on the method by which the artefact was virtualized: a sine qua non condition for real world applications. In order to achieve rotation invariance, the Eigen vectors of the tensor of inertia are calculated. Once normalized, the unit vectors define a unique reference frame, which is independent on the pose and the scale of the corresponding artefact: the so-called Eigen frame. The unit vectors are identified by their corresponding Eigen values. The descriptor is based on the concept of a cord. A cord is a vector that originates from the centre of mass of the artefact and that terminates on a given tri-centre. The coordinates of the cords are calculated in the Eigen reference frame in cosine coordinates. The cosine coordinates consist of two cosine directions and a spherical radius. The cosine directions are defined in relation with the two unit vectors associated with the smallest Eigen values i.e. the direction along witch the artefact presents the maximum spatial extension. In other words, the cosine directions are the angles between the cords and the unit vectors. The radius of the cords are normalized relatively to the median distance in between the tri-centres and the centre of mass in order to be scale invariant. It should be noticed that the normalization is not performed relatively to the maximum distance in between the tri-centres and the centre of mass in order to achieve robustness against outliers or extraordinary tri-centres. From that point of view, the median is more efficient that the average. The cords are also weighted in terms of the area of the corresponding triangles; the later being normalized in terms of the total area of the artefact. The statistical distribution of the cords is described in terms of three histograms. The first histogram described the distribution of the cosine directions associated to the unit vector associated with the smallest Eigen value, the second one describe the distribution of the cosine directions associated with the unit vector associated with the second smallest Eigen value ant the third one describe the distribution of the normalized spherical radius as defined in the previous paragraph. The three histograms together constitutes the shape index of the corresponding artefact.

2.3. Application to the EROS Database

The C2RMF is a pioneer in applying new technologies in the field of Cultural Heritage. The activities of the C2RMF in the field of High Resolution imaging for Cultural Heritage started in 1989 with the high quality digitization of large size transparencies (photos and X-ray plates) via the Thomson-Broadcast flatbed scanner developed for the NARCISSE European project. Then we proceeded with the acquisition of direct digital imaging, by panoramic views of objects, by direct 3D acquisition of the surface of paintings and objects, by 3D reconstruction from panoramic views and by multispectral imaging from ultraviolet to infrared allowing us to reconstruct the colour for any illuminant.

All these techniques give us an enormous amount of data and information to organize and to exploit. This information is focused on scientific and technical data.

The EROS system is organized in several parts: the storage back-end, the relational database, the image server, the

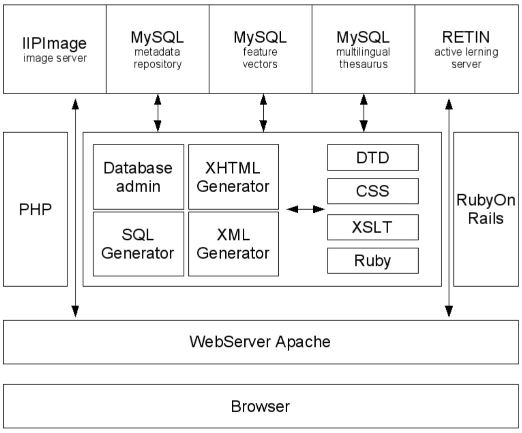

Figure 1: *EROS database organization.*

middle-ware and the web server. The data are stored on 15TB HP RAID 5 hard disk racks managed by a file server and consist of:

- meta-data related to 65,000 works of art, 200,000 high resolution images, 10,000 reports, 170,000 analysis, analytical reports, restoration reports, 6,600 conservation surveys, the chemical, structural, isotopic and molecular quantitative and qualitative analytical results and published papers;
- high definition digital images (some of them are gigapixel images);
- feature vectors for 2D and 3D image content recognition for automatic classification and image category retrieval (for different engines).

The EROS system is an Open Source project available under the GNU Public License (GPL). It is based on powerful and industry-leading free software.

In the following examples we compare the results obtained using the meta-data with the ones using “query by content”.

Finding 2D images from 3D models.

A snapshot of the 3D model of a Chalcidian amphora is used to query the database (Figure 2).

Our content-based recognition algorithm allows us to retrieve in the first screen the 36 2D images (from a panoramic view) of the same vase and then similar Greek vases. As 3D acquisition techniques are improving, more and more 3D models will be produced in place of high quality 2D images. As this algorithm is able to compare 3D models with 2D images, our existing image database will continues to be useful.

Finding the overview using a detail.

Lost pictures and slides are sometimes registered with wrong reference number. In this case content-based recognition can help us to retrieve the overview of the work of art.

Figure 2: *Chalcidian amphora - Louvre Museum, Paris, inv. E795. This amphora was made during the Archaic Greek Period (620-480 B.C.) and was found in the South of Italy*

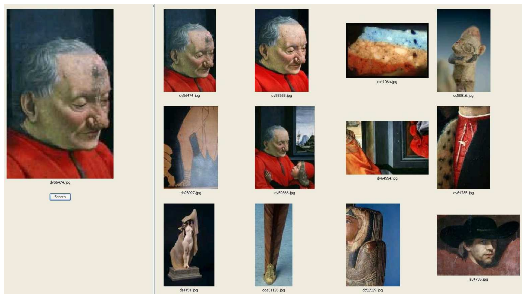

Figure 3: *PORTRAIT OF AN OLD MAN WITH A YOUNG BOY - GHIRLANDAIO Domenico (1449-1494), Louvre Museum, Paris, inv. RF266*

For example the detail showed in the left side of Figure 3 was compared to the contents in the database. Similar images made at various periods of time under different experimental conditions are retrieved first as well as the overview of the painting.

Style recognition.

Figure 4: *SITTER SEEN FROM THE FRONT - SEURAT Georges-Pierre (1859-1891), Louvre Museum, Paris, inv. RF1947-13*

When the meta-data related to the image content are not indexed, our algorithm is very useful in retrieving similar paintings having the same pictorial style. For example images characteristic of the "pointillism" style, which is the painting technique of, for example, Georges-Pierre Seurat, can be automatically retrieved using a detail of one of his paintings (Figure 4).

The first results are images representing part of the same painting and then images of paintings having the same pictorial style.

Figurine classification by type.

Figure 5: *ANADYOMENE VENUS - PRISCUS, France, Moulins, Anne de Beaujeu Museum, inv. 5.3.20*

5.500 Gallo-Roman figurines produced in workshops in France between 40-300 a.C. are stored in the EROS database. 500 3D models and several thousand images were compared. In Figure 5 a 2D image of a mould of "Venus", characteristic of the "Anadyomene" type, is used as a reference. Content-based recognition applied to flat images (Figure 5) gives results of groupings of similar object moulds and then figurines issued from these moulds. This example is typical of situations in which one wants to retrieve a certain group of pictures that are very similar but that do not necessarily correspond to the same artefact. In the following examples a 3D model of a statue (Figure 6) of a "Prudish Venus" and after a 3D model of a mould (Figure 7) of an "Anadyomene Venus" are used. We obtain an impressive level of coherence in the results.

The system was able to retrieve very similar Venuses irrespective of their orientation in space. Shape is a powerful retrieval paradigm for 3D models in Cultural Heritage.

Content-based recognition can be used for semi-automatic classification of ceramic production presenting similar artefact and can be used also to link signed moulds to figurines for attributing the production of an antique workshop.

Robustness of the algorithm.

A test of robustness in detecting images both before and after restoration was made using a shroud.

It appears that the content-based recognition algorithm is able to retrieve images corresponding to the same object in

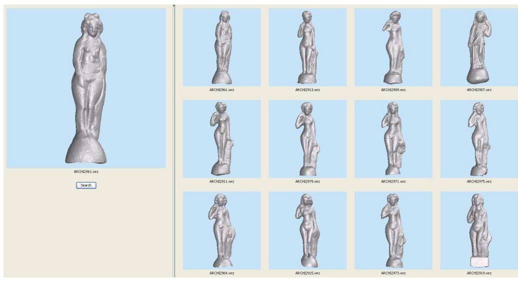

Figure 6: *PRUDISH VENUS - France, Moulins, Anne de Beaujeu Museum, inv. 5.7.6*

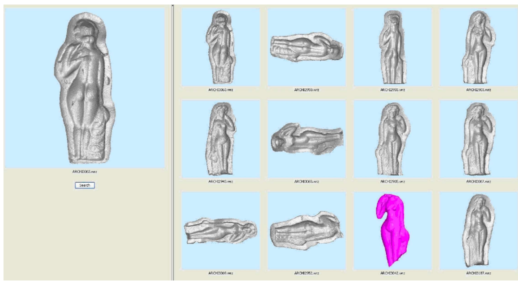

Figure 7: *ANADYOMENE VENUS - France, Moulins, Anne de Beaujeu Museum, inv. 5.3.33*

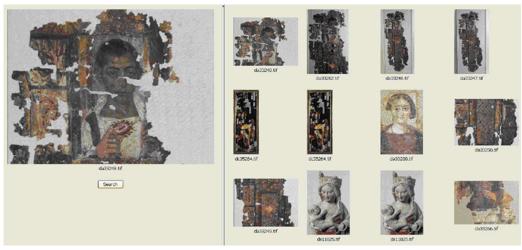

Figure 8: *Shroud - Louvre Museum, Paris, inv. AF6482. Made around 200-299 A.D. in Egypt*

different states of conservation (Figure 8). It means that in addition to the pictorial information, for example the painting, there is a significant amount of information that is related to the deteriorated textile. It is relatively easy for a human being to make an abstraction of the deterioration information and to solely concentrate on the pictorial one. For a content-based retrieval system, it is extremely difficult to handle such a situation. This is because the system does not know a priori which information is original and which is related to the deterioration. Nevertheless, we managed to retrieve many views (partial and complete) of the work of art. This demonstrates that the algorithm is maintaining a good balance between colour information, which here corresponds to the pictorial information, and textural information, which corresponds to the deterioration.

3. Conclusions

In this paper, we have presented novel indexing algorithms developed by the National Research Council of Canada for 2D and 3D digital data for Cultural Heritage. In particular, we have applied the proposed approaches to the heterogeneous EROS Database of the C2RMF.

Our results have shown the efficiency of our algorithms. In many situations, content-based retrieval has proved itself to be, not only a complement to text-based retrieval, but as a sine qua non condition for efficient retrieval. The retrieval of the "pointillist" paintings from a detail as start point is a spectacular example of such a situation. In any case, the synergy between text-based and content-based searching should be exploited to the maximum.

At the moment, around 150,000 images have been indexed at low resolution (1,000x1,000 pixels), 14,000 at high resolution (up to 12,000x8,000) and around 300 3D models (from 30,000 to 3,000,000 vertices). The feature vectors at low resolution have been calculated on a high-end laptop while the indexes for the high-resolution paintings have been calculated in Paris on the C2RMF server (this operation took about 5 hours on both machines). The fact that 150,000 images can be indexed on a laptop and that the query, on the same laptop, takes between 1 to 3 seconds, shows that the algorithms are well optimized.

All the indexes were calculated offline. An evaluation of the system has shown that paintings and artefacts should be indexed automatically as soon as they are store in the database. This would ensure that a content-based index is attached to each item as soon as it becomes available in the database. We are currently working on a grid-computer architecture in order to be able to index a massive amount of ultra-high resolution 2D and 3D. By distributing the load of many nodes, it will be possible to increase substantially the performance of the system. This task will be facilitated by the fact that the approach is well-suited for parallelization.

References

[ALPP05] AITKEN G., LAHANIER C., PILLAY R., PITZALIS, D.: EROS : An Open Source Database For Museum Conservation Restoration. In *Preprints de la 14ème réunion triennale du Comitè International pour la Conservation ICOM-CC*, 2005, pp. 15–23. 2

[ALPP06] AITKEN G., LAHANIER C., PILLAY R., PITZALIS, D.: Database Management and Innovative Applications for Imaging within Museum Laboratories. emph7th European Commission Conference "SAUVEUR", Prague, Czech Republic, 2006 2

[AML*05] ADDIS M., MARTINEZ K., LEWIS P., STEVENSON J., GIORGINI F.: New ways to search, navigate and use multimedia museum collections over the web. In *Proceedings of Museums and the Web* (2005). 2

[FP02] FEKETE, J., PLAISANT, C.: Interactive Information Visualization of a Million Items. *IEEE Symposium on Information Visualization*, Boston, October 2002 2

[FTT05] FELIPE J. C., TRAINA A. J. M., TRAINA C. J.: Global warp metric distance: Boosting content-based image retrieval through histograms. *ism 0* (2005), pp. 295–302. 2

[GBT*02] GODIN G., BERALDIN J. A., TAYLOR J., COURNOYER L., RIOUX M., EL-HAKIM S., BARIBEAU R., BLAIS F., BOULANGER P., DOMEY J., PICARD M.: Active optical 3d imaging for heritage applications. *IEEE Comput. Graph. Appl. 22*, 5 (2002), pp. 24–36. 1

[IJL*05] IYER N., JAYANTI S., LOU K., KALYANARAMAN Y., RAMANI K.: Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design 37*, 5 (April 2005), pp. 509–530. 3

[LAC*04] LAHANIER C., AITKEN G., COTTE P., PILLAY R., PITZALIS, D.: Digital capture of archive, objects and paintings for research and conservation. In *EVA* (2004). 1

[LDP02] LAHANIER C., DUFRESNE J. L., PILLAY R.: Acohir: a system for 3d digitization applied to conservation and research imaging. In *The 13th Triennial Meeting, ICOM, J&J* (2002). 1

[LFJ*05] LAHANIER C., FEIHL O., JEANLIN M., PITZALIS D., SCHMITT F.: 3D modelling of archaeological objects for conservation, visualisation, colour and shape characterisation, details comparison In *Preprints de la 14ème réunion triennale du Comitè International pour la Conservation ICOM-CC*, 2005, pp. 43–51. 1

[PPLA05] PILLAY R., PITZALIS D., LAHANIER C., AITKEN G.: Eros: New development in digital archiving for research in conservation. In *EVA* (2005), pp. 192–197. 2

[TS06] TAE P. K., SHIK M. Y.: Object-based image retrieval using dominant color pairs between adjacent regions. In *Computational Science and its Applications - ICCSA 3984* (2006), pp. 404–411. 2

[TV04] TANGELDER J. W. H., VELTKAMP R. C.: A survey of content based 3d shape retrieval methods. In *SMI '04: Proceedings of the Shape Modeling International 2004 (SMI'04)* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 145–156. 3