



NRC Publications Archive Archives des publications du CNRC

De l'application des technologies 3-D dans le cadre du patrimoine culturel

Paquet, Eric; Viktor, H.L.

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

Annales des telecommunications, 2005

NRC Publications Record / Notice d'Archives des publications de CNRC:

<https://nrc-publications.canada.ca/eng/view/object/?id=ef441831-8357-42e8-96d8-68aa2a98b7b8>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=ef441831-8357-42e8-96d8-68aa2a98b7b8>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.





National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC - CNRC

On the Application of 3-D Technologies to the Framework of Cultural Heritage / De l'application des technologies 3-D dans le cadre du patrimoine culture *

Paquet, E., Viktor, H.L.
2005

* published in the *Annales des télécommunications*. 2005. NRC 48249.

Copyright 2005 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables
from this report, provided that the source of such material is fully acknowledged.

On the application of 3-D technologies to the framework of cultural heritage

De l'application des technologies 3-D dans le cadre du patrimoine culturel

E. Paquet^a and H. L. Viktor^b

^a Visual Information Technology, National Research Council, M50 Montreal Road, Ottawa, K1A 0R6, Canada – eric.paquet@nrc-cnrc.ca

^b SITE, University of Ottawa, 800 King Edward Ave., Ottawa, K1N 6N5, Canada – hlviktor@site.uottawa.ca

Key words: 3-D, acquisition, content-based, cultural heritage, data mining, imaging, indexing, modelling, retrieval, three-dimensional, virtual collections, visualisation

Abstract:

Unrestricted access to both historical and archaeological sites is highly desirable from both a research and a cultural perspective. With the recent developments in 3-D scanner technologies and photogrammetric techniques, it is now possible to acquire and create accurate models of such sites. Through the process of virtualisation, numerous 3-D collections are created that need to be visualised, indexed and eventually characterized. This paper presents a mobile virtual environment designed for the visualization of photorealistic high-resolution virtualised cultural heritage scenes and artefacts. In addition, it is shown how it is possible to describe and index the texture and the 3-D shape of these artefacts with compact support feature vectors. A data mining system, based on these vectors, allows the characterization and the exploration of the 3-D collections, through cluster analysis.

Résumé :

Un libre accès aux sites historiques et archéologiques est hautement désirable tant du point de vue de la recherche que de la culture. Avec les récents développements en numérisation 3-D ainsi qu'en photogrammétrie, il est maintenant possible d'acquérir et de créer des modèles précis et représentatifs de tels sites. À travers le processus de virtualisation, de nombreuses collections 3-D sont ainsi créées et doivent par la suite être visualisées, indexées et éventuellement caractérisées. Cet article décrit un environnement virtuel mobile pour la visualisation haute résolution photo-réaliste de sites et d'artéfacts patrimoniaux. De plus, on montre qu'il est possible de décrire et d'indexer la texture et la géométrie 3-D de tels artéfacts à l'aide de vecteurs signature compacts. Un système d'exploration des données, basé sur ces vecteurs, permet de caractériser et d'explorer les collections 3-D en analysant les amas sous-jacents.

I. Introduction

The public interest toward historical, cultural and archaeological sites and artefacts has significantly increased over the past decade. An equivalent phenomenon has been observed among specialists and scholars. Such phenomenon, desirable from a cultural point of view, has introduced a new problematic: many sites are extensively visited and many artefacts are manipulated by a very large number of people. The consequences have been dramatic: a gradual degradation of numerous sites has been observed and countless artefacts have been permanently damaged or lost [1].

Several curators have reached the conclusion that the sole solution to ensure the conservation of the artefacts is to restrict their access to the minimum possible level. This, in turn, has generated many questions. Who should be allowed to access the sites? Under which conditions should this access be granted? For instance, if the duration of the visit is too short and if the conditions are too constraining, the visitor cannot appreciate the site and the scholar is limited in terms of the feasibility of the research.

A proposed solution has been to virtualize such sites and artefacts, i.e. to create an "identical" digital copy of the original sites [2]. Initially, it seems that virtualisation is the panacea to the above-mentioned difficulties. A virtual model provided an

unlimited access, does not suffer from degradation and can be accessed from virtually any location. Nevertheless, one has to recognise that such an approach conveys new challenges. For instance, can we really certify that the virtualised site conforms to the original? The answer is twofold. With the current techniques, it is possible to achieve a high degree of realism. However, it is still difficult to provide the scholar with an accurate representation both from a chromatic and geometric point of view. For a valuable study of this aspect, the reader is referred to [3].

Consequently, the accuracy of virtualised sites can still be considered as an open issue. For this reason, scholars should have access to the raw data in addition to the final model, in order to inspect the calibration files and to provide an exhaustive list of the equipment and software utilised for the acquisition and creation of the model. Furthermore, the scholar should be provided with such information as the precision and the accuracy of the model in order to be able to estimate the level of detail to which the model can be analysed.

Despite of tremendous improvements, the acquisition and creation of a model still involve a significant amount of manual intervention and savoir-faire. Nevertheless, if the acquisition pipeline is well designed [4], it is possible to acquire and create a large amount of models. It is conceivable that in the near future, such models will be so numerous that the creation of

databases will be justified and needed. A certain number of issues then arise. One should be able to access the database, to search for a particular item and to formulate specific queries in order to retrieve artefacts and sites of interest. In addition, the retrieval system should be designed in such a way that it facilitates comparative studies of sites and artefacts.

Once the content of interest has been identified, it must be visualised. The visualisation process should be designed in such a way that it provides a realistic view of the artefact under consideration. In order to achieve such a high degree of realism, the visualisation system should provide depth perception, high spatial resolution, high refresh rate, immersion, navigation, manipulation and a faithful reproduction of the colours and geometry of the original [5].

This paper is organised as follow. Firstly, a cost effective stereo visualisation system is described. Then, it is shows how the composition of an image and the shape of an artefact can be indexed. A retrieval system is presented. It is shown how to combine this system with cluster analysis in order to perform comparative studies. Finally, an integrated approach or framework for virtual collections is presented. This framework encompasses the acquisition and the creation of the models, virtual model documentation, content indexation, retrieval, comparative studies and visualisation.

II. Visualisation of virtual collections

II.1 Architecture of the system

In this section, we describe a cost-effective approach for the stereo visualisation of virtual sites. When a site is virtualised, it is possible to make it available to a large variety of people, as far as the intellectual property protection allows.

While most users are content to visualise a site with a standard computer, specialists and scholars are not. Most of the time, scholars need access to high quality pictures of sites. Despite the fact that the pictures provide a high-resolution representation of the visual appearance of the site, they provide limited and ambiguous information about the three-dimensional shape. Furthermore, the pictures are taken from a very limited subset of viewpoints, which are not necessarily the ones required by the scholar. This is why, among many reasons, archaeological sites need to be accessed. The situation is rather different if the scholars can visualise the site in three dimensions and navigate within the scene. This can be achieved with dynamic stereo visualisation.

Most commercial stereo systems suffer either from high-cost or from poor performances. In both cases, the access to the data is limited: in the first case, because a small number of users can afford the system, while in the second case, the system can only display low-resolution versions of the original data. For these reasons, we have created a cost-effective stereo visualisation system suitable for the visualisation of large amount of complex three-dimensional data. The system is made of off-the-shelf components and is characterised by its scalability, its distributed architecture and its portability. Let us review its architecture.

The system is formed of two interconnected laptops and two DLP video projectors. The first computer, the sender, receives the events from the user and renders the right-eye view. The second computer, the receiver, is synchronised with the sender

and renders the left-eye view. Each computer is attached to an ultra-compact DLP video projector. The left and right-eye views are simultaneously projected though polarised filters on a special screen that maintains polarisation. The polarisation is circular right for the right-eye view and circular left for the left-eye view. The circular polarisation insures that the orientation of the user's head does not introduce cross talk in between the right and the left view, which would deteriorate the stereo effect. The scene can be visualised in three dimensions with the corresponding passive stereo glasses.



Figure 1. Visualisation of the Crypt of *Santa Cristina* in *Capignano Salentino*, Italy with the virtual theatre. Details of a fresco on a column.

Visualisation de la crypte de *Santa Cristina* située dans le *Capignano Salentino* en Italie à l'aide du théâtre virtuel. Détails d'une fresque située sur une colonne.

The software is written in Java/Java 3-D and the same byte code can run transparently on OpenGL and DirectX. The synchronization is ensured as follows. Before it renders a frame, the sender updates its navigation parameters from the events generated by the user. It sends the navigation parameters through a firewire to the receiver, which renders the left-eye view of the scene. Meanwhile, the sender renders the right-eye view of the scene. Then the sender sends a synchronisation signal to the receiver and both computers swap their respective buffer. In order to obtain a higher resolution, an arbitrary number of receivers can be connected to the sender. Thus, it is possible to render a scene at ultra-high resolution by simply adding more computers and projectors: each computer-projector pair being responsible for the rendering of a subset of the scene; the so-called tiled wall display.

Currently, the system runs on Dell Precision M60 laptops with a NVIDIA® Quadro FX Go700™ 4XAGP graphics card with 128 MB of texture memory. The selected DLP video projectors are the NEC LT260K: their brightness is 2100 ANSI lumens and their native resolution is 1024x768. These projectors have a proprietary technology that corrects rapidly and efficiently the horizontal and vertical keystones, which make the alignment relatively straightforward even in the case of multiple displays. The selected polarised filters are the 3M HNCP37. They have been chosen because their light transmission curve is relatively uniform over a wide range of visual wavelengths, which means that the colour distortion is reduced to a minimum. This is an important aspect to ensure the faithfulness toward the original.

II.2 Experimental results

This section discusses two heritage sites that has been visualised using our system as described above. Firstly, the Crypt of *Santa Cristina* is a 9th century Byzantine crypt situated in *Capignano Salentino* in the South of Italy. It has an irregular shape and contains one of the most ancient Byzantine frescoes signed and dated. For instance, *Theophylact* painted the Christ and the Annunciation in 959 A.D. The dimensions are about 16.5 x 10.0 x 2.5 m.

An Italo-Canadian team from SIBA, University of Lecce and Visual Information Technology, National Research Council virtualised the site with a range laser scanner and photogrammetric techniques [4]. The crypt was scanned with a MENSIS SOISIC™ range laser scanner. The spatial sampling was 5 mm on the surface of the walls and the depth uncertainty was 0.8 mm. The crypt was scanned in various sections of about 2.5 m wide each. Texture information was acquired separately with a high-resolution digital camera Nikon D1x at a resolution of 3008 x 1960 pixels. The sections were aligned together with the ICP algorithm and a global alignment was performed on the complete model. These operations were carried out with a commercial software package named Polyworks™. The crypt has subsequently been visualised with our system. Figure 1 shows a column of the crypt as visualised with our system.



Figure 2. Visualisation of the *Torre dell'Aquila* in Trento, Italy with the virtual theatre. Partial view of the Summer Cycle.

Visualisation de la *Torre dell'Aquila* située à Trento en Italie à l'aide du théâtre virtuel. Vue partielle des fresques du cycle de l'été.

The "Palace of the Good Council" or *Castello del Buon Consiglio* was the symbol of the temporal power of the Prince

Bishop of Trento; now part of Italy. The actual palace consists of many sections from the medieval and renaissance periods: one of the best known is the "Tower of the Eagle" or *Torre dell'Aquila*. The *Torre dell'Aquila* contains a magnificent cycle of frescoes known as *Dei Mesi* i.e. "of the Months" painted by an anonymous Bohemian painter in the 15th century. They depict medieval life, month by month, comparing the richness and splendour of the court with the simple life in the country.

The interior of the tower was virtualised with photogrammetric techniques. Pictures were acquired at high-resolution by an Italo-Canadian team from IRST – ITC and VIT - NRC with a digital camera and a model was created with photogrammetric techniques with a commercial software package called ShapeCapture™ [6]. The tower has also been visualised with our system. Figure 2 shows a section of a wall corresponding to the so-called "Summer Cycle" while Figure 3 shows details of the ceiling.

Both sites are currently visualised in Italy and Canada with our system: one virtual theatre has been deployed at SIBA (Lecce, Italy), one at IRST – ITC (Trento, Italy) and the other one at VIT in our laboratory. A fourth system is currently dedicated to the mobile visualisation of industrial design at INDACO, the School of Industrial Design of the Politecnico di Milano (Milan, Italy).



Figure 3. Visualisation of the *Torre dell'Aquila* in Trento, Italy with the virtual theatre. Details of the ceiling.

Visualisation de la *Torre dell'Aquila* située à Trento en Italie à l'aide du théâtre virtuel. Détails du plafond à caisson.

III. Exploration of virtual collections

III.1 Content-based retrieval of images and texture

This section presents a new algorithm for the indexation and retrieval of images. Pictures and images are of the utmost importance in virtual collections. They are (and will remain in a foreseeable future) the easiest, fastest and most economical mean for creating virtual collections. Furthermore, most three-dimensional models are covered with textures. The textures constitute an important visual descriptor for the model under consideration and convey essential historical, artistic and archaeological information. For instance, the most important information about the Crypt of *Santa Cristina* described above comes from the frescoes, which correspond to the textures.

This phenomenon is even more evident in the case of the *Torre dell'Aquila*.

Images are difficult to describe. They convey a large amount of complex and ambiguous information. The ambiguity is due to the fact that an image is a bidimensional projection of the three-dimensional world and by 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 [7]. For the above-mentioned reasons, we propose a statistical approach in which the overall composition of the image is described in an abstract manner.

We now depict our algorithm. The colour distribution of each image is described in terms of hue and saturation. This colour space imitates many characteristics of the human visual system. The hue corresponds to our intuition of colour e.g. red, green or blue while saturation corresponds to the colour strength e.g. light red or deep red.

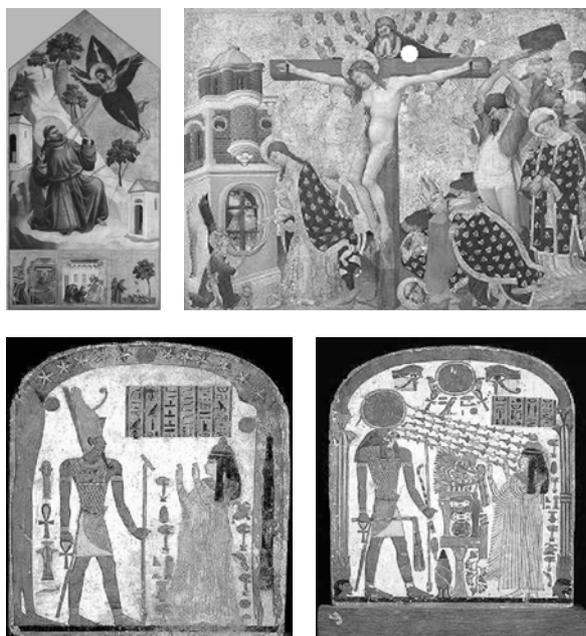


Figure 4. Retrieval of paintings from a database containing 1.000 elements. The top illustrations show the retrieval of a crucifixion while the bottom illustrations show the retrieval of an Egyptian tablet.

Recherche de tableaux dans une banque de données contenant 1.000 éléments. Les illustrations supérieures montrent les résultats obtenus pour une crucifixion tandis que les illustrations inférieures montrent les résultats obtenus pour une tablette égyptienne.

Next, a set of points is sampled from the image. A quasi-random sequence generates the points. The choice of such a sequence is justified by the fact that no particular or a priori underlying structure or composition is assumed for the image. In the present implementation, the Sobol sequence has been selected. Such a sequence allows sampling the image at various resolutions while filling the gaps in the previously generated distribution. Each point of this sequence becomes the centre of a small rectangular window on the image. For each centre position, the pixels inside the corresponding window are extracted and the associated hue and saturation images are

calculated. The statistical distribution of the colours within the window is characterized by a bidimensional histogram. The first dimension of this histogram 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. Intuitively, such a histogram can be assimilated to the palette (as understood in art) associated with the visual content of the window.

This bidimensional 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. The displacement of the window over the image insures that not only global information about the later is accumulated but structural information is extracted as well. Here, structural information should not be understood as a set of spatial relations in between segmented areas, but as the repetition on various locations of small and similar patterns that characterise the composition of the image. Finally, each set of three histograms is converted into a compact descriptor or index. Such an index is a quantized representation of the bidimensional histograms. In the current approach, the proportions are quantized on four channels or bins while the colours (hue and saturation) are quantized on six channels. The indexed is stored as a compact binary representation and has a typical size of 200 bytes.

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. This is why that algorithm is robust against occlusion, composition, partial view and viewpoint. Nevertheless, this approach provides a good level of discrimination.

A better understanding of the approach can be obtained by considering Figure 4. The top illustrations show the results obtained for the retrieval of a crucifixion from a database of 1.000 paintings and frescoes. Both paintings are from the same period and have been painted with similar techniques e.g. with a golden background. This example illustrates very well the concept of structural information. A careful observation of these two paintings shows that both present similar local patterns despite the fact that the spatial relationships in between them are quite different.

The bottom illustrations show the retrieval of an Egyptian wooden tablet from the same database. Both images have the same composition and the colour palette is similar. This represents an ideal case for the algorithm because it is precisely based on the above mentioned characteristics. Initially we thought that the two frescoes were part of the same temple. We found later, by seeing them in a museum that they belong in fact to the same wooden tablet. If we would have utilised a text-based query, we would have probably never found these images because of our wrong assumption.

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

III.2 Content-based retrieval of 3-D artefacts

This section presents a new algorithm for the indexation and retrieval of three-dimensional artefacts [8]. The indexation of three-dimensional artefacts differs fundamentally from the indexation of images. 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 virtualised geometry and the physical geometry of the artefact. 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 artefact 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 artefact. From the proto-artefact, 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 artefacts 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 3-D) is the index. We now describe our algorithm for three-dimensional artefact description. We assume that each artefact has been modelled with a mesh. This is a non-restrictive representation for virtualised artefact 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 artefact 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 vertices are normalised relatively to the position of its centre of mass. A translation invariant representation is then achieved. Translation invariance is important since a priori we do not know the location of the artefact. 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 utilise the vertices 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 normalised 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 virtualised: a sine qua non condition for real world applications. Under certain assumptions, the area of the triangles is related to the local curvature: the smaller the area, the higher the curvature. Such a hypothesis is valid if the number of acquired points is related to the complexity of the local structure and constitutes a sine qua non condition for any realistic shape acquisition.

In order to achieve rotation invariance, the Eigen vectors of the tensor of inertia are calculated. Once normalised, 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. It is very common to encounter axes which are orientated along the shortest and the longest dimensions of the artefact. For instance, for an Ionic column, one axis would be orientated along the rotation axis of the column while the other one would be orientated along a radial direction.

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 which 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 normalised 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 normalisation 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 than the average. The cords are also weighted in terms of the area of the corresponding triangles; the later being normalised 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 histogram described the distribution of the cosine directions associated with the unit vector associated with the second smallest Eigen value. The third histogram described the distribution of the normalised spherical radius as defined in the previous paragraph. The ensemble of the three histograms constitutes the shape index of the corresponding artefact.

A concrete example is provided in the next section.

III.3 Characterization of virtual collections

This section shows how to perform a comparative study on a collection by using the apparatus developed in the previous two sections. We saw that shape constitutes a universal language. This is obvious in art and architecture. Indeed, numerous styles can be characterized by the shape of their associated architectural elements. For instance, let us consider the three classical orders of columns: Ionic, Doric and Corinthian. Columns of the same order are of course never identical but their shape is sufficiently similar so that it is relatively straightforward to identify them. The same can be said about cupolas, Roman arcs and Gothic arcs; the list is literally endless.

While studying a collection, a scholar usually studies particular artefacts but he may as well want to compare similar artefacts in order to determine their common characteristics, their discrepancies or even the temporal evolution of a particular style. In order to be able to perform these actions, the scholar must be able to group or cluster similar artefacts within a common class.



Figure 5. Creation of a palette of fragments for a broken vase. The top view corresponds to the central section of the vase while the bottom view corresponds to the border.

Création d'une palette de fragments à partir des tessons d'un vase. Les figures supérieures correspondent à la panse du vase tandis que celles du bas sont associées au bord du vase.

The same approach can be followed in the reconstitution of broken artefacts. Most of the time, similar fragments occupy adjacent regions on the artefact or belong to the same component: for instance a vase ear. Consequently, the palette facilitates the work of the scholar by providing him with a pre-classified ensemble of fragments into clusters. In practice this classification is performed manually: a time consuming approach. With the proposed method, it is possible to perform this classification semi-automatically. We first consider the general theory and then we apply the results to the creation of a palette of fragments for the reconstitution of a broken vase.

Our system can find, in a collection, similar pictures and three-dimensional artefacts. We have seen that by following an iterative or spiral approach, the user can converge to the item of interest. Let us review this process from a more fundamental point of view. An index can be visualised as a point in an N -dimensional space where N depends on the number of channels. When the user chooses a prototype, he chooses a point or a seed

in this space. Then, the search engine determines the closest points to the seed, with the proximity being defined according to a metric such as the Euclidian distance. It should be remembered that a point corresponds to an artefact in the collection. In this way, the nearest-neighbourhood approach is used to identify clusters of similar objects.

If the outcome of the query corresponds to the foreseen results, it means that the seed point is approximately located near the centre of the cluster. If it is not the case, it means that the seed point is situated on the outskirts of the cluster. In that case, the user selects the best candidate from the closest M points obtained from the previous iteration and reiterates the process. After a few iterations, the process converges to a point near the centre of the cluster. This point and its neighbourhood constitute the cluster and the corresponding artefacts form the class.

We have applied this approach for the creation of a geometrical palette of similar fragments from a collection of 448 fragments from a broken vase (2000 AD!). We were able to group similar pieces such as the ear, the rim and the central region. Some experimental results are shown in Figure 5. The top view shows two fragments of the central section that were grouped together while the bottom view shows two fragments of the rim.

If it is not possible to process the palette manually because the number of segments is too high ($\sim 100,000$), cluster analysis [9] should be considered. We refer the reader to reference [10] for an analysis of this important issue and limit our discussion to one algorithm of particular relevance. A cluster is a collection of objects that are similar to one another and are dissimilar to the objects in other clusters. The goal of clustering is to find inter-cluster similarity and intra-cluster dissimilarity, through the discovery of a hidden pattern that gives meaningful groups (clusters) of objects.

From the analysis performed in [10] we found that the WaveCluster algorithm [11] is suitable for image and 3-D descriptors. WaveCluster is based on the wavelet transform and the grid-based approach. A wavelet transform is a signal processing technique that decomposes a signal into different resolution sub-bands. WaveCluster divides the range of values associated with each dimension into a finite set of equal length intervals. The ensemble of all discrete dimensions forms a discrete multidimensional space. The smallest volumetric element of this space is called a voxel. WaveCluster assigns the continuous values associated with each dimension to the corresponding voxel in the multidimensional space. If more than one value is attributed to a voxel, their sum is accumulated. WaveCluster applies discrete wavelet transform to the multidimensional discrete space in order to find connected components or clusters in the transformed wavelet space. WaveCluster works well with high dimensional data sets, such as the image and 3-D descriptors. The quality of the output clusters are high, it successfully handle noises and is able to find irregular shape clusters with no a-priori knowledge. In order to use WaveCluster, an estimation of the resolution, i.e. the number of discrete channels associated with each dimension, must be provided as well as a basis function for the wavelet transform.

IV. An integrated approach to virtual collections management

In the previous sections, we have presented some considerations on virtualisation, a cost-effective stereo visualisation system for three-dimensional complex data, an indexation and retrieval system for images and three-dimensional artefacts based on composition and shape and, an approach for characterising a collection in terms of clusters and classes. In this section we present an integrated approach to virtual collections management. The proposed approach consists of five steps: virtualisation, documentation, indexation, retrieval, characterisation and finally, visualisation.

The first step is well known and thousands of scientific papers have been devoted to the subject. Although virtualisation is and will remain a major research area, it is our opinion that documentation has been neglected. Most scholars assume that virtualised models are conformed to the original ones within the precision of the apparatus utilised to virtualize them. It has been shown that although a model might be precise, it is not necessarily accurate [3]. Precision is related to the reproducibility of the measured data while accuracy is related to the conformity of the measured data with the original model. It is difficult in an extra-laboratory situation to determine if the measured data are accurate or not. Unfortunately, this is the most common situation. For that reason, in addition to the virtualised model, additional information should be provided. The provided information should at least include, the type of apparatus utilised to virtualize the artefact, resolution, precision, time-stamped calibration, time-stamped raw data, time-stamped temperature and humidity, etc. The availability of such data would assist the scholar to assert the validity of the measured data.

Indexation turns out to be important when the number of virtualised artefact becomes significant. Our aptitude to retrieve pertinent information is primarily determined by the quality of our indexes. An index can be a textual description of the artefact, a set of measurements, a set of keywords and metadata or, an innovative technique like composition and shape description: the so-called content-based approach. While some fields like archaeology have developed a systematic procedure for indexation, one has to recognise that it is not the case for most fields. Even worst, in most cases, there is no information at all associated with the artefacts. A content-based approach can provide a partial solution to this problem since it automatically indexes a collection. The content-based indexes provide a description of the artefacts in terms of composition and shape. A search engine then provides a fast and efficient access to the collection. The content-based approach does not replace classical indexation but it complements it and provides an alternative if the later is absent.

As we have seen earlier, characterisation is intimately related to retrieval. The combination of classical documentation with content-based indexes is of particular interest. Classical information can act as a filter on the collection, which then can be searched with a content-based approach. Some results in relation with this approach are reported in [12]. Once the artefacts of interests have been retrieved, that must be visualised. We have presented a cost-effective stereo visual system for the visualisation of complex three-dimensional artefacts and scenes. One of the advantages of this system is that it can be rapidly deployed on any working site and may therefore be utilised as a terrain and operational device.

V. Conclusions

The creation of virtual collections of cultural heritage sites provides both experts and novices with unrestricted access to these places of interest. The stereo environment, as described here, provides a cost-effective system to visualize, search and further explore such virtualised collections. The system is mobile, scalable and is able to visualise large amounts of complex three-dimensional data, as illustrated by the use thereof to visualise a number of cultural heritage sites in Italy. Experimental results show that use of content-based retrieval, together with traditional information retrieval techniques, offer the user the ability to accurately characterise virtual collections. Here, the “query by example” or “query by prototype” paradigm allows for the clustering of similar pictures and three-dimensional artefacts. Through the use of this technology, novel insights into previously restricted historical and archaeological sites can be obtained through comparative studies, virtual model documentation and visualisation.

VI. References

- [1] KUMAN (S.) et al., Digital Preservation of Ancient Cuneiform Tablets Using 3-D Scanning, *Fourth International Conference on 3-D Digital Imaging and Modeling*, pp. 326-333, Oct. 2003.
- [2] EL-HAKIM (S. F.), BERALDIN (J.-A.) et al., Effective 3D Modeling of Heritage Sites, *Fourth International Conference on 3D Digital Imaging and Modeling*, pp. 302-309, Oct. 2003.
- [3] GUIDI (G.) et al., Accuracy Verification and Enhancement in 3D Modeling: Application to Donatello’s Maddalena, *Fourth International Conference on 3D Digital Imaging and Modeling*, pp. 334-341, Oct. 2003.
- [4] BERALDIN (J.-A.), VALZANO (V.) et al., Virtualizing a Byzantine crypt: challenge and impact, *Videometrics VII*, pp.137-147, Jan. 2003.
- [5] GAIANI (M.), Strategie di rappresentazione digitale: modelli per la storia e la conservazione dei beni architettonici e ambientali, *Centro di Ricerche Informatiche per i Beni Culturali*, **10**, pp. 47-69, 2000.
- [6] HAKIM (S. L.), GONZO (L.) et al., Visualization of Highly Textured Surfaces, *4th International Symposium of Virtual Reality, Archaeology and Intelligent Cultural Heritage*, pp. 1-9, Nov. 2003
- [7] TSEKERIDOU (S.) and PITAS (I.), Audio-Visual Content Analysis for Content-Based Video Indexing, *IEEE Multimedia Computing and Systems*, **1**, pp. 667-672, Jun. 1999.
- [8] PAQUET (E.) and RIOUX (M.), Nefertiti: a query by content system for three-dimensional model and image databases management, *Image and Vision Computing*, **17**, pp. 157-166, 1999.
- [9] HAN (J.) and KAMBER (M.), Data Mining: Concepts and Techniques, *Morgan Kaufman Publishers*, Boston, 2001.
- [10] PAQUET (E.), VIKTOR (H. L.) et al., Exploring anthropometric data through cluster analysis, *SAE Digital Human Modeling for Design & Engineering*, CD, Jun. 2004.
- [11] AGRAWAL (R.) et al., Automatic Subspace Clustering of High Dimensional Data for Data Mining Application,

Proceedings of the ACM SIGMOD Conference, pp. 94-100, Jun. 1998.

[12] PAQUET (E.) and RIOUX (M.), Anthropometric Visual Data Mining: a Content-based Approach, *International Ergonomics Association XVth Triennial Congress*, pp. 78-91, Aug. 2003.