

NRC Publications Archive Archives des publications du CNRC

Visualisation, Exploration and Characterization of Virtual Collections 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:

*XXth Congress of the International Society for Photogrammetry and Remote
Sensing [Proceedings], 2004*

NRC Publications Archive Record / Notice des Archives des publications du CNRC :
<https://nrc-publications.canada.ca/eng/view/object/?id=446eec11-3a1e-4a80-9021-74d6b7d933d0>
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=446eec11-3a1e-4a80-9021-74d6b7d933d0>

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

*Visualisation, Exploration and Characterization of Virtual Collections **

Paquet, E., and Viktor, H. L.
July 2004

* published in XXth Congress of the International Society for Photogrammetry and Remote Sensing. July 12-23, 2004. Istanbul, Turkey. NRC 46569.

Copyright 2004 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.

VISUALISATION, EXPLORATION AND CHARACTERIZATION OF VIRTUAL COLLECTIONS

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.ca

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

Commission VI, WG V/2^a

KEY WORDS: Virtual Reality, Visualisation, Three-dimensional, Mobile, Data Mining, Retrieval, Query

ABSTRACT:

Unrestricted access to both historical and archaeological sites is highly desirable from both a research and a cultural perspective. However, due to security and preservation considerations, access is becoming more and more restricted and subject to various conditions. With the recent developments in 3D scanner technologies and photogrammetric techniques, it is now possible to acquire and create accurate models of such sites. Through the process of virtualisation, numerous virtual collections are created that need to be visualised, searched and eventually characterized. This paper presents a mobile virtual environment designed for the visualization of photorealistic high-resolution virtualised scenes and artefacts. The mobile virtual environment also includes a component for retrieving artefacts from virtual collections. This stereo virtual environment is portable and can be easily and rapidly deployed at any suitable location, for instance an archaeological site. The architecture and the implementation of the mobile virtual environment are described. This environment is characterized by a massively asynchronous architecture that optimises the rendering performances by distributing the calculations over various graphical processing units. A request broker insures the synchronization among the various components of the system. The performance of the system is illustrated through multiple examples of the visualisation of virtualised cultural heritage sites. In addition, it is shown how it is possible to describe the geometry of the artefacts by representing them with compact support feature vectors. A recurrent data mining system, based on these vectors, is presented. This system allows the characterization and exploration of the collection, through cluster analysis. The system employs the “query by example” paradigm and the knowledge of the expert in a recurrent approach, in order to identify clusters of artefacts. The virtual environment is subsequently utilised in order to perform visual data mining on the clusters, as identified during data mining, and to characterize and further explore the clusters by defining archetypes.

1. INTRODUCTION

The public interest toward historical, cultural and archaeological sites and artefacts has been increasing significantly 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 over visited and many artefacts are over manipulated. 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 feasible of the research.

A proposed solution has been to virtualise 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 opened issue. For this reason, scholars should have access to the raw data in addition to the final model, in order to the calibration files and to 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 a 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, 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, the creation of the models, virtual model documentation, content indexation, retrieval, comparative studies and visualisation.

2. VISUALISATION OF VIRTUAL COLLECTIONS

2.1 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 it.

While most users are content to visualise a site with a standard computer, specialists and scholars are not. Most of the time, scholars have 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 they can visualise the site in three dimensions and navigate within the scene. This can be achieved with dynamic stereo visualisation.

Most of the 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 limited number of users can afford the system, while in the second case, because the system can only display a low-resolution version of the 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 through 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.

The software is written in Java/Java 3D and the same byte code can run transparently on OpenGL and DirectX. The synchronization is ensured as follow. 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 keystone, 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 faithfulness toward the original.

2.2 Experimental results

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 MENSİ 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 been visualised with our system. Figure 1 shows a column of the crypt as visualised with our system.



Figure 2. Visualisation of the Torre dei Mesi in Trento, Italy with the virtual theatre. Partial view of the Summer Cycle.

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 called ShapeCapture™ [6]. The tower has 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 dei Mesi in Trento, Italy with the virtual theatre. Details of the ceiling.

3. EXPLORATION OF VIRTUAL COLLECTIONS

3.1 Content-based retrieval of images

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 comes from the frescoes, which correspond to the textures. This phenomenon is even more evident in the case of la Torre dell’Aquila.

Images are difficult to describe. They convey a large amount of complex and ambiguous information. The ambiguity comes from 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.

Then, 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 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. This bidimensional histogram is computed and accumulated for each point of the sequence i.e. that 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 it is 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 be 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. 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. A demonstration can be found in [8] with a database of more than 2100 images.

3.2 Content-based retrieval of 3D artefacts

This section presents a new algorithm for the indexation and retrieval of three-dimensional artefacts [9]. 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 in 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 in between the two systems 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 follow. 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. Then, the tensor of inertia of the artefact is calculated. This tensor is a 3×3 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 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.

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.

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 demonstration with a database of more than 4400 three-dimensional models can be seen at [8].

3.3 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 4. 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.

The same approach can be followed in the reconstitution of broken artefacts. Here, clustering (or cluster analysis) a data mining technique that deals with the extraction of the implicit knowledge, data relationship or other patterns not explicitly stored, by grouping or classifying related records together, is performed [10]. 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. For example, the first step in the reconstitution of a broken artefact is to create a palette of similar fragments. 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, when a cluster has been constructed, 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 palette 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. Here, the Euclidian metric was used to create the clusters of similar pieces. Some results are illustrated in Figure 4. The top view shows two fragments of the central section that were grouped together while the bottom view shows two fragments of the border. A demonstration can be found in [8].

4. 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 virtualise 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 virtualise 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.

5. 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.

6. REFERENCES

1. Kuman, S. et al., 2003. Digital Preservation of Ancient Cuneiform Tablets Using 3D Scanning. In: *Fourth International Conference on 3-D Digital Imaging and Modeling*, Ottawa, Canada, pp. 326-333.
2. El-Hakim, S. F., Beraldin, J.-A. et al., 2003. Effective 3D Modeling of Heritage Sites. In: *Fourth International Conference on 3-D Digital Imaging and Modeling*, Ottawa, Canada, pp. 302-309.
3. Guidi, G. et al., 2003. Accuracy Verification and Enhancement in 3D Modeling: Application to Donatello's Maddalena. In: *Fourth International Conference on 3-D Digital Imaging and Modeling*, Ottawa, Canada, pp. 334-341.
4. Beraldin, J.-A., Valzano, V. et al., 2003. Virtualizing a Byzantine crypt: challenge and impact. In: *Videometrics VII*, Santa Clara, USA, pp.137-147.
5. Gaiani, M., 2000. 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.
6. Hakim, S. L., Gonzo, L. et al., 2003. Visualization of Highly Textured Surfaces. In: *4th International Symposium of Virtual Reality, Archaeology and Intelligent Cultural Heritage*, Brighton, United Kingdom, pp. 1-9.
7. Tsekeridou, S. and Pitas, I., 1999. Audio-Visual Content Analysis for Content-Based Video Indexing. In: *IEEE Multimedia Systems*, Florence, Italy, Vol. 1, pp. 667-672.
8. Paquet, E., “Content-based Indexing and Retrieval of 3D and Image Databases”, Ottawa, Canada. <http://www.cleopatra.nrc.ca> (accessed 7 April 2004)
9. Paquet, E. and Rioux, M., 1999. Nefertiti: a query by content system for three-dimensional model and image databases management. *Image and Vision Computing*, 17, pp. 157-166.
10. Han, J. and Kamber M., 2001. *Data Mining: Concepts and Techniques*. Morgan Kaufman Publishers, Boston.
11. Paquet, E. and Rioux, M., 2003. Anthropometric Visual Data Mining: a Content-based Approach. In: *International Ergonomics Association XVth Triennial Congress*, Seoul, Korea, pp. 78-91.