The virtual museum: virtualisation of real historical environments and artifacts and three-dimensional shape-based searching

Paquet, Eric; El-Hakim, Sabry; Beraldin, Jean-Angelo; Peters, Shawn

Publisher's version / Version de l'éditeur:
https://doi.org/10.1007/978-1-4471-0337-0_20
Virtual and augmented architecture (VAA'01): proceedings of the International Symposium on Virtual and Augmented Architecture (VAA'01), Trinity College, Dublin, 21-22 June 2001, pp. 183-193, 2001

National Research Council Canada
Conseil national de recherches Canada

Institute for Information Technology
Institut de technologie de l'information

NRC-CNRC

The Virtual Museum:virtualisation of real historical environments and artifacts and three-dimensional shape-based searching. *

Paquet, E., El-Hakim, S., Beraldin, J.-A., and Peters, S.

June 2001

* published in: Proceedings of the International Symposium on Virtual and Augmented Architectures (VAA'01), Dublin, Ireland. June 21-22, 2001. pp. 183-194. NRC 44913.

THE VIRTUAL MUSEUM: VIRTUALISATION OF REAL HISTORICAL ENVIRONMENTS AND ARTEFACTS AND THREE-DIMENSIONAL SHAPE-BASED SEARCHING

Eric Paquet, Sabry El-Hakim, Angelo Beraldin and Shawn Peters

Visual Information Technology, National Research Council, M-50 Montreal Road
Ottawa, Ontario, K1A 0R6, Canada
eric.paquet@nrc.ca, sabry.el-hakim@nrc.ca, angelo.beraldin@nrc.ca, shawn.peters@nrc.ca

ABSTRACT

An intuitive technique for managing three-dimensional objects by shape is presented. This technique is applied to the Virtual Museum in the framework of heritage applications. The Virtual Museum is a virtual reality-based museum that allows the public to visualise, compare, manipulate and search the collection with an intuitive query by shape paradigm. The artefacts and the museum are virtualised with laser scanners and photogrammetric techniques. A high degree of realism can be attained by those techniques and the context in which the artefacts were discovered or were used can be recreated.

1. INTRODUCTION

The use of three-dimensional objects is relatively new in Internet-based heritage applications. Most heritage web sites rely principally on text and images. Nevertheless, the use of three-dimensional artefacts and scenes can drastically improve those sites by providing more realism, immersion, context and interaction. Furthermore, three-dimensional information can be used in order to search more efficiently the museum's catalogue. The virtualisation of historical sites and artefacts is interesting from many point of views. It allows a wider access: a virtual artefact can be accessed from anywhere at any time without restrictions. The artefact can be manipulated and viewed from any angle which, is not the case in a real museum. The virtualised artefact can also be viewed and studied simultaneously by an unlimited number of persons. Virtualisation also allows the access of fragile, closed, destroyed and remote sites. If the site or artefact is scanned with high precision and resolution then the virtual representation can be used as a reference version of the real artefact if the latest is damaged, lost, stolen or destroyed.

It is difficult to navigate through large collections. Most visitors do not have the historical background to navigate efficiently through a large museum and even the most seasoned expert becomes rapidly overwhelmed by the number of items. Text-based searching can help since most artefacts have been catalogued. But text-based search requires some knowledge about the vocabulary used in the corresponding field and it is language dependent. Language independence is important if the site is intended to have an international scope. Furthermore, some concepts like shape and appearance are very difficult to be explained with words alone. This is where content-based searching can help. Content-based uses three-dimensional shape and appearance in order to describe the artefacts. Shape and colour are very intuitive and can be very efficient to search historical databases. Many styles are characterised by their shape and or colour e.g. most Greek vase of a given period have the same colour and a Corinthian columns have a unique three-dimensional shape.

In this paper, a technique for searching three-dimensional object databases based on shape is first presented. It is explained how this technique can favour the whole museum-visitor interaction process. Finally a dedicated environment for three-dimensional virtual visit is presented: The Virtual Museum. The Virtual Museum supports realistic and fast rendering, extensive manipulation, reconfiguration, searching capability and virtualisation of real physical environments.

2. SEARCHING IN THREE-DIMENSIONAL SHAPES

The Visual Information Technology Group from the National Research Council of Canada has developed a system that describes the shape of a large set of three-dimensional objects in a totally automatic fashion. A descriptor is associated with each object. The descriptor is an abstract representation of their shape. The descriptor is very compact in size: it has typically 120 bytes. The size does not depend on the dimensions of the original three-dimensional object.

The descriptor is independent on the orientation of the original object. In order to achieve invariance, a reference frame that is orientation independent is defined. This reference frame is utilised in order to compute all the descriptors. All the calculations are performed by assuming that the objects are modelled as a set of connected triangles called tri-mesh. This is not a restrictive assumption because any mesh can be converted to a triangular mesh and any cloud of points can also be converted to such a representation. It can also be demonstrated that for any parametric surface representation like non-uniform rational B-splines or NURBS, a corresponding triangular mesh representation can be defined. Most of the time, our system can perform those conversions automatically.

The reference frame is obtained by computing the principal axis of the tensor of inertia. The principal axes correspond to the eigen vectors of the tensor of inertia. It can be shown that the normalised principal axes are translation, scale and rotation invariant. The possible mirror symmetry of the tensor of inertia is broken by a statistical analysis of the mass distribution.

The shape description is based on the concept of the cord vector. A cord is defined as a vector that goes from the centre of mass of a three-dimensional object to the centre of mass of a given triangle. The cord is not a unit vector since it has a length. The statistical distribution of the cord orientations can be represented by three histograms: two for the angles and one for the radius. This radius histogram is scale-dependent but it can be made scale-independent, by normalising the radius.

In order to process the data two basic problems need to be solved: the choice of the metric [1] and the number of channels or bins in each histogram. In order to determine the outcome of the search, the closest histograms to the reference histogram has to be found in the multidimensional space. A recent study [2] has shown that the Hamming and the Euclidean distance are the sole metrics that should be used. Any other metric is meaningless as far as similarity is concerned. This is because, from a statistical point of view, the points become equidistant if another metric is used. Those studies have also shown that the number of bins should be kept as small as possible. If the number of bins become too large, the density of points in the solution space increases and the discrimination is consequently reduced. Not only the discrimination is lost but also the probability to obtain false trues becomes very high. It is possible to reduce this undesirable effect by reducing the number of bins; but then some information is lost and the sampling can become inadequate. Although it is not clear at the moment, those results may indicate that the use of metrics may not be the best way to compare histograms. A histogram can be seen as a distribution. In such a case the metric can be replaced by a hypothesis test that determines if the two distributions are similar within a confidence interval. Clustering could also be a good approach since it does not directly based on metrics.

Actually the process can be made more robust by using a Bayesian analysis [3]. Assuming nothing is known about the objects, it is reasonable to say that our a priori knowledge about each bin can be summarised by saying that each bin is equiprobable and that only the average per bin is known. This is the most basic assumption that can be made. By using the maximum entropy principle it is possible to show, under those assumptions, that each channel of the histogram has a Poisson distribution. The variance of the signal is then given by the second statistical moment of this distribution. The bins are statistically independent. The signal in each bin is not known a priori which means that the prior has to be uniform. If the Baye's theorem is applied taking into account those assumptions it can be shown that the log of the probability is, up to a constant, proportional to the Chi-square. The Chi-square is given by the sum of the square of the differences between the bins of the reference histogram and the bins of the unknown object weighted by the inverse of the sum of their variance. In other words, it is possible to assign a level of confidence to each bin. If more a priori information is known, it is possible to refine the model.

3. VIRTUAL REALITY FOR HERITAGE APPLICATIONS: THE VIRTUAL MUSEUM

Visiting a museum is by nature an unorganised process: visitors look, touch, compare, and move in what seems to be a random walk. In an ideal case, a visit would be based on two paradigms that are absent in most if not all heritage web sites: comparison and query by example. In such a visit, the virtual curator shows some artifacts of interest to the visitor. If the visitor is not interested by those items, the virtual curator shows him additional artifacts that can be either similar or related.

Figure 1: Reconfiguration of the museum. A room representing contemporary furniture is shown. In the bottom view, the sofas, the tables, the light and the plant have been substituted.

If the choice becomes too difficult, the visitor considers all the items of interest and tries to reach a decision by comparing them on a relative basis. The first situation uses the query by example paradigm: the visitor provides an example and asks the virtual curator to find similar and related artefact. The second situation reflects the fact that for most peoples, it is easier to choose on a relative basis than on an absolute basis: that means that the artefact should not be isolated but placed in context.

For all those reasons The Virtual Museum has been created. The Virtual Museum is searchable and reconfigurable. The algorithms described in section 2 perform the search. The outcome of the search can be visualised through a substitution of the reference object or through a reconfiguration of the whole museum. In the first case the results are displayed sequentially. The visitor selects an object and asks for a similar or related object. The virtual curator determines the closest match and displays it by replacing the reference object with the closest match. All subsequent results are displayed accordingly

An application of this paradigm is shown in figure 1 for a collection of virtualised contemporary furniture. The top view shows an initial configuration for a living room while in the bottom view shows a room that has been reconfigured: the sofas, the tables, the lamps and the plan have been changed. This paradigm is very useful if the visitor wants to recreate various combinations of artefacts that were in effect at the time of their use. The fact that the pieces of furniture are displayed within a living room provides the context to which the visitor can refer.

The visitor can also reconfigure the whole museum with the sole outcome of the search: meaning that all other artefacts are removed from the scene. All the artefacts of interest are displayed within the museum. While the first paradigm allows interclass comparison, the second paradigm allows intraclass comparison. Again one may be looking for a Greek vase. If absolute visualisation is used, it might be very difficult to determine which one is the most interesting for him but if relative visualisation is used it is much easier. Relative visualisation is the paradigm that is the most commonly used in comparative studies. It is very useful both for professional and non-professional user. Of course, the visitor can bounce back and forth between the substitution and the reconfiguration paradigm. Furthermore, a division into collections allows the visitor to select a subset of the artefacts according to his own interests.

The Virtual Museum combines realistic rendering, multimedia integration, efficient and intuitive navigation and content-based searching capability. The mouse and the keyboard control the navigation. The mouse controls the direction of navigation and the inclination of the head while the keyboard controls the displacements. Many physical parameters are taken into account like inertia and collisions. The lighting is designed in order to favour the visualisation and to create the required atmosphere. Some shadows are pre-calculated in order to provide a more realistic rendering. The museum per se can be created with a computer or can be obtained from the virtualisation or digitalisation of a real scene. In the first case, the museum is made out of modules. Those modules can be easily modified in order to fit any particular requirement.

Virtualisation is by far the most innovative solution for creating sites [4]. High-resolution digital pictures of the scene or museum are taken. From those pictures a three-dimensional representation of the scene is reconstructed by using photogrammetry. A digital camera is first calibrated on an object of known dimensions. Once the calibration is completed, pictures of the scene of interest are taken with the same camera settings. It is also very important to take those pictures under the same lighting conditions in order to favour a seamless integration. The scene can be an outdoor scene, a building, a room or a full museum. Certain rules must apply while taking the pictures: there must be a common sub-region between the pictures and each picture must be taken from a distinct viewpoint. The configuration of the camera positions and angles, known as network design in photogrammetry [5], should be planned to increase the geometric accuracy of the reconstruction and reduce error propagation. For example the farther the distance or baseline between the viewpoints relative to the stand off distance, combined with a convergent angle, the better the quality of the reconstruction.

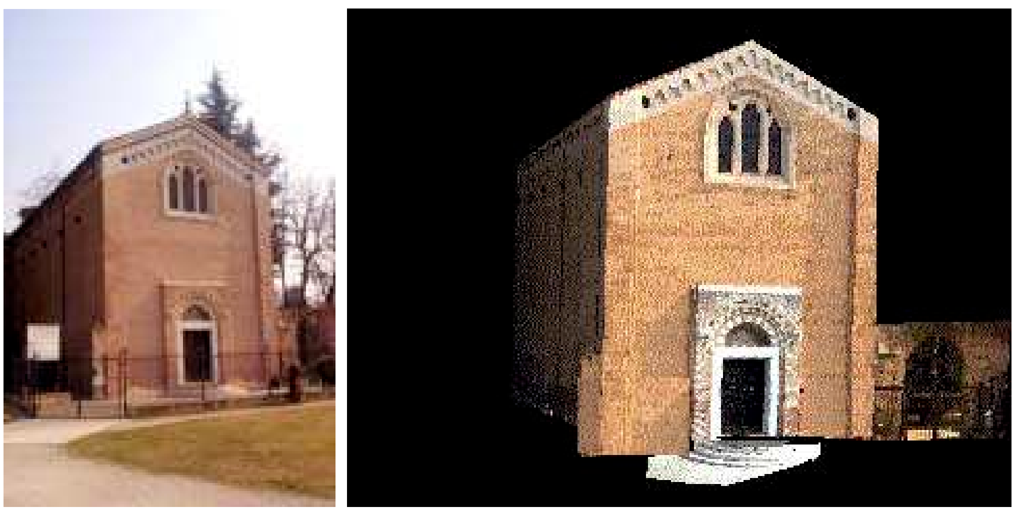

Figure 2: Virtualisation of the outside of the Chapel of the Scrovegni situated in Padova, Italy with photogrammetry.

The number of pictures needed depends on the complexity of the scene. While a few pictures are enough for a simple scene, many pictures may be required for a complicated one. Also, single images may be used if the scene is composed of known shapes such as planes or cylinders. The first step in the reconstruction process is to localise, with high precision, common points between the pictures. A photogrammetric algorithm uses those points and the calibration in order to reconstruct a three-dimensional model of the scene [4]. The number of points determines the resolution of the reconstructed geometry. Most of the time a triangular mesh represents the model on which are mapped one or more texture maps: the texture maps being subsets of the original pictures. The virtualisation is carried out on software developed at NRC. Photogrammetry has many advantages over traditional computer generated models: it provides highly accurate and realistic-looking models with non-repetitive texture. Those features increase the level of realism and the visitor confidence. Two examples of virtualisation of sites with photogrammetry are shown in figure 2 to 4. Figure 2 shows the reconstruction of the outside of Chapel of the Scrovegni situated in the city of Padova in Italy. Figure 3 shows the reconstruction of Constantine's Basilica in the Forum in Rome. Figure 4 shows an exposition of Canadian Native art within a stone room recreated with photogrammetry.

Figure 3: Virtualisation of Constantine's Basilica in the Roman Forum with photogrammetry.

The rendering is very important. The quality of the pictures on heritage sites is usually very high. Visitors are accustomed to this quality and expect a similar level for three-dimensional objects. One of the most viable solutions is texture mapping. In this technique, one or many pictures are mapped on the three-dimensional object. The texture can be obtained from a digital camera or from our polychromatic three-dimensional scanner based on synchronised scanning [6]. This innovative approach that is based upon triangulation allows very large fields of view with small triangulation angles and that without compromising precision. With smaller triangulation angles, a reduction of shadow effects is inherently achieved. The intent is to synchronise the projection of the laser spot with its detection. The instantaneous field of view of the position detector follows the spot as it scans the scene. The focal length of the lens is therefore related only to the desired depth of field or measurement range and not to the field of view. Implementation of this triangulation technique by an auto-synchronised scanner approach yields a considerable reduction in the optical head size compared to the conventional triangulation method. Furthermore, with optical synchronisation an improvement in ambient light immunity due to a small instantaneous field of view (long focal lengths) and a reduction of speckle noise (through spatial filtering) are achieved.

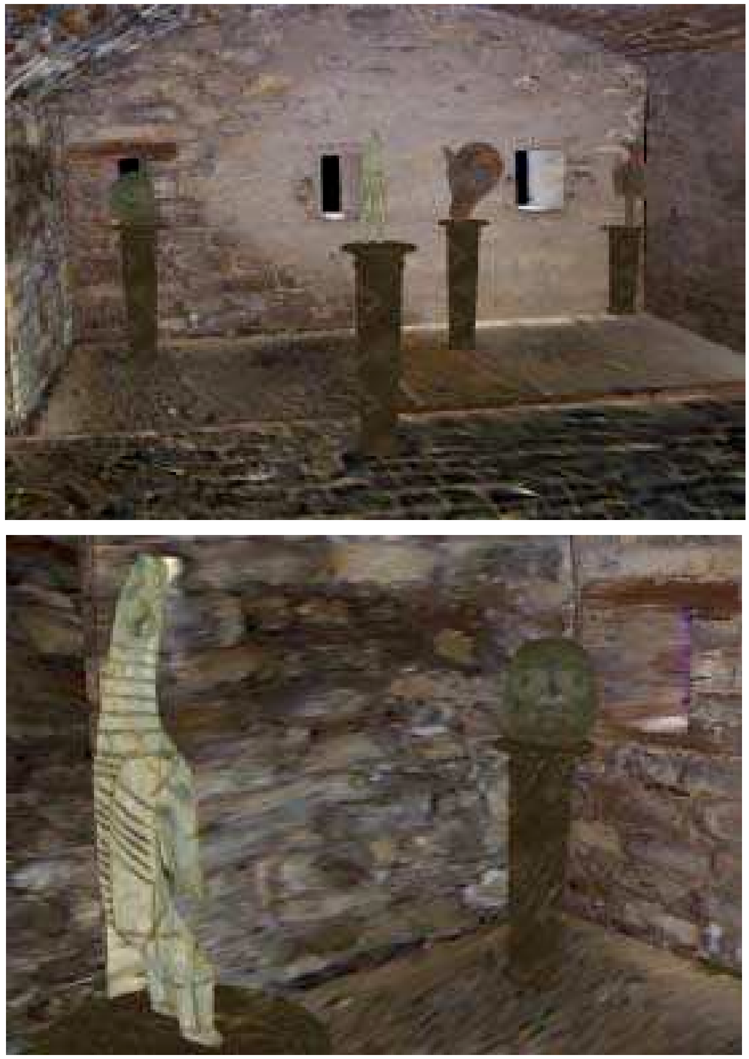

Figure 4: Canadian Native sculptures in a stone room. The sculptures were scanned with NRC's colour laser scanner and the room reconstructed with photogrammetry.

Colour is added to this auto-synchronised range camera by coupling a polychromatic laser (e.g. RGB laser) to the optical fibre that brings the light source to the 3D camera. On the scene, the laser light appears white. At the collection, colour separation is made on a single linear position sensor using a collinear dispersive optical element close to the collecting lens. The brightness signals measured by the position sensor at the three laser wavelengths depend on the absolute power of the laser, the distance from the camera to the surface element being measured, the orientation of that surface element relative to camera, and the physical properties of the surface itself. The goal of colour correction is to obtain unique values for the colour components that depend only on the physical properties, namely the reflectance, intrinsic to the surface element. Because the laser is intense and its characteristics do not fluctuate with time, the colours obtained are stable (e.g. at the three wavelengths used) and independent on the surrounding lighting conditions. Detailed information on the colour correction algorithms can be found in [7]. Other mathematical models to describe this phenomenon for digital cameras have been proposed in the literature [8]. Multimedia integration is also very important [9].

Inside the Cappella degli Scrovegni (Padova, Italy see 4) in the apse, there is an altar dedicated to the Virgin Mary. There stand three statues signed by Giovanni Pisano (circa 1305). The 3D model of the sculpture of the Madonna and Child (central one) was created with multiple range images taken at various locations all around a sculpture by. The views have enough overlap between them to find the registration and to merge them together. Innovmetric Software does the generation of the models. The final model is created after the merging process (removal of overlaps) or after compressing the polygons to a desired resolution. Figure 4 presents the whole statue, which contains approximately 1,500,000 polygons. All the range images for the sculpture, the vaults and the Altar table (not shown) were acquired in a single one-day session (10 h long). Finally, the model of the statue was created in a separate 3-day session of 6 h per day [10].

Figure 5: Cappella degli Scrovegni, Madonna and Child, G. Pisano (circa 1305):
a) portable biris system used to scan the sculpture, b) site in the Cappella degli Scrovegni in Padova, and, c) completed 3D model of the statue shown from a virtual point of view not available to tourists visiting the Chapel. Height is about 1.29 m.

The Virtual Museum can be made more attractive and more useful by the integration of multimedia. One may want to access the web page corresponding to an object or a video providing additional information about the artefact like excavation or a reconstitution of the period in which the artefact was in use. The user selects an object and then presses on the HTML or video button in order to access the corresponding media. The video can be streamed or downloaded. The HTML page is particularly adapted to provide text information and images.

One of the problems of three-dimensional virtual site is the time required to download the objects. Texture maps can be compressed by using the JPEG standard or the wavelet transform. The shape can also be compressed. In the Virtual Museum the following technique is used: the object is converted from a triangular mesh representation to a triangular strip representation. That means that the object is peal from top to bottom in a spiral fashion in order to generate a set of continuous strips of triangles. Then, instead of using the absolute location of the vertices, the relative location is used. Finally the relative position is coded on a restricted number of bits, saved in a binary format and zipped. Another approach is to cache the decor and the most commonly used objects on the local computer.

The Virtual Museum is coded with Java 3D. There are many reasons for that. Java 3D can run on any platform without modification. Java 3D is built on top of two API: the OpenGL (or Mesa for Linux) and DirectX: that means that it is possible to take advantage of both API without having to rewrite the code. Java 3D can be executed very rapidly: this is because Java 2 code is executed and not interpreted. Java 2 code is compiled on the fly by a compiler called just in time compiler or JIT. Because the JIT is process optimised and platform independent class files, compilation can be performed in almost no time. The Java 3D API is also very sophisticated and can handle complicated problems like synchronisation of events, collisions, high quality rendering and optimisation.

4. CONCLUSIONS

This paper discussed an intuitive technique for managing three-dimensional objects by shape in the context of Virtual Museum and especially in the framework of heritage applications. The Virtual Museum is a virtual reality-based museum that will empower the public to visualise, compare, manipulate and search the collection with we believe an intuitive query by shape paradigm. Virtualising artefacts and the museum with laser scanners and photogrammetric techniques has shown promising results.

The potential of 3D vision for applications in heritage or as an input to virtualised reality environments was discussed. A high degree of realism can be attained by those techniques and the context in which the artefacts were discovered or were used can be recreated. Real world acquisition and modelling is now possible. Technological advances are such that difficulties are more of a logistical nature than technology per se. Many techniques exist to digitise small objects with both a high-resolution 3D surface and a view independent surface texture with perfect registration between them. Models of large objects, structures and environments are possible but we are convinced that the combination of the current technique with other methods that are being explored in our laboratory and others around the world will be required. In the case of virtual and virtualised environments offer the possibility to expand our abilities in planning, creating and experiencing new surroundings, e.g. virtual tourism on sites that have been closed to the general public.

5. REFERENCES

[1] Y. Liu and F Dellaert, "A Classification Based Similarity Metric for 3D Image Retrieval", CVPR, pp. 800-805 (1998).

[2] K. Chakrabarti and S. Mehrotra, "Local Dimensionality Reduction: A New Approach to Indexing High Dimensional Spaces", VLDB, 89-100 (2000).

[3] J. H. Yi and D. M. Chelberg, "Model-Based 3D Object Recognition Using Bayesian Indexing", Computer Vision and Image Understanding 69, pp. 87-105 (1998).

[4] S .F. El-Hakim, "A practical approach to creating precise and detailed 3D models from single and multiple views." International Archives of Photogrammetry and Remote Sensing, Volume 33, Part B5A, Commission V, 122-129, Amsterdam, July 16-23 (2000).

[5] C.S. Fraser, "Network design considerations for non-topographic photogrammetry", Photogrammetric Engineering and Remote Sensing, 50(8); 1115-1126, 1994.

[6] M. Rioux, "Laser range finder based on synchronized scanners", Applied Optics, 23, 3837-3844, 1984.

[7] Baribeau, R., Rioux, M., Godin, G., 1992. Color reflectance modeling using a polychromatic laser range sensor. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 14 (2), 263-269.

[8] Rushmeier, H., Bernardini, F., 1999. Computing consistent normals and colors from photometric data 3-D Digital Imaging and Modeling. Proceedings Second International Conference on 3D digital imaging and modelling, Ottawa, Canada, pp. 99 –108.

[9] A. Biancardi and V. Moccia, "Integrating Applications into Interactive Virtual Environments", VISUAL'99, 703-710 (1999).

[10] Beraldin, J.-A., Blais, F., Cournoyer, L., Rioux, M., El-Hakim, S.H., Rodella, R., Bernier, F., Harrison, N., 1999. 3D Digital Imaging and Modeling, Proceedings Second International Conference on 3D digital imaging and modelling, Ottawa, Canada, pp. 34-43.