SENSORS AND ALGORITHMS FOR THE CONSTRUCTION OF DIGITAL 3-D COLOUR MODELS OF REAL OBJECTS

Marc Soucy[*], *Guy Godin*[†], *Réjean Baribeau*[‡], *François Blais*[†], *and Marc Rioux*[†]

[*] InnovMetric Software Inc.
2065, Charest ouest, Suite 218
Ste-Foy, Québec, Canada G1N 2G1
URL: http://www.innovmetric.com

[†] Visual Information Technology Group
National Research Council of Canada
Ottawa, Canada K1A 0R6
URL: http://www.vit.iit.nrc.ca

[‡] Analytical Research Services
Canadian Conservation Institute
1030, Innes Road
Ottawa, Canada K1A 0M5

ABSTRACT

This paper describes sensors and algorithms developed and used for the creation of coloured 3-D triangular meshes from a set of range images with registered colour measurements. The objective is the creation of a digital model of a real object that is compatible with established computer graphics techniques. Triangular meshes are the fundamental rendering primitive supported by most high performance graphics workstations. The surface colour of objects can be represented by attributing reflectance values to the vertices of the mesh. Such a geometric and reflectance model provides a uniform and general representation for sculptured surfaces with non-uniform colouring.

1. INTRODUCTION

This paper describes a suite of technologies (sensors and algorithms) and their application for the construction of 3-D colour models of real objects. First, a range sensing technology that allows simultaneous capture of range and colour is described. Intrinsic reflectance properties are then recovered from the measured colours and the range image, through the inversion of an illumination model. Since more than one view is usually required for most objects, it is necessary to recover the rigid transformation that brings the range data from a camera-centered coordinate system to a single common system: a registration method has been developed to achieve this goal. A key step in the process is the construction of a non-redundant triangular mesh that represents the original data at full resolution, with an intrinsic colour value at each vertex. The resulting models can then be decimated by compressing the triangulation within a preestablished error threshold. The colour information of the removed vertices is then transformed into a texture map to be applied over the remaining triangles. This second form of models is particularly suited for dynamic viewing, such as interactive manipulation or virtual reality. The paper ends with a discussion of a successful demonstration, in which this suite of technologies was used to model and display museum artifacts.

2. SENSOR

The range sensor used to acquire colour reflectance of an object is based on the auto-synchronized laser scanning technique developed at the National Research Council of Canada [4]. This system makes use of a double-sided scanning mirror and, for colour acquisition, of a mixture of red, green, and blue laser beams. Laser light is projected on the object from the first side of the scanning mirror. It is diffusely reflected by the object, captured by the second side of the scanning mirror and imaged on a CCD photo-detector after spectral separation through a prism. The position and intensity of the peaks formed by the three primary red, green, and blue wavelengths of the white laser beam are recorded. This results in three separate range and intensity images that are merged together into a single 3-D RGB image. As a single laser beam is used for the simultaneous measurement of colour and range, 3-D and colour data are in perfect registration. The use of laser illumination also accounts for the stability and repeatability of colour measurements.

3. ALGORITHMS

Researchers at InnovMetric Software, at the National Research Council of Canada, and at the Canadian Conservation Institute, have jointly developed a 3-D modelling methodology for constructing coloured 3-D triangulated meshes from a set of 3-D RGB images, as depicted in Figure 1. The first three processing steps, enclosed in a dotted box, form the acquisition loop. The 3-D RGB images are acquired one by one, until all the object's surface has been sampled. A semi-automated software tool is used to interactively align the different surface meshes, allowing the human operator to rapidly detect areas that have not yet been measured by the sensor. The five processing steps following the acquisition loop form the 3-D modelling sequence. With the exception of the optional Manual Editing step, all these modelling operations are completely automated processes that are controlled by a very small number of intuitive parameters.

0-7803-3258-X/96/$5.00 © 1996 IEEE

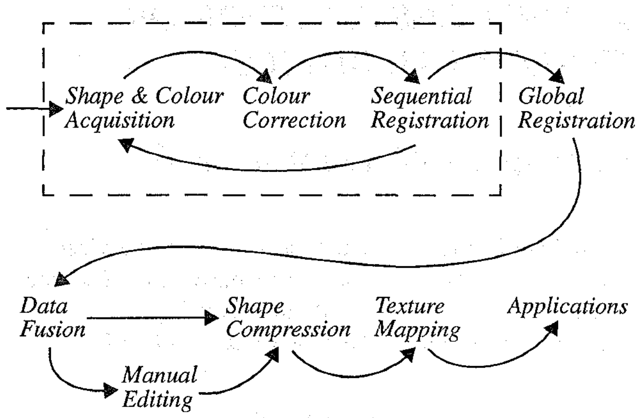

Figure 1 This figure illustrates the methodology used in constructing coloured 3-D triangulated meshes of an object. The dotted box forms the acquisition loop, while the other modules form the 3-D modelling sequence.

3.1 Colour Correction

The brightness signals measured by the CCD at the three laser wavelengths depend on the absolute power of the laser, the distance from 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. To achieve this, the sensor is first calibrated using a white diffusing target that is scanned at various depths in the working volume. This establishes the spatial brightness sensitivity of the sensor. The optical geometry of the camera is also modelled so that for any given surface element, one can calculate the angle between the projected laser beam and the surface, as well as the angle at which it is detected. Reflectance values are obtained by comparing the signal from surface elements to that from a white standard kept near the border of the scene, taking into account a reflectance model as well as all the calibration data. Because the colour information obtained this way is intrinsic to surface elements, colour data collected from multiple viewpoints can be integrated seamlessly. More information on the colour correction algorithms can be found in [1].

3.2 Registration of Multiple 3-D Images

Since a single 3-D image can only partially describe the surface of an object, it is necessary to acquire several 3-D images in order to entirely measure an object. Each 3-D image has its own coordinate system, related to the sensor's position and orientation in the 3-D world. As a result, a processing step is needed in which all image coordinate systems are transformed into a unique coordinate system for the purpose of object modelling. The operation of estimating the rigid inter-frame transformations between a set of 3-D images is called registration. We propose an innovative two-step registration technique that does not rely on the calibration of a mechanical system, and allows the user to freely move the object in front of the 3-D sensor.

Both the sequential and global registration steps are based on an automated least-squares algorithm that aims at minimizing the 3-D distance between surface overlaps in a set of 3-D images [3]. In the acquisition loop, the human operator uses this least-squares algorithm to compute a very good approximation of the global registration of the set of images. When a new 3-D image is processed, the operator will begin by interactively aligning the image surface with respect to the previously registered images. The least-squares algorithm is then used to accurately register the new image while all previously registered images are locked. Registering a new image usually takes between 1 and 2 minutes. The operator then rotates and translates the acquired surface data, in order to detect areas where additional scans are needed. The Acquistion-Registration steps are repeated until all the object's surface has been measured.

The sequential alignment procedure described above provides a very good approximation of the optimal registration of a set of 3-D images. This approximation is not accurate enough, however, to generate a 3-D model from the measured 3-D data. It is well known that this kind of sequential procedure suffers from the phenomenon of cumulative error. Once all 3-D images have been acquired, their registration is globally optimized by applying the iterative least-squares process to the whole set of 3-D images, in a parallel manner. For this purpose, only one 3-D image is kept fixed, and all other images are allowed to move. A Maximum Distance threshold controls the maximum normal distance between two matched surface locations, and is typically set to 10 times the standard deviation of the sensor noise.

3.3 Integration of a Set of 3-D RGB Images

Once all 3-D RGB images have been transformed into a common coordinate system, they are merged by a general integration technique based on the reparameterization of the canonic subsets of the Venn diagram [5][6]. This step is identified by the Data Fusion block in Figure 1. The output of the integration algorithm consists of a high-resolution coloured 3-D triangulated mesh in which RGB colours have been computed for each mesh vertex. Vertices are uniformly distributed over the surface described by the integrated mesh and the vertex density is determined by the operator. When a surface area is measured by more than one 3-D image, the shape and colour measurements provided by several images are merged using a weighted average [6].

3.4 Shape Compression and Texture Mapping

Texture mapping is a well-known computer graphics technique used to provide the illusion of high levels of details on models of moderate geometric complexity by the application of colour patterns onto coarser geometric structures. Because this mode is usually faster than the display of full-resolution geometry, it is particularly useful for dynamic displays where rendering speed is more important than detail accuracy. In this paper, we demonstrate new shape compression and texture mapping algorithms that enable the creation of highly realistic and compact triangulated representations.

The implemented shape compression algorithm is based on previous work [7]. A sequential optimization process iteratively removes vertices from the high-resolution triangulation, always minimizing the retriangulation error. The vertices used to generate compressed triangulations are thus part of the original set of vertices. The operator controls the compression process by specifying a 3-D tolerance level, defined as the maximum 3-D distance between the original model vertices and the compressed surface representation. Once the triangulation error crosses a specified tolerance level, the compressed representation is saved in a file. Two major improvements have been brought to this original algorithm. First, the time complexity has almost been reduced to $O(n)$, which makes the compression of large models more tractable. Second, the new compression technique maintains a mapping between the original vertices and the compressed model. Once a model has been compressed, all removed vertices are attributed the barycentric coordinates (u,v,w) of their projection on the larger triangles of the remaining model. Thus each triangle of the compressed model has a colour triplet at each vertex, as well as a variable number of removed vertices mapped onto its planar surface, each carrying three barycentric coordinates (u,v,w) as well as a colour triplet (RGB). The colour information carried by the removed vertices is then used to generate a high-resolution texture map for a compact triangulated mesh.

Most 3-D graphics workstations have the possibility of handling (either in hardware or software) texture-mapped models for which a texture is parametrically mapped onto geometric primitives. In the approach described here, texture can be used as a way to represent, at least partially, the colour of vertices that have been removed by the geometric compression process. The texture map is organized as a rectangular array of colour values. In the cases where texture mapping is implemented in hardware, it provides very fast and efficient rendering. As the amount of texture space is relatively limited (in the order of a few Mbytes even on high-end workstations), making efficient use of this valuable resource is a key concern. Since the triangles forming the compressed model may vary widely in size and angles, and the model's topology is not restricted, the mapping of all the model's triangles is not a straightforward process. An algorithm developed by Soucy and Godin takes as input the compressed model with the barycentric coordinates of removed vertices and generates a rectangular texture map of a size chosen to fit the graphics hardware. The algorithm operates by first tessellating the rectangular map into triangles of sizes related to the 3-D length of the compressed model's triangles, in order to ensure an approximately uniform texture pixel density on the surface of the model. The original model vertices are then mapped onto the texture map, and their RGB values are assigned to the nearest texture element. Finally, a modified fast distance transform algorithm is used to assign RGB values to empty texture elements whose RGB values has not been set by a colour 3-D vertex. In order to maintain the visual quality, texture continuity is enforced between triangles that are adjacent on the model.

4. EXPERIMENTAL RESULTS

Figure 2 clearly demonstrates the high visual impact of texture-mapped models. A coloured 3-D triangulated mesh of a wooden duck was built from 12 three-dimensional RGB images. The images were corrected, registered, and integrated, yielding an integrated mesh containing 166278 triangles. The resulting high-resolution coloured triangulated model is shown on top of Figure 2. The shape compression algorithm was then applied to this model for creating a compact 1000 triangle representation, and an associated 512x512 texture map was computed. The texture-mapped model is shown on the bottom of Figure 2. The differences between the top and bottom models are almost imperceptible. An additional figure shows a wireframe representation of the compressed model (Figure 3).

5. APPLICATIONS

The process described here allows the creation of digital models for practically any object that can be imaged by active range sensing. The coloured or textured triangular mesh is a generic representation that can handle complex geometry and colour patterns. Among the main motivations behind the work described here is the creation of digital 3-D colour models of museum artifacts that can be stored on disk or transmitted over communications networks. These models can then be displayed and examined using a graphics workstation. Computer graphics techniques can be used to simulate particular lighting conditions, since the intrinsic colours recovered in the step described in Section 3.1 correspond to the reflectance data required by lighting models.

As a joint effort by National Research Council of Canada, InnovMetric Software, and the Canadian Conservation Institute, an actual demonstration of this technology was held at the Canadian Museum of Civilizations in Hull (Quebec), as part of the "Mothers of Time: Seven Palaeolithic Figurines from the Louis Alexandre Jullien Collection" exhibition. Palaeolithic stone and ivory figurines of sizes varying from 3 to 7 cm were digitized using the polychromatic 3-D colour sensor, and the models were prepared using the methodology described above. A graphics workstation connected to a large screen was installed in the gallery where the figurines were on display. The visitors were able to manipulate the models and examine particular details of the objects. When moved by the user or a predefined animation sequence, compressed texture-mapped models were used, whereas static viewing displayed the full resolution models. The availability of a digital model also permitted operations that would have been unthinkable on real artifacts. For example, the figurines could be rendered as if they were made of a pure white material simply by replacing the measured colour values by a "pure white" value; shading of the model then enhanced geometric details, such as tool marks, that would otherwise be drowned in the colour texture.

6. ACKNOWLEDGEMENTS

The authors would like to thank Luc Cournoyer, J.-Angelo Beraldin, Denis Laurendeau, and Sabry El-Hakim for their participation in the development of these technologies. Special thanks to Jacques Domey for the support of the whole project.

7. REFERENCES

[1] R. Baribeau, M. Rioux, G. Godin, "Colour reflectance modelling using a polychromatic laser range sensor", in *IEEE Trans. on PAMI*, vol. 14, no.2, pp. 263-269, 1992.

[2] F. Blais, M. Rioux, "Real-time numerical peak detector", in *Signal Processing*, vol. 11, pp. 145-155, 1986.

[3] H. Gagnon, M. Soucy, R. Bergevin, D. Laurendeau, "Registration of multiple range views for automatic 3-D model building", in *Proc. of IEEE Conf. on CVPR*, Seattle, Was., pp. 581-586, June 21-23 1994.

[4] M. Rioux, "Laser range finder based on synchronized scanners", *Applied Optics*, vol. 23, no. 21, pp. 3837-3844, Nov. 1984.

[5] M. Soucy, D. Laurendeau, "A dynamic integration algorithm to model surfaces from multiple range views", in *Machine Vision and Applications*, 8:53-62, 1995.

[6] M. Soucy, D. Laurendeau, "A general surface approach to the integration of a set of range views", in *IEEE Trans. on PAMI*, vol. 17, no.4, pp. 344-358, April 1995.

[7] M. Soucy, D. Laurendeau, "Multi-resolution surface modeling based on hierarchical triangulation", in *Computer Vision and Image Understanding*, vol. 63, no. 1, pp. 1-14, January 1996.

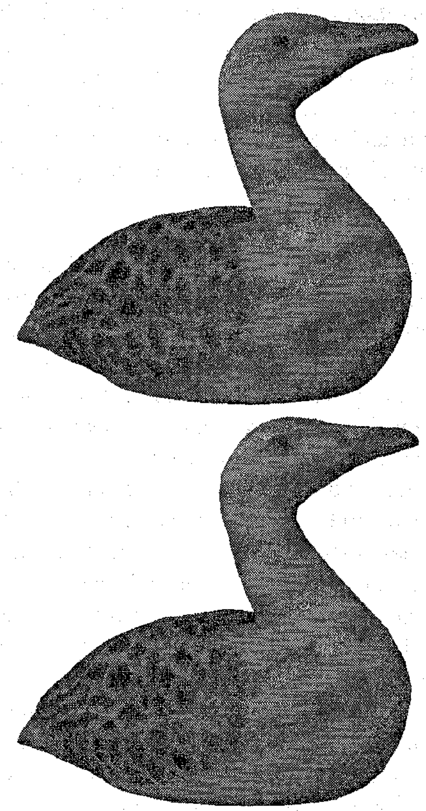

Figure 2 The top image is a high-resolution coloured 3-D mesh of a duck containing 166278 triangles and 83141 vertices. The bottom image shows the texture-mapped representation of the duck that only contains 1000 triangles.

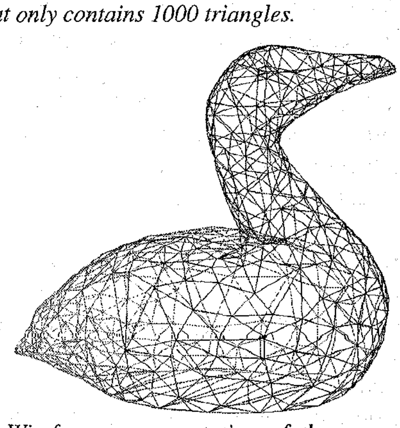

Figure 3 Wireframe representation of the compressed duck model made of 1000 triangles.