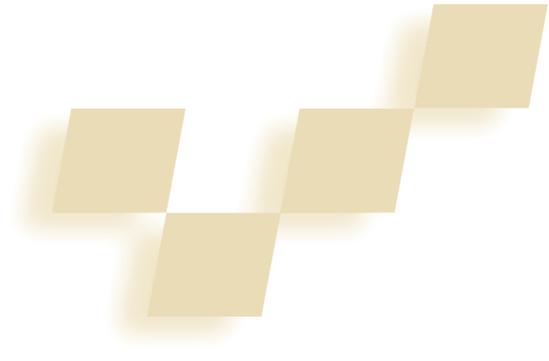


# Building a Digital Model of Michelangelo's *Florentine Pietà*



Fausto Bernardini, Holly Rushmeier, Ioana M. Martin, Joshua Mittleman, and Gabriel Taubin  
IBM T. J. Watson Research Center

**B**ecause scanning devices are less expensive and easier to use than they were just a few years ago, a wide range of applications employ 3D scanning technology. As a result, various organizations can now produce models of cultural artifacts and works of art.

Members of the National Research Council of Canada, developers of high-accuracy scanning equipment, have applied their technology to scanning paintings, sculptures, and archaeological sites. Recent work emphasizes the importance of portable, reliable equipment that researchers can easily deploy at the scanning site.<sup>1</sup> Zheng and Zhong scanned archaeological relics in cooperation with the Museum of Terra Cotta Warriors and Horses, China.<sup>2</sup> Their goals included creating a database of information about the excavation site and testing and employing virtual restoration techniques. Recently, Marc Levoy and a team from Stanford University scanned many of Michelangelo's sculptures,<sup>3</sup> including the 5-m tall

*David* in the Galleria dell'Accademia. They used several types of scanners, including a high-resolution laser triangulation system mounted on a custom-made mechanical gantry and a time-of-flight long-range sensor. The large quantity of data collected should have a major impact in future development of shape reconstruction algorithms. Numerous other projects have been conducted or are currently underway.

In this article, we describe a recent project to acquire and build a 3D model of the Michelangelo's *Florentine Pietà*. Figure 1 (next page) shows a photograph of the *Pietà* and an image of our model. The work we describe here is unique in that an art historian, not a technologist, conceived and specified the project. Our goal wasn't simply to produce the statue's model but also to provide the art historian with material and tools to enable him to answer his own research questions. The project gave us

the opportunity to explore the value of 3D scanning and visualization in a nontechnical discipline, art history. The project's second goal was to develop scanning technology accessible to other cultural heritage projects both in terms of cost and usability. Such technology could potentially be used in widespread commercial applications, such as e-commerce, where equipment cost must be minimal.

## Overview

Jack Wasserman, professor emeritus of art history at Temple University in Philadelphia, had been studying Michelangelo's *Florentine Pietà* for several years. He intended primarily to document all aspects of this important work and its history for future researchers, and secondarily to investigate his own theories on Michelangelo's composition. He had used high-quality traditional photography, x-ray, and ultraviolet light studies, as well as researching the complex history and symbolism of the statue and its analysis by past art historians.

Although it's not clear that a 3D model would be useful in studying every sculpture, Wasserman felt that this new technology was especially well suited to studying the *Pietà*.<sup>4</sup>

Accounts from Michelangelo's contemporaries tell us that the artist intended the *Florentine Pietà* to serve as his own tomb monument. Beginning late in his life, in the 1550s, he carved four larger-than-life figures from a single block of marble. The Christ figure in the center rests across the lap of the Virgin Mary, supported on the left by Mary Magdalene. Behind and above, supporting the statue of Christ, is a figure believed to represent Nicodemus and to bear Michelangelo's face. At some point, Michelangelo decided to break off parts of the statue. He then abandoned it. Shortly before his death he permitted one of his students, Tiberio Calcagni, to repair the statue. Calcagni reattached several pieces to the statue and partially finished the figure of the Mary Magdalene. Thus, what we see today is, in a sense, a composite of Michelangelo's work and his student's: the original design damaged, repaired, and overlaid by later work.

The unique aspects of the history of this statue make

We present the methodology we used to acquire the data and construct a computer model of Michelangelo's *Florentine Pietà* with enough detail and accuracy to make it useful in scientific studies.

1 A photograph of Michelangelo's Florentine Pietà (left). A synthetic picture from our 3D computer model (right).



it a promising candidate for using 3D scanning technology. It's important for art historians to view the statue in the environment Michelangelo intended, examine it without the pieces Michelangelo removed, and analyze the detailed toolmarks in the unfinished portion of the work. Furthermore, the statue's complex geometry limits what can be done with traditional techniques. A camera can't capture certain views of the statue because the statue itself or the walls of the room where it stands interfere with proper camera placement.

**Scanning system and methodology**

Three-dimensional scanning technology is evolving rapidly. Current sensors use a number of techniques such as laser triangulation, laser time-of-flight, passive stereo vision, and structured light projection to sample surfaces. Typical considerations in choosing the most appropriate scanning technology include target accuracy, surface reflectance characteristics, and cost.

**Design considerations**

Scanning a large statue in a museum poses a number of constraints in designing the scanning system and process. In our case, the small size of the room in which the statue is displayed limited scanner size and the distance from the scanner to the statue. We didn't have per-

mission to work or leave any equipment visible around the statue when the museum was open to the public. Therefore, we needed a system that we could easily set up and remove at the beginning and end of each evening scanning session. The irreplaceable nature of the piece restricted contact to a minimum and required us to ensure that we operated the scanner system safely. The complex shape of the group of figures in the Pietà required the ability to freely position the sensor to access the marble surface's recessed parts.

We had a limited budget for buying noncomputer equipment and a limited amount of time for design and customization. These constraints led us to consider a small, portable structured-light system rather than a more expensive laser triangulation scanner. By this choice, we sacrificed geometric resolution, which we would have had to recover with a supplementary system.

Our main technical requirements were dictated by the nature, resolution, and accuracy of the data needed to address Wasserman's needs. The goal was to obtain data to allow realistic rendering of the synthetic model. The statue is 2.25 meters tall, and we wanted to capture its shape and surface details, such as cracks and toolmarks, on the scale of 1 to 2 mm. Besides geometry, we were interested in capturing the surface's reflectance properties. We therefore needed to achieve submillimeter accuracy in measurements.

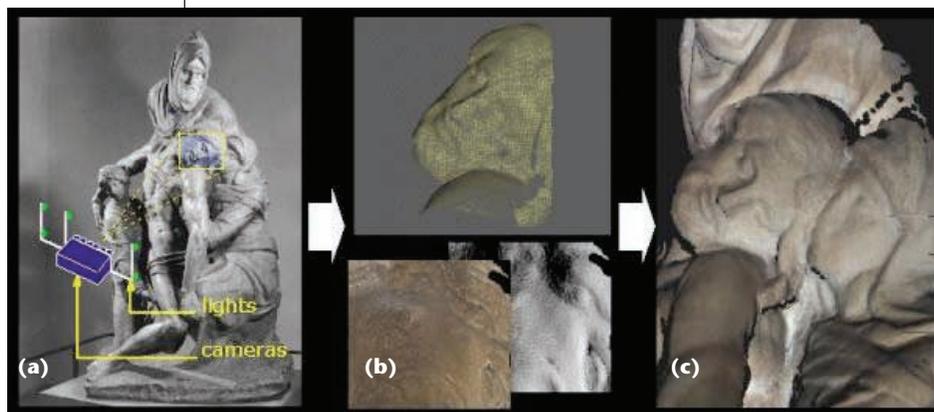
Capturing a large object at such fine resolution entails a number of difficulties, especially under the less-than-ideal conditions outside a lab. Issues of repeatability and precision make scanners based on moving parts expensive to build and difficult to transport and operate. Sub-surface scattering of laser light in marble limits the accuracy that laser triangulation systems can achieve. We decided to use a system that could capture a small portion of the surface from a single position, acquire a large number of overlapping scans, and rely on software registration to integrate the results.

The amount of data we needed to represent the surface at such a fine level of detail presents additional problems. For example, we can't store, process, and visualize a triangle mesh of hundreds of millions or billions of triangles on current PCs or midrange workstations. Because we aimed to make the results accessible to a

wide audience, we decided to represent shape as a triangle mesh with a resolution of a few millimeters and to store additional fine geometric details as a normals map. Plus, we could store reflectance values as RGB image maps. Having thus chosen the final representation of our model, we avoided a great deal of intermediate computation by designing a system that captures data directly in that format.

**Scanning**

Figure 2 shows a schematic of our 3D capture methodology. Our scanner is based on a multibaseline stereo system, supplemented by a



2 Our 3D capture methodology. (a) Take multiple digital photos. (b) Compute each scan's surface shape, color, and details. (c) Align and merge the scans into a single model.

photometric system. The scanner, visible in Figures 3a and 3b, is a customized version of the Virtuoso ShapeCamera. A photographic flash projects a pattern of vertical stripes on the subject. At the same time, six black-and-white digital cameras photograph the illuminated area from different angles. Figure 3c shows a sample stripe image. An additional digital color camera provides a texture image. A multiview stereo algorithm,<sup>5</sup> part of the software suite provided with the Virtuoso system, computes a triangle mesh approximating the scanned area. In our scanning conditions, each scan typically covered a 20 cm by 20 cm area and comprised on average about 10,000 measured points. The typical intersample distance for these scans is about 2 mm. In tests we conducted on reference objects, we were able to accurately measure a depth of 0.1 mm for a single scan.

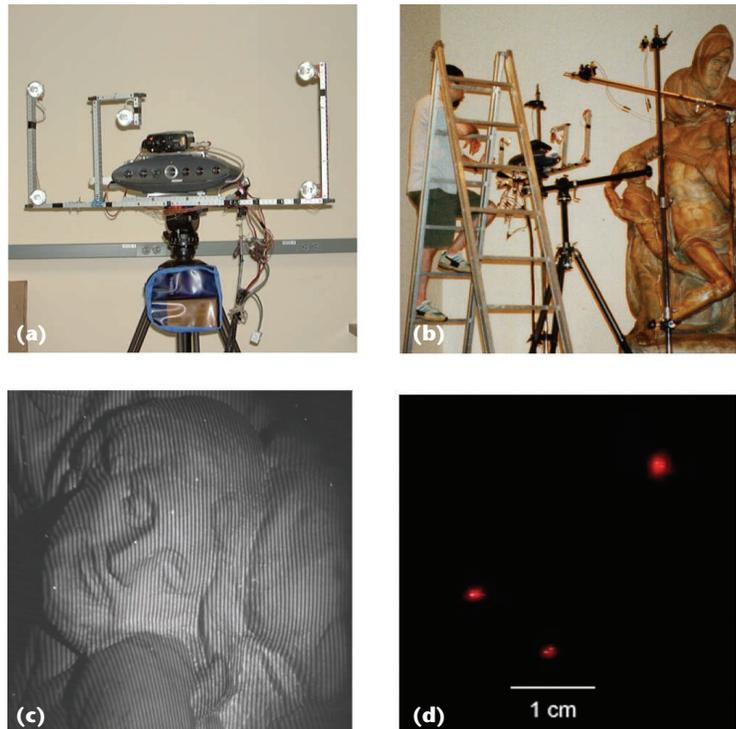
The photometric system (Figures 3a and 3b) consists of five light sources and the built-in color camera, plus some control electronics. For each camera pose, we took five additional color pictures, each with one of the five light sources, while all other lights were turned off. We also used low-power laser sources to project red dots onto the statue (shown mounted on light stands in Figure 3b). The laser projectors that we used each generated an  $11 \times 11$  grid of rays. From a distance of about 1 meter, they produced an irregular pattern of red dots on the statue, with an average spacing of 2 to 4 cm. For each pose, we took a picture of the dots (with all other light sources turned off) to help align overlapping meshes (see Figure 3d). The color pictures have a resolution of  $1280 \times 960$  pixels, with 24 bits of RGB color per pixel. Typically we had a 0.5-mm intersample distance in the color images. As a result, we computed reflectance and surface normals from these pictures at a resolution about four times greater than the underlying geometry.

Our initial design included a magnetic tracker to record an approximate estimate of the camera position and orientation with respect to a global frame of reference. We hoped to use this pose estimate to provide a starting point for our software registration process. We used a Polhemus system, fitted with the long-range source to provide an electromagnetic field large enough to cover our work volume. We attached a sensor at the tip of a 40-cm plastic rod, rigidly secured to the camera body. Unfortunately, we quickly discovered that metallic material in the room, including our own equipment, distorted the field to the point of making measurements useless. We also had initially planned to use additional hardware and software to remotely control the scanner

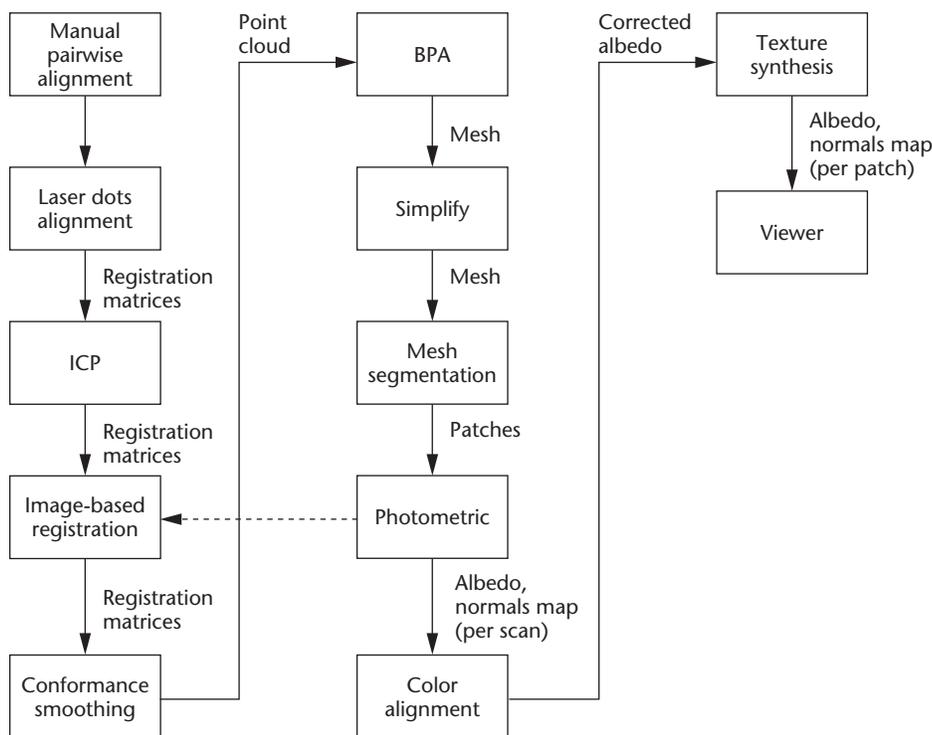
and facilitate data transfer operations. However, to keep setup and teardown time to a minimum, we simplified our system considerably.

Our streamlined procedure consisted of the following steps. We positioned the large photographic tripod and secured the scanner to it. Then, we placed the five laser projectors on three light stands to cover the area to be scanned in one session with a grid of laser dots. We captured the data by placing the scanner at one area to be covered and shot one set of pictures, then we moved the scanner across the target area to take successive overlapping picture sets, covering the region with a regular pattern of scanned tiles. We kept track of the approximate area covered by each picture set on paper diagrams of the statue. We estimated that we had enough overlap by comparing previews on the scanner display. We moved the scanner conservatively, about 10 cm between shots, to ensure that we had enough data. We processed the stripe pictures during the day, before the next evening's scanning session, to make sure that we had acquired enough data and not left any holes. One person operated the scanner, while another acted as supervisor to ensure that the proper standoff distance was respected and safety rules followed. Wasserman was present during the entire process to provide input on his priorities.

We did a preliminary scan of the statue (without the photometric system) in February 1998, spending five six-hour evenings and four full days in the museum. We repeated the scan using the photometric system during



**3** (a) The scanner used in our project. We added the five-light photometric system to a Virtuoso ShapeCamera. (b) The scanner in use in the museum. The stands visible in the picture held the laser projectors. (c) Detail of one of the six stripe images simultaneously taken by the scanner. (d) Detail of the laser dots projected on the statue, as captured by the color camera mounted on the scanner.



4 The reconstruction pipeline.

two one-week visits, one in June 1998 and one in July 1999. The total time spent doing the final scanning was about 90 hours over 14 days, including the equipment setup and teardown each day. It took about 800 scans to cover the whole statue. The raw data consisted of 4,800 640 × 480 pixel, 8-bit gray-scale stripe pictures and 4,800 coregistered 1280 × 960 pixel, 24-bit RGB color images. Stored in lossless compressed format, the raw data occupies 3 Gbytes of storage.

In retrospect, we believe that we chose the right scanning technology, although with additional planning and design we could have built a more efficient system. The main bottlenecks in the process were the scanner’s relatively long cycle time, the small area covered by each scan, and the offline processing of data. The time required to complete the acquisition and local storage of one set of images was about 2 minutes, and we required about the same amount of time to process one set of striped images to obtain a triangle mesh. Tracking the camera pose would have saved us from doing the pairwise manual alignment of the scans that provided a starting point for our registration algorithms.

The main advantage of our scanning system, besides meeting the requirements of our original design, is that it potentially provides a starting point for future development of a low-cost system built out of commodity components. If it’s augmented with reliable tracking, fast capture, and high-resolution cameras, it could lead to a system for real-time scanning of large objects.

### Reconstruction pipeline

The acquired raw data consisted of roughly 800 scans, each consisting of six black-and-white stripe

images and six color images. We used the Virtuoso Developer software to compute triangle meshes for each single scan from the six stripe images. From this point on, we applied several algorithms to the data to build the final model. We registered the individual scans together based on matching geometric and image features. We remeshed the resulting point cloud to obtain a seamless geometric model. We extracted the color and detail information from five of the color images and reorganized them in the form of normals and reflectance maps. Figure 4 illustrates the sequence of steps.

### Registration

We start with a pairwise, manual, approximate alignment of the meshes, obtained interactively by selecting three or more matching features on overlapping pairs of scans. We used the diagrams recorded during scanning to identify sequences of overlapping scans and constructed a tree of pairwise

alignments that spans the whole set of scans. We needed to do this initial manual alignment because the tracking hardware we intended to use didn’t perform satisfactorily in the museum. We progressively refined the alignment using several registration algorithms that used geometric and image information at increasing resolution levels.

For each scan, we found the red dots in the image taken with the laser projectors on and mapped these image points back onto the scan geometry. Given the initial manual alignment, we searched in the neighborhood of each laser point for matching points in overlapping scans, adding additional consistency constraints to prune false matches. We improved the registration by minimizing the sum of square distances between matching points using Besl and McKay’s method.<sup>6</sup> We then ran several iterations of an *n*-scan iterated closest point (ICP) algorithm<sup>7</sup> to further reduce the registration error. Additional details of our geometry-based alignment appear elsewhere.<sup>8</sup>

To refine the geometry-based alignment obtained with the ICP algorithm, we applied an image-based registration method that considers additional information contained in the high-resolution reflectance maps computed for each scan (see the following section). We used a combination of smoothing, edge-detection, and thresholding operations for selecting candidate points in feature-rich areas of each image. We conducted a correlation-based search in images associated with overlapping scans to find corresponding points. We subsequently back projected the resulting pairs onto the scans and derived a rigid transformation that minimizes distances in a least-squares sense. See Bernardini, Mar-

tin, and Rushmeier<sup>9</sup> for additional details of the image-based registration phase.

We also attempted to reduce scanner line-of-sight errors by computing more accurate estimates of true surface points from multiple overlapping scans while filtering out small high-frequency components (which the photometric system can better capture). We call this process *conformance smoothing*. Figure 5 shows an example of the successive alignment steps on three test scans of Nicodemus' face. The finer grain of the color variations after the conform step indicates that the shape of the overlapping scans are nearly the same. Experiments showed we could improve the registration error if we accounted for the line-of-sight error during the alignment. We're experimenting with alternating iterations of registration and conformance smoothing to obtain a better final alignment.

We didn't have equipment to accurately measure large distances between points on the statue (for example, from the top of Nicodemus' head to a point on the base). Therefore, we were unable to quantitatively state the global accuracy of the alignment. We discuss the validation of our results using 2D projections and photographs in the section "Validating and using the model."

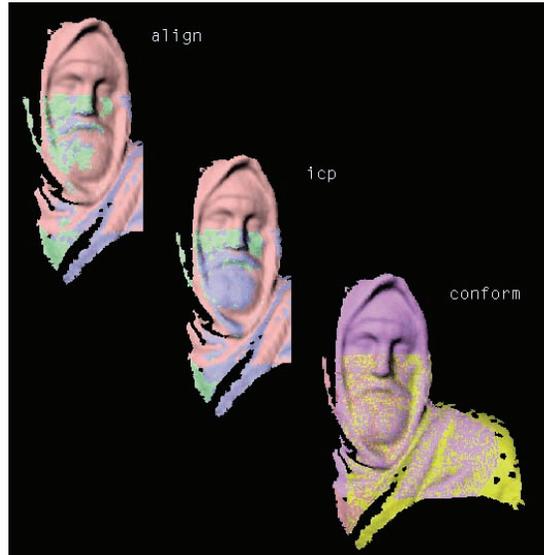
### Meshing

The result of the alignment and conformance processing is a large set of estimated surface samples. This point cloud has nonuniform density because the number of overlapping scans varies from one part of the surface to another and because the density within each scan varies locally depending on the angle at which the scanner saw the surface. However, except for areas that the scanner couldn't reach, the sampling density is usually larger than necessary to recover the surface's shape to a reasonable approximation. We designed our system to acquire geometry with an average intersample distance of 2 mm. Note that this spatial resolution is independent of the accuracy in measuring point position. The scanner we used has a precision of 0.1 mm in computing depth for each sample point.

The ball-pivoting algorithm (BPA)<sup>10</sup> computes a triangle mesh interpolating the point cloud, using a region-growing approach. Our implementation of the BPA handles large data sets in a memory-efficient way, by processing input data in slices.

The *Pietà* data consist of 800 scans containing a total of 7.2 million points. We processed the data in slices of 10 cm, using ball radii of 1.5, 3, and 6 mm. The BPA ran in 30 minutes on a Pentium II PC using 180 Mbytes of memory and output a 14-million triangle mesh.

We applied a mesh simplification algorithm to generate a hierarchy of models at different resolutions. We found that conventional, in-core, simplification algorithms couldn't handle the large mesh generated from our data. We computed simplified models by breaking up the mesh into smaller, manageable pieces. We then applied a traditional, high-quality simplification algorithm,<sup>11</sup> leaving the boundary of each piece untouched and stitched the resulting simplified pieces together. In a successive pass, we broke up the mesh along different edges, so that we could simplify the previous bound-



5 Subsequent steps of the registration of three test scans. Each scan shows a different color. In the conform step, the scans have been slightly deformed (shown with a changed color) to compensate for residual registration and measurement error.

aries. We can repeat the process as many times as needed. Eventually, the simplified mesh is small enough to be further processed in a single pass by the in-core algorithm. We expect memory-efficient simplification algorithms to become a hot topic of research as capture methods improve and large models become widespread.

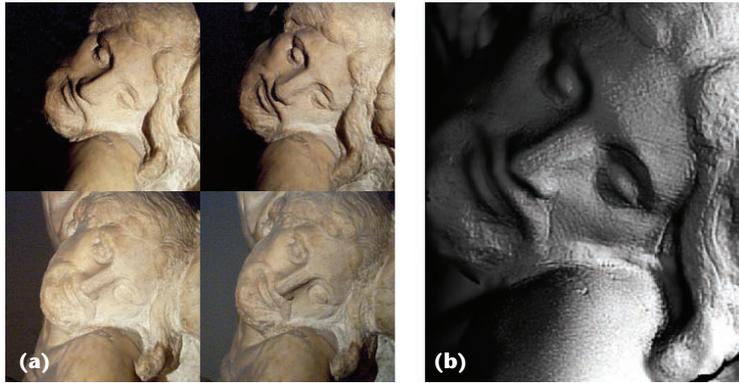
### Details and color

The mesh produced using the Virtuoso camera has a spatial resolution of approximately 2 mm, which is adequate for studying the statue's proportions from various viewpoints. However, it's inadequate for studying the small-scale tool marks. To capture data at a higher spatial resolution, we exploited the fact that the Virtuoso includes a color camera that produces images with a resolution of about 0.5 mm per pixel. We computed detail at this pixel resolution using a photometric stereo system built around the Virtuoso.

Figure 3 shows our photometric system. Given three images of a surface lit by three different light sources in known positions, we can solve a set of simultaneous equations for the surface normals corresponding to the points visible at each pixel in the image. Given the normal at each pixel, we can compute the relative reflectance, or albedo, at each pixel for the red, green, and blue bands. We used five light sources rather than three because in any given image a point may be in a shadow or a specular highlight. Figure 6a shows four typical images obtained from a single camera position. Figure 6b shows the resulting normals (lit from the side). For further detail of the physical system design, see Rushmeier et al.<sup>12</sup>

To compensate for possible errors in the photometric normals calculations, we used data from the 2-mm resolution mesh to compute the direction and relative distance to each point visible in each image and to estimate the relative light source intensity in the neighborhood of each pixel from each of the five lights. To compensate for scan-to-scan color variations, we performed a color registration analogous to the geometric registration of scans. We found corresponding points in all overlapping color albedo maps and then found a least-squares solu-

**6** (a) Color images taken with four of the five light sources. (b) Synthetic picture computed using the surface normals obtained with the photometric system.



tion for scaling factors for each of the color channels in each of the images to obtain the best color match on corresponding points. Rushmeier and Bernardini<sup>13</sup> discuss additional details of adjustments made using the underlying mesh and color registration.

### Texture synthesis

We partitioned the triangle mesh into height-field patches with a simple region-growing heuristic. For each patch, an orthogonal projection in the direction that maximizes the projected patch area defines a mapping between geometry and corresponding textures.

The texture synthesis process computes surface normals and reflectance maps as weighted combinations of corresponding values in all the overlapping images. Weights are assigned to take into account the degree of confidence in each pixel value, based on distance to the camera and viewing angle. Because weight maps correspond to scans and not to patches, transitions across patch boundaries aren't visible. Also, since the weights for each scan decrease with distance to the scan border, scan-to-scan boundaries aren't visible.

In our implementation, we streamlined computations by presampling the patch geometry and loading values from all maps simultaneously. Occlusions are handled elegantly by comparing depth values in precomputed depth buffers. Image and geometric information is loaded on demand to allow for processing of large data sets that don't fit in memory. We provide additional details regarding our image-based registration and texture synthesis algorithms elsewhere.<sup>9</sup>

### Validating and using the model

A large digital model isn't useful to art historians. As a result, we needed to derive an assortment of presentations of the data suited to Wasserman's needs, which in some cases required new techniques. Before developing other results from our model we needed to validate its accuracy to Wasserman's satisfaction. His test was that images derived from our model must correlate well with the high-quality photographs he had commissioned from a professional photographer.

To perform the validation, we selected features in digitized versions of Wasserman's photographs and found the corresponding 3D coordinates of those points on our model. We then used Tsai's calibration methodology<sup>14</sup> to compute camera parameters to gen-

erate a synthetic image from the same viewpoint. We couldn't estimate the lighting conditions in the commissioned photographs. To address the effect of lighting, we also matched camera viewpoints for images we took with a digital camera for which we knew the flash location. Initially, we computed images with geometry alone and found that including the surface albedo was essential to perceiving the proportions in the synthetic image.

### Overview of results

Wasserman's research questions defined our primary goals for the *Pietà* project. We shaped our presentation of the results to fit his needs. We developed a plan to fulfill his requirements by delivering a variety of results, including

- precisely defined views,
- impossible views,
- embedding the statue in virtual environments,
- providing precise measurements,
- modifying the statue, and
- providing an interactive viewer.

To answer certain questions about Michelangelo's composition, Wasserman wanted to see the statue from physically or practically impossible points of view. These included views from directly above the statue to reveal details of the composition not normally visible (Figure 7c, next page) and from various angles at a height below the base of the statue to illustrate it as it would have appeared in the context Michelangelo originally intended. We also recreated some of the settings in which the *Pietà* stood during its history, using 3D models and animations to illustrate the statue's visual impact in these various environments (Figures 7e and 7f). To reconstruct the tomb and garden settings shown in these figures, Wasserman provided drawings and images of similar environments and some crude dimensions. Accurately modeling the environments required a number of variations of each environment that Wasserman evaluated against his understanding of the historical record.

### Editing the model

The ability to modify our model of the statue provided Wasserman with opportunities to study it in ways otherwise impossible. Using the 3D model, we reconstructed the statue with Christ's missing left leg, approximating its appearance before Michelangelo broke it. We removed the pieces that Calcagni reattached, illustrating the statue as it may have appeared without his efforts. Figure 7g shows this second modification. In the figure, some of the surfaces that would be occluded by the limbs removed are now visible. Internal areas that revealed where the marble was broken are colored a flat gray. We also separated the four figures that make up the statue so that they can be examined in isolation.



Identifying the pieces that Michelangelo removed is itself a problem. Four major pieces were removed from the statue (three arms and a portion of a leg that was never replaced). The three pieces that were reattached were each composed of a set of smaller fragments, so it's not obvious what was removed. Based on his own close study of the physical work and the x-rays he had commissioned, Wasserman sketched on electronic images of the statue the various lines where he believed the breaks were made.

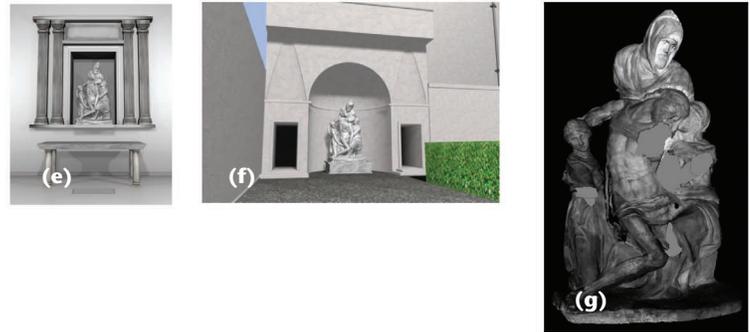
Directly editing a large model with millions of vertices isn't feasible, particularly because our triangle mesh doesn't have sufficient resolution to model the breakage exactly as Wasserman wanted. We tried two methods of editing the model. First, we tried painting each of the color images associated with the individual scans to precisely mark which parts belonged to the removed sections. This approach had some problems, since hand marking didn't give pixelwise identical locations for all of the breaks across the various scans. However, given the painted images, we could automatically segment the statue by simply removing vertices that were painted with the broken color. This simple computation was useful while we were producing early versions of the model (before all the data was acquired, added, and tightly aligned) to give Wasserman an indication of what results to expect.

For the final model, we crudely segmented the model by defining bounding boxes enclosing the broken segments. We then identified individual patches containing portions of the cracks for editing. While this approach would be tedious to repeat many times (the cracks extend over many different patches), it was the most reliable approach for the final model.

We also used the painting to separate the four figures in the model. This task is more sensitive to the problem of ambiguous identification across scans because no precise lines on the statue define the figures. Wasserman wanted the figures separated because they revealed shapes and relationships, like the relative position of Mary Magdalene's hands, that we can't observe from the solid statue. While we were able to achieve the segmentation of the statue we needed for this study, our experience indicates that detailed editing of high-resolution models is an area that requires additional research.

### Interactive viewer

To enable Wasserman to study the statue on his own computer, we designed a viewer that could run on a PC



**7** (a) Black-and-white rendering of the model. (b) Close-up view of the model. (c) Bird's eye view of the model. (d) Interactive viewer. (e) Synthetic image in a niche above the tomb. (f) Synthetic image in a garden. (g) The statue without the pieces Michelangelo removed.

(Figure 7d). The combination of a large data set and a slow computer required special attention to the trade-offs between speed and quality and between usability and flexibility. Our target audience consisted of unsophisticated computer users unaccustomed to the navigation paradigms common to interactive 3D graphics. We found that we needed to radically simplify the controls to provide a fast learning curve and then adapt them to our users' abilities and interests. Maintaining interactivity was essential, so that users could easily grasp the function of the navigation controls and compensate for their inevitable limitations.

We had two objectives in designing an intuitive interface: maintaining a frame of reference when zooming in on detail and providing clear, separate controls for altering views and lighting. Our viewer presents users with a simplified model of the statue around which they can navigate interactively and lets them render a full-resolution image of a selected detail. The simplified model, only 1 percent the complexity of the full model, acts as a map to the more detailed model. Users can select an area of interest, using simple 3D navigation controls to reach the desired view. We chose a navigation paradigm in which the camera orbits around a user-selected center at a fixed distance and zooms in and out along a radius. Users can select a new center by picking or dragging the image plane to the view plane.

To examine a region in more detail, users frame a section of the scene and render the desired image from a database of the full-detail model. This step currently can take a few minutes on a laptop computer. We enhanced the resulting 2D image to let users interactively vary the light-

ing by dragging the mouse across the window. Many details invisible with the light in one position appear when the light moves. Users can thus understand fine-scale structures more clearly. We designed this technique to model a method that we observed Wasserman using on the statue. He moved a small flashlight slowly back and forth across the statue to highlight small surface irregularities. The virtual light editor produces similar effects.

The viewer was useful for isolating portions of the model and rendering high-resolution close-ups of sections of interest. We had hoped that the viewer would also be helpful to Wasserman in evaluating the appearance of the statue from various views, to develop a theory of exactly how Michelangelo intended the statue to be viewed. The simplified model in the viewer proved inadequate for this interactive study. When we simplified the model to the extent that allowed interactive viewing, it still looked like a good representation to the casual observer. However, for Wasserman's in-depth study of specific proportions, the simplified model wasn't accurate enough.

The current viewer still needs improvement. It's far from satisfactory as a tool for art historians. We need to do a great deal of work to find intuitive methods for non-technical users to interact with 3D data, especially when viewing large data sets that we must simplify to allow interactive display. Part of the problem is rendering speed. Incorporating ideas such as point-based rendering<sup>15</sup> or multipass texturing now available on inexpensive graphics cards would improve this aspect of our system.

### Conclusions

This project has demonstrated that 3D scanning and graphics can be useful tools for art historians, and by extension for similar studies in archaeology, architecture, and other disciplines where detailed examination and manipulation of artifacts is necessary but not feasible in reality. Wasserman noted three aspects of the virtual model that were especially useful to art historians:

- Having a virtual model of the statue lets researchers examine the statue at their leisure, to discover details they hadn't noticed in the time they spent with the statue and to verify their recollections and theories.
- The ability to control lighting precisely lets researchers see the statue as it might have appeared in environments outside the museum and to highlight small details not easily visible.
- Measuring the statue freely and precisely lets historians factor out subjective perception and better understand the artist's use of perspective and composition.

Our technical experience shows that it's possible to build a detailed computer model of a large object from a large collection of small individual scans, which we can acquire with a relatively inexpensive device. Our plans for future work include studying improved algorithms for accurately registering multiple scans and developing a handheld, real-time scanning system. For more information about our work, visit <http://www.research.ibm.com/pieta>.

The final conclusions Wasserman draws from the digital model will be in his book published by Princeton University Press, which will include more of our results on a CD-ROM. Kiosks are now (or have been) presenting our work at a number of museums throughout the world (including sites in Asia, Australia, South America, Europe, and the US). ■

### Acknowledgments

We express our appreciation to Claudio Silva and André Guezic for their contributions to this project, the Museo dell'Opera del Duomo in Florence for their assistance in our study of the statue and Jack Wasserman for his inspiration and for creating the opportunity for our work.

### References

1. J.-A. Beraldin et al., "Portable Digital 3D Imaging System for Remote Sites," *Proc. 1998 IEEE Int'l Symp. Circuits and Systems*, IEEE Press, Piscataway, N.J., 1998, pp. 488-493.
2. J. Zheng and L. Zhong, "Virtual Recovery of Excavated Relics," *IEEE Computer Graphics and Applications*, vol. 19, no. 3, May/June 1999, pp. 6-11.
3. M. Levoy, "The Digital Michelangelo Project," *Proc. 2nd Int'l Conf. 3D Digital Imaging and Modeling*, IEEE Press, Piscataway, N.J., 1999, pp. 2-13.
4. J. Abouaf, "The Florentine Pietà: Can Visualization Solve the 450-Year-Old Mystery?," *IEEE Computer Graphics and Applications*, vol. 19, no. 1, Jan./Feb. 1999, pp. 6-10.
5. C.L. Zitnick and J.A. Webb, *Multi-Baseline Stereo Using Surface Extraction*, tech. rep. CMU-CS-96-196, School of Computer Science, Carnegie Mellon Univ., Pittsburgh, Penn., Nov. 1996.
6. P.J. Besl and N.D. McKay, "A Method for Registration of 3D Shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, Feb. 1992, pp. 239-256.
7. R. Bergevin et al., "Towards a General Multi-View Registration Technique," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 5, May 1996, pp. 540-547.
8. F. Bernardini and H. Rushmeier, "Strategies for Registering Range Images from Unknown Camera Positions," *Three-Dimensional Image Capture and Applications III* (Proc. SPIE), vol. 3958, Jan. 2000, pp. 200-206.
9. F. Bernardini, I.M. Martin, and H. Rushmeier, "High-Quality Texture Reconstruction from Multiple Scans," *IEEE Trans. Visualization and Computer Graphics*, Oct.-Dec. 2001, vol. 7, no. 4, pp. 318-332.
10. F. Bernardini et al., "The Ball-Pivoting Algorithm for Surface Reconstruction," *IEEE Trans. Visualization and Computer Graphics*, vol. 5, no. 4, Oct.-Dec. 1999, pp. 349-359.
11. A. Guézic, "Locally Toleranced Surface Simplification," *IEEE Trans. Visualization and Computer Graphics*, vol. 5, no. 2, Apr.-June 1999, pp. 168-189.
12. H. Rushmeier et al., "Acquiring Input for Rendering at Appropriate Level of Detail: Digitizing a Pietà," *Proc. 9th Eurographics Workshop on Rendering*, Eurographics Assoc., Aire-la-Ville, Switzerland, 1998, pp. 81-92.
13. H. Rushmeier and F. Bernardini, "Computing Consistent Normals and Colors from Photometric Data," *Proc. 2nd Int'l Conf. 3D Digital Imaging and Modeling*, IEEE Press, Piscataway, N.J., 1999, pp. 99-108.

14. R.Y. Tsai, "An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision," *Proc. IEEE Computer Vision and Pattern Recognition (CVPR 86)*, IEEE Press, Piscataway, N.J., 1986, pp. 364-374.
15. S. Rusinkiewicz and M. Levoy, "Qsplat: A Multiresolution Point Rendering System for Large Meshes," *Computer Graphics (Proc. Siggraph 2000)*, ACM Press, New York, 2000, pp. 343-252.



**Joshua Mittleman** is a programmer at the IBM T.J. Watson Research Center, specializing in simplification, level-of-detail control, and other algorithms applied to rendering complex models. He has an MS in computer science from Rutgers University.



**Fausto Bernardini** is a research staff member and manager of the Visual and Geometric Computing group at the IBM T.J. Watson Research Center. His research interests include interactive visualization of large models, 3D scanning and reconstruction, and geometric modeling. He has been one of the principal investigators in a research effort to build a digital model of Michelangelo's Florentine Pietà. He received an electrical engineering degree from the University of La Sapienza, Rome, Italy, in 1990, and a PhD in computer science from Purdue University in 1996.



**Gabriel Taubin** is a research staff member and former manager of the Visual and Geometric Computing Group at the IBM T.J. Watson Research Center. He is known for his work on 3D geometry compression and progressive transmission of polygonal models, the signal processing approach for polygonal meshes, algebraic curve and surface rasterization, and algebraic curves and surfaces in computer vision. He received the IBM Research 1998 Computer Science Best Paper Award for his paper on geometry compression through topological surgery. He earned a PhD in electrical engineering from Brown University and an MS in mathematics from the University of Buenos Aires, Argentina.



**Holly Rushmeier** is a research staff member at the IBM T.J. Watson Research Center. Her research interests include data visualization, rendering algorithms, and acquisition of input data for computer graphics image synthesis. She received a BS, MS, and PhD in mechanical engineering from Cornell University in 1977, 1986, and 1988, respectively. In 1990, she was selected as a US National Science Foundation Presidential Young Investigator. She has served as papers chair or cochair for the ACM Siggraph conference, the IEEE Visualization conference, and the Eurographics Rendering Workshop. From 1996 to 1999, she was editor in chief of ACM Transactions on Graphics.

Readers may contact Fausto Bernadini at the IBM T.J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, email [fausto@us.ibm.com](mailto:fausto@us.ibm.com).

For further information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.



**Ioana M. Martin** is a research staff member in the Visual Technologies Department at the IBM T.J. Watson Research Center. Her research interests include interactive computer graphics, visualization, and image processing. She has developed and implemented methods for processing and rendering large data sets, including progressive refinement and multiresolution techniques with applications to structural biology. She has worked on various projects including modeling with subdivision surfaces, adaptive delivery of 3D models over networks, 3D data compression, and model reconstruction. She received a PhD in computer science from Purdue University in 1996.