Michael J. Potel
http://www.wildcrest.com

The Florentine Pietà: Can Visualization Solve the 450-Year-Old Mystery?

Jeffrey Abouaf

Michelangelo carved his "Florentine Pietà" (Figure 1), the second of three Pietà statues, from a single block of marble while in his 70s. The piece includes posed figures of Christ in the arms of the Virgin Mary, Nicodemus (properly Joseph of Arimathea) looking on, and Christ touching Mary Magdelene. The statue was thought to be for the artist's tomb.

Two years after he had stopped working on the unfinished piece, Michelangelo attacked the statue with a sledgehammer. Why? Only the interruption of a servant prevented its complete destruction. The servant disposed of the statute. Subsequently, a second artist repaired some of the damage and completed some unfinished sections. What portions represent the original work, the restoration, and new work by the second artist? What parts were obscured by the subsequent sculptor, and which were reworked?

Renowned art historian Jack Wasserman hopes to settle 400 years of debate—and raise new issues—by bringing together contemporary scholars and technologies to literally "view" the work in a completely new way. In 1997, he enlisted researchers at IBM's Thomas J. Watson Research Center to create a high-resolution virtual model of the sculpture. The team planned to use new technologies in 3D digital photography, database management, compression, and detail visualization to enable examination and analysis of the source work in ways not possible or permitted using the original sculpture.

IBM funded the project, in which data collection now nears completion. This data and the team's results contribute to Wasserman's upcoming collection of essays (Princeton University Press, scheduled for 2000). This book and accompanying CD-ROM will document the statue factually in as many ways as possible. The discussions will address the Florentine Pietà mystery from many perspectives—historical, theological, aesthetic, and scientific—based on papers and supporting data from scholars, theologians, scientists, and visualization researchers. IBM and Wasserman plan a press conference and preview at the Morgan Library, New York City (December 1999 as of this writing) accompanied by an exhibit and interactive kiosk showing some final images from the project.

The statue

Although this work did not enjoy permanent installation after its creation, the Florentine Pietà's significance was recognized and protected by the few responsible for its custody. Michelangelo carved the piece some 50 years after having carved his most famous Pietà, located in Rome. The first reference to it appeared when Michelangelo's artist friend Giorgio Vasari cited the Florentine Pietà in his book, *Le vite de' più architetti, pittori, et scultori italiani* in 1550:

> It is impossible to speak of its beauty and its sorrow, of the grieving and sad faces of them all, especially of the afflicted Mother. Let it suffice: I tell you it is a rare thing, and one of the most laborious works that he has yet done... He intends to give the Deposition from the Cross to some church, and to be buried at the foot of the altar where it is placed.

Between 1555 and 1556 Michelangelo broke off parts of the statue, specifically, the Christ figure's left arm and leg and part of the right arm; the Virgin's left arm; and the Magdelene's right arm.

Thereafter the statue became the property of Francesco Bandini, who had it finished and restored by Michelangelo's pupil and young friend, Tiberio Calcagni. Calcagni reattached all but Christ's left leg. His style, especially on the Magdalene, differs noticeably and seems incongruous with Michelangelo's; the master's face is believed to be reflected in the sad face of Joseph.

On Bandini's death, the statue passed to his son, and in 1674 to Cosimo III, Graduke of Tuscany, who had it installed at San Lorenzo, Florence. In 1721 it was transferred to Santa Maria del Fiore (Duomo) and installed behind the high altar in 1722. In 1933 it was moved to the Chapel of St. Andrew in the left tribune of the Duomo. In 1946 it was brought to the Accademy delle Belle Arti for cleaning, then reinstalled in the Chapel of St. Andrew in the Duomo in 1948. In 1980 the statue was exhibited briefly in Santo Stefano al Ponte, Florence, before being installed in the Museum of the Opera of Santa Maria del Fiore in 1981, in the room it occupies today.

The statue evaluated

The polished details and classical concept of the early Rome Pietà give way in this work to a rough dramatic style, whose broken angular lines speak to the scene's pathos. The four figures are portrayed as a group, with Nicodemus lowering Christ from the cross into his moth-

0272-1716/99/$10.00 © 1999 IEEE

er Mary's lap. Many believe Nicodemus' mournful face to be that of the aged artist. The figures intertwine, but are carved in unusual proportions, individually and in relation to each other.

Many 16th century discussions on the topic of ideal proportions in sculpture focused on understanding Michelangelo's style and perspective. To modern viewers, the piece looks different based on the viewing angle: proportions look incorrect viewed head-on, but appear properly foreshortened when the statue is elevated and seen from below.

Wasserman finds the proportion and details of the second Pietà curious; some parts seem excessively elongated, while others seem suspiciously small:

> The ability to stand each figure of this Pietà up straight without distorting the dimensions and proportions would provide valuable insight into the question of what Michelangelo's proportions were like, his general concept of proportions, how he meant the work to be viewed at his tomb site, and, perhaps, his intent in taking a hammer to it.

Does Michelangelo's destruction and abandonment of the statue acknowledge a mistake, or merely a change of mind or purpose? Was it in response to flaws discovered in the stone or various accidents? Some believe the sculptor's purpose changed from an original desire to be buried at the foot of the statue in Santa Maria Magglore in Rome to having his final resting place in Florence. One account claims Michelangelo attempted to destroy the flawed work, while another suggests it was merely abandoned and sold. Were limbs removed for reasons other than the statue's destruction? The historical accounts come from third-party opinions and interpretations, not from the artist's statements. Maybe a virtual statue can help us find the answers to these questions.

The team

Jack Wasserman is Professor Emeritus of Art History at Temple University with a focus on Italian renaissance art. He became interested in Michelangelo's Florentine Pietà after studying artist Jacopo Pontormo's portrayal of the Virgin Mary mourning the body of Christ. In 1996, with the help of Peter Rockwell, a noted stone-carver and the son of painter Norman Rockwell, Wasserman began study of the Pietà.

Computer scientist Jacob Ukelson is general manager of the Applications and Solutions Technologies department at the IBM T.J. Watson Research Center. The team's manager, Gabriel Taubin, heads the Visual and Geometric Computing Group and is a specialist in geometric and image-based computation. Taubin's staff researchers include Fausto Bernardini, Joshua Mittleman, Andre Gueziec, and Holly Rushmeier. (Additional information is available at IBM's sites, http://www.pl.ibm.com/news/pieta.html and http://www.ibm.com/Press/media.nsf under "Pieta.")

The equipment

Visual Interface's (VI) six-lens Virtuoso camera uses structured light and multi-baseline stereo techniques to capture 3D patches composed of surface points ("shape images"), each containing between 10,000 and 20,000 *X,Y,Z* coordinates. Texture and 24-bit color information about the statue were captured by a Kodak DC120 (later a DC420) digital color camera mounted on top of the Virtuoso. VI's algorithm determines relative positions of points in the shape images to the surface, reconstructs the image as a triangular mesh with UV indices, then maps the texture data onto the mesh.

1 Michelangelo's Florentine Pietà (1550)

As described by VI President Jon Webb, each lens is a fixed focus, with one charge-coupled device (CCD) per lens, and the lenses positioned 20 mm apart. The field-of-view (FOV) is 25 cm at a 50-cm distance and 50 cm at 150 cm. The key to the 11-pound system's 3D capture capability is proper lens alignment—accurate to 1 mm at workspace level, which translates to less than 1 micron adjustment error at the camera head. VI performs and locks this calibration. Each lens captures a gray-scale image. An internal flash projects a vertical stripe on the object, which the software uses for registration in aligning the six images and the texture map.

VI's capture and stitching algorithm translates the six captures to one 3D image and can stitch up to 10 images together. The unit interfaces with a standard PC (using a PC card) and outputs data in several common file formats (Open Inventor format for this project). IBM used VI's calibration and capture software without modification, but asked VI to modify the camera and provide calibration information for the special lighting array deployed to capture texture data.

IBM used a different Virtuoso camera on each of two trips to photograph the statue. VI calibrated the six CCDs and color camera prior to each trip; the special lighting array was calibrated to the camera on site (see below).

2 Reconstructions of part of the statue relit with four different lighting specifications. Only the simple color and surface texture appear in this example.

3 IBM researchers on site in the museum with the statue.

IBM purchased the first shipping model of the Virtuoso camera in December 1997 at an estimated price of $25,000. (Further information on Virtuoso and other Visual Interface products can be found at http://www.visint.com.)

The project

The Florentine Pietà is wonderfully suited for study using virtual modeling and digital technologies because of its unfinished state, damage, partial restoration, and the physical limitations imposed by the site and the work itself. Researchers can't move or disassemble the work, must keep physical contact to a minimum, and must make do with available lighting conditions.

The virtual model, on the other hand, can be seen in whole or in part; from any angle, distance, or perspective; at any scale; in any light; and at any location. The virtual Pieta can be studied as if installed at the originally intended burial site in Rome. Virtually removing the reattached parts (the arms of the Christ figure, the Virgin's left arm, and the Magdelene's right arm) may reveal what lies underneath, as Michelangelo saw it. Figure 2 provides examples of this flexibility.

Lights, cameras in action

Digitizing the Florentine Pietà for a near-perfect replica required more than 1,000 individual digital shape photographs of this 2.5-meter, 10-ton statue. Each scan was performed at a distance of 750 to 1,000 mm from the surface. The raw source data consists of tens of gigabytes—mostly photometric texture data. The relative positions of the shape photographs with respect to the sculpture must be determined to assemble all of them into a single, accurate, digital 3D representation of the work.

To Taubin, not only was the task at hand of artistic interest, but

> It presented technological challenges in the area of data collection and assembly into a three-dimensional model that has the desired properties, and using it to render very accurate images under different lighting conditions, from different points of view. And for all these processes, the volume of data that we have to handle carries the problem beyond the scope of existing techniques.

The Virtuoso cameras needed no repair or recalibration during the shoots. However, the marble surface produced reflections and specular artifacts that generated errors in the texture data. They could manually adjust the texture camera's exposure to compensate for on-site light variations, but this wasn't sufficient. The IBM team developed a five-light lighting array for photometric texture capture from different specified lighting angles. They mounted the lights and camera on an aluminum pole bracket. (Figure 3 shows team members on site with the camera assembly and statue.)

Instead of a single texture map, for each capture position the team took five photometric exposures with the color camera, each using a different light (at a different angle from the camera). The researchers determined that they needed three exposures to align the image according to the surface normals and to filter out reflection, specularity, and noise artifacts caused by the marble surface. (If this were a white wall, two images would be necessary; given a mirrored surface, the technique could not work.) They discard the two exposures deviating most from the norm and combine the remaining three to generate a single accurate texture and bump map representation, which they then apply to the 3D mesh.

The team used a Colortron II to color calibrate the statue surface. Although the photometric data is in RGB format, there is no assurance these values are correct. The Colortron II was originally developed for the Macintosh by Light Source in the San Francisco Bay area, and later acquired by X-Rite (http://www.x-rite.com) for use in desktop publishing and prepress. It is a hand-held object with a lens at one end, held against the surface to be measured. It projects a light onto the surface and samples absolute RGB values for use in calibrating other devices.

The lights are identical except for their position on the frame. Although VI calibrated the camera as a unit, it had to be calibrated to the lighting array. This took two steps: Using a Faro digitizing arm, the Virtuoso took a shape image of points on a predefined target. These coordinates were then entered as data. The Faro arm digitized the target points, aligning the Faro coordinate system with the Virtuoso system. Then the team used the Faro arm to locate three points on the Virtuoso camera body and to measure the distance to each of the light positions.

The raw data from each scan consists of one shape image scan (3D accurate to 2 mm) and five photometric 24-bit images, each 1280 × 960 resolution (2D accurate to 0.25 mm). Because of the irregular patch shape, each mesh is smaller than the dimensions of the combined texture map—about 700 × 700. Nonetheless, there will be around 620 final patches made from more than 1,000 scans.

4 **These images show a preliminary reconstruction of the statue in wireframe. The red dots are the laser spots projected onto the statue and used to align the small individual meshes.**

The next problem was how to register and put the images together. Here the team used stand-alone red lasers to project registration spots onto the statue surfaces. The lasers remained in place for each day's shooting while the Virtuoso was moved from scan to scan. The museum required they be taken down during business hours; the lasers were installed in a different position for the following evening's shoot. Figure 4 shows a preliminary reconstruction of the statue in wire frame. The red dots are the laser spots projected onto the statue, used to align the small individual meshes to combine them into a larger mesh representing the full statue. Note how they aid registration on the reconstructed mesh.

5 **The geometry underlying the images in Figure 2 and the underlying preliminary resolution of the mesh.**

For the final image, the photometric texture and bump maps at 0.25 mm resolution, with the 3D data at 2mm, will yield a virtual model resolved to 1 mm. This translates to a final database likely between 5 and 10 Gbytes. Part of IBM's research will involve developing technologies to manage, distribute, and dynamically represent accurate and efficient imagery from this data set. Specifically, this means new compression and dynamic level-of-detail (LOD) technologies for transmitting and viewing models at a continuous increasing and decreasing resolution, whether stand-alone or over a network. All of these solutions are proprietary to IBM.

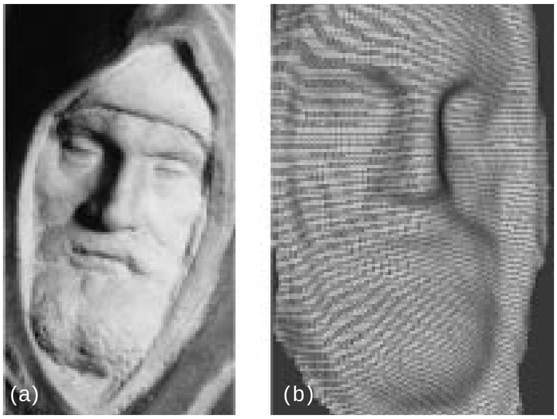

6 **A reference photograph of the Nicodemus head (a). The Nicodemus head in wireframe, synthesized from the composited 3D photographs (b).**

Facing the challenges

Physical access was limited to off-hours (nights and weekends). The statue could not be moved, and no evidence of the team's presence was permitted. Rushmeier explained, "We worked at night and processed data during the day, and our only chance to see Florence was when we went to the hardware stores to get odds and ends." This was in January, in a museum with no heat.

With no guarantees that the camera could capture the entire surface of the sculpture, researchers took 1,000 overlapping photographs. Not until returning to the IBM lab and combining all of the digital patches on computer could the team locate gaps. They found holes, Rushmeier recalls, which necessitated a second trip to Florence that summer to capture additional data. (Wasserman suggests they may return for a third trip.) Figure 1 illustrates the reason: the statue's geometric complexity, the overall size of the piece and range of length scales, and the intertwined limbs of the figures create many small spaces that are difficult to measure with many existing 3D digitizing devices. Typically, gaps in data showed up at a cavity between figures or objects, most often limbs and fingers. Likewise, grabbing texture data of the backside of a limb required particular care, such as positioning a mirror behind the object and photographing the reflection.

Each trip to Florence required 10 days of round-the-clock work. Figure 5 shows the resolution of the preliminary reconstruction. This "crude" model can be augmented by small-scale details, measured by processing multiple color pictures taken under controlled lighting conditions. Even without the finer details, the model contains millions of points and triangles.

IBM and Wasserman expect to continue working with the data set beyond the year 2000. IBM has the opportunity to investigate a complete data set it helped create, and Wasserman will need years to interpret the imagery to come.

Figure 6a is a reference photograph showing a detail of the Nicodemus head. Figure 6b shows the head in wire-

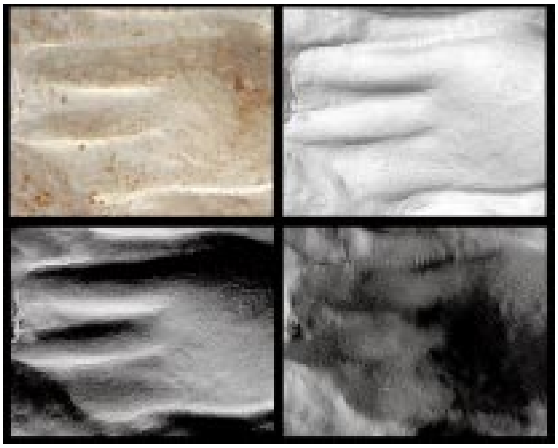

7 The types of color maps used for texture data to be applied to the mesh, in this image the Magdalene's left hand.

8 Comparison views of the Nicodemus head with regular color mapping (left) and the synthesized texture mapping techniques developed by the IBM team (right).

frame, synthesized from the composited 3D photographs; Figure 7 illustrates the types of texture data (diffuse maps and bump maps) to be applied to the mesh—in this image, the Magdelene's left hand. Figure 8 provides comparison views of the Nicodemus head. Note the differences in how fabric and facial hair look in these images.

The expectations

Clearly the technologists and the historian have enjoyed learning from each other. Wasserman has only positive observations about this first experience with technology and a corporate partner. In applying digital technology to the Florentine Pietà, Wassermann generally focuses on five advantages of the technology, with a specific purpose for each:

1. *Convenience and interaction*. Viewing the statue in the computer lets people see the work at a realistic resolution in ways not otherwise available. The data can be presented as an interactive experience, as a self-directed tour from perspectives not otherwise available, or as selected parts in isolation.
2. *Precision*. Recording the exact dimensions and proportions of the whole and of its parts in natural and 3D form results in a replica accurate to 1 mm. Wasserman seeks to uncover facts unknown to historians and intends this study to become a major part of the historical record. For example, he has had the Pietà x-rayed, which revealed placement of metallic pins. The presence of a pin settled one debate—whether one limb had been removed and reattached or merely exhibited a crack in the stone.
3. *Problem solving*. Exploding the figure visually, raising the figure to various heights, or viewing the statue gradually while moving from one side to the other gives the perspective the artist had when carving the piece. For example, Christ's torso and right arm are proportioned too long, which appears odd when viewed at eye level. However, by elevating the virtual model in the computer to the height it might have occupied had it been installed above an altar, the right and left arms appear of equal length and the torso realistically foreshortened.
4. *Documentation* (printout and CD-ROM). Illustrating the exploded parts of the figure alongside the remaining block or illustrating the desired views and details of the statue give insight into the work in progress and the special qualities of individual figures.
5. *Stimulate new research*. Although the project began in 1997, Wasserman is just now beginning to see some of the final imagery. The entirety of the IBM visualization won't be available for many months, and Wasserman's interpretations will come after that. IBM's interest revolves around having a complete data set in which they controlled every aspect of its preparation.

The future

IBM calls this project the most extensive technological study ever done on a single work of art. And the project appears to be "win-win" for the historians and technologists. From the historical perspective, I hope the IBM researchers captured more data than they can analyze with their current technology, allowing future technologies to delve deeper, yielding interpretations currently not possible.

The sheer amount of the data offers plenty of fodder for current research goals, for example developing compression technology for sharing a multi-gigabyte data set like this across a network or the Internet using existing PC technology. Or, developing combined 3D and texture data capture and filtering that can separate specular, reflective, refractive, bump, and transparent properties on the fly or capture live subjects in motion at the same detail.

Taubin sees some of these methods as applicable more generally to the problem of digitizing very large real-world objects like large works of art or even architecture. This might also benefit general applications in the graphical representation of 3D objects on the Internet, from displaying art work to allowing the manufacturing industry to produce online catalogues with 3D displays of their products. The ability to examine an object in exquisite detail even without access to the original is an exciting possibility. In the long term, we should expect consumer-grade hardware capable of generating this kind of imagery. What we do with it is then up to us. ■

Acknowledgment

My special thanks to Holly Rushmeier for her efforts in coordinating and providing information for this article.