Sculpting Virtual Reality

3-D models offer new ways of seeing art

By DAMARIS CHRISTENSEN

In 1555, Michelangelo took a hammer and broke his unfinished sculpture of Mary Magdalene, the Virgin Mary, Nicodemus, and the body of Christ into 15 separate pieces. A contemporary of the temperamental Renaissance man repaired the broken statue, although the left leg of Christ was still missing.

It still isn't clear what prompted Michelangelo to mutilate his work, but art historian Jack Wasserman of Temple University in Philadelphia hopes that a new digital representation of the sculpture will offer him some clues. Already, thanks to the work of a team at the IBM Thomas J. Watson Research Center in Yorktown Heights, N.Y., he has been able to see what the main part of the statue would have looked like before the broken pieces were replaced.

Michelangelo may have originally intended the sculpture, now known as the "Florentine Pietà," to be mounted above his grave. The pietà now sits at eye level in the Museum of the Opera del Duomo in Florence, not far from where the artist is buried.

Digital imaging will allow art historians to see the statue as it looked before Michelangelo's fit of fury, says Wasserman. Additionally, once the researchers create a digital model, a computer program should be able to separate and animate the bent and intertwined human forms, he says.

"The ability to stand each figure of this pietà up straight without distorting the dimensions and proportions will provide valuable insight into the question of what Michelangelo's proportions were like, his general concept of proportions, [and] how he meant the work to be viewed at his tomb site," says Wasserman.

Digital models enable art historians to view sculptures from different heights and angles and in different lighting, he explains. As more become available, researchers and possibly museum visitors could use three-dimensional models to look at sculptures in ways that are impossible in works' current settings.

Such models are better than photographs at showing the proportions within a figure and between figures in a sculpture, he says. Such graphics may be able to restore on the computer screen the

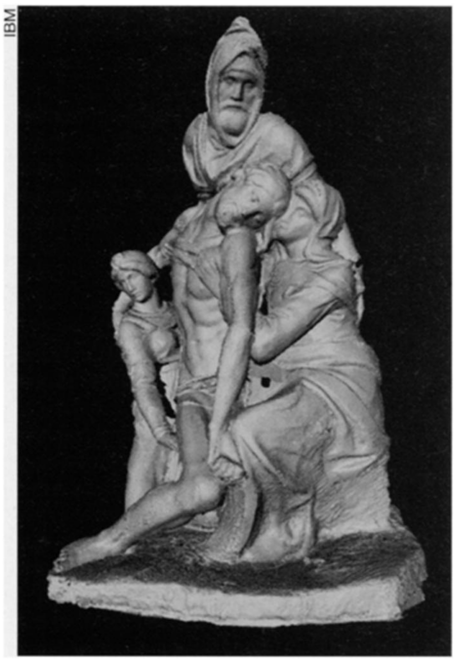
IBM

The IBM team's still incomplete digital model of the "Florentine Pietà."

damaged areas of artworks. Moreover, digital models may bring a fuller sense of distant masterpieces to people than two-dimensional photographs can.

Building an accurate, 3-D model of a sculpture in virtual reality, however, is nearly as arduous as carving away the original stone. Researchers must make thousands of measurements and align hundreds of photographs to place every fragment of the artwork in space. The steps required to convert these data into an image on a computer screen multiply the complexity of the task.

The more detailed the model is, the more computer memory it takes up and the more difficult it is to display, rotate, and modify the virtual sculpture.

The idea of digitizing statues is not entirely new. Practical techniques for 3-D scanning have been around for years, and several museums have successfully copied their works of art to the computer screen. However, before the work of this IBM team and that of a group from Stanford University, which spent the past year digitizing many of Michelangelo's other sculptures, no one had captured a large statue with enough precision to serve as a primary resource for scientific work, says Stanford's Marc Levoy.

Wasserman originally approached IBM about scanning the "Florentine Pietà" because he found that detailed photography was insufficient to answer his questions. By manipulating the model on his own personal computer, he wanted to study the statue at scales ranging from meters to millimeters.

To help Wasserman, the IBM team needed to capture large amounts of data efficiently and render accurate images from various points of view with several lighting combinations. "The volume of data we have to handle carries the problem beyond the scope of existing techniques," says IBM's Gabriel Taubin.

Still, the IBM team worked as much as possible with available technology. They adapted a black-and-white camera system that plastic surgeons have used to take 3-D pictures of patients' faces. The system projects a pattern of stripes onto the face or sculpture. It then takes six pictures simultaneously from different positions.

A computer algorithm next uses the stereo images to compute the 3-D shape of the object, based on how parts of the picture shift from one stereo image to another. Because the 3-D camera system scans an area of only about 10 by 10 inches, the researchers must patch many scans together to model the entire sculpture. The stripes help the computer match corresponding parts of the pictures.

The first step in converting the many scans into a complete model is to take points measured at 2-millimeter intervals on each surface and combine them into a

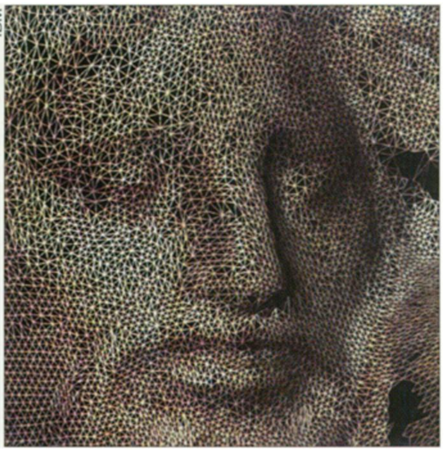

IBM

A close-up of a single scan of Christ's face, showing lines drawn between the measured points.

3-D cloud of points. The next step is to connect the points into triangles, essentially building a 3-D mosaic of some 14 million tiles. Further calculations smooth the surface of the statue and—based on the color images taken with a different camera—add information on the statue's color and texture.

Incorporating both the black-and-white and color data, the computer model reconstructed details of the "Florentine Pietà" with a resolution of about half a millimeter—enough for Wasserman to see the tool marks left by Michelangelo.

Wasserman's computer can't handle a detailed 3-D model of the entire sculpture, so the researchers created a program in which he first sees a simplified model that he can view from any angle. He can then choose to see small sections in detail.

Despite the detail of the IBM scanning, this model of Michelangelo's "Florentine Pietà" is not as precise as models produced at Stanford University. Levoy and his colleagues spent the past year imaging Michelangelo's "David" and five other statues in the Galleria dell'Accademia in Florence, four statues in the Medici Chapel in Florence, and about a thousand pieces of a stone map in Rome.

"I wanted to create 3-D models larger [with more detail] than anyone else had ever created," says Levoy. The challenge was not only to help art historians understand the works but also to stretch the limits of current 3-D imaging and model-building technology.

Levoy's team used a custom-built scanner to collect the data. The scanner projects onto an object a 6-inch line of laser light. When viewed from an angle, the line reflects the curves of the surface it covers. By tracking some 200 to 300 points as the line sweeps across the surface, a computer records an extremely detailed 3-D image. Like Taubin's group, Levoy's team must repeat their scans many times to cover the entire surface of a statue.

Unlike the IBM team, the Stanford group has designed their computer program to connect the dots in each small image. This technique is less likely than that using a cloud of points to incorrectly show two surfaces as being connected when they are not, Levoy says.

Algorithms that the researchers developed combine these 3-D meshes into a single seamless object made up of tiny triangles, each a quarter of a millimeter on a side. For the triangles to be clearly visible to the unaided eye, the image would have to be so large that just one of David's eyes would cover a magazine page. Levoy's digital model of Michelangelo's "David" contains about 2 billion of these minuscule triangles.

Levoy can use standard computer-graphics techniques to illuminate each surface of the computer image as though the statue were reflecting a single stationary light source, giving his geometric model the look of a black-and-white photograph.

Levoy's team also took color photographs of the statue under different lighting conditions. They were then able to calculate the color of each point on the statue's surface, independently of how it was lit when they photographed it. This color data, combined with the geometric data, allow the researchers to generate new images of the statue under any virtual lighting conditions they like.

Levoy says that 3-D imaging of marble statues has probably reached the limit of current technology. Highly polished, fine marble causes the tight laser line to blur very slightly, explains Levoy. At high resolution—less than a tenth of a millimeter or so—the scattering of the laser light from individual marble crystals creates tiny bumps on the 3-D model. To compensate for this, the computer program rejects oddly scattered data points.

Over the next few months, Levoy and his colleagues expect to finish compiling their models of several of Michelangelo's sculptures. Developing programs to display these large, complicated models will pose many technical problems, he says.

Although he has used his data to create low-resolution images of entire sculptures, "no computer in the world is big enough to view an entire statue at once" with a resolution of one-quarter of a millimeter, Levoy says. Portions of the artwork, however, can be examined at that level of detail.

Such precise measure of a sculpture's surface is impossible to achieve by any other means. "It's very hard to get exact measurements with a ruler," Wasserman notes.

The precision of digital models is valuable beyond looking at fine details and tool marks. Levoy's team discovered that Michelangelo's "David" was 3 feet taller than the 14 feet 3 inches most textbooks—and even the guidebooks sold at the gallery—reported. It isn't clear where that error originated.

The careful measurements of digital models offer art historians the opportunity to closely examine size relationships within and between characters in a sculp-

Paul Debevec/University of California, Berkeley

Michelangelo's "David," the 3-D scanner, and members of the Stanford team.

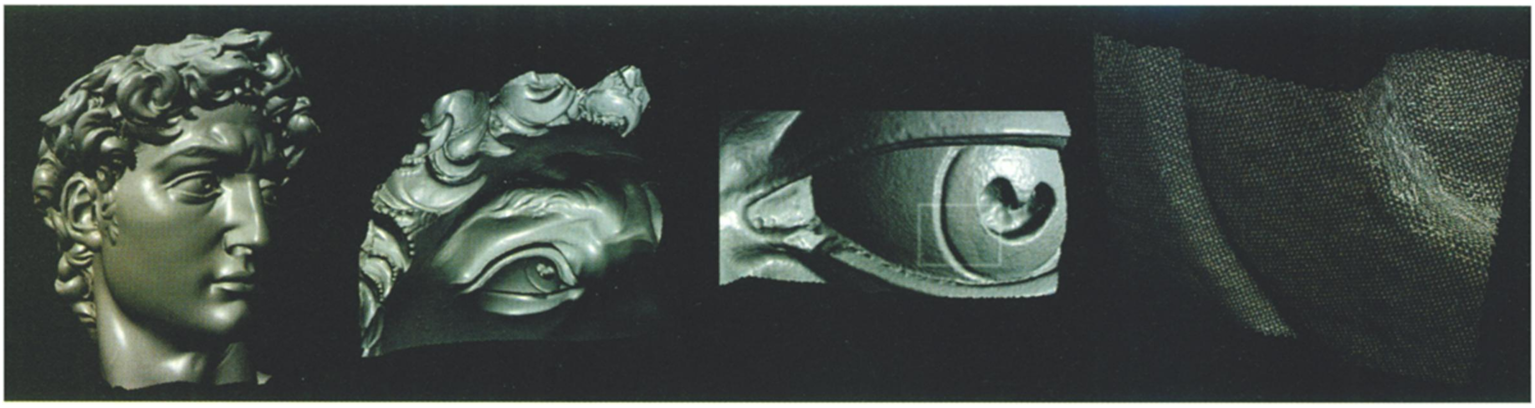

Images of "David." Left to right: A model of the head shown at 2-millimeter resolution; a 1-mm model of the right eye; a 0.25-mm model of the eye; and the triangle mesh obtained by scanning a portion of the eye several times from different directions. Levoy/Stanford University

ture. The proportions of Michelangelo's sculptures are often unrealistic. With a digital model—unlike a photograph or the actual artwork—researchers can separate characters and line them up. Then, they can compare a figure's proportions with those of other characters in the same sculpture, other sculptures, or an idealized human.

Michelangelo may have played with the proportions in his sculpture to make the piece look more realistic from the viewer's perspective. He designed some of his statues to be seen from below, some to be viewed from many sides, and others to be seen from just one angle. Digital models enable art historians to change the perspective from which they view the statue, a task that might otherwise be impractical. Actually hoisting the heavy marble would be complicated and expensive and would risk damaging the priceless art.

For example, Michelangelo probably expected the 10-ton "Florentine Pietà" to sit in a niche about 10 feet off the ground, Wasserman says. So he used the IBM model to simulate what the statue would look like if it could be raised. He concludes that the proportions of the characters would seem less curious if they were seen from below.

Imaging techniques such as those developed by these two teams could have more mundane applications. "Three-D technology is finding more widespread use, as faster computers and better hardware make it more accessible. We are trying to push the technology to the point where it becomes very easy to put rich, three-dimensional models on a Web site," says Josh Mittleman of the IBM team. Merchandisers with online catalogues, for example, might find that capability useful.

The techniques might also aid the entertainment industry in its perpetual quest for bigger, better, more complicated ways of showing imaginary creatures, points out Levoy. He also suggests that accurate 3-D imaging would be useful for architects remodeling or putting additions onto buildings. They could scan the existing structure to check whether it matches the original blueprint before using it as a basis for new plans.

Finally, he says, art museums could use 3-D models to create reproductions of their masterpieces that are more accurate than any photograph or drawing. Difficulties in enforcing digital copyrights may complicate that practice, Levoy says. Currently, it's not clear who would control and collect royalties from such images.

Those issues of copyright are all the more important because 3-D model building is not cheap. The Stanford team spent about $2 million to digitize 10 statues, the interiors of two buildings designed by Michelangelo, and the Roman map.

"One of the magical things about [such detailed 3-D modeling] is that it allows you to see things in a way you just couldn't before," says Erik Brisson of Boston University. "People can see sculptures more accurately than they can in a photograph; they can rearrange lighting and pieces to see things they couldn't see in real life; they can compare the sculpture with others in ways they couldn't before."

Carving these 3-D models out of virtual space is far from routine, says Brisson. It could be as long as a decade before the computing power to deal with these complicated computer models is readily available. Yet, these researchers suggest, it's increasingly clear that this technology one day will shape not only the way people view art but also the way they interact with their world. □