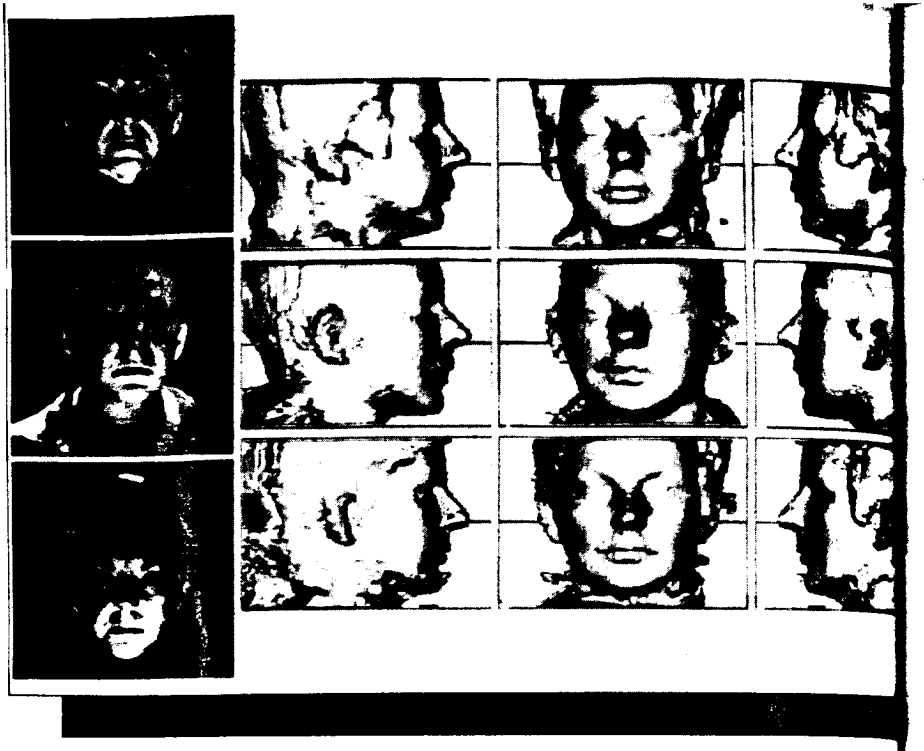

Facial Surface Scanner

Michael W. Vannier, Tom Pilgram, Gulab Bhatia, and Barry Brunsden
Washington University School of Medicine

Paul Commean
Cencit

We modified Cencit's optical noncontact 3D range surface digitizer to help us plan and evaluate facial plastic surgery.

A design team at Cencit developed a noncontact 3D digitizing system to acquire, process, display, and replicate the surface of the human head. Key requirements were all-around coverage of the complex head surface, accuracy and surface quality, a data acquisition time of less than one second (the approximate time a person can remain motionless and expressionless), and automated operation, processing, and object replication. The designers also wanted easy operation and operational safety in a medical clinical environment. The resulting design is unique in its combination of complex 3D surface coverage, accuracy, speed, and ease of use through fully automatic operating and 3D processing.[1,2] For this reason, we chose to modify the Cencit digitizer to meet our specific medical application, facial plastic surgery.

Other researchers have developed several different technical approaches for active, optical, noncontact range sensing of complex 3D digitization surfaces. Their techniques include laser moiré, holographic methods, and patterned light. Paul Besl[3] of the GM Research Laboratory recently reviewed 3D optical active ranging systems, used primarily for industrial machine vision.

0272-17-16/91/1100-0072$01.00 ©1991 IEEE

One aspect of the Cencit system makes it distinctive: the integration of multiple stationary sensors, which you can arrange to cover complex contoured surfaces. Another benefit of the approach is its digitization speed—less than one second for data acquisition. Processing and display requires less than 10 minutes on a Silicon Graphics Personal Iris 4D/20-GT workstation and less than two minutes on the more powerful 4D/240-GTX workstation. The Cencit team developed algorithms to enable automatic processing without operator intervention. Applications for the system include biomedicine, consumer portrait sculpture, and anthropometric studies. We modified the system to assess the facial changes possible with and resulting from plastic surgery.

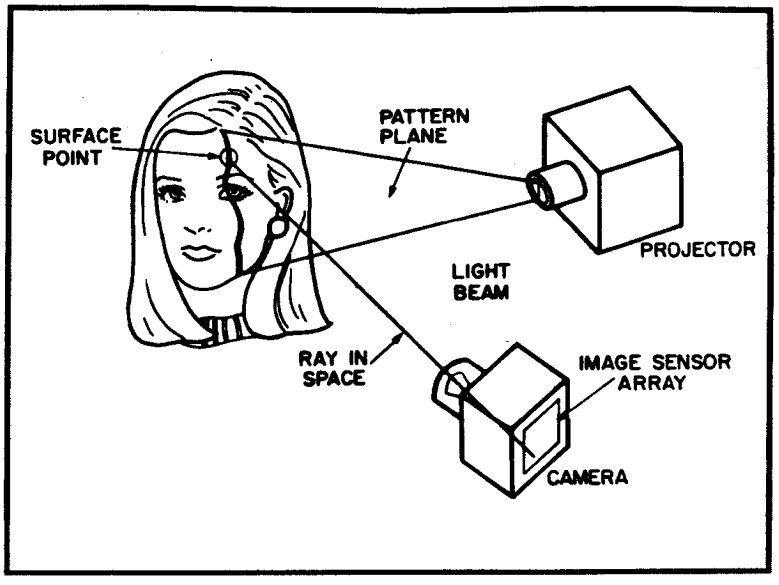

Figure 1. Methodology for determining 3D points in space. To identify this 3D point in space, we can use simple algebra. If we know the equation of the pattern plane and that the equation of the ray in space intersects the pattern plane, we can find the point of intersection. Drawing a line from the pixel location found in the image sensor array through the camera lens center to the subject determines the ray in space. Identifying the pattern number that produced the profile tells us the pattern plane.

Design concept for the 3D digitizer

The design team chose to use structured incoherent light to project a predetermined pattern of light onto the subject, viewing it with an area imager offset from the projector.[4] This offset is necessary to create a triangulation baseline. You determine positions of contours on the subject's surface by solving for the intersections of the known projected pattern surface and the rays passing through the lens of the imaging sensor onto its imaging plane. Knowing the positions, orientations, and other parameters of the projector and imaging sensor and observing the imaged intersection of the projected pattern with the subject's surface, you can find the solution.

The system employs a stationary, multiple-sensor fixed geometry, illustrated in Figure 1, rather than using a single moving-sensor approach. The designers arranged the sensors to cover the surface in overlapping segments for several reasons. First, with no mechanical motion required of either the sensors or the subject, they avoided the hazards typically caused by quickly moving devices: excessive mechanical resonances and vibrations, deflection caused by accelerations and centrifugal forces, and the problems of bearing play, maintenance, adjustment, and wear. They also avoided the expense, weight, and safety concerns involved with moving masses at high speeds in close proximity to patients. Second, they chose multiple sensors for their flexibility in positioning to reach portions of the surface perhaps not viewable by other methods, such as a single sensor rotating about the subject in a single horizontal plane. Without the constraints imposed by a motion path, you can position the stationary sensors to meet the application's needs. You can also select the number of sensors based on the surface's complexity, thus matching the system to the problem.

Given their choice of stationary rather than moving area sensors, the designers had to address a number of issues. To achieve the speed requirement, as well as to reduce the amount of image memory and processing needed, they decided that each sensor should digitize a surface segment, not just a single profile line as in the past. Another important design problem involved the number and arrangement of sensors needed to cover the entire surface of the human head. To successfully integrate multiple surface segments, you must obtain segments accurate enough that any two segments when joined produce unnoticeable seams, or merges. This imposes a far more stringent accuracy requirement upon each sensor than is the case for a single moving sensor, because the single moving sensor generates only one surface. Thus, in the single-sensor case nominal inaccuracies go unnoticed.

The problem of digitizing a complete surface segment (rather than only one contour at a time) from a given sensor position presented a number of problems that the team solved uniquely. The benefits of the solution are substantial, with a dimensional improvement in imaging and processing efficiency. This is one key to the digitizing speed achieved.

Because many light profiles are projected at once, each image contains many contour lines. Together these lines define a surface segment. Historically, this approach has encountered the problem of uniquely identifying each separate contour in the sensed image. This identification is necessary to correctly solve for the 3D surface. The concept employed to solve this problem, illustrated in Figure 1, involves using a sequence of several patterns, each including a portion of the total number

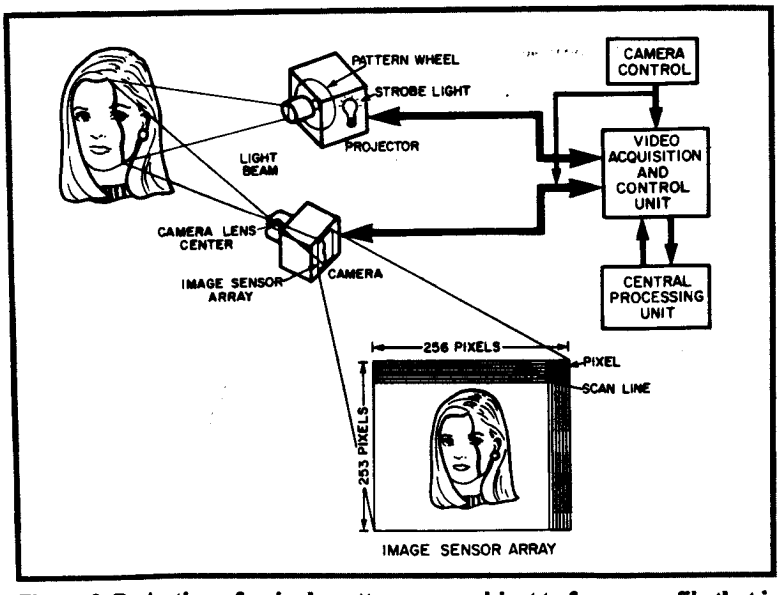

Figure 2. Projection of a single pattern on a subject to form a profile that is captured by the camera image sensor array.

of profiles. When interleaved, these profiles describe the complete surface. The key to identifying the individual contours lies in the interleaving pattern, which is coded so that you can uniquely identify the contours in subsequent processing (see Figure 2).

A further advantage arises from projecting a sequence of patterns, each containing a portion of the profiles: You can space the projected profiles widely enough to make them separable in the sensed images while providing dense surface coverage when you interleave the sequence of images.

System operation

The Cencit team found that they could digitize the surface of the human head by combining overlapping surface segments from a total of six sensor positions. They space three sensors circumferentially around and slightly above the subject, with three more interleaved among them but slightly below the subject's head. The sensors thus form a zig-zag pattern around the subject (see Figures 2 and 3). This provides coverage of areas (such as eye recesses and under the chin and nose) that a single sensor restricted to motion in a plane could not "see." With this configuration they found that most places on the surface of the head were covered by two or more sensors, thus providing a substantial amount of overlapping coverage. This assures a sufficiently complete digitized surface for the variety of subjects encountered.

Our application required a fully enclosed standalone system. A functional enclosure, shown schematically in Figure 4, provides rigid mounting points for the cameras and projectors. Rigid mounting of the stationary sensor elements provides long-term accuracy and infrequent need for calibration (typically, two or three times a year in a commercial environment). When the system does need recalibration, we can accomplish it with parameter estimation algorithms that process a known reference object.

A "sensor" consists of a pattern projector and a solid-state video camera. The projectors are sequenced by a module called the Video Acquisition and Control Unit. The operator initiates a digitizing session with a hand-held controller that, through a small embedded host computer, begins sequencing of the projectors and acquisition of the video images. For the system described here, this takes less than one second, during which the subject must remain still.

Upon completion of the video acquisition, the images are normally downloaded to a streaming tape for transport to a central processing facility. There—in the case of portrait sculpture, for example—the system processes the image data to compute the 3D surface, then replicates it on a standard numerically controlled milling machine or other reproductive device. Alternatively, you can process the images directly using a computer or workstation interfaced to the embedded host computer, as we did for our modified system. Using a Silicon Graphics 4D/340-VGX workstation, in less than two minutes you can digitize a subject and display on the workstation a shaded polygon model of the processed 3D surface.

The digitizing, 3D processing, and tool path generation are automatic, requiring no human intervention. You can process groups of digitized subjects in unattended batch mode, producing models if desired, ready for 3D graphics display or replication. The system achieves automatic operation through processing algorithms developed to perform all operations that otherwise would require interactive manipulation on a graphics

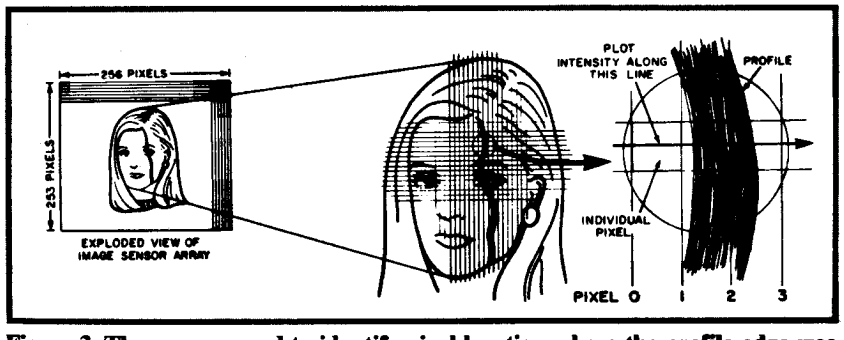

Figure 3. The process used to identify pixel location where the profile edge was found in the image sensor array. Light intensity profiles are evaluated in each pair of an image sequence to identify and locate local surface variations.

workstation. These algorithms rely heavily on statistical estimation as well as image processing.

In many medical and industrial applications, often practitioners must mark the subject to be digitized with reference points that are carried through the digitization process and displayed on the 3D surface model. For example, in digitizing a subject for medical orthotics, the technician must find and mark on the patient the locations of underlying bony prominences, then show them on the surface displayed for the orthotist.[5] The Cencit system can accommodate this and other special applications that require specific information and measurements in association with the digitized 3D surface.

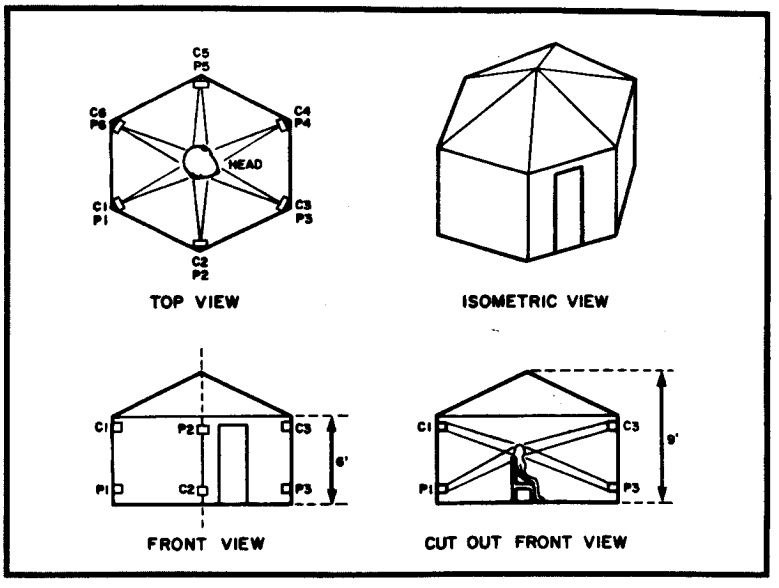

Figure 4. The scanning apparatus is contained in a hexagonal chamber that serves to provide structural integrity, exclude ambient light, and house much of the system electronics. Six camera-projector pairs are located about the subject at different elevations. These units operate in synchrony under the control of the Video Acquisition Unit.

3D digitization procedure

The projector contains a set of coded circular bar patterns for projection onto the subject. These patterns are captured in the camera image, mensurated, and tagged to identify each projector profile with the profiles as observed in the camera (see Figures 2, 3, and 5). A selected set of points belonging to a profile on the image plane is subjected to *2D to 3D solution*. The 2D to 3D solution refers to an analytical procedure whereby a point on the image plane (2D point) is translated to a point in space (3D point). In this procedure the circular projected profile is projected using the calibrated parameters of the projector, making a cone in space (see Figure 6). A point of the corresponding profile on the image plane is used along with calibrated camera parameters to form a ray in space. The intersection of this ray with the cone gives the 3D point in space. This procedure repeats for all points of all profiles to produce a set of 3D points lying on the surface of the scanned subject.

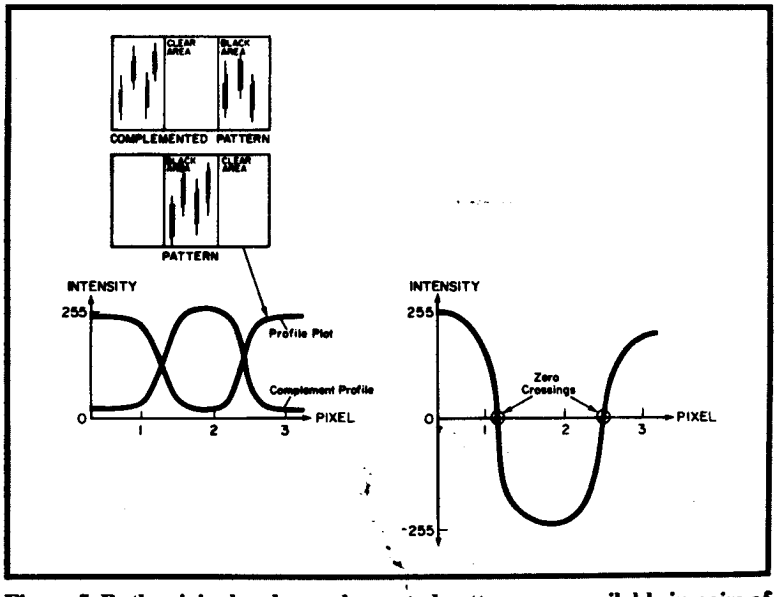

Figure 5. Both original and complemented patterns are available in pairs of images from the 144-frame sequence. By plotting the intensity profiles in each of these images (bottom left), we can determine the location of surface patterns. Summing the paired profiles, we see the intensity plot as a function of pixel location (bottom right). The zero crossings, interpolated to subpixel precision, provide an accurate, reproducible means of locating surface points in 2D. Combining multiple 2D profiles from pairs of adjacent cameras gives an accurate 3D surface estimate.

The determination of 3D points in space follows from the sampling geometry. Once you find the pixel location for a given pattern, you can solve the ray equation and pattern plane equation simultaneously to find the 3D point of intersection.

The method resamples the 3D points onto a uniform Cartesian or cylindrical grid. The location of each grid point is influenced by the weighting of each nonuniform point within a specified distance of the grid point. The method then sums and averages them to give a final value.

The pattern number identification for determining the matching pattern plane equation is an important practical issue. Since every pattern line in every set of patterns is uniquely identifiable, the following combination of observables makes the profiles (pattern lines) distinguishable:

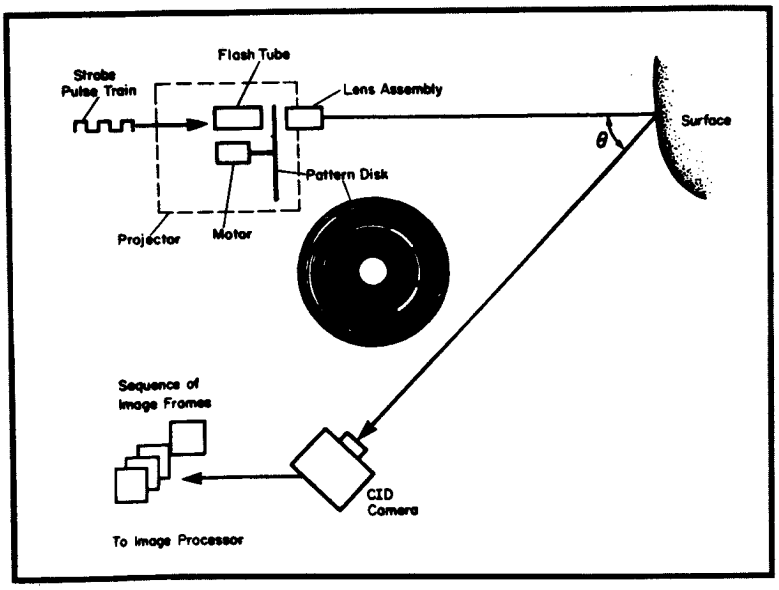

Figure 6. The camera-projector geometry. A pulsed flash-tube projector illuminates the object surface with eight different patterns of light and dark lines stored in different octants of a rotating pattern disk. A charge-injection-device camera views the surface at a fixed orientation, q, and synchronously samples 256×253 element frames. We use adjacent camera and projector pairs together so that the eight patterns and six projectors are viewed by three cameras each to form $8 \times 6 \times 3 = 144$ frames.

1. Identify in which set of patterns a profile lies, by knowing which pattern numbers correspond to the image sequence number.
2. Identify the direction of the profile boundary edge, that is, whether the profile edge progresses from dark to light or light to dark.
3. Identify the type of area, that is, light or dark, that lies at the corresponding physical position in each of the remaining three patterns.
4. Identify the sequence of pattern numbers corresponding to the sequence of imaged and mensurated profiles.

The process of mensuration and 2D to 3D solution is carried out for all camera-projector pairs (see Figure 4) to form surface patches as "seen" by these pairs. The system then combines the surface patches using 3D transformations and 3D resampling to form a complete surface representation of the scanned subject (see Figure 7).

The data from six cameras has a substantial amount of overlap. To achieve a seamless merging of this data, our method transforms each camera's data from its local coordinate system to a global coordinate system. It then resamples this data onto a uniform grid that uses groups of four adjacent points in a linear interpolation method. This procedure repeats for each camera. Following the resampling, we again have a substantial amount of data overlap, handled in our modified system by a constrained averaging of the overlap data. You can fill any holes (missing data) appearing in the surfaces by applying the resampling procedure at every point within the hole and using the four nearest points, one in each quadrant.

Image production

In our modified system, the 3D data set produced by the Cencit scanner is resampled in the form of a cylindrical grid (see Figure 8). The grid consists of 256 slices, each containing 512 radial data points equally spaced in azimuth. This data set contains holes and missing data segments, from regions obscured from the cameras in the surface digitzing process or those with low reflectance.

We transform the data set into a voxel format to use it with Analyze[6] software (see Figure 9). The voxel data set is a $256 \times 256 \times 160$ binary volume. In other words, the total volume consists of ap-

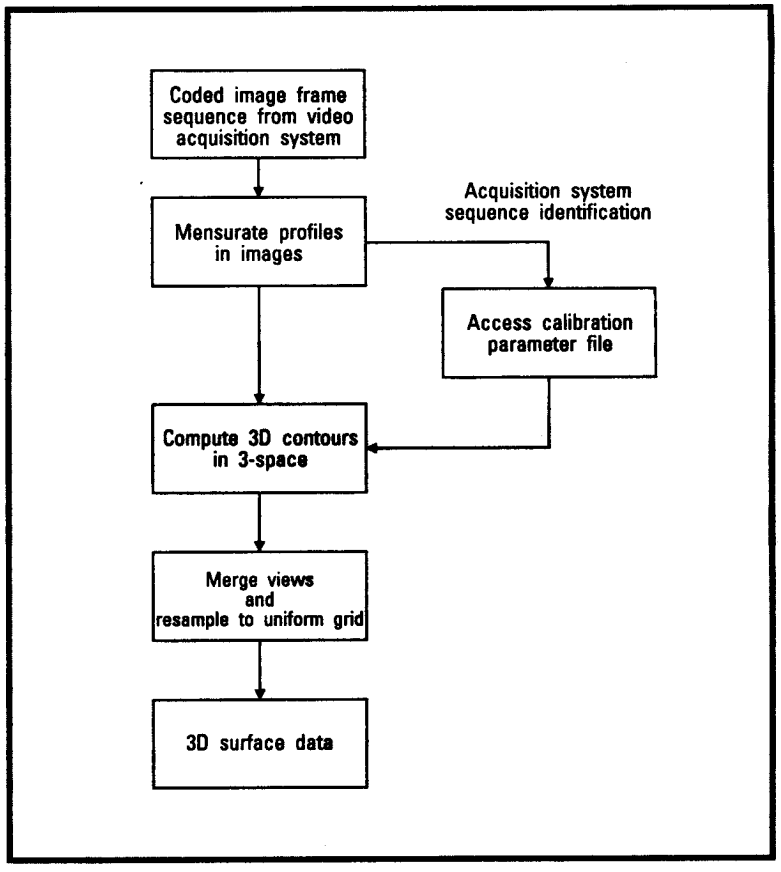

Figure 7. Data processing scheme for reconstruction of 3D surface coordinates from 2D image sequence.

proximately 10.5 million identical cubes, with the presence of the surface within the volume defined by a binary value for each cube. If the surface passes through a cube, its value is 1; if it does not, the value is 0.

We scale surfaces to make their sizes consistent with other images of the same type of subject. We do this by giving each data set an identical voxel size for a given image type. When we have consistent image sizes, we can interpolate images to fill in as many of the missing data points as possible.

Quality and accuracy

To evaluate the quality and accuracy of the digitized data and resulting images, we tested the digitization process and the scanner. Our findings follow.

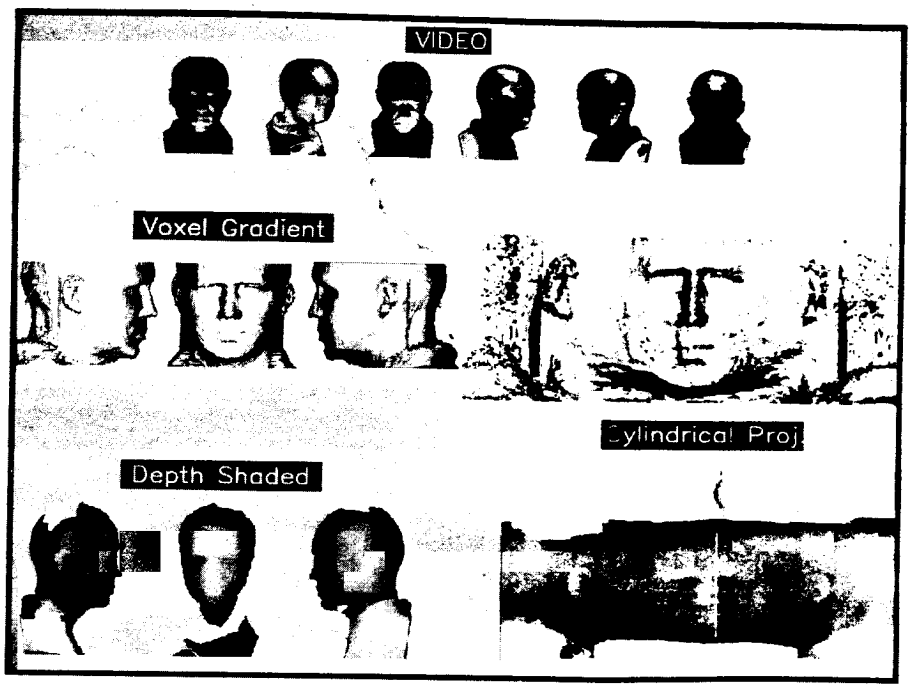

Figure 8. Video images (six from a set of 144) are shown at the top, one from each of the six cameras. We can represent these as isotropic voxels and render them using a voxel gradient in orthographic projections (lateral and frontal view—middle left) or cylindrical maps (middle right). We can render the reconstructed 3D surface data as orthographic (frontal and lateral—bottom left) or cylindrical views.

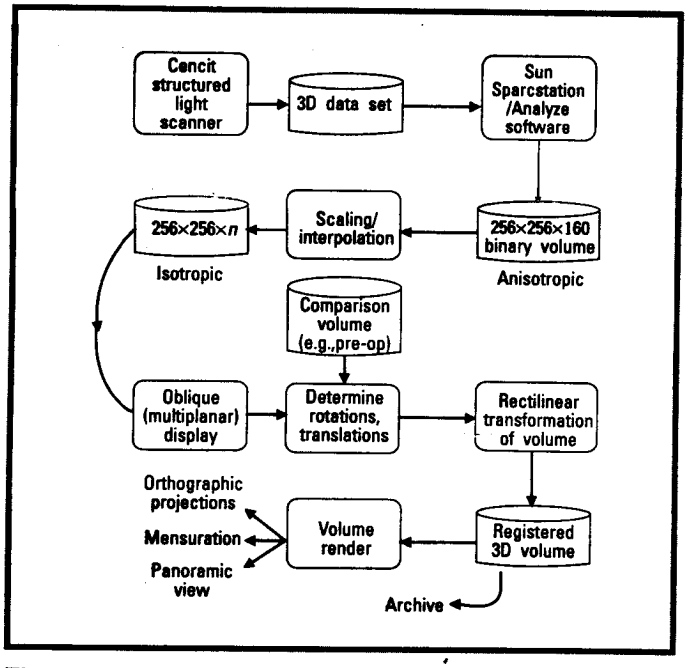

Figure 9. 3D data set processing. The Cencit scanner produces a 3D data set consisting of surface coordinates. We transfer these data via Ethernet to a Sun Sparcstation for processing with the Analyze software system from the Mayo Clinic. A binary volume of 256 × 256 × 160 anisotropic voxels is computed from the original 3D irregularly spaced surface coordinates. We scale and interpolate these data to isotropic voxels at 256 × 257 × *n* resolution. We use the multiplanar oblique reconstruction tool in Analyze to determine the translations and registrations needed to register the sampled data set with a previously stored reference volume. This might be a pre-op volume used in comparison to a post-op result, for example. A rectilinear transformation produces a registered 3D data volume that we can archive and volume render as needed.

Digitization accuracy

We found the accuracy of the digitization process to be on the order of 0.01 inch or 0.25 mm, as assessed by several different methods. Measurements made on known reference objects indicated errors of this magnitude. Calibration error residuals indicate a similar error magnitude. Finally, since all sensor pairs are calibrated separately, the error seen in overlapping data from different sensor elements provides a good indicator of error.

Image accuracy

We tested the quantitative accuracy of the image by comparing three images of a plastic surgery patient (see Figure 10). The images were made a few hours before (pre-op), 24 hours after (immediate post-op), and two weeks after surgery (late post-op). Surgery consisted of a browlift, facelift, nose trim, and chin implant.

We chose several standard surface anatomic points or landmarks on the face, based on our expected ability to relocate them consistently on this patient and on different

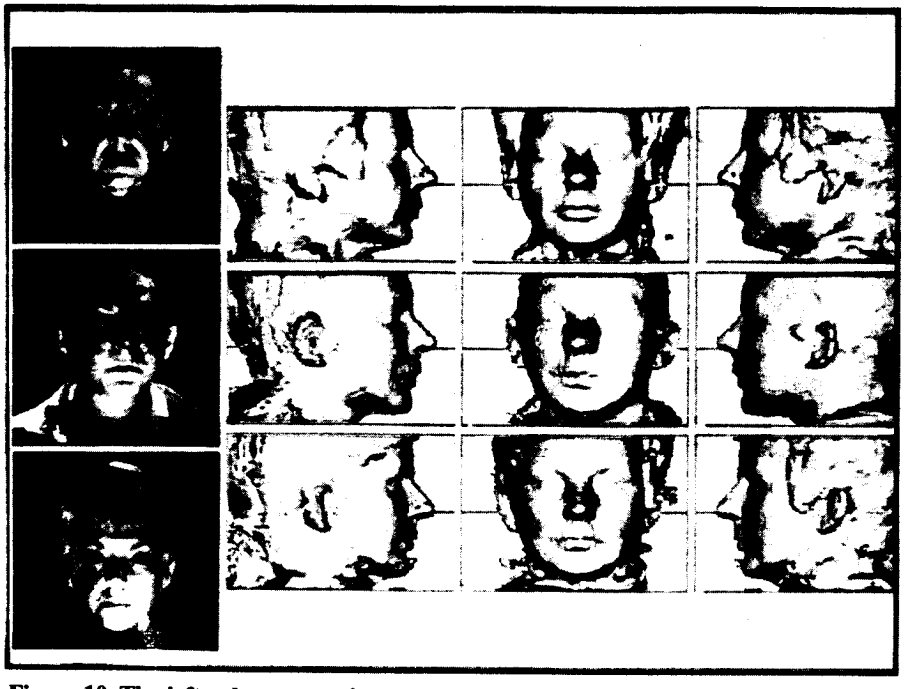

Figure 10. The left column contains raw unprocessed 256 × 256 video images. The top row is pre-op, the middle row immediately post-op, and the bottom row late post-op. The right three columns show voxel gradient volume rendering of the Cencit facial surface data. Cylindrical surface data was converted to 256 × 256 × 156 × 1 bit, where $x + y = 1.27$ mm and $z = 1.6$ mm. This produced a set of contour slices. One-voxel-thick contours do not produce suitable volume rendered images, so we added a one-voxel thickness to the inside of the contour. This allowed adequate volume rendering.

patients (see Figure 11). To test our ability to locate points consistently from one image to another, we located the points on the facial midline twice, once from a right 45 degree angle and once from a left 45 degree angle. Because the points were all on the same image, no registration error exists to consider. The only source of error arises from the operator's inability to perfectly locate anatomical markers on the image.

The size of the marker location error depends on the anatomical point being located. The menton (the bottom of the chin; see Figure 11, point 6c) is the most difficult point to locate, particularly in the horizontal (x) dimension. The location error for the menton is smaller in the depth (y) dimension and comparable to other anatomical points in the vertical (z) dimension. Locating the labiale superius (the center of the upper lip; see Figure 11, point 5c) also produces some error, but not as much as with the menton. The horizontal error, again largest, was roughly comparable to other anatomical regions in the y and z dimensions.

The size of the error was generally largest in the horizontal dimension, regardless of anatomical point. With all anatomical points and all three stages (pre-op, immediate post-op, and late post-op) included, the mean horizontal error measured 3.0 voxels, or 3.8 mm (1 voxel = 1.27 mm). The mean depth error was 1.3 voxels, or 1.7 mm (1 voxel = 1.27 mm), and the mean vertical error was 1.2 voxels, or 1.9 mm (1 voxel = 1.60 mm). Producing an image with greater voxel density would probably help reduce the error, both by making it easier to locate comparable anatomical points and by reducing the physical dimensions of each voxel, thereby reducing the consequence of a one-voxel error.

The mean error, excluding the horizontal dimension, is about one voxel. The size of the horizontal error probably results from the way we displayed the image to test repeatability. Rotating an image 90 degrees horizontally will have the greatest effect on our ability to locate points horizontally. Although the most practical way to do a repeatability test, this exaggerates the amount of error. The typical error from images processed in this fashion probably measures slightly more than one voxel, or a bit less than 2 mm.

To compare the location of anatomic points in space, we must register the images as closely as possible. A number of factors complicate this problem. For one thing, the angle of the subject's head usually changes during the scan. These are not simple rotations, because the change in each dimension moves around a different center. In addition, it is difficult to pick good registration points, because the rotation alters the surface description. Three good registration points independent of the surface—thus immune to alteration—would allow exact registration. With a typical data set, however, truly correct registration is impossible. More elaborate procedures, although probably more accurate than simple ones, will still have some level of error. Worst of all, we cannot know exactly the magnitude and direction of these errors.

We used a simple registration procedure. The otobasion inferius (the point where the earlobe joins the face; see Figure 11, point 1r) served as our reference point because, of all points on the face, its position seemed likely to be the least altered by surgery. Also, it is probably the easiest to locate exactly on different images. We registered right and left side measurements by adding or subtracting the change in position from each measurement on the appropriate side. We registered midline measurements by adding or subtracting the mean of the right and left side changes. This registration procedure probably compensates quite well for simple position changes, reasonably well for lateral head tilt, somewhat less well for horizontal rotation, and not very well for vertical tilt.

The failure to compensate for vertical tilt made an additional registration procedure necessary when examining vertical change along the profile. We noticed an apparent upward tilt of the profile in the immediate post-op image. The second registration was done in only the vertical dimension. We used the nasion (the junction of the nose and forehead; see Figure 11, point 2c) as the landmark because the close conformity of the skin to a pronounced underlying structure at that point makes its location the least likely of points on the profile to be affected by surgery. We adjusted the values of both post-op images so their values at the nasion were identical to the pre-op measure.

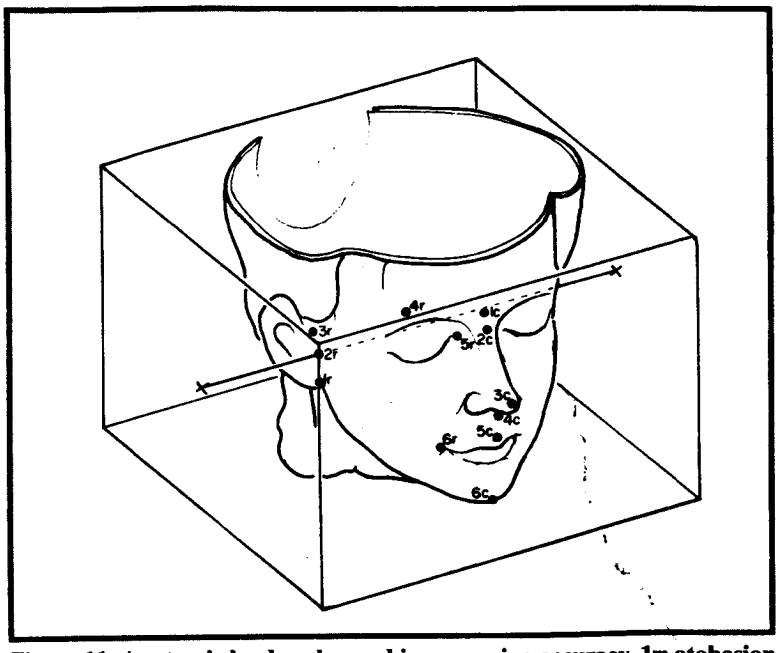

Figure 11. Anatomic landmarks used in measuring accuracy. 1r: otobasion inferius; 2r: *t* region; 3r: preaurale; 4r: superciliare; 5r: endocanthion; 6r: cheilion; 1c: gonion; 2c: nasion; 3c: pronasale; 4c: subnasale; 5c: labiale superius; 6c: menton.

Medically relevant results

The data show two clear cases of simple edema or facial swelling. The horizontal locations of the preaurale (the junction of the upper front part of the ear and the face; see Figure 11, point 3r) and the superciliare (the point on the eyebrow where the forehead joins the temple; see Figure 11, point 4r) change in such a way that the total width of the head at those points increases in the immediate post-op measurement. Then, in the late post-op measurement it decreases to approximately the original width. The results are more consistent for the preaurale than for the superciliare.

The results show slight registration errors, probably due to horizontal rotation, but this has no effect on the measurement of total width. We would expect some point location errors, but not large enough to be responsible for these changes in width. Also, both the direction and magnitude of the changes are consistent with the expected physiological effects of this type of surgery.

At least two points on the profile conform to a pattern of surgical change initially modified by edema. First, the vertical location of the pronasale (the tip of the nose; see Figure 11, point 3c) moves markedly upward in the immediate post-op image, then slightly further upward in the late post-op image. Because the patient's nose was shortened and reshaped, this is the most likely cause of the change in vertical location. The slight vertical difference between the immediate and late post-op images probably results from edema, which was present immediately following surgery and disappeared before the later image was made.

Second, the vertical location of the menton (the bottom of the chin; see Figure 11, point 6c) moves markedly upward between the immediate and late post-op images, though there is no change between the pre-op and immediate post-op measurements. The skin under the subject's chin was tightened as part of a general facelift, and this is almost certainly the cause of the upward change in the location of the menton. The absence of a change between the pre-op and immediate post-op images probably results from edema completely masking the surgical change.

The gonion (the center of the eyebrows; see Figure 11, point 1c) also shows some vertical change. Because it is a difficult point to locate precisely, the changes might result from location error. However, the pattern of first upward and then downward gonion movement, ending up slightly above the original position, is consistent with a facelift where edema initially exaggerated the amount of skin tightening.

The positions of the subnasale (the point centered just below the nose; see Figure 11, point 4c) and labiale superius (the center of the upper lip; see Figure 11, point 5c) also show changes. These likely result from slight differences in the way the patient held her mouth during the different imaging sessions. However, the direction and size of the changes is also consistent with both a slight tightening of the skin and a shortening of the nose.

Conclusions

We adapted the Cencit scanner, developed for facial portrait sculpture, to use as a medical imaging system. We applied its special capabilities—rapid, safe, noncontact 3D measurements in a form you can display and manipulate on a computer graphics workstation—to the quantitative assessment of facial plastic

surgery. Our results demonstrate that the Cencit system's accuracy is adequate for quantitative studies of facial surfaces.

We continue to pursue our investigations on several fronts. Our current work focuses on increasing the accuracy of facial surface measurements. Improved registration is one of the most important needs, so we are exploring complex algorithms that use the entire facial surface in the registration process. Since location of anatomical points on different images also constitutes an important source of error, we are looking into ways to describe portions of the face with mathematical models. This would let us locate anatomical points more objectively, on the basis of quantitative measures, rather than subjectively, as we do now. These improvements will greatly increase the system's usefulness in facial research applications.

The Cencit system might also prove useful in surgical planning. Currently, it provides a way to record and view a 3D facial surface image acquired noninvasively. This assists planning more than ordinary photographs. If the system could modify images in real time, as many engineering CAD/CAM systems do, the surgeon could easily experiment with alternatives. Alterations viewed in combination would increase the surgeon's ability to foretell the cumulative aesthetic effect of multiple subtle modifications. The patient could view the potential outcomes as well, and become a better informed participant in the decision. Planning for facial plastic surgery can thus become a more thorough and interactive process. ❑

Acknowledgments

This work was supported in part by the State of Missouri, Missouri Research Assistance Act.

We wish to thank Michel Morin and Universal Vision Partners for their support and continued encouragement. John R. Grindon was responsible for the initial system concept and design. Technical discussions with Arjun Godhwani of Southern Illinois University at Edwardsville are gratefully appreciated. Clinical application of the system for facial plastic surgical studies was performed with Leroy V. Young and Jane Riolo of the Division of Plastic Surgery at Washington University Medical Center and Barnes Hospital. The Analyze software system was provided by Richard Robb and Dennis Hanson of the Mayo Biodynamics Research Unit in Rochester, Minnesota. We appreciate suggestions regarding the manuscript presentation by Ronald Walkup. Manuscript preparation by Mary M. Akin is gratefully acknowledged.

References

1. J.R. Grindon, *Optical Means for Making Measurements of Surface Contours*, U.S. Patent No. 4,846,577, issued July 11, 1989.
2. J.R. Grindon, *Means for Projecting Patterns of Light*, U.S. Patent No. 4,871,256, issued Oct. 3, 1989.
3. P.J. Besl, "Active Optical Range Imaging Sensors," *Machine Vision and Applications*, Vol. 1, 1988, pp. 127-152.
4. J.R. Grindon, "Noncontact 3D Surface Digitization of the Human Head," *NCGA Conf. Proc.*, Vol. 1, April 1989, pp. 132-141.
5. G. Bhatia, A. Godhwani, and J.R. Grindon, "Optimal 3D Surface Metrology—Localization of Fiducial Points," *IEEE Proc. Southeastcon 91*, Vol. 2, 1991, pp. 925-930.
6. R.A. Robb et al., "Analyze: A Comprehensive Operator-Interactive Software Package for Multidimensional Medical Image Display and Anlaysis," *Computerized Medical Imaging and Graphics*, Vol. 13, 1989, pp. 433-454.

Michael W. Vannier is presently a professor in the department of radiology at the Mallinckrodt Institute of Radiology at the Washington University School of Medicine in St. Louis, Missouri. His primary interests are research in radiology and 3D imaging.

Vannier graduated from the University of Kentucky School of Medicine and completed a diagnostic radiology residency at the Mallinckrodt Institute of Radiology. He holds degrees in engineering from University of Kentucky and Colorado State University, and was a student at Harvard University and the Massachusetts Institute of Technology.

Paul Commean is a senior engineer at Cencit in St. Louis, Missouri, where his primary work interest is the development and production of 3D surface scanner equipment. From 1982 to 1985 he designed and integrated automatic test equipment for the F-18 Flight Control Computer of McDonnell Douglas in St. Louis.

Commean graduated from the Georgia Institute of Technology in 1982 with a bachelor's degree in electrical engineering.

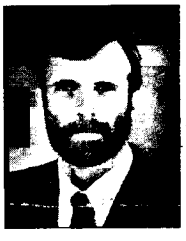

Thomas Pilgram is a research associate at the Mallinckrodt Institute of Radiology. His research there has concentrated on diagnostic performance and its measurement. From 1983 to 1985, Pilgram was funded by the New York Zoological Society for a study of the effect of the international ivory trade on the African elephant population. From 1985 to 1988 he held academic appointments in anthropology at the University of California, Berkeley and Washington University, St. Louis.

Pilgram received a BA degree from the University of California, San Diego in 1974, and a PhD degree from the University of California, Berkeley in 1982, both in anthropology.

Gulab Bhatia is presently a research engineer working in the School of Medicine at Washington University. He is interested in 3D imaging and computer graphics. From 1987 to 1990, he worked as a senior engineer in charge of software and algorithm development for Cencit, developing 3D scanner systems.

Bhatia received his BSEE from Birla Institute of Technology and Science, India in 1982 and his MSEE from Southern Illinois University at Edwardsville in 1987.

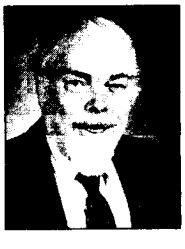

Barry Brunsden joined the Mallinckrodt Institute of Radiology in 1990 and has been involved chiefly with 3D imaging and measurement methods. He has been involved in 3D imaging and measurement since he worked in a physics lab of the Ontario Cancer Foundation in the late 1950s. From 1960 to 1989 he was employed at the University of Chicago developing instrumentation for gamma, *x*-ray, ultraviolet, and visible light measurements and images in the areas of biophysics, electrophysiology, cardiology, and nuclear medicine. He was responsible for developing quality assurance procedures and digital image processing techniques in nuclear medicine.

Readers may contact the authors at the Mallinckrodt Institute of Radiology, Washington University School of Medicine, 510 S. Kingshighway Blvd., St. Louis, MO 63110.