

Gulf and Caribbean Research

Volume 13 | Issue 1

January 2001

Accurate 3-D Morphological Measurement Using a Structured-Light Range Sensor

Robert L. Cromwell

Purdue University

Stuart G. Poss

Gulf Coast Research Laboratory

DOI: 10.18785/gcr.1301.02

Follow this and additional works at: <http://aquila.usm.edu/gcr>



Part of the [Marine Biology Commons](#)

Recommended Citation

Cromwell, R. L. and S. G. Poss. 2001. Accurate 3-D Morphological Measurement Using a Structured-Light Range Sensor. *Gulf and Caribbean Research* 13 (1): 11-28.

Retrieved from <http://aquila.usm.edu/gcr/vol13/iss1/2>

This Article is brought to you for free and open access by The Aquila Digital Community. It has been accepted for inclusion in *Gulf and Caribbean Research* by an authorized editor of The Aquila Digital Community. For more information, please contact Joshua.Cromwell@usm.edu.

ACCURATE 3-D MORPHOLOGICAL MEASUREMENT USING A STRUCTURED-LIGHT RANGE SENSOR

Robert L. Cromwell¹ and Stuart G. Poss²

¹*Robot Vision Laboratory, School of Electrical and Computer Engineering, Purdue University West Lafayette, IN 47907-1285, USA E-mail: cromwell@ecn.purdue.edu*

²*Gulf Coast Research Laboratory Museum, Institute of Marine Sciences, The University of Southern Mississippi, P.O. Box 7000, Ocean Springs, MS 39566-7000, USA, E-mail: Stuart.Poss@usm.edu*

ABSTRACT A single-plane structured light range-sensor was tested to establish its usefulness in acquiring 3-D measurements of fish skulls. Twenty-one distances among 22 landmark points for each of 12 neurocrania of the scorpaenid fish *Neomerinthe hemingway* were taken with digital calipers, with a video-based 2-D imaging system widely used in systematic studies, and with a single-plane structured-light range sensor of inexpensive and simple design. Measures taken by 3-D sensor are highly correlated with those obtained from caliper measurement ($r = 0.9995$, $P \ll 0.001$), with a precision ranging from 0.08–0.43 mm. Like caliper-based measurements, they are less strongly correlated with measurements derived from projected video-imaging. Most skulls were scanned in just over 3 minutes each. Range maps, typically establishing the (x,y,z) coordinates of more than 75,000 points per scan, can be obtained in about 40–50 CPU seconds using software running on multiple platforms. Sensor data taken from different views can be merged to build a more complete 3-D reconstruction. System design, calibration, and use are discussed. By eliminating error due to perspective effects inherent in measuring from projected video images, such sensors hold considerable promise in quantifying biological shape in 3-D for comparative and functional studies.

When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge of it is of a meager and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced it to the stage of science.

Sir William Thomson, Lord Kelvin (1824–1907)

INTRODUCTION

Since measurement of anatomical dimensions provides an extremely important means to distinguish among taxa, it is central to the practice of taxonomy and systematics. Measurement is essential to studies of growth and development, as well as comparative investigations of function. Mensuration also forms a necessary element in the advancement of theoretical systematics by providing empirical data needed to refute hypotheses. Measurable differences, although seldom actually taken, are also implicit in establishing morphological character states for cladistic study and in providing identifications that unify morphological and molecular systematics.

Perhaps because of their ubiquity, much attention has been directed at what measures might be most appropriately taken and how they can be evaluated. Much less attention has been given to the study of how measurements may be acquired more accurately and precisely, with numerous publications often failing to indicate what acquisition methods were used. Given the unification and standardization measurement provides science (Klein 1974, Wise 1994), the constant need to obtain

more and better data, and recent advances in computer imaging, it is useful to investigate alternative approaches of measurement.

In taxonomy and systematics, measurements of macroscopic objects have been taken primarily with vernier or dial calipers. In the past, such measurements were transcribed from data sheets and then analyzed. When linked via RS-232 interfaces, electronic calipers now provide a means of acquiring measurement data directly without the need of a separate data entry step. Although this has resulted in data capture rates roughly 5 to 10 times faster than traditional methods, a relatively limited number of measurements can be captured by this means. Consequently, calipers are seldom used to study quantitatively many complex shapes, such as those observed in fish bones. The development of digitizers has permitted morphometricians to use more readily projected images to measure specimens (e.g., Currents et al. 1989, Reilly 1990, Brooks 1991, Hastings 1991), with the use of video-digitizers becoming increasingly common (e.g., Fink 1987, White and Prentice 1988, White et al. 1988, Lindberg 1990, Ray 1990, Sanfilippo and Riedel 1990, Ehlinger 1991, Douglas 1993, Fink 1993, Meacham 1993, McElroy and Kornfield 1993, Newton

and Kendrick 1993, Wimberger 1993, Kaiser et al. 1994, Robinson and Wilson 1995, Zelditch and Fink 1995, Zelditch et al. 1995, McElroy and Douglas 1995). Applications of such 2-D techniques can increase the rate of data acquisition an additional one or 2 orders of magnitude. However, for strongly 3-dimensional objects for which measures can extend out of the plane of the projected image, error resulting from measurement of a foreshortened projected image can be significant. Measurement in multiple orientations and repeated calibration may be required to correct for error due to foreshortening. Although some (e.g., Corner et al. 1992, Richtsmeier et al. 1993) have employed 3-D digitizers to circumvent this problem, most existing tools provide only a limited number of data points for larger objects. Bookstein et al. (1991) successfully demonstrated agreement between measures of the human skull taken from lateral and posteroanterior projections of 3-D cephalograms produced using computed tomography and caliper measures of the same landmarks.

Machine vision offers a variety of sensing methods for accurately and more comprehensively measuring in 3-D. Confocal microscopy provides 3-D measurement for suitably small objects (e.g., Foster et al. 1990, Birkmann and Lundin 1996). Computed Tomography (CT), Nuclear Magnetic Resonance Imaging (NMRI), and Positron Emission Tomography (PET) scanning techniques are powerful, in wide use for medical applications, and have considerable potential for use in measuring biological materials (e.g., Conroy and Vannier 1984, Zangrel and Schultze 1989, Schultze 1991a,b, Kalvin et al. 1995). However, costs associated with these sensors make their use impractical for most systematists. Consequently, development of low-cost imaging systems capable of accurate measurement will likely remain focused on systems utilizing visible light, at least in the near future.

Typically (x,y,z) locations, visible on a target object, are gathered as an array of measured points. This 2-dimensional array of (x,y,z) positions is commonly called a "range map". Range maps may be passively or actively produced. Active methods require special lighting to illuminate the object. Passive methods do not. Without such lighting, range maps are often sparse. Light from a given point on the target may not reach the sensor or it may result in ambiguous values. Nonetheless, one passive method, stereopsis, has seen wide use (Kaufman 1964, Julesz and Miller 1975, Mayhew and Frisby 1976, Yakimovsky and Cunningham 1978, Marr and Poggio 1980, Baker and Binford 1981, Grimson 1981, Mayhew and Frisby 1981, Barnard and Fischler 1982, Marr 1982,

Boyer and Kak 1988). However, stereopsis requires matching points between "left-eye" and "right-eye" images, an ambiguous task without tightly controlled or structured lighting. The resultant range maps can be noisy as well as sparse. Automated stereopsis thus produces significant measurement error. Manual methods reduce ambiguity, but require much work per measured point. Other passive techniques, such as range-from-focus and shape-from-shading methods (Horn 1975) also provide sparse maps, and they provide relative position only.

Active sensing methods provide more complete range maps. Laser radar probes a target with a collimated beam (e.g., Caulfield et al. 1977, Lewis and Johnston 1977, Nitzan et al. 1977, Riggs et al. 1986). Time of flight is measured for pulsed systems, while FM or AM systems modulate the frequency or amplitude of the transmitted beam and compare the transmitted signal to the return signal to determine distance. However, system cost and complexity are high. Because beams used in such systems dwell on any given point on the target for only a very brief moment, power must be increased to provide sufficient lighting. As a result, such systems are not necessarily eye-safe and are relatively expensive. Flying-dot systems can gather good range data (Rioux 1986, Blais et al. 1988), but again system cost and complexity are high for designs that gather data quickly.

Structured light is a relatively simple active sensing method. Although complex designs involving multiple light stripes (Yeung and Lawrence 1986, Boyer and Kak, 1987) or Gray-coded binary patterns (Inokuchi et al. 1984) are possible, simple sensors that produce range maps of high quality can be built from inexpensive components. A single-plane structured light range sensor is simply a projector, which produces a plane of light, and a camera, which views the intersection of that light plane with the target. The system shown in Figure 1 provides an example of one possible arrangement. When the light plane from the projector intersects a staircased target it forms a light stripe that the camera sees as a series of line segments. With a monochromatic light source and a bandpass filter on the camera, the image contains only those line segments. It is easy to see that the shape of the target defines the general shape of the camera's image. Less obvious, but critical to this sensing method, is that if the optical characteristics of the camera and the geometric relationships of the camera and projector are known, then the (row, column) location of an illuminated point in the image defines a unique (x,y,z) location on the target.

A single-plane structured light sensor is really a 2-D sensor—it gathers 3-D data by assembling a series of 2-D slices. The sensor (or target) is moved so that the light plane sweeps across the volume of interest. Assembling data for individual slices yields a 2-dimensional array of (x,y,z) measures. Each digitized video frame, corresponding to a 2-D slice, yields one column of (x,y,z) values. In our experiments we moved the sensor on an overhead track; moving the target on a stage would yield identical results.

Single-plane structured light range-sensing technology has been widely employed in industrial robotics applications (e.g., Shirai 1972, Agin and Binford 1976, Posdamer and Altschuler 1982, Sato et al. 1982, Yang and Kak 1986, Kak et al. 1987, Wang et al. 1987, Hu and Stockman 1989, Hutchinson et al. 1988, Hutchinson et al. 1989, Hutchinson and Kak 1988, Hutchinson and Kak 1989, Kak et al. 1988, Chen and Kak 1989, Kim and Kak 1991, LaValle and Hutchinson 1991, Cromwell 1992, Cromwell 1993, Wang et al. 1994, Grewe and Kak 1995). However, its use for measuring biological objects has been limited (McLeod 1991).

In this paper, we describe the design, calibration, and use of a 3-D sensor, and the steps involved in processing sensor data for presentation and investigation. We then compare measurements between landmark points on osteological materials using data obtained from the sensor against those gathered using methods already familiar to most systematists. Such comparison provides an estimate of their relative accuracy and a demonstration of their potential usefulness in systematic studies.




Figure 1. Range-data gathering with a structured light sensor. C = camera position; P = laser projector; IP = image plane; LS = light stripe.

METHODS

Our approach is to scan a target with a low-cost laser range finder to produce a dense array of (x,y,z) measures. A set of images is then displayed on a workstation. As the user selects landmarks of interest, 3-D distances between landmarks are displayed. The overall process may briefly be described as: 1) calibrate the sensor; 2) sweep the sensor across the volume of interest, producing a 2-D array of initial sensor data; 3) convert the sensor data into a 2-D array of (x,y,z) measures; 4) calculate a 2-D array of local surface orientation from the (x,y,z) data; 5) resample the (x,y,z) and local surface orientation arrays to ensure that the displayed images will have appropriate aspect ratios when displayed with square pixels, if necessary; 6) produce images to present the (x,y,z) data in a variety of forms; 7) present those images to the user for specification of landmark points from among the approximately 75,000 to 100,000 measured points; and 8) derive distance measures between landmarks and report these values to the user.

Sensor Design Considerations

There are 2 primary criteria for designing a structured light range sensor. The first is the field of view. The volume within the camera's field of view beyond the image plane is the "viewing frustum". The intersection of the frustum with the light plane generated by the projector forms a trapezoid. Points within this trapezoid can be simultaneously illuminated and viewed, thereby defining the field of view of the sensor. This imposes a limit on the cross-sectional dimensions of the target object.

The second criterion is resolution, or more appropriately, the spatial quantization. The number of rows and columns in the digitized image limit the possible (x,y,z) measurements to a finite set. The spacing of the corresponding points within the trapezoidal field of view of the range sensor determines the achievable resolution.

Spatial quantization is a function of several variables, including field of view. It is defined in terms of the 3-D distance between points corresponding to adjacent image locations. The 3-D distance between adjacent points on a single row of the image is of particular importance. As illustrated in Figure 2, spatial quantization is a complicated function of the following parameters: 1) the baseline distance, B the perpendicular distance from the light plane to the camera; 2) the toe-in angle, Φ —the angle between the camera's optical axis and the baseline; 3) the horizontal angular field of view of the camera, θ_h —the angle between the left and right

planes of the viewing frustum; 4) the vertical angular field of view of the camera, θ_v —the angle between the top and bottom planes of the viewing frustum; 5) the number of rows in the digitized image, N_R ; and 6) the number of columns in the digitized image, N_C .

Formally, spatial quantization has row and column components:

$$\delta_R(r, c) = \left| \bar{X}(r, c) - \bar{X}(r+1, c) \right| \quad (1.1)$$

and

$$\delta_C(r, c) = \left| \bar{X}(r, c) - \bar{X}(r, c+1) \right| \quad (1.2)$$

where,

$$\bar{X}(r, c) = \begin{bmatrix} x(r, c) \\ y(r, c) \\ z(r, c) \end{bmatrix} = \begin{bmatrix} B \tan(\Phi + \psi_h(c)) \\ 0 \\ \frac{B \cos(\psi_h(c))}{\cos(\Phi + \psi_h(c))} \tan(\psi_v(r)) \end{bmatrix} \quad (1.3)$$

and where,

$$\psi_h(c) = \tan^{-1} \left[\left(1 - 2 \frac{c-1}{N_C - 1} \right) \tan \left(\frac{\theta_h}{2} \right) \right] \quad (1.4)$$

$$\psi_v(r) = \tan^{-1} \left[\left(1 - 2 \frac{r-1}{N_R - 1} \right) \tan \left(\frac{\theta_v}{2} \right) \right] \quad (1.5)$$

A complete derivation of spatial sampling measures, including some sensor design guidelines, is presented in Cromwell (1992). Because $\delta_R(r, c)$ and $\delta_C(r, c)$ can not be readily used to determine the sensor




Figure 2. Geometric relationships among sensor elements that determine the level of spatial quantization for a given sensor design. B = baseline distance, Φ = toe-in angle, θ_h = horizontal angular field of view of the camera, θ_v = vertical angular field of view of camera.

design parameters (B , Φ , θ_h , θ_v , N_R , and N_C), sensor design tends to be an iterative process. The best values for these parameters for a particular class of target size and shape are obtained through experimentation.

As for physical construction, Figure 3 depicts a compact sensor design. For this experiment, we used a Pulnix TM-540 camera and a series of lenses, selected and positioned depending on the size and orientation of target skulls. The lens used for each of the scans is given in Table 1. Lateral and dorsal views were taken with the sensor position at location L1 of Figure 3. Ventral views required a top-down orientation (location L2). A wider-angle lens was used for larger targets. The sensor-to-target distances were adjusted so that the target just fit within the field of view. The Pulnix camera has a 510×492 pixel CCD array and produces an RS-170 video signal. This signal was sampled to produce a 512×480 image using an Imaging Technologies FG-100-V video digitizer attached to a Sun 3/280. The light source was a low-cost 0.5 mW HeNe laser producing collimated red light at a wavelength of 632.8 nm.

To spread the collimated beam into a plane we used a 3-element lens system. The beam first passed through a pair of cylindrical lenses separated by 7 mm. The first lens had a focal length of 1.2 mm. The second “lens” was simply a 4 mm glass rod, carefully aligned to present a section free from surface and internal defects. To focus the width of the light stripe to about 0.3–0.5 mm, the beam was passed through a biconvex lens with a focal length of 30 cm separated from the preceding lens by 20 mm.

The sensor can be arbitrarily positioned along a rail. We adjusted the position to view the target from the




Figure 3. Sensor design used in this study. B = ball screw carriage; C = camera; G = gantry; L₁ = position 1; L₂ = position 2; P = projector; R = rail; S = stepper motor.

TABLE 1

Spatial Sampling Granularity For All Datasets² Camera views (dorsal, lateral, and ventral) are indicated in parentheses by D, L, and V respectively, with the camera position (position 1—toward end of arm, position 2—directly overhead) as specified in Figure 3.

Dataset Name	Spatial Quantization, mm		Typical Lens
	Minimum	Maximum	
GCRL26643 (D1)0.2115	0.6479	0.3379	75mm@f/5.6
GCRL26666 (D1)0.1661	0.2579	0.2038	75mm@f/8
GCRL26668 (D1)0.0655	0.1006	0.0799	75mm@f/8
GCRL26669 (D1)0.2326	0.6212	0.3539	75mm@f/5.6
GCRL26670 (D1)0.1526	0.2769	0.1997	75mm@f/8
GCRL26671 (D1)0.1471	0.2789	0.1963	75mm@f/8
GCRL26672 (D1)0.1400	0.2954	0.1946	75mm@f/8
GCRL26674 (D1)0.2068	0.6662	0.3356	75mm@f/5.6
GCRL26676 (D1)0.1432	0.2940	0.1967	75mm@f/8
GCRL26677 (D1)0.2021	0.6709	0.3307	75mm@f/5.6
GCRL26678 (D1)0.1269	0.2102	0.1602	75mm@f/8
GCRL26689 (D1)0.1772	0.7227	0.3105	75mm@f/5.6
GCRL26643 (L1)0.1605	0.2650	0.2025	75mm@f/8
GCRL26666 (L1)0.1524	0.2727	0.1986	75mm@f/8
GCRL26668 (L1)0.0677	0.0968	0.0800	75mm@f/8
GCRL26669 (L1)0.1546	0.2742	0.2010	75mm@f/8
GCRL26670 (L1)0.1374	0.2960	0.1929	75mm@f/8
GCRL26671 (L1)0.1342	0.2994	0.1909	75mm@f/8
GCRL26672 (L1)0.1580	0.2641	0.2001	75mm@f/8
GCRL26674 (L1)0.1595	0.2690	0.2030	75mm@f/8
GCRL26676 (L1)0.1466	0.2812	0.1965	75mm@f/8
GCRL26677 (L1)0.1602	0.2987	0.2122	75mm@f/8
GCRL26678 (L1)0.1456	0.2095	0.1728	75mm@f/8
GCRL26689 (L1)0.1432	0.3391	0.2081	75mm@f/8
GCRL26643 (V2)0.1786	1.9185	0.4154	25mm@f/4
GCRL26666 (V2)0.1171	0.4121	0.1970	75mm@f/8
GCRL26668 (V2)0.0962	0.1786	0.1278	75mm@f/8
GCRL26669 (V2)0.1644	3.2031	0.4286	25mm@f/4
GCRL26670 (V2)0.1471	0.3079	0.2047	75mm@f/8
GCRL26671 (V2)0.1452	0.2976	0.2006	75mm@f/8
GCRL26672 (V2)0.1535	0.2879	0.2049	75mm@f/8
GCRL26674 (V2)0.1786	1.9443	0.4161	25mm@f/4
GCRL26676 (V2)0.1432	0.3171	0.2045	75mm@f/8
GCRL26677 (V)0.2094	0.9245	0.3781	25mm@f/4
GCRL26678 (V2)0.1364	0.3390	0.2035	75mm@f/8
GCRL26689 (V2)0.1953	1.1991	0.3943	25mm@f/4

desired direction and to attain an adequate field of view. The rail was mounted to a ball-screw carriage, with the long axis of the rail perpendicular to the carriage's direction of travel. The ball screw was mounted above the work table, and was driven by a stepper motor that received drive signals from a serial port of the Sun 3/280.

Data Gathering

Operation is quite simple and takes under 4 minutes per scan. The video digitizer grabs a 512×480 frame, the computer processes the resulting image to extract (row, column) positions of illuminated points, and the stepper motor moves the sensor. Movement (translation) is perpendicular to the projected light plane, and the distance




Figure 4. Arbitrary single image of light stripe.

moved per step is the total target length divided by the number of slices. This continues until the desired number of slices have been digitized. Each slice contributes a 480-point column of data to the range map.

For processing a single image, consider the sensor image in of an arbitrary “slice” in Figure 4. Due to orientation of the camera relative to the light plane, the light stripe will be roughly perpendicular to the scan lines in the image. Some rows may have no illuminated points. We describe the shape of the illuminated stripe segment(s) in a frame by a series of numbers, one per row, where zero indicates that no point in that row was illuminated and any number other than zero indicates the column location of the illuminated point. For a 512×480 image, this generates a sequence of 480 numbers in the range 0 through 512 inclusive. In the terminology of machine vision, this is frequently called “offset data”, as it expresses the offset of the light stripe from the left margin of the image.

Offset data is stored in a 2-D array. Each column of the array represents a slice derived from one digitized video frame. The row position within a column is the row position from the associated frame. The value stored at offset $[i,j]$ is thus the offset data from the j 'th row of the i 'th frame. If zero is stored, the point is invalid—there was no measurement made for row j of frame i . If a non-zero offset distance d is stored, then a 3-D position was measured. It has been stored in terms of an image-plane (row, column) position equal to (j,d) .

The data structure just described assumes an ideal image. An actual video image will present a somewhat fuzzy stripe, more than one pixel wide, against a background that is not entirely dark. To deal with real images, we have achieved acceptable results with the following procedure. As shown diagrammatically in Figure 5A, for each scan line we find the brightest pixel. If that pixel is not above some absolute threshold, T_A (which may depend on light source, lens speed, and stand-off distance), then store zero. If that pixel is at least as bright as T_A , then we test pixels at an outlier distance w on either side of the peak. If they are not dimmer than T_R , a percentage of the intensity of the brightest pixel, store zero. As shown in Figure 5B, this avoids artifacts caused where the light plane grazes a nearly coincident surface, keeping only isolated peaks.

To reduce fuzziness caused by vibration, which may result from using long focal length lenses positioned toward the end of the rail, short time delays were added between frames during sensor movement. To limit noise introduced by low light performance of the Pulnix camera, we averaged 4 digitized frames for each slice when using long focal length lenses.




Figure 5. Measured Intensity Values Across Light Stripe. A = Valid peak; T_A = absolute threshold; T_R = relative threshold; W = outlier distance; B = Invalid peak resulting from grazing.

The distance to these outlier pixels and the absolute threshold of the brightest pixel are specific to a particular sensor design. For the sensor used in these experiments, with pixels digitized to take intensity values from zero through 255, we used an absolute threshold of 20, a relative threshold of 0.66 times the peak intensity, and an outlier distance of 10 pixels. If the pixel in question meets the criteria associated with the absolute and relative thresholds, then its column position is stored in the offset data. We also save the intensity of the original point. This means that we have intensity data that is registered to (x,y,z) location—the utility of this is explained in a later section. The intensity is stored in the same order as the offset data. While the offset data encodes (x,y,z) in a somewhat indirect way, it is possible to plot the offset data to produce a composite light-stripe image (Figure 6). This image strongly suggests the 3-D shape of the surface. It provides immediate feedback to the user regarding the quality of the measurements.

We store both offset and intensity as ASCII values, for maximum portability. The data are run-length-encoded, so that stretches of missing data are replaced by their lengths. The file has a header that includes the calibration matrix, sensor use characteristics, date of capture, the name of the user, and a user-specified comment field, which allows us to record information about the specimen. Our comment field included the catalog number of the specimen and other relevant information about the scan.

Calibration

The sensor records provides only (row, column) information. Measurement requires (x,y,z) locations. Sensor calibration provides a mapping from image plane coordinates—(row, column), or (r,c) —into world coordinates— (x,y,z) . It might seem appropriate to calibrate the sensor by trigonometric analysis, as done in equation 1.3 to define $\bar{X}(r,c)$. However, this is not the case, as the optical and geometric characteristics are difficult to measure and small errors have profound effects. The best method is to solve analytically for a calibration matrix, using known (x,y,z) and (r,c) pairs. In other words, image a calibration target with features in known (x,y,z) positions, measure the resulting (r,c) image locations, and solve for the relationship providing the mapping from image coordinates (r,c) to world coordinates (x,y,z) .

Our calibration target was a small panel with 6 pins protruding from known locations. We use the calibration method described in Chen and Kak (1987) to map the image plane onto the light plane. This yields a calibration matrix T , where:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} a \\ d \\ b \\ d \\ c \\ d \end{bmatrix} \quad \text{where} \quad \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = T \begin{bmatrix} r \\ c \\ 1 \end{bmatrix} \quad (2.1)$$




Figure 6. Unrectified, composite light-stripe image (GCRL 26643 in dorsal view).

or:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \frac{T_{11}r + T_{12}c + T_{13}}{T_{41}r + T_{42}c + T_{43}} \\ \frac{T_{21}r + T_{22}c + T_{23}}{T_{41}r + T_{42}c + T_{43}} \\ \frac{T_{31}r + T_{32}c + T_{33}}{T_{41}r + T_{42}c + T_{43}} \end{bmatrix} \quad (2.2)$$

Each pair of image and world coordinates yields 3 equations in twelve unknowns, the elements of T . A set of 4 calibration points would be adequate to solve for T . Our problem is overdetermined, so we form the normal equations as described in Chen and Kak (1987), and find the linear least-squares solution as described in sections 2.0–2.1 of Press et al. (1988).

Given T , we would like to calculate bounds on its spatial quantization. The finest and coarsest sampling will both occur along the margins of the image, as long as the image does not include the point on the light plane that is closest to the camera. The error to consider is a one-pixel error in column position for a given row, as that is the error associated with incorrectly choosing the peak in a row of the image. So, for each border pixel, we calculate the 3-D distance between the (x,y,z) coordinates corresponding to that pixel and its neighbor to the left or right. A typical error is that same measure carried out for a pixel in the region of the frame where illuminated points typically appear. The calculation of this error is critical, as error in peak detection and selection is the source of practically all error in establishing (x,y,z) . Sub-pixel peak detection and selection provides a means of improving spatial quantization. We recalibrated the sensor before every scan of every target to minimize error. The typical and bounding spatial sampling values for all datasets are shown in Table 1. Note that these values provide the bounds on the best possible accuracy for a given configuration. Typical best-case calibration errors, $e(T)$, were in the range 0.05–0.2 mm. If $e(T)$ is within or below the spatial quantization (as it always was in this study), that provides high confidence in the accuracy of the calibration matrix T .

Applying the calibration matrix to the offset data yields a range map \vec{X} , a 2-D array of 3-D vectors. $\vec{X}[i,j]$ is the vector (x,y,z) stored in column i , row j of this array. This corresponds to the (x,y,z) position obtained from the j 'th row of the i 'th digitized frame. We then add in the slice-to-slice translations. We translate the sensor along the y -axis, as the projected light defines the x - z

plane. The translation per step is just enough to cover the long dimension of the target in 256 slices. If each step is s millimeters, then every (x,y,z) vector stored in column i has $(0, s \cdot i, 0)$ added to it. In other words, for every i in the range $[0,256]$ and each j in $[0,480]$, if point $[i,j]$ was valid, we make the assignment:

$$\vec{X}[i, j] = \vec{X}[i, j] + \begin{bmatrix} 0 \\ s \cdot i \\ 0 \end{bmatrix} \quad (2.3)$$

By “valid”, recall that we mean those points for which we extract an offset value and thus are able to calculate (x,y,z) . This is typically about 50–70% of the points per scan.

Self-Occlusion

The only measurable points are those that can be simultaneously illuminated with the laser and viewed with the camera. If the light source and detector were coincident, as is the case with laser radar, there would be no missing data. Of course, for such a sensor the light stripe would be a straight line in every image and thus reveal no 3-D information.

With the camera outside the light plane, as it must be, there is a set of points that can be viewed but not illuminated, and a set that can be illuminated but not viewed. This is what is referred to as self-occlusion. Consider a smooth spherical target. Almost 50% of the surface can be viewed, and almost 50% can be illuminated, but the overlap is perhaps only 30–40% of the total surface area. If the target has protruding structures, those structures may cast 2 “shadows”. One shadowed region is the set of points that cannot be illuminated because an occluding structure blocks the light plane. These illumination-occluded points will lie in the same slice, thus the same range map column, as the occluding surfaces. The other shadowed region is the set of points that cannot be viewed because an occluding structure blocks the camera’s line of sight. These view-occluded points will not necessarily lie in the same slice, but for a “left-eyed” sensor configuration like ours, with sensor translation to the right, they appear to the right of the occluding structure in the range map. Many points near an occluding structure will belong to both occlusion sets.

Processing the Range Map

The above section describes how offset data can be captured and converted to a 2-D array of (x,y,z) vectors.

The next step is the calculation of the local surface normals. A detailed description of the calculation of surface normals and local curvatures from a range map is given by Yang and Kak (1986). Simply stated, \bar{X}_u and \bar{X}_v are 2-dimensional arrays of 3-D vectors, just as \bar{X} is a 2-dimensional array of (x,y,z) locations. $\bar{X}_u[i,j]$ and $\bar{X}_v[i,j]$ are the partial derivatives of the range map \bar{X} with respect to row and column position, respectively and where “*” implies convolution:

$$\bar{X}_u = \bar{X} * \begin{bmatrix} \frac{-1}{12} & \frac{-4}{12} & \frac{-1}{12} \\ 0 & 0 & 0 \\ \frac{1}{12} & \frac{4}{12} & \frac{1}{12} \end{bmatrix} \quad (2.4)$$

$$\bar{X}_v = \bar{X} * \begin{bmatrix} \frac{-1}{12} & 0 & \frac{-1}{12} \\ \frac{-4}{12} & 0 & \frac{4}{12} \\ \frac{1}{12} & 0 & \frac{1}{12} \end{bmatrix} \quad (2.5)$$

The surface normal vector at the location [i,j] is thus:

$$\bar{S}[i,j] = \frac{\bar{X}_u[i,j] \times \bar{X}_v[i,j]}{|\bar{X}_u[i,j] \times \bar{X}_v[i,j]|} \quad (2.6)$$

where, \times is the cross product operator.

This means that \bar{S} would be defined only at points that have 8 valid neighbors. If a point is valid, but lacks a completely valid neighborhood, the surface normals are examined at its 8 neighbors. If a surface normal could be calculated at one or more neighbors, then a neighboring surface normal vector is assigned. If there is a choice of more than one neighboring surface normal vector, then we select the one for which the point in question best fits the plane of the corresponding 3×3 neighborhood.

Image Rectification

If we simply display the range map, coding intensity or hue for 3-D vectors (as discussed in detail below), then depending on the sensor field of view and the sensor translation per slice, we may or may not have an image that appears correct when displayed with square pixels. The (x,y,z) vectors will be correct, but the row and column dimensions of the image may be such that it has the appearance of being compressed or stretched. This can be corrected by resampling the range map through a step we refer to as “rectification”.

The range map is first transformed so that the average surface normal vector points in the direction (0,0,1). Since the range map is filled with one column per measured slice, points lying along one column are at a constant value of x. Therefore, the derivative dx/di describes the step per slice, and dx/dj is zero (where i,j are indices within the range map corresponding to the ith row in the jth original frame). After the transformation, dy/di is essentially zero, dy/dj is essentially some




Figure 7. Reconstruction of Measured Target. A = Registered intensity data; B = Z-value coded as intensity; C = Rendered image D = Encoding of position and orientation.

non-zero constant, and both dz/di and dz/dj vary with the 3-D shape of the measured surfaces. We resample the range map to make dx/dj and dy/di roughly equal, which ensures that images displayed with square pixels present the expected aspect ratio. The surface normal map is then recalculated for the resampled \vec{X} .

Presentation and Visualization

Given a range map and a map of surface normal vectors, we can produce a rendered image. This rendered image displays the intensities we would expect if the 3-D surface were lit by a point light source, ignoring shadow effects. Given a location of a point on a surface, the expected intensity is a function of the angles between 3 vectors: 1) the local surface normal vector, 2) the vector to the hypothesized viewpoint, and 3) the vector to the hypothesized light source. The precise function used is defined by the lighting model (see Foley and van Dam 1984, Ballard and Brown 1982, Marr 1982 for specific algorithms). We used a combination of Marr and Minnaert lighting models, positioning the light source location to illuminate the surface from above our left shoulder. Note that this is rendering only, not raytracing, which is far more computationally expensive. Rendering alone presents adequately realistic imagery for landmark selection.

We now produce 4 images. The first, (Figure 7A) is the registered intensity data, and it resembles a photograph of the target. The second (Figure 7B) encodes z as intensity. The third (Figure 7C) is the rendered image. The fourth (Figure 7D) encodes both position and orientation. We can use hue on a computer monitor to encode z , with blue furthest from the sensor and red closest (Cromwell, 1997). The intensity is that of the rendered image. The combination of these 4 images guarantees that any measured feature will be visible in at least one of them. Suture lines, for instance, are too small to show up in the range map but they frequently appear in the registered intensity image. Many uniform regions within the intensity image have complex 3-D shapes revealed by other images.

Data Extraction

The above steps produce a 2×2 panel of 4 images. This panel is displayed on a workstation screen, and the user selects points. We have developed a simple interface in the X Window System (TM) that allows the user to accurately select points for measurement using a mouse. The corresponding (x,y,z) location follows immediately from the $[i,j]$ index of the selected image point. The user

can extract positions and orientations of single points, distances between selected points, or cross sections along arbitrary cutting planes. The selection can be done on any of the 4 images presented, so a pair of points does not have to be visible in a single image.

Registration, Range-map Merging, and Raytracing

Thus far, we have limited discussion to a single translational scan of the target. In addition to the small regions of self-occlusion described earlier, this sensing mode restricts the range map to less than 50% of the target. An obvious next step is to integrate range maps gathered from a variety of directions to fully image the object in 3-D. To do this we must solve several problems.

First consider the problem of merging 2 range maps taken from views perhaps 90° apart. Let us refer to these range maps as \vec{X}_1 and \vec{X}_2 . We have thus far used rows and columns of such range maps to define an appropriate grid for presentation. In merging range maps it is convenient to use the points as vertices in a mesh, representing the surface as one or more contiguous regions of triangular facets that tessellate the surface. If the target were roughly spherical, we would expect about half of each range map to cover a common area. Thus, half of \vec{X}_1 would form a mesh roughly coincident with the mesh formed by \vec{X}_2 . We say "roughly coincident" because spatial quantization and measurement error will prevent exact coincidence. Even in the absence of measurement error, \vec{X}_1 and \vec{X}_2 will, in all probability, sample the 3-D spaces slightly differently with individual points not falling exactly at the same locations.

Consider the union of these range maps, which forms a more complete model of the object and can be expressed as,

$$\vec{X}_u = \vec{X}_1 \cup \vec{X}_2 \quad (2.7)$$

with \vec{X}_1 and \vec{X}_2 transformed so that their coordinate frames are coincident. \vec{X}_u could be formed in 3 ways: 1) by moving the target mechanically between scans, with sufficient precision to provide the reverse translation; 2) by a human expert picking landmarks visible in both range maps to derive the needed transformation; 3) by a procedure in which corresponding landmarks in the 2 range maps are identified automatically, without human intervention. Presently, only the first 2 approaches are practical and we have used both. After the union, \vec{X}_u now describes roughly 75% of the object. However, approximately one-half of \vec{X}_u is redundant.

A second major problem is that we can no longer display the range map using the indices as a grid. Instead, we must resort to raytracing to create a realistic image. Since our measured points are described by several sets of points gathered from rectangularly sampled grids, for each range map we form a triangular mesh. We thus have a (rather long) list of triangular facets described by the (x,y,z) positions of their vertices. To perform raytracing we follow techniques described in detail by Glassner (1989). For each pixel in the image, raytracing determines which, if any, facet will appear and how it will be rendered. If the 2 range maps are of high quality and if the registration between the maps is accurate, the region they share will appear as a finely-mixed combination of the 2 meshes. However, it is possible for many facets of one range map to be completely obscured by facets of the other. In such cases, there may be facets that are completely inside the surface, so that they would not appear from any exterior viewpoint. Although it would be difficult to establish a list of such facets analytically, it would be relatively easy to generate a number of raytraced views from canonical positions and then search the union range map for facets that appear in none of these views. Figure 8 shows the results of multi-scan registration and merging using 3 views (dorsal, lateral, and ventral).

For many applications it may be reasonable to delete such facets from later consideration, given the com-

putational expense of raytracing. If they were visible only from very limited regions of space, their deletion would have little effect on overall accuracy or completeness. Nonetheless, more work needs to be done to facilitate landmark selection and measurement from raytraced imagery.

Computer Hardware Used

As previously mentioned, a Sun 3/280 was used to drive the scanner and digitize the images. The sensor used was a general-purpose machine vision tool, and was not specially designed for biological applications. In fact, it was designed for far coarser sampling of larger targets. The use of the Sun was due to the arrangement of available hardware rather than to any specialization. Since then, we have implemented the system on a PC platform running the Linux operating system. Readers interested in constructing such a system may contact the senior author.

Data analysis, from offset data to presentation images, was done on both Sun (4/390S and Sparc 1000) and PC (K6-333 MHz CPU, 32 MB RAM) platforms. On both Sparc platforms, a 256×480 range map can be processed and presented in 40 to 50 CPU seconds. On the PC platform, range map processing and presentation takes only 10 CPU seconds.




Figure 8. Composite reconstruction produced through the union of 3 “canonical” views (dorsal, lateral, and ventral) of GCRL 26666. Note the reduction in self-occlusion as compared to Figure 7.




Figure 9. Selected landmark points and distances (in white) used to compare sensor-derived measures with those taken using caliper and projected-video imaging. A = Lateral view; B = Ventral view. 1) anteriormost median tip of median ethmoid, 2) posteriormost ventral margin of basioccipital, 3) lateralmost extension of median ethmoid on left side [4 on right side], 5) lateralmost point on lateral extension lateral ethmoid immediately anterior to palatine socket, [6 on right], latero-anterior most point on lateral extension of lateral ethmoid immediately posterior to palatine socket [8 on right side], 9) lateralmost point of ventroposterior margin of lateral ethmoid at anterior rim of orbit [10 on right side], 11) tip of left preorbital spine [12 on right side], 13) tip of supraorbital spine [14 on right side], 15) anterolateral tip of left dermosphenotic [16 on right side], 17) posteriormost tip of left posttemporal spine [18 on right side], 19) tip of left posttemporal spine, 20) tip of left nuchal spine, 21) tip of left pterotic spine, 22) tip of anteriorprojection of left frontal. Landmark points 19–2 not consistently visible or damaged in some specimens and not used in computations shown in Figure 10.

Specimen Preparation and Caliper and 2-D Measurement

Specimens of the scorpaenid fish *Neomerinthe hemingway* were taken near the mouth of the Mississippi River and frozen for subsequent skeletonization. Catalog numbers for these specimens are provided in Table 1. This species was selected because of its relatively large size and availability. Partially frozen specimens were autoclaved for 6 to 12 minutes at 123°C, depending on

size, to facilitate removal of muscle and other tissues without disarticulating the skulls. Posttemporal bones of some specimens were subsequently reattached. To minimize reflections that can result in beam deflection and image loss and to insure that adequate light is more uniformly returned off more translucent regions, all skulls were spray painted with a flat white enamel prior to measurement of any kind.




Figure 10. Comparison of methods of measurement of inter-landmark distances between landmarks seen in Figure 9 (shown as lines). Closed circles = correlation between range-sensor and caliper-based measurements; Open circles = correlation between projected video and caliper based measures. Open circles denoted by A are measures of distances between points 2 and 15 and 2 and 16; Open circles denoted by B are measures of distances between points 9 and 11 and 10 and 12. Note divergence due to foreshortening in these measurements. Circles may represent more than one measurement because of overlap.

Upon skeletonization, caliper based measures were taken by one of us (SGP) to the nearest 0.1 mm using a set of Fowler digital calipers connected to microcomputer using the convenient and flexible DataQ™ interface software. Landmark points (Figure 9) were selected (again by SGP) to insure a reasonable sampling of the 3-dimensional conformation of the skulls and to provide data on dimensions important to comparison with other scorpionfishes. Twenty-one measured distances among these landmark points were used to compare measuring techniques. The skulls were then subsequently measured for the same landmark points using the MorphoSys video digitization (Meacham 1993) using an Olympus 35 mm lens at the maximum possible working distances (50 to 70 cm) to minimize parallax, while providing a sufficiently large and clear image. Because some of the chosen measurements can not be visualized in a single plane, measurements were obtained from lateral, dorsal, and ventral projections.

Comparison with Caliper and Projected Image Video Measurement

Data extracted from 3-D sensor based measurements and data taken for the same inter-landmark points using dial calipers were very nearly identical, with $r = 0.9995$

and $P \ll 0.001$ (Figure 10). The corresponding correlation for inter-landmark distances measured by sensor and projected video-imaging is somewhat lower ($r = 0.9934$ and varies depending primarily upon the orientation used for a given measurement). Four measures that, when taken in ventral view exhibit notable foreshortening are primarily responsible for the difference between projected video-based measurement and either caliper and sensor based methods (distance between points 2 and 15, 2-16, 9-11 and 10-12). The first 2 distances are relatively distant, but extend simultaneously from the broadest dorso-lateral part of the skull (dermosphenotic) to the postero-ventral margin of the medial skull axis (basisphenoid). The second pair of measurements (posterior margin of lateral ethmoid process to tip of preocular spine) appear proximate in ventral view, but extend strongly dorso-ventrally, when viewed laterally.

DISCUSSION

Although our system was not specifically designed to measure biological objects, our results indicate it produces abundant data within the range commonly used in comparative morphometrics. Data storage costs increase roughly linearly with an increase in the number of pixels measured. However, this is not prohibitive as compression can result in considerable savings. Our offset data files ranged from 69-158KB compressed and about 3.2 MB uncompressed. Processing time is also roughly linearly related to the number of points in the range map. Increases in the number of optical slices per scan, as well as increases in pixel density per digitized video frame, could be accommodated without exorbitant computational expense. For a system costing only about \$2,000, excluding the workstation and labor costs, we believe our results demonstrate the practicality of single-plane structured light measurement of osteological materials.

At sharp edges, finite beam width and optical bloom limit achievable accuracy. We have coated our targets with a matted surface, carefully selected camera optics, and used appropriate lens speed to minimize ambiguity and loss of resolution. For this reason, we limited our inquiry to relatively large, non-specular objects that are relatively hard, opaque, and dry. Nonetheless, significant scope for improvement at relatively low cost exists and additional experimentation will likely improve accuracy and the range of objects that may be examined. Since these experiments were conducted, we have built a more capable sensor capable of gathering (x,y,z) data

with a precision of 0.2–0.05 mm at about the same cost as the original sensor.

For measures of the kind described here, error due to perspective effects can be significant when measuring from projected 2-D imagery. Although MorphoSys was designed to measure plant leaf shapes, use of projected measurement of more 3-dimensional objects requires greater circumspection. Studies utilizing such measures to analyze shape must account for potential errors arising from perspective effects. Our results indicate that use of a single-plane structured-light range sensor can essentially eliminate errors due to perspective effects.

In our experiments we used only 21 of the more than 75,000 available points with estimated (x,y,z) coordinates per scan. Although these are sufficient to demonstrate the relative accuracy and considerable potential of this approach, we believe that future research will be usefully focused on developing automated methods capable of analyzing more numerous and more informative elements of the data set, thereby better characterizing subtle shape differences among objects. Such studies may lead to more objective methods of establishing landmarks among a neighborhood of potentially corresponding points. Such studies may also lead to more precise understanding of mechanisms, selection, and constraints that lead to and control shape differences. They may also aid in developing a better understanding how humans recognize taxonomically useful characters. However used, one of the greatest strengths of single-plane structured light sensing methods is that 3-D data from the rest of the scanned object remains readily available for display and a variety of additional new measurements become possible without subsequent rescanning of the original object.

Like 3-D data obtained from tomographic methods using other energy sources, these data can be archived and distributed electronically to provide relatively complete representations and to assure the same bases of comparison. See <http://rv14.ecn.purdue.edu/~cromwell/3d.html> and <http://lionfish.ims.usm.edu/~musweb/lasermesure.html>. Although limited to surface features, single-plane structured light imaging can secure data at relatively limited cost, unlike CT, PET, or NMRI imaging. By facilitating quantitative characterization and thus common understanding among scientists with a diverse range of perspectives, objects of common interest may be more accurately studied.

ACKNOWLEDGMENTS

The authors wish to express thanks to the following individuals for their assistance in aspects of this project: A.C. Kak, Director of Robot Vision Laboratory, Department of Electrical and Computer Engineering, Purdue University, who graciously made facilities of his Laboratory available to us; B. Rohr of the National Marine Fisheries Service, Pascagoula Laboratory, who kindly collected many of the specimens used in this study; C. Aadland, formerly of the Gulf Coast Research Laboratory (GCRL), who collected specimens and greatly assisted in preparing the skeletons, as did S.Y. Poss; D. Rebarchik also of GCRL, who provided us the use of the autoclave. We thank S. Blask, F. Bookstein, S. Hutchinson, and C. Meacham, for comments on the manuscript. This project was funded in part by NSF BSR 8705373 awarded to SGP.

LITERATURE CITED

- Agin, G. and T. Binford. 1976. Computer description of curved objects. *Institute of Electrical and Electronics Engineers Transactions. Computers* 25:439–449.
- Ballard, D.H. and C.M. Brown. 1982. *Computer Vision*. Prentice Hall, Englewood, NJ. 523 p.
- Baker, H.H. and T.O. Binford. 1981. Depth from edge and intensity based stereo. *Institute of Electrical and Electronics Engineers Trans. Pattern Analysis and Machine Intelligence* 2:333–340.
- Barnard, S.T. and M.A. Fischler. 1982. Computational stereo. *Computing Surveys of the Association for Computing Machinery* 14:553–572.
- Birkmann, H. and R.F. Lundin. 1996. Confocal microscopy: potential applications in micropaleontology. *Journal of Paleontology* 70:1084–1087.
- Blais, F., M. Rioux, J. Domey, and J.A. Beraldin. 1988. A very compact real-time 3-D range sensor for mobile robot Applications. *Proceedings of Society of Photo-optical and Instrumentation Engineers, Mobile Robots III* 1007:330–338.
- Bookstein, F.L., B. Grayson, C.B. Cutting, H.C. Kim, and J. G. McCarthy. 1991. Landmarks in three dimensions: reconstructions from cephalograms versus direct observation. *American Journal of Orthodontics and Dentofacial Orthopedics* 100:133–140.
- Boyer, K.L. and A.C. Kak. 1987. Color-encoded structured Light for rapid active ranging. *Institute of Electrical and Electronics Engineers Transactions on Pattern Analysis and Machine Intelligence* 9:14–28.
- Boyer, K.L. and A.C. Kak. 1988. Structural stereopsis for 3D vision. *Institute of Electrical and Electronics Engineers Transactions On Pattern Analysis and Machine Intelligence* 10:144–166.
- Brooks, M.J. 1991. The ontogeny of sexual dimorphism: quantitative models and a case study in labrisomid blennies (Teleostei: Paraclinus). *Systematic Zoology* 40:271–283.

- Caulfield, H.J., T. Hirschfeld, J.M. Weinberg, and R.E. Herron. 1977. Laser Stereometry. Proceedings Institute of Electrical and Electronics Engineers 65:84–88.
- Chen, C.H. and A.C. Kak. 1987. Modeling and Calibration of a Structured Light Scanner for 3-D Robot Vision. Proceedings Institute of Electrical and Electronics Engineers, International Conference on Robotics and Automation. p. 807–815.
- Chen, C.H. and A.C. Kak. 1989. A robot vision system for recognizing 3-D objects in low-order polynomial time. Institute of Electrical and Electronics Engineers Transactions on System, Man, and Cybernetics 19:1535–1563.
- Conroy, G.C. and M.W. Vannier. 1984. Noninvasive three-dimensional computer imaging of matrix-filled fossil skulls by high-resolution computed tomography. Science 226:456–458.
- Corner, B.D., S. Lele, and J.T. Richtsmeier. 1992. Measuring precision of three-dimensional landmark data. Quantitative Anthropology 3:347–359.
- Cromwell, R.L. 1992. Tyro - a robot vision system that can learn from observations of its environment. Unpublished Ph.D. Thesis, Purdue University. 402 p.
- Cromwell, R.L. 1993. Sensors and Processors Enable Robots to See and Understand. Laser Focus World, March 1993, Pennwell Publications, Nashua, NH, p. 67–78
- Cromwell, R.L. 1997. Laser scanner digitizes fish skeletons. Vision System Design 2:9–10.
- Currents, K.P., C.S. Sharpe, R. Hjort, C.B. Schreck, and H.W. Li. 1989. Effects of different feeding regimes on the morphometrics of chinook salmon (*Oncorhynchus tshawytscha*) and rainbow trout (*O. mykiss*). Copeia 1989:689–695.
- Douglas, M.E. 1993. Analysis of sexual dimorphism in an endangered cyprinid fish (*Gila cypha* Miller) using video image technology. Copeia 1993:334–343.
- Ehlinger, T.J. 1991. Allometry and analysis of morphometric variation in the bluegill, *Lepomis macrochirus*. Copeia 1991:347–357.
- Fink, W.A. 1987. Video digitizer: a system for systematic biologists. Curator 30:63–72.
- Fink, W.A. 1993. Revision of the piranha genus *Pygocentrus* (Teleostei, Characiformes). Copeia 1993:665–687.
- Foley, J.D. and A. Van Damm. 1984. Fundamentals of Interactive Computer Graphics. Addison-Wesley, Reading, MA, 664 p.
- Foster, B., E. Williams, R.J. Witmer, and K.M. Piel. 1990. Confocal microscopy: a new technique for imaging microorganisms and morphology in three-dimensions. Palynology (Abstracts) 14:212.
- Glassner, A.S. 1989. An Introduction to Ray Tracing. A.S. Glassner, ed. Academic Press, London, 225 p.
- Grewe, L. and A.C. Kak. 1995. Interactive learning of a multi-attribute hash table classifier for fast object recognition. Computer Vision and Image Understanding 61:387–416.
- Grimson, W.E.L. 1981. From Images to Surfaces: A Computational Study of the Human Early Visual System. MIT Press, Cambridge, MA, 274 p.
- Horn, B.K.P. 1975. Obtaining Shape from Shading Information. In: P.H. Winston, ed. The Psychology of Computer Vision, McGraw Hill, New York, NY, p. 115–155.
- Hu, G. and G. Stockman. 1989. 3-D surface solution using structured light and constraint propagation. Institute of Electrical and Electronics Engineers Transactions Pattern Analysis and Machine Intelligence 11:390–402.
- Hutchinson, S.A., R.L. Cromwell, and A.C. Kak. 1988. Planning sensing strategies in a robotwork cell with multi-sensor capabilities. Proceedings, 1988 Institute of Electrical and Electronics Engineers International Conference on Robotics and Automation, Philadelphia, PA. April 24–29, 1988, p. 1068–1075.
- Hutchinson, S.A., R.L. Cromwell, and A.C. Kak. 1989. Applying uncertainty reasoning to model based object recognition. Proceedings, Institute of Electrical and Electronics Engineers Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, CA, June 4–8, 1989, p. 541–548.
- Hutchinson, S.A. and A.C. Kak. 1989. Applying uncertainty reasoning to planning sensing strategies in a robot work cell with multi-sensor capabilities. Proceedings, Institute of Electrical and Electronics Engineers Symposium on Intelligent Control, Albany NY, September 25–26, 1989, p. 129–134.
- Hutchinson, S.A. and A.C. Kak. 1989. Planning sensing strategies in a robot work cell with multiple-sensor capabilities. Institute of Electrical and Electronics Engineers Transactions on Robotics and Automation 5:765–783.
- Inokuchi, S., K. Sato, and F. Matsuda. 1984. Range-imaging for 3-D object recognition. Proceedings, International Conference on Pattern Recognition, Montreal, PQ, Canada, November 14–17, 1988, p. 806–808.
- Julesz, B. and J.E. Miller. 1975. Independent spatial-frequency-tuned channels in binocular fusion and rivalry. Perception 4:125–143.
- Kaiser, H., J.D. Hardy, and D.M. Green. 1994. Taxonomic status of Caribbean and South American frogs currently ascribed to *Eleutherodactylus urichi* (Anura: Leptodactylidae). Copeia 1994:780–796.
- Kak, A.C., A.J. Vayda, R.L. Cromwell, W.Y. Kim, and C.H. Chen. 1987. Knowledge-based robotics. Proceedings, 1987 Institute of Electrical and Electronics Engineers International Conference on Robotics and Automation, Raleigh, NC, March 31–April 3, 1987, p. 637–646.
- Kak, A.C., A.J. Vayda, R.L. Cromwell, W.Y. Kim, and C.H. Chen. 1988. Knowledge-based robotics. International Journal of Production Research 26:707–734.
- Kaufman, L. 1964. On the nature of binocular disparity. American Journal of Psychology 77:393–402.
- Kalvin, A.D., D. Dean, and J.J. Hublin. 1995. Reconstruction of human fossils. Institute of Electrical and Electronics Engineers Computer Graphics and Applications 15:12–15.
- Kim, W.Y. and A.C. Kak. 1991. 3-D object recognition using bipartite matching embedded in discrete relaxation. Institute of Electrical and Electronics Engineers Transactions On Pattern Analysis and Machine Intelligence 13:224–251.
- Klein, H.A. 1974. The Science of Measurement. A Historical Survey. Dover Publications, New York, NY, 736 p.

- LaValle, S.M. and S.A. Hutchinson. 1991. Considering multiple surface hypotheses in a Bayesian hierarchy. Proceedings, SPIE Conference on Stochastic and Neural Methods in Signal Processing, Image Processing, and Computer Vision, San Diego, CA, July 24–26, 1991, p. 1–15.
- Lewis, R.A. and A.R. Johnston. 1977. A scanning laser rangefinder for a robotic vehicle. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, IJCAI Inc., Cambridge, MA, August 22–25, 1977, p. 762–768.
- Lindberg, D. 1990. Morphometrics and the systematics of marine plant limpets (Mollusca: Patellogastropoda). In: F.J. Rohlf and F.L. Bookstein, eds. Proceedings of the Michigan Morphometrics Workshop, University of Michigan Museum of Zoology Special Publication No.2, Ann Arbor, MI, May 16–28, 1988, p. 301–310
- Marr, D. 1982. Vision. A Computational Investigation into Human Representation and Processing of Visual Information. W. H. Freeman and Co., New York, NY, 397 p.
- Marr, D. and T. Poggio. 1980. Theory of human stereopsis. Proceedings of the Royal Society of London 207:187–217.
- Mayhew, J.E.W. and J.P. Frisby. 1976. Rivalrous texture stereograms. Nature 264:53–56.
- Mayhew, J.E.W. and J.P. Frisby. 1981. Psychological and computational studies towards a theory of human stereopsis. Artificial Intelligence 17:349–385.
- McElroy, D.M. and M.E. Douglas. 1995. Patterns of morphological variation among endangered populations of *Gila robusta* and *Gila cypha* (Teleostei: Cyprinidae) in the Upper Colorado River Basin. Copeia 1995:636–649.
- McElroy, D. M. and I. Kornfield. 1993. Novel jaw morphology in hybrids between *Pseudotropheus zebra* and *Labeotropheus fuelleborni* (Teleostei: Cichlidae) from Lake Malawai, Africa. Copeia 1993:933–945.
- McLeod, A. 1991. Illumination Techniques for 3-D Machine Vision. In: L.F. Pau and R. Olafsson, eds. Quality Control by Computer Vision, Marcel Dekker, New York, NY, p. 151–182.
- Meacham, C.A. 1993. MorphoSys: An interactive machine vision program for acquisition of morphometric data. In: R. Fortuner, ed. Advances in Computer Methods for Systematic Biology. Artificial Intelligence, Databases, Computer Vision, Johns Hopkins University Press, Baltimore, MD, September 9–14, 1990, p. 393–402.
- Newton, G. and B. Kendrick. 1993. Image processing in fungal taxonomy and identification. In: R. Fortuner, ed. Advances in Computer Methods for Systematic Biology. Artificial Intelligence, Databases, Computer Vision, Johns Hopkins Univ. Press, Baltimore, MD, p. 403–412.
- Nitzan, D., A.E. Brian, and R.O. Duda. 1977. The measurement and use of registered reflectance and range data in scene analysis. Proceedings Institute of Electrical and Electronics Engineers 65:206–220.
- Posdamer, J. and M. Altschuler. 1982. Surface measurement by space-encoded projected beam systems. Computer Graphics and Image Processing 18:1–17.
- Press, W.H., B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. 1988. Numerical Recipes in C. The Art of Scientific Computing. Cambridge University Press, Cambridge, MA, 735 p.
- Ray, T. 1990. Application of eigenshape analysis to second order leaf shape ontogeny in *Syngonium podophyllum* (Araceae). In: F.J. Rohlf and F.L. Bookstein, eds. Proceedings of the Michigan Morphometrics Workshop, University Michigan Museum Zoology, Special Publication No. 2, Ann Arbor, MI, May 16–28, 1988, p. 201–213.
- Reilly, S. 1990. Comparative Ontogeny of Cranial Shape in Salamanders Using Resistant fit Theta Rho Analysis. In: F.J. Rohlf and F.L. Bookstein, eds. Proceedings of the Michigan Morphometrics Workshop, University Michigan Museum Zoology Special Publication No. 2, Ann Arbor, MI, May 16–28, 1988, p. 311–321.
- Richtsmeier, J.T., B.D. Corner, H.M. Gausz, J.M. Cheverud, and S.E. Danahey. 1993. The role of postnatal growth pattern in the production of facial morphology. Systematic Biology 42:307–330.
- Riggs, A., L. Tomko, K. Wesolowicz, and C. Jacobus. 1986. Design of a laser radar range imaging sensor for postal applications. In: Proceedings of the Second United States Postal Service Advanced Technology Conference, Washington, DC, October 21–23, 1986, p. 240–264.
- Rioux, M. 1986. Laser Range-Finder Based on Synchronized Scanners. In: A. Pugh, ed. Robot Sensors, Vol. 1 Vision. IFS Ltd., London, p. 175–190.
- Robinson, B.W. and D.S. Wilson. 1995. Experimentally induced morphological diversity in Trinidadian guppies (*Poecilia reticulata*). Copeia 1995:294–305.
- Sanfilippo, A. and W.R. Riedel. 1990. Homology, pattern, and morphometric framework in cyrtid radiolarians. In: F.J. Rohlf and F.L. Bookstein, eds. Proceedings of the Michigan Morphometrics Workshop, University Michigan Museum Zoology Special Publication No. 2, Ann Arbor, MI, May 16–28, 1988, p. 345–362.
- Sato, Y., H. Kitagawa, and H. Fujita. 1982. Shape measurement of curved objects using multiple slit-ray projections. Institute of Electrical and Electronics Engineers Transactions. Pattern Analysis and Machine Intelligence 4:641–646.
- Schultze, H.P. 1991a. Computed tomography and magnetic resonance imaging studies of *Latimeria chalumnae*. Environmental Biology of Fishes 32:159–181.
- Schultze, H.P. 1991b. CT scan reconstruction of the palate region of *Latimeria chalumnae*. Environmental Biology of Fishes 32:183–192.
- Shirai, Y. 1972. A context sensitive line finder for recognition of polyhedra. Artificial Intelligence 4:243–250.
- Wang, Y.F., A. Mitiche, and J.K. Aggarwal. 1987. Computation of surface orientation and structure of objects using grid coding. Institute of Electrical and Electronics Engineers Transactions Pattern Analysis and Machine Intelligence 9:129–137.
- Wang, S., R. Cromwell, A.C. Kak, I. Kimura, and M. Osada. 1994. Model-based vision for robotic manipulation of twisted tubular parts: using affine transformations and heuristic search. Proceedings, 1994 Institute of Electrical and Electronics Engineers International Conference on Robotics and Automation, San Diego, CA, May 8–13, 1994, p. 208–215.
- White, R.J. and H.C. Prentice. 1988. Comparison of shape description methods for biological outlines. In: H.H. Bock, ed. Classification and Related Methods of Data Analysis, Amsterdam: Elsevier, p. 395–402.

3-D MORPHOLOGICAL MEASUREMENT

- White, R.J., H.C. Prentice, and T. Verwijst. 1988. Automated image acquisition and morphometric description. *Canadian Journal Botany* 66:450–459.
- Wimberger, P.H. 1993. Effects of vitamin C deficiency on body shape and skull osteology in *Geophagus brasiliensis*: implications for interpretation of morphological plasticity. *Copeia* 1993:343–351.
- Wise, M.N. 1994. *The Values of Precision*. Princeton Workshop in the History of Science, Princeton University Press, Princeton, NJ, 372 p.
- Yakimovsky, Y. and R. Cunningham. 1978. A system for extracting three-dimensional measurements from a stereo pair of TV cameras. *Computer Graphics and Image Processing* 7:195–210.
- Yang, H.S. and A.C. Kak. 1986. Determination of the identity position, and orientation of the topmost object in a pile. *Computer Vision, Graphics, and Image Processing* 36:229–255.
- Yeung, K.K. and P.D. Lawrence. 1986. A low cost three-dimensional vision system using space-encoded spot projections. *SPIE Optics, Illumination, and Image Sensing for Machine Vision*, Vol. 763, Cambridge, MA, October 30–31, 1986, p.160–172.
- Zangrel, R. and H.P. Schultze. 1989. X-radiographic techniques and applications. In: R.M.Feldmann, R.E. Chapman and J.T. Hannibal, eds. *Paleotechniques*, Paleontological Society, Special Studies, Vol. 4, Knoxville, TN, p. 165–178.
- Zelditch, M.L. and W.L. Fink. 1995. Developmental integration and complexity of body growth in a piranha, *Pygocentrus nattereri* (Teleostei: Ostariophysi). *Journal of Morphology* 223:341–355.
- Zelditch, M.L., W.L. Fink, and D.L. Swiderski. 1995. Morphometrics, homology, and phylogenetics: quantified characters as synapomorphies. *Systematic Biology* 44:179–189.

