

Digital Hammurabi: Design and Development of a 3D Scanner for Cuneiform Tablets

Daniel V. Hahn^{*a}, Donald D. Duncan^a, Kevin C. Baldwin^a,
Jonathon D. Cohen^b, and Budirijanto Purnomo^b

^aJohns Hopkins University Applied Physics Lab, 11100 Johns Hopkins Road, Laurel, MD 20723

^bJohns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218

ABSTRACT

Cuneiform is an ancient form of writing in which wooden reeds were used to impress shapes upon moist clay tablets. Upon drying, the tablets preserved the written script with remarkable accuracy and durability. There are currently hundreds of thousands of cuneiform tablets spread throughout the world in both museums and private collections. The global scale of these artifacts presents several problems for scholars who wish to study them. It may be difficult or impossible to obtain access to a given collection. In addition, photographic records of the tablets many times prove to be inadequate for proper examination. Photographs lack the ability to alter the lighting conditions and view direction. As a solution to these problems, we describe a 3D scanner capable of acquiring the shape, color, and reflectance of a tablet as a complete 3D object. This data set could then be stored in an online library and manipulated by suitable rendering software that would allow a user to specify any view direction and lighting condition. The scanner utilizes a camera and telecentric lens to acquire images of the tablet under varying controlled illumination conditions. Image data are processed using photometric stereo and structured light techniques to determine the tablet shape; color information is reconstructed from primary color monochrome image data. The scanned surface is sampled at 26.8 μm lateral spacing and the height information is calculated on a much smaller scale. Scans of adjacent tablet sides are registered together to form a 3D surface model.

Keywords: cuneiform, 3D scanner, photometric stereo, structured light, shape reconstruction

1. INTRODUCTION

Cuneiform tablets vary from the size of a human torso to the size of a quarter. Scholars estimate that some characters, even in well preserved tablets, contain features as small as 50 μm ; this imposes a rather stringent resolution requirement on the cuneiform scanner. A small sample tablet is shown in Fig. 1. This tablet is deemed to be in excellent condition and is used as the test specimen for scanner development. Hereafter, “the tablet” refers to this test specimen.

Several technologies were evaluated as potential scanning solutions; among them were a tri-color laser scanner¹, a laser-line scanner, and conoscopic holography². Each of these technologies relies on laser technology as the illumination source and each has related problems. The tri-color laser scanner and laser line scanner have an inherent trade-off between lateral resolution and depth of field. To achieve the depth of field necessary to scan the entire tablet face in a single pass, the lateral resolution and height accuracy fall below acceptable levels. The conoscopic technique fell short because of its sensitivity to multiple surface reflections of the laser light; V-shaped grooves appeared W-shaped. Because of these problems, laser technologies were abandoned and an imaging system developed that utilizes incoherent illumination and digital camera technology.

2. PROTOTYPE SCANNER

A diagram of the prototype scanner is shown in Fig. 2. The scanner uses fixed camera (Lumenera Lu120) and tablet positions to maintain image registration while varying the lighting conditions by rotating the source about the tablet and by projecting different illumination colors and patterns (the polar angle of the source is fixed). This technique maintains

* daniel.hahn@jhuapl.edu; phone 443.778.7404; fax 443.778.6779

exact registration between the tablet and camera. An InFocus LP120 digital projector was selected as the light source as it provides excellent illumination uniformity, can easily project custom patterns, and, through the use of an additional lens, provides adequate collimation. The lens placed in front of the projector is chosen to approximate a telecentric configuration when used in combination with the output lens of the projector. An Edmund Optics 0.25X telecentric lens is attached to the camera to magnify the field of view of each pixel by a factor of four ($26.8 \mu\text{m}$ by $26.8 \mu\text{m}$). Sharp image focus is obtained by attaching an ND filter to the telecentric lens so that the iris of the lens remains open. A V-block mounted on top of an elevation stage is used to position the tablet within the focal range of the telecentric lens. Color information is obtained by illuminating the tablet with solid primary colors over various azimuthal angles. Shape information is obtained in two ways, the method of photometric stereo³ and a form of structured light analysis (of which there are many^{4,5}). Each image is pre-processed upon data acquisition to correct for camera noise and non-linearity.



Figure 1: Test specimen.




Figure 2: Prototype scanner design.

The method of photometric stereo is used to obtain a surface normal map of the tablet. This is accomplished by acquiring images over various azimuthal angles under collimated white illumination. The brightness of each pixel in each image is dependent upon the illumination, view, and normal directions as well as the bi-directional reflectance distribution function (BRDF) of the surface. Given the data and known illumination and view directions, the normal map and reflectance are estimated. Normal data resulting from photometric stereo analysis can be integrated over small areas to obtain good estimates of the surface height. Unfortunately, the normal map does not form a conservative surface and small errors accumulate when integration is attempted over larger areas; the data are locally accurate but suffer larger scale inaccuracies.

The particular structured light technique implemented projects a series of 1D sinusoidal patterns onto the tablet at a fixed polar and various azimuthal angles. Four patterns, each out of phase with one another by 90° , are projected for each of a series of iteratively doubled frequencies starting with only one quadrant of a sine wave over the entire projector array and ending with 128 periods. The finest resolution projects each sinusoidal cycle over a lateral distance of approximately 0.6 mm. Each of the images of different phase for a single frequency is used to determine an absolute phase that is unaffected by variations in surface reflectance. This processing is performed via the Carré technique of phase-measurement interferometry⁶. The resulting images are compared to images of a flat white background to calculate the phase difference and corresponding relative tablet height. The resulting height data, although sampled at the same resolution as the normal data, are inherently lower in resolution. This results in the opposite characteristics of the normal

data – globally accurate but of low resolution. Together, however, the two analysis techniques form a synergistic data set that contains all information necessary to construct an accurate surface map of the tablet.

3. DATA PROCESSING

3.1 Photometric Stereo

The technique of photometric stereo is used to calculate the surface normal map of the tablet. The main premise of the technique is that a surface will appear brighter when the illumination direction converges towards the surface normal. This concept is illustrated in Fig. 3, which shows two images acquired under opposite azimuthal illumination directions. The image on the left is of the tablet being illuminated from the left; the image on the right of right illumination. As can be seen, sections of the tablet which are sloped toward the left appear bright in the left image but dark in the right image. Likewise, rightward slopes are brighter in the image on the right.




Figure 3: Identical views of a tablet under varying illumination directions. The left image is being illuminated from the left, the right image from the right.

Treating the problem with more mathematical rigor, the intensity values of the point (x, y) for a series of images acquired under uniform and collimated illumination are written as³

$$\bar{I} = Q\bar{N}\bar{n}, \quad (1)$$

where Q is the reflectance of the point, \bar{N} is a matrix which describes the directions of incident illumination, and \bar{n} is the surface normal at (x, y) . This equation assumes a Lambertian BRDF. Although only three images are required to uniquely invert Eq. 1, more are used and a least squares approach taken to reduce error and to account for shadowed facets. Defining the z -axis to point downward from the camera towards the tablet, Eq. 1 becomes

$$\begin{bmatrix} I_1 \\ \vdots \\ I_K \end{bmatrix} = Q \begin{bmatrix} \sin(\theta) \sin(\phi_1) & \sin(\theta) \cos(\phi_1) & -\cos(\theta) \\ \vdots & \vdots & \vdots \\ \sin(\theta) \sin(\phi_K) & \sin(\theta) \cos(\phi_K) & -\cos(\theta) \end{bmatrix} \begin{bmatrix} n_x \\ n_y \\ n_z \end{bmatrix}, \quad (2)$$

where θ is the polar angle and ϕ is the azimuthal angle. The least squares solution is

$$Q\bar{n} = \left(\left(\begin{bmatrix} \bar{N} \\ \bar{N} \end{bmatrix} \right)^T \begin{bmatrix} \bar{N} \\ \bar{N} \end{bmatrix} \right)^{-1} \begin{bmatrix} \bar{N} \\ \bar{N} \end{bmatrix}^T \bar{I}. \quad (3)$$

Note that the values used for \bar{I} are background corrected image values; these are obtained by dividing the tablet images by the corresponding background images.

As previously noted, the normal map resulting from this approach does not form a conservative surface due to the nature of the point-by-point calculations. Integration from the normal map to a height field is path-dependent and results in unrealistic shapes when performed on a global scale. To counter these problems, structured light data are incorporated into the final surface determination.

3.2 Structured Light

The basic premise of the structured light technique employed is to measure the phase shift of a sinusoidal pattern projected onto the tablet versus onto a flat background. The resulting phase difference is proportional to the relative tablet height where the constant of proportionality is determined by applying the technique to a flat object of known height. There are three main problems with this approach. First, each projection angle results in some of the object features being shadowed. This problem is easily resolved by using multiple projection angles and statistical analysis to intelligently select an appropriate final value of the phase difference at the point (x, y) . Any remaining “holes” in the data are filled when the data is combined with the normal map to construct the final surface.

The second problem with this approach is that illumination non-uniformities, camera noise, and variations in surface reflectance and orientation make it difficult to accurately measure the phase. This is illustrated in Fig. 4, which shows raw structured light images of the background (left) and the tablet (right) along with vertical line profiles through the centers of the images. As can be seen from the background data, it is difficult to construct a perfect sinusoid even with a flat target surface. When viewing a textured object such as the tablet, changes in surface reflectance and orientation mask the sinusoidal profile and make it impossible to accurately measure phase. Use of the Carré technique of phase-measurement interferometry solves this problem as it does not depend on local reflectance or illumination level. This technique requires that four images of differing phase shifts be acquired. An absolute value of the phase is then calculated via the relation⁶

$$\varphi = \tan^{-1} \left[\frac{\sqrt{(I_1 - I_4 + I_2 - I_3)(3(I_2 - I_3) - (I_1 - I_4))}}{I_2 + I_3 - I_1 - I_4} \right]. \quad (4)$$

This equation, however, is not the final solution; the resulting phase is in fact ambiguous due to the range of the inverse tangent function (φ is bound to $\pm\pi/2$). The value of φ depends upon the order in which the intensity values I_k are input to Eq. 4. In addition, the wrapping of the inverse tangent function causes alternating periods of the phase to switch from ascending to descending values; this in turn wrecks havoc when attempting to calculate the phase difference between object and image data. Resolution of these problems requires that the intensity values be input in a consistent order amongst all points (x, y) . Since determining this order requires full calculation of the phase four times, it is easier to choose a consistent phase value from amongst the four calculated values. In particular, the second positive value is chosen by applying the selection algorithm

$$\Phi = \varphi_1 \times (\varphi_4 > 0) + \sum_{k=2}^4 \varphi_k \times (\varphi_{k-1} > 0), \quad (5)$$

where φ_1 through φ_4 are calculated by varying the order of the input intensity values in Eq. 4. The necessity of this operation is illustrated in Fig. 5, which compares φ and Φ calculated from the same set of raw background data. While φ alternates between ascending and descending slopes and has a range of $\pm\pi/2$, Φ is always ascending and ranges from 0 to $\pi/2$. This range limitation is the negative consequence of implementing the selection algorithm. It is for this reason that absolute certainty of the phase of a given point (x, y) requires that the period of the lowest frequency sinusoid be four times the width of the projected image (that only 1/4 of the cycle be projected).

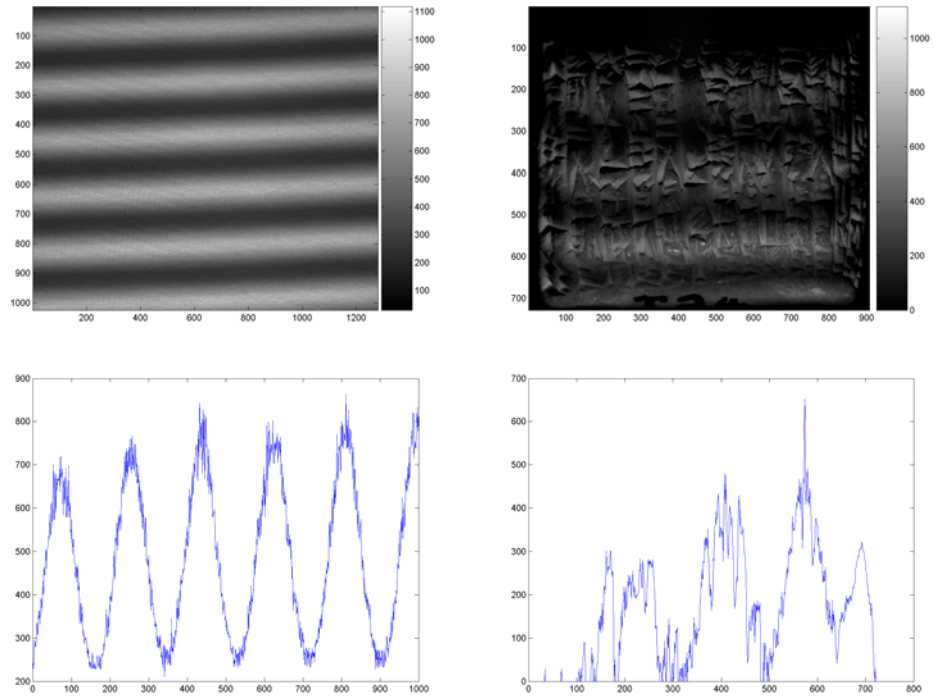


Fig. 4. Raw structured light data of the background (left) and tablet (right). The lower plots are center vertical line profiles.

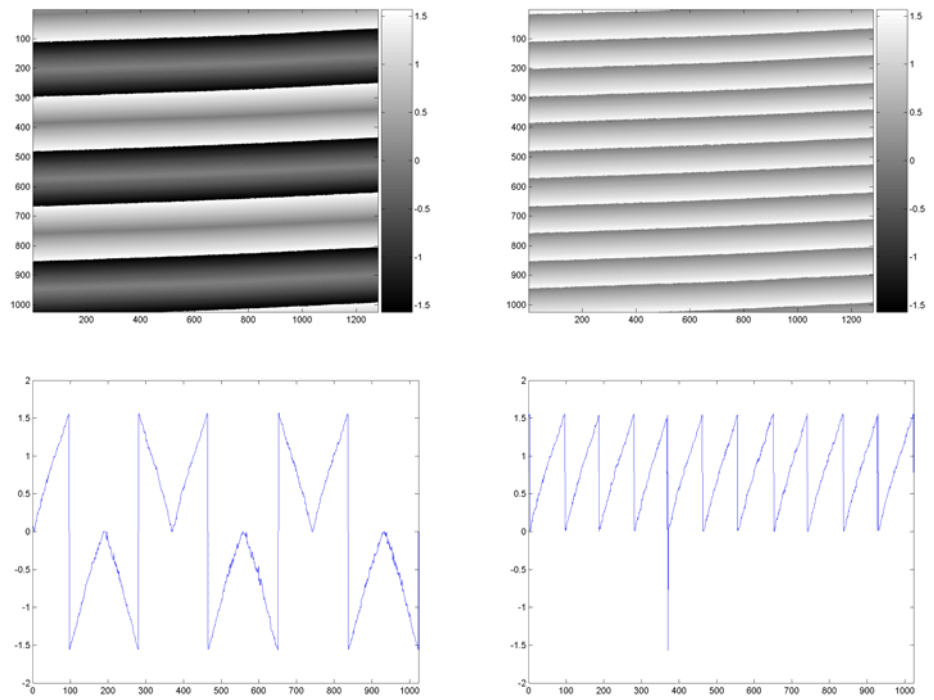


Fig. 5. Comparison of φ (left) and Φ (right) on a set of background data of the same projection frequency. The lower plots are center vertical line profiles.

This limitation on the projected sinusoid leads to the third and final problem associated with the implemented structured light technique. In short, the greater the period, the greater the measurement error. Fortunately, countering this problem is much more straightforward than the last. An iterative approach is taken in which the frequency of the projected sinusoid is doubled and the resulting phase used to refine the original value. Looking at the solution from the opposite perspective, the highest resolution sinusoid is used to determine the phase and the iteratively frequency-halved sinusoids are used to resolve the $\pm n\pi/2$ ambiguities.

The shortcoming of this structured light technique, or rather, its implementation, is its low resolution. The projector has a resolution of 1024x768, which overfills the area viewed by the camera (1280x1024 resolution). This results in a noisy, oversampled, and low resolution surface (in comparison to the normal map). However, the benefit of this technique is its high level of global accuracy, which is unattainable by the method of photometric stereo.

3.3 Final Surface Determination

Each of the previously described measurement approaches has shortcomings that prevent it from being a standalone solution to the scanning needs of the application. Together, however, they compose a complementary data set that contains all information necessary to construct an accurate surface map of the tablet. The normal map resulting from photometric stereo analysis does not form a conservative surface and integration of the data yields global shape inaccuracies. The resolution of the normal data, however, is excellent. Structured light measurements, on the other hand, provide globally accurate height information that is inherently consistent but low in resolution. An iterative minimization algorithm was therefore designed to combine the data sets in such a way as to take advantage of the benefits of each and to discount the drawbacks.

Two main constraints are incorporated into the algorithm. The first minimizes the error between the slope of the final surface and the normal map on a point-by-point basis⁷, thereby taking advantage of the high resolution of the normal data and avoiding problems due to large-scale integration. The second constraint minimizes the relative height difference between the final surface and a 5x5 median filtered structured light height map. This constraint uses the global accuracy of the height data while removing effects due to isolated noisy data points. A complete description of the algorithm follows.

The height of the tablet surface is updated according to the rule

$$h(n+1) = h(n) + ((1-\lambda)\delta h_{PMS} + \lambda\delta h_{SL}). \quad (6)$$

In this equation, δh_{SL} is the difference between the 5x5 median filtered height, h_{SL5} , and the surface height,

$$\delta h_{SL} = h_{SL5} - h(n); \quad (7)$$

λ is a weighting factor bound to the interval [0,0.5],

$$\lambda = \begin{cases} (\delta h_{SL}/25 \text{ um})^2 / 2; & \delta h_{SL} < 25 \text{ um} \\ 1/2; & \text{otherwise} \end{cases}; \quad (8)$$

and δh_{PMS} is the height error calculated by comparing the shape of the current surface to the normal data,

$$\delta h_{PMS}(x, y) = \frac{\chi}{4} [\delta S_x(x-1, y) - \delta S_x(x+1, y) + \delta S_y(x, y-1) - \delta S_y(x, y+1)], \quad (9)$$

where χ is the length of an image pixel (26.8 μm) and $\delta \bar{S}$ is the slope error,

$$\delta \bar{S}(x, y) = \bar{S}(x, y) - \bar{S}_{PMS}(x, y). \quad (10)$$

$\bar{S}(x, y)$ is the slope as calculated from the surface height,

$$S_x(x, y) = \frac{h(x-1, y) - h(x+1, y)}{2\chi}; \quad S_y(x, y) = \frac{h(x, y-1) - h(x, y+1)}{2\chi}, \quad (11)$$

and $\bar{S}_{PMS}(x, y)$ is the slope measured by photometric stereo analysis,

$$\bar{S}_{PMS}(x, y) = -\frac{n_x}{n_z} \hat{x} - \frac{n_y}{n_z} \hat{y}. \quad (12)$$

The initial guess, $h(0)$, used in the algorithm is a 4x4 block-integrated surface (the x and y -slope maps are combined and locally integrated using the Fried algorithm⁸) where the shape of each block is determined by integration of the normal data. The center-height of each block is set to the average height over the region as measured by structured light analysis. An average height adjustment of less than 1/100 of the pixel size (0.268 μm) is used as the exit criterion for the algorithm, with the added restraint that at least 10 iterations be performed.

3.4 Tablet Construction and Visualization

Adjacent tablet scans are typically acquired at 60° view increments by manually repositioning the tablet with the V-block mount along the two major axis of the tablet. A total of ten scans is required to image the entire artifact. Overlapping areas of the data are used to register the scans together for display; the end result mimics a rigid body merging of adjacent “faces” of the tablet. Viewing software, which was written to display the registered data, allows the user to set any desired view and lighting direction, as well as to adjust other shading parameters such as accessibility, curvature, and depth-based shading⁹.

4. RESULTS

The tablet was scanned using the methods described hereto. Meshed surface maps of a 2.68 mm by 2.68 mm cross-section (100 x 100 pixels) of the “front” of the tablet are shown in Fig. 6; the left mesh shows the structured light height map while the right mesh depicts the final surface. These figures substantiate the claim that the minimization algorithm preserves the global height information resident in the structured light data while discounting the local noise.

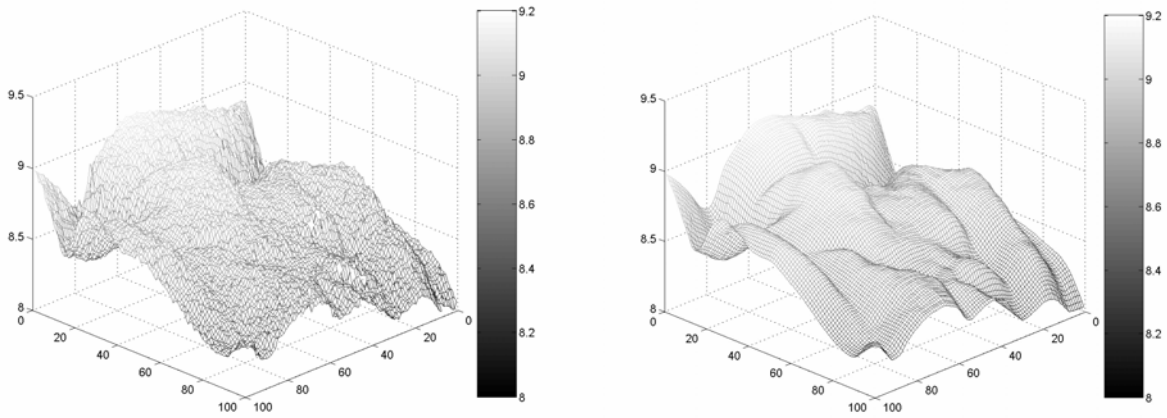


Figure 6: Meshed surface maps of a 2.68 mm by 2.68 mm cross-section of the tablet showing the structured light height map (left) and final surface (right).




Figure 7: X- (top) and y-components (bottom) of the normal vectors over a 2.68 mm by 2.68 mm cross-section of the tablet as measured by the method of photometric stereo (left) and computed from the final surface (right).

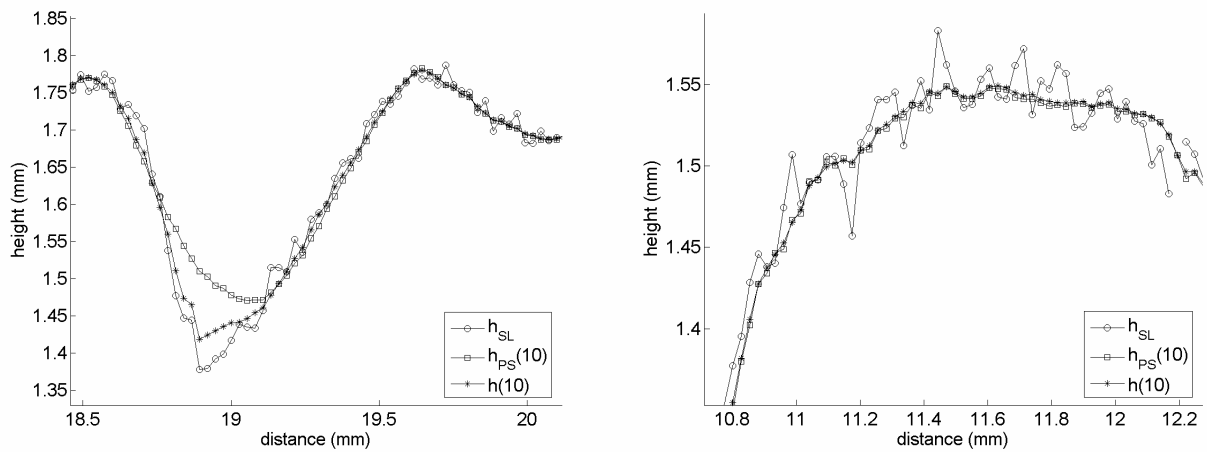


Figure 8: Height profiles of the tablet. The circles represent the structured light height map. A local integration of the normal data is shown with squares. The stars are the final surface (10 iterations).




Figure 9a: Photograph of the tablet.




Figure 9b: Rendering of the tablet from approximately the same view direction as the photograph in Fig. 9a.

A comparison of the normal vectors of the final surface to those measured by the method of photometric stereo over the same area is shown in Fig. 7. Overall, the slope information is preserved well. In areas of steep slopes, however, the final surface exhibits a slightly steeper slope than the measured data. This is because the minimization algorithm adjusts the final surface to more closely match the known height, thereby avoiding excessive smoothing of genuine structure.

Height profiles of tablet data are shown in Fig. 8. Circles represent the structured light height map and squares a local integration of the normal data. The stars are the final surface after 10 iterations. Both the normal integration and the final surface suppress the noise of the height data. However, integration is inaccurate with respect to the genuine structure of the tablet in comparison to the minimization algorithm in areas of steep slopes. This is evident in the center valley in the left plot in Fig. 8. Here, a sharp groove was detected in the structured light data but smoothed over by the integration. The final surface, on the other hand, comes within approximately 50 μm of the groove depth as measured by structured light analysis.

A photograph of the tablet under ambient lighting is shown in Fig. 9a and a rendering from approximately the same view direction and with the light source towards the right is shown in Fig. 9b. The position of the light source was chosen to accentuate the features of the tablet in order to demonstrate the utility of having a 3D surface model compared to photographic records. The surface model can be rotated to any orientation and the light source placed in any position so that the best possible view of a given tablet feature may be obtained. The rendering matches the photograph's character for cuneiform character and also maintains the gross shape of the tablet. This figure pair also points out one of the distinct features of a rendering versus a photo. Photos inevitably display a finite depth of field in which some features are in sharp focus and others are blurred. This is not the case for a rendering which has an inherent infinite depth of field.

Another rendering is shown in Fig. 10. In this rendering, the view direction matches that of the photograph in Fig. 1, which compares the tablet size to that of a quarter. This clearly shows the ability of the 3D model and viewing software to display zoomed-in views of the tablet.

5. CONCLUSIONS

Overall, the prototype scanner does an excellent job of determining the surface shape of the tablet. It acquires data at 26.8 μm x - and y -sample intervals over an area of approximately 34.3 mm by 27.4 mm. The scanner uses off-the-shelf hardware components, thereby minimizing the system cost and allowing for easy expansion and scalability. The resulting final surface is both globally accurate in accordance with height information as measured by a structured light technique and locally accurate in accordance with slope information obtained by the method of photometric stereo.

Scans of the various faces of the tablet have been registered together to form a complete 3D surface model of the tablet. This model and the viewing software allow for examination capabilities that far surpass photographic records. In addition, the software provides capabilities unattainable by inspection of the artifact itself. These include the ability to adjust the exact location of the light source, to incorporate depth-based shading effects on the surface, to adjust surface reflectance characteristics, and to zoom in on any region of the tablet from any view direction.

Future work will seek to determine the absolute accuracy of the measurement techniques and the shape of the final surface. This can be accomplished by scanning objects of precisely known shape. Several functional and algorithmic improvements will also be evaluated. In the short-term, these include utilizing a graycode scheme⁵ to speed up acquisition and processing of structured light data, developing a means of determining the BRDF as well as the surface normal from photometric stereo methods, and composing a 3D background cube to more accurately correct for illumination non-uniformities. Long-term scanner modifications include the use of additional cameras to accelerate scan times of complete tablets, the incorporation of x - and y -translation stages to scan larger tablets, and the construction of a custom digital projector to improve resolution and accelerate data acquisition speed.

In 3D visualization, future work will investigate feature detection and feature-based registration techniques to improve automation of the registration and reconstruction processes and to merge the scans together to form a single surface as opposed to a number of registered surfaces. These algorithms could potentially evolve to the point where they become capable of performing a virtual assembly on fragmented, broken, and/or damaged artifacts. Finally, development of a 3D

viewer that incorporates view-dependent, progressive level of detail will be undertaken. Such a viewer can selectively refine the detail of tablet model in those areas the reader wishes to examine more closely.



Figure 10: Rendering of the tablet from approximately the same view direction as the photograph in Fig. 1. The distance from the top to the bottom of this rendering is approximately the diameter of a quarter.

ACKNOWLEDGEMENTS

This work has been supported by NSF Medium ITR IIS-0205586. In addition, the authors would like to thank Dean Snyder of Johns Hopkins University for his tireless efforts seeking and obtaining funding.

REFERENCES

1. Rioux, M., "Laser range finder based on synchronized scanners," *Appl. Opt.*, V. 23, pp. 3837 – 3844, 1984.
2. ConoLine camera, 75mm focal length, Optimet Optical Metrology, Limited, P.O.B. 45021 Jerusalem 91450.
3. Woodham, R. J., "Photometric method for determining surface orientation from multiple images," *Optical Engineering*, Vol. 19, No. 1, pp. 139-144, 1980.
4. Pan, J., Huang, P. S., and Chiang, F., "Color-coded binary fringe projection technique for 3-D shape measurement", *Optical Engineering*, Vol. 44, No. 2, 2005.
5. Sansoni, G., Carocci, M., and Rodella, R., "Three-dimensional vision based on a combination of gray-code and phase-shift light projection: analysis and compensation of the systematic errors", *Appl. Opt.*, V. 38, No. 31, pp. 6565-6573, 1999.
6. Gåsvisk, K. J., *Optical Metrology*, Second Edition, John Wiley & Sons, pp. 254-260, 1995.
7. Smith, D. J., and Bors, A. G., "Height Estimation from Vector Fields of Surface Normals," *IEEE DSP 2002*, pp. 1031-1034, 2002.
8. Barchers, J. D., Fried, D. L., and Link, J. D., "Evaluation of the performance of Hartman sensors in strong scintillation," *Appl. Opt.*, V. 41, pp. 1012-1021, 2002.
9. Cohen, J., Duncan, D., Snyder, D., Cooper, J., Kumar. S., Hahn, D., Chen, Y., Purnomo, B., and Graettinger, J., "iClay: Digitizing Cuneiform," *The 5th International Symposium on Virtual Reality, Archaeology, and Cultural Heritage*, December, 2004.