Digital Preservation of Ancient Cuneiform Tablets Using 3D-Scanning*

Subodh Kumar   Dean Snyder   Donald Duncan   Jonathan Cohen   Jerry Cooper

Johns Hopkins University†
Baltimore, MD, USA

Abstract

Cuneiform is the world's oldest known writing system. Ancient scribes impressed reed styluses onto damp clay to write the approximately 900 different logographic, syllabic and taxographic signs. These clay tablets were preserved in the ruins of ancient Near Eastern cities for millennia, and today sit in museum storerooms around the world, where they are subject to degradation and are available for study only to scholars who can visit the collections in person.

Cuneiform documents exhibit three dimensional writing on three dimensional surfaces and we seek to provide accurate, high resolution 3D models of these tablets for scholars' use in their research and for digital preservation of these unique historical artifacts. We describe our project and the requirements for constructing and publishing large digital libraries of these ancient tablets. We also present visualization requirements and algorithms for producing them.

1 Introduction

The earliest cuneiform texts appeared in Mesopotamia around 3200 B.C., and the last native cuneiform texts were written around 75 A.D. Over the centuries of its use, cuneiform script evolved from its more expansive early pictorial forms to the later more condensed, abstract forms, in which the signs bear little resemblance to the original symbols. Since the decipherment of Babylonian cuneiform some 150 years ago, museums have accumulated hundreds of thousands of tablets, written in most of the major languages of the Ancient Near East - Sumerian, Akkadian (Babylonian and Assyrian), Eblaite, Hittite, Persian, Hurrian, Elamite, and Ugaritic. These texts include genres as varied as mythology and mathematics, law codes and beer

Figure 1. 3D Scan of the Obverse and Right Side of a "micro-tablet" in the Johns Hopkins Archaeological Collection. This tiny tablet dates to around 2000 BC, and records rations of beer, bread, garlic, oil and soap supplied to eight imperial messengers in the service of the Third Dynasty of Ur. Notice the rounded shape of the tablet, the three dimensional impressions of the wedges in the clay, and the writing streams extending onto the side of the tablet.

recipes. In most cases these documents are the earliest exemplars of their genres, and cuneiformists have been able to make unique and valuable contributions to the study of such modern disciplines as history, linguistics, mathematics and science.

In spite of continued great interest in mankind's earliest documents, it has been estimated that only about 1/10 of the extant cuneiform texts have been read even once in modern times. There are various reasons for this: the complex Sumero/Akkadian script system is inherently difficult to learn; there is, as yet, no standard computer en-

coding for cuneiform; there are only a few hundred qualified cuneiformists in the world; the pedagogical tools are, in many cases, non-optimal; and access to the widely distributed tablets is expensive, time-consuming, and, due to the vagaries of politics, becoming increasingly difficult. By creating digital replicas of the tablets we seek to create an *accessible archive* that will not erode.

1.1 Justification

Cuneiform, as written by ancient scribes, is, by its very nature, a three dimensional writing system:

1. most of the media are three dimensional – rounded clay tablets;

2. the characters are three dimensional – stylus tracings and wedge impressions in wet clay;

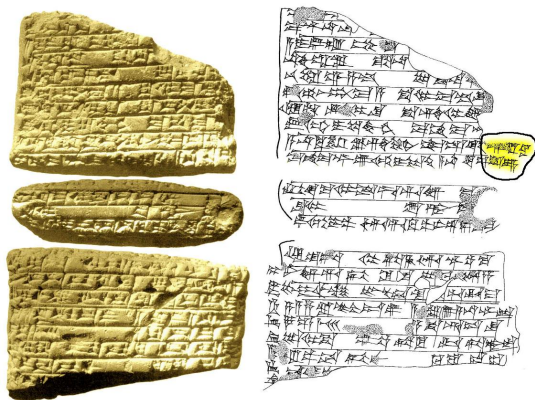3. the writing streams are three dimensional – sentences often run over the edges and onto the backs of tablets.

Figure 2. Photograph and Autograph of an Old Babylonian Text. The highlighted signs on the hand copy run off the right edge of the tablet and are not visible in the 2D photographs. "Sargon, the Lion", Joan Goodnick Westenholz, 1997, Legends of the Kings of Akkade

Heretofore cuneiformists have used two main techniques to represent and archive cuneiform documents graphically - 2D photography and hand-drawn copies, or autographs.

A photograph of a tablet has the advantage of conveying a fairly accurate 2D visual representation of a tablet's layout. However, due to the multi-tiered three dimensionality of cuneiform documents (wedge impressions, round tablets, and over-the-edge writing) one photograph is not enough. Several photographs are needed, taken at different angles, with different lighting, and at different magnifications, in order to convey enough useful 2D information to enable the

collation of a single tablet. This is obviously an onerous burden for research and publication. Furthermore, even if done successfully, these photographs of a single tablet do not contain views from *every* direction, something that can be generated from a high resolution 3D scan.

The primary advantage of a hand copy, or autograph, is also its primary disadvantage – it provides a representation of an author's interpretation of what signs are on a tablet. Manually drawing autographs is a laborious, time consuming, error-prone, and highly subjective process requiring direct access to the tablets. Cuneiformists must apply for travel grants to visit the tablet collections in London, Berlin, Istanbul, etc. And, though the resulting autographs have the advantage of recording a scholar's interpretation of difficult to read signs, both the quality of the interpretation and the quality of the drawing vary widely, and disputed readings are common. In order to verify disputed readings, cuneiformists must apply for additional travel grants to inspect the tablets once again. The entire process is obviously slow, delicate, expensive, tedious, and, in the end, unproductive. Moreover, unlike photographs, autographs, by definition, are useless for collation, and are practically useless for paleography.

In order to compensate for the deficiencies inherent in both techniques, lately we are seeing more and more cuneiform publications using both 2D photographs and autographs. But the combination is still inadequate for tablet collation. It is no wonder then that we are also seeing a number of recent forays into 3D surface scanning of cuneiform tablets, including by our Digital Hammurabi project introduced by Lee and Snyder [30].

Accurate, detailed, and efficient 3D visualization will enable the virtual "autopsy" of cuneiform tablets and will revolutionize cuneiform studies, not only by making the world's tablet collections broadly available, but also by limiting physical contact with these valuable and unique ancient artifacts, while at the same time providing redundant archival copies of the originals. It will have the happy side effect of dramatically reducing the need for expensive travel grants.

1.2 Previous work

Image-based re-lighting [17, 18] is an interesting approach to generate novel views of a tablet from only a few photographic images. However these are only 'interpolated' views and lose much of the accuracy. Furthermore, knowing the real geometry can help us do geometric analysis and fragment matching.

Levoy and Anderson [4] developed a parametrization of the 3D tablet scans onto planes and used it to "unwrap" the tablets on a sheet of paper. This has great potential for cuneiform print publication. They used 3D data acquired

by the NRC technology [7, 23].

The problem of model reconstruction from scanned data is well addressed in the literature. Mencl et. al. [21] categorize the methods into those using *surface warping, spatial subdivision, distance fields, or front growing.*

Warping-based reconstruction methods deform a known initial surface to fit the input point set [27, 25]. Spatial subdivision based techniques partition the data into cells. For each cell that the surface (or the enclosed volume) passes, these techniques approximate or interpolate a surface [14, 11, 6]. Guo et. al. [13] use visibility algorithms and Teichmann et. al. [26] use density scaling and anisotropic shaping to improve the results of reconstruction using α-shapes. Volume subdivision schemes decompose space into volumetric cells and remove ones that are found to be 'outside' the sampled surface. The final surface is reconstructed from the remaining cells [10, 29, 1]. In particular, Amenta et. al. [1, 2] use Voronoi filtering approach in 3D to construct the *crust* [3] of the sample points. They provide theoretical guarantees on the topology of their reconstructed mesh assuming "sufficiently dense" sampling. These algorithms are too CPU and memory intensive, however.

Another class of algorithms use a distance field that gives the shortest distance from any point to the surface. The surface passes through the zeroes of this distance field [15, 5]. These schemes are able to process large data, albeit sometimes at the cost of accuracy. Surface growing algorithms [10, 19, 20, 8, 12] construct the surface incrementally by connecting more samples to the set of triangles generated so far. Among the more accurate and efficient methods in this class is the the Ball pivoting algorithm [8]. Given an unstructured point set and a radius ρ, it finds triangles ensuring that a ball of radius ρ containing its vertices does not contain any other.

None of these algorithms take advantage of the structure in the scanned data, which [28] does by triangulating each scan separately and zippering the overlapping scans in a post-processing step. Our scheme is similar in spirit, but much simpler and hence more efficient.

2 Scanning Requirements

Cuneiform tablets are precious artifacts and require safe and minimal handling. In addition, they typically cannot be removed from their collections to be brought to a scanning facility. Therefore a cuneiform scanner must be portable and easily calibrated in the field (a non-controlled environment).

Apart from the usual requirements for precision, stability, and repeatability, the speed and automation of scanning is paramount. There are more than approximately 300,000, mostly palm-sized, tablets in various collections around the

Figure 3. This enlargement from one of the tablet in our collection gives some indication of the current resolutions attainable with laser triangulation - ca. 50 micrometers, or 500 dpi. In a 2mm long wedge one can count approximately 40 data points; this yields about a 50 micrometers. We estimate our resolution requirement to be approximately 25 micrometers to capture all details.

world. A scan time of even 2 to 3 minutes per tablet implies that it will take several scanner-years to scan all of them.

Related to scan times, we must ensure that the scan and model reconstruction process is not labor-intensive. Due to the requirement of limited handling and high automation we must implement *conclusive scanning*: it is crucial to collect all the required data in the first attempt in order to reduce scan time, overhead, and tablet re-handling. Lastly, in order to truly preserve the characteristics of the tablets, we require high resolution and accuracy.

In order to perform virtual tablet autopsy, cuneiformists must be able to differentiate stylus marks from erosion, damage, and stray marks. They must also be able to distinguish fine markings of rolled seals and even detect unambiguously the "damming" that results when one wedge, overwriting another, pushes up a little of the clay. Although we are still actively engaged in the requirements establishing phase of the project, preliminary tests suggest a target of at least 25 micrometer (or ca. 1000 dpi) resolution will be adequate for our application.

3 Scanning Procedure

We have used laser triangulation based scanners in our experiments. All images presented in this paper have been generated from models scanned using the Canadian National Research Council scanner [23]. A conceptual diagram of this scanner is shown in figure 4. The object is
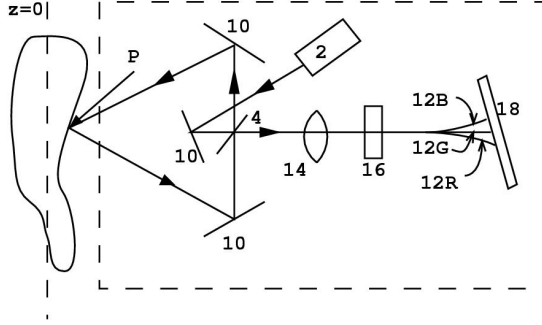
Figure 4. Canadian National Research Council RGB scanner conceptual diagram (reproduced from [22])

illuminated with red, green, and blue lines from a laser (2) that is directed to the point P via a scanning galvanometer mirror (4) and turning flat (10) at the top of the figure. The imaging system (10) views the point P through the back of the scanning galvanometer mirror and the other turning flat (10) at the bottom of the figure. An image of the spot P forms on the linear array detector (18). The dispersive element (16) produces three such separate images corresponding to the red, green, and blue wavelengths from the laser. In operation, the galvanometer mirror dithers back and forth producing a spot that scans across the object. First consider a flat object located at the reference plane. After the twin offsetting reflections of the light from the front and back of the galvanometer mirror, as the spot moves across the reference plane, the image on the detector remains stationary. For object heights departing from the reference plane, the position of the image moves in proportion to these heights. Thus, determination of the position of the focused spots on the detector array allows estimation of the height of the object above the reference plane.

This scanner has a nominal resolution of 50 micrometers.

3.1 Scan-planning

Even though the topology of Cuneiform tablets is simple, there are ample variations in the height to cause under-sampling in many regions. Most reconstruction algorithms generate holes or incorrect triangles in the under-sampled areas. Others fill the holes with interpolated geometry. We must avoid under-sampling in the first place in order to truly preserve the tablets and allow full "autopsy". Hence, we perform automatic detection of holes to help guide the operator in planning the next scan. This requires that we merge multiple scans quickly, online, to detect the remaining holes. We use a fast and automatic registration process as described in the next subsection to detect these holes.

3.2 Registration

Since not all of a tablet is visible from any single position, multiple scans are required, each producing a range image. These images must be combined to produce the final 3D model. We assume that the range image is specified as depth values on an XY grid. Although not all scanners directly produce grids, they can be rearranged in a 2D array, based upon the order of scanning or simply by resampling. We exploit this information in our system. One important aspect of our application is the large number of tablets that require scanning. As a result, the common strategy of registering scans by first manually aligning the individual range images is too labor intensive. At the same time, to eliminate the risk of damage, we must not apply any physical fiducial marks on the objects, nor may we use any clamping devices to hold these priceless artifacts. Objects must be carefully handled by hand only.

Hence we have devised a scheme of approximate fiducial markers near the object using pointing pins attached to flexible arms. In this scheme we allow both the scanner and the object to move, but limit the number object re-positioning. The fiducial pins are not attached to the tablet but are moved with it. Each pin is associated with a point on the rim of the tablet by the operator and each time the object is moved, the pins are repositioned to stay close to their associated points. If a pin occludes the object from the scanner, it is removed for that scan. For scanning tablets, it can be shown that a set of fifteen pins is sufficient to see at least eight non-occluding pins from any view around a cube. These eight pins provide sufficient redundant equations to obtain a close registration between scans. Only a few further iterations of the modified iterated closest points (ICP) algorithm [9, 28] achieves excellent registration.

We are further investigating methods for accurate registration using only the surface geometry. We need to run the registration algorithm only in instances when the tablet needs to be physically lifted and repositioned. Since we scan at a high resolution, it is inefficient and unnecessary to consider all points for registration of individual range images. It is enough to identify edges and pass the points on an edge and those in their immediate neighborhood to the ICP algorithm. Furthermore, like [28], we prevent the boundary points as well as geometrically distant points from matching and use normal estimates to assign weights to each point. Formally, using a technique similar to [28], we minimize

$$E = \sum_{i=1}^{n} w[i]|M[i] - \mathbf{R}(M'[i]C') - \mathbf{T}|^2,$$

where $M[i]$ and $M'[i]$ are corresponding samples from mesh M and M'. C' is the centroid of mesh M'. $w[i]$ is the accuracy of point $M'[i]$. We use $w[i] = Average\{\frac{1}{d_{ik}}\}$,
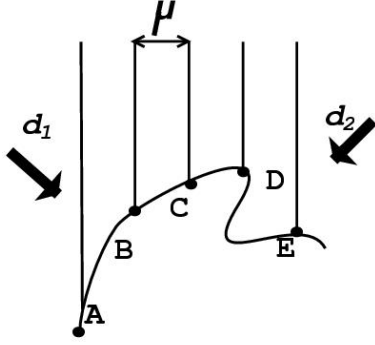
Figure 5. Detection of under-sampled areas dues to slope (as in A to B) or due to occlusion (as in D to E)

Figure 6. Highlighting areas of unreliable data (See the electronic version for color.)

where d_{ik} is the distance between point $M[i]$ and neighbor $M[k]$. We consider a non-boundary point that is closer to its neighbors more reliable. In practice, the distance may simply be approximated by the difference in the depth values. Note that high depth disparity implies large slope or missing data, both sources of error. $\mathbf{T}$ is the difference of centroids of M and M': $C - C'$ [16]. $\mathbf{R}$ is derived from the cross-covariance matrix of centroid-adjusted pairs of points. $\mathbf{R}$ as a quaternion is the eigenvector corresponding to the largest eigenvalue of the element matrix of this cross-covariance matrix [16].

3.3 Under-sampling detection

Recall that a scanner may scan at a fixed lateral resolution in the range space, call it μ. We do not have the same density of samples on the manifold. For example, in Figure 5 points may be arbitrarily far apart on the surface, either due to a large inclination of the surface (e.g., points A and B) or due to occlusion (e.g., points D and E). We require that the samples on the surface be close. We conservatively determine that more samples are needed in any part of the surface where inter-sample distance on the surface is larger than 3μ, thus discarding data when the slope of the surface is larger than $70°$. We compute and mark all points bounding an under-sampled region. Currently this data is visualized with the marked points highlighted (as shown in Figure 6). This allows the operator to determine where to place the scanner next. In the future, we plan to use direct robot-controlled placement of the scanner (e.g., directions d_1 and d_2 in Figure 5) to minimize the number of scans.

Since under-sampling is a scan time operation, it is important to detect it efficiently. Triangular topology reconstruction can aid this detection but triangulation needs time in the order of a minute. Hence, we have chosen to determine under-sampling directly using the sample points. Although, it is straightforward to perform this in the scanner
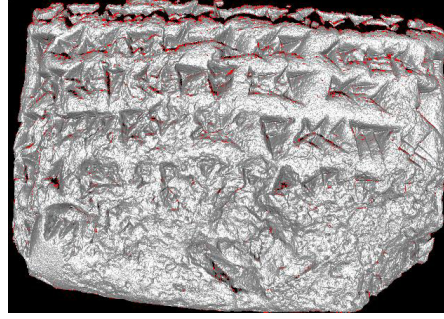
space, one must find holes not simply in each scan but the combined set of samples found so far. This requires registration of the scans with each other. our registration algorithm takes only a few seconds and hence this is possible. In fact, our under-sampling detection algorithm is robust enough to work well even with imprecise registration. The algorithm works iteratively, processing each scan as it is generated:

1. Compute holes within the new range image. This step simply considers the neighboring X and Y pixels, and compares their depth values to determine the distances. A pixel with any adjacent pixel missing or having depth disparity is marked *potentially sparse* (see Figure 7).

2. A statistical approach is needed. If there is true under-sampling in a region (not simply scanner noise), several points must be marked potentially sparse. If more than half the adjacent points of a point are potentially sparse, we mark the point *sparse*.

3. As new range images are generated, we register them to the rest of the registered model, matching points from different scans based on their distances. If a sparse point, P, matches P' that is non-flagged, the sparse flag of P is removed.

4. We generate a visualization of holes to help plan subsequent scanning by coloring the sparse vertices in red (Figure 6). The goal of our 'conclusive scanning' is that every sparse point in a given range image must correspond to a non-sparse point in some other range image.

4 Triangulation

Unorganized point triangulation has been a popular research area. Most algorithms are either too memory intensive to enable the triangulation of many million points;
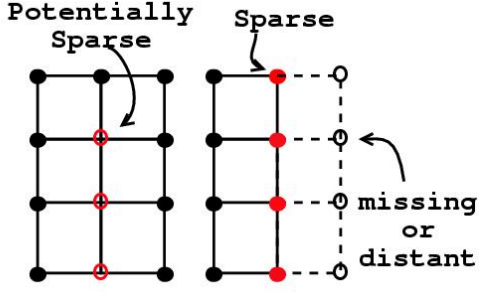
Figure 7. Sparse points are locally discovered. For a point to be marked sparse, it must be a part of a sparse region.

Figure 8. Range image triangulation respecting ridges

others do not guarantee triangle quality. Another common assumption in most algorithms is that the underlying surface is smooth, which is untrue in our application. We need to capture the normal discontinuities, i.e., cuneiform wedge boundaries with high fidelity. In our algorithm, we take advantage of the fact that scanners do indeed produce "organized" points at mostly regular spacing. We also exploit the knowledge that our models are 2-manifolds.

Triangulations related to Delaunay triangulations are commonly applied for constructing the topology. Some operate volumetrically [1] but are impractical for large data size. Others [12] operate in some 2D projection of the points but must find the correct plane to project to. Delaunay triangulations are known to minimize the integral of square of gradient across the surface [24]. This is appropriate for smooth surfaces, but not for our case, when wedges should not be smoothed.

Due to our high resolution scans, it is sufficient to use the plane of the range image to perform triangulations most of the time, when sufficient data is available. We only employ the more expensive topology construction at the boundaries of the images and in sparse regions. At the same time we avoid the geometric complexities of zippering [28]. Our main premise is that any given point on the surface has a primary scan: the scanner view that corresponds to the angle of view closest to that point's normal. This view is used for the triangulation. Any other views of the point are only used to refine the triangulation based on the primary view.

We start by noting that the 2D Delaunay triangulation of a uniform grid of points is simply the quadrilaterals formed by joining the adjacent points as shown in Figure 8. Although, our grids are not necessarily rectilinear or uniform, they are close to it. (We only join non-sparse points to their neighbors in this step.) There are two choices for the diagonal. We choose the one that represents the underlying ridge or valley on the surface the most. In other words, if CD has a larger gradient change from C to D than AB does from A to B (see Figure 8), we choose AB.
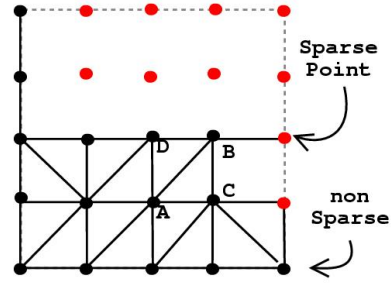
Triangulating each range image thus takes time proportional to the number of samples. Moreover, multiple scans do not need to be in the memory together. We do require any overlapping scans to be together in the memory for the stitching step.

We stitch the triangulation of overlapping range images as follows: Recall that every sparse point in range image must have a corresponding point in some other image that is non sparse.

We describe here the algorithm for stitching only two meshes, although the same process is repeated for merging more than two meshes together. In order to stitch mesh M and M', we must determine the location of each resulting 3D point and its topology, i.e., other points it is joined to by the triangulation. Assume point P in mesh M corresponds to point P' in mesh M'. Call them a *pair*. In general, many points in a mesh M are too far from their closest point in mesh M' and hence do not have a pair. These points and their triangles are retained in the final model.

To determine the final location of a point P that is more reliable than its pair P', we use a weighted average. We first compute the projection P_p of P on mesh M'. The final location of P is given by $P \frac{R[P]}{R[P]+R[P']} + P_p \frac{R[P']}{R[P]+R[P']}$

To find the topology of a point, P, we need to determine its *link*, $\{P_i\}$: the ordered list of vertices around it (and connected to it). In the example in Figure 9, the link of point P is the list $ABDE$. If the point P is more reliable than P', we start with the topology of P and merge with it the topology of P' and *vice-versa*. (If P is a boundary point, we always choose P'. Note that the *conclusive scanning* condition requires that both P and P' may not be boundary points.)

For the example in Figure 9(a), consider point pairs P and P', P being the more reliable of the two. We call a point P *conforming*, if the link of P matches the link of P'. In other words, if the list of points $\{P_i\}$ form the link of P, $\{P_i'\}$ form the link of P'.

If the links of P and P' are conforming, we just retain the link of P in the final mesh. Otherwise, there are two
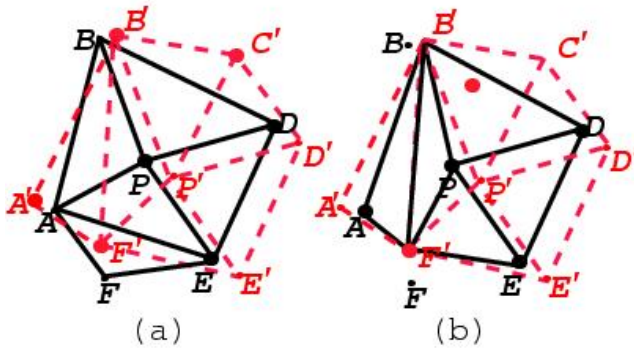
Figure 9. Triangulation stitching: M is shown in black, M' in red (gray). (a) shows the two meshes, (b) shows them merged.

cases to consider:

1. Point F' is in the link of P' but, F is not in the link of P: If F' is more reliable than F, we designate PF a *weak* edge, otherwise P and F are not connected. The edge PF is *associated* with mesh M (in which P and F are not connected).

2. Point A is in the link of P but A' is not in the link of P': If A' is the more reliable vertex, PA' is marked *weak*, otherwise AP is connected. Edge PA' is *associated* with M' (in which A' and P' are not connected).

In the second pass we process weak edges. If edge PF is listed as a weak edge, we project all the edges connected to P and F on the scanner plane of the associated mesh. Due to our distance restriction, this turns out to be a single edge: the edge that forms the base of two triangles that have P and F as their apexes. In our example, the projected edges are, respectively, $B'F'$ for weak edge AP and AE for edge PF. If a weak edge intersects a non-weak edge (as AP intersects $B'F'$, it is deleted. If a weak edge intersects only another weak edge, we choose the one that has a lower change of gradient between its end points.

Recall that the scanner data usually has gaps we choose to not flag as holes in the *potentially sparse* areas, where only few data items are missing. We currently fill those values by interpolation if no other scan covers the area. If larger holes still remain in the data, our triangulation will leave holes in these areas also, as one of its basic assumptions fails. Due to our efficient algorithm, more than a million samples can be triangulated in about a minute. With further enhancements, we believe it can be performed online and thus improve the on-line registration without fiducials.

5 Conclusion

We have embarked on an ambitious project to enable in-the-field, highly accurate and high-resolution scanning of hundreds of thousand cuneiform tablets. This requires that the scanner be light, fast and largely automatic. As a part of this project, we are also developing methods for scan-planning and automated scanner placement. Accurate and automatic registration of separate range images and final model reconstruction is also important. We have developed early versions of these algorithms and continue to improve them. Efficient, high quality visualization of the resulting models, while not addressed in this paper, is also important. We have found so far that conclusive scanning is a tough challenge. Interpolating neighboring values to fill small holes works to an extent. We are developing statistical algorithms to determine when the missing samples may, in fact, not be necessary and hence ignored.

6 Acknowledgments

References

[1] N. Amenta, M. Bern, and M. Kamvysselis. A new voronoi-based surface reconstruction algorithm. In *Proc. ACM Siggraph Annual Series*, pages 415–421, 1998.

[2] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th ACM Sym. on Comp. Geom.*, pages 213–222, 2000.

[3] N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19(2-3):127–153, 2001. (special issue on surface reconstruction).

[4] S. Anderson and M. Levoy. Unwrapping and visualizing cuneiform tablets. *IEEE Computer Graphics and Applications*, 22(6):82–88, 2002.

[5] M. L. B. Curless. A volumetric method for building complex models from range images. In *Proc. ACM Siggraph Annual Series*, pages 303–312, 1996.

[6] C. Bajaj, F. Bernardini, and G. Xu. Automatic reconstruction of surfaces and scalar fields from 3d scans. In *Proc. ACM Siggraph Annual Series*, pages 109–118, 1995.

[7] J.-A. Beraldin, F. Blais, P. Boulanger, L. Cournoyer, J. Domey, S. F. El-Hakim, G. Godin, M. Rioux, and J. Taylor. Real world modelling through high resolution digital 3d imaging of objects and structures. *ISPRS Journal of Photogrammetry & Remote Sensing*, 55:230–250, 2000.
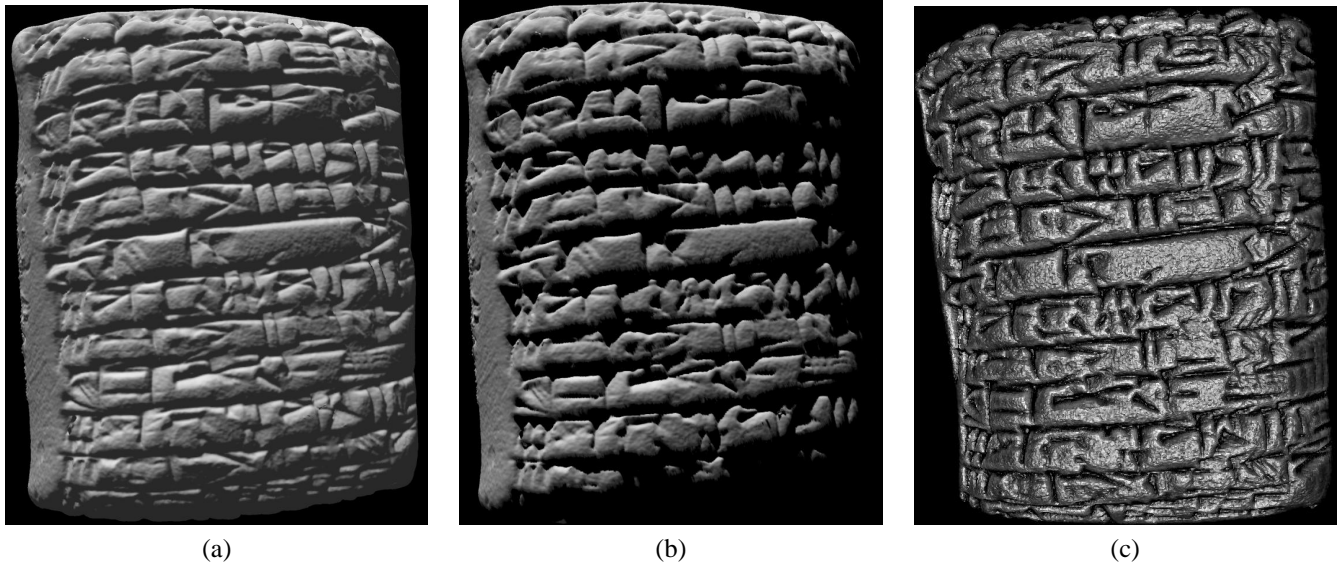
Figure 10. Tablet visualization. (a) uses Gouraud Shading. (b) adds shadows. Note the enhancement of wedge details. (c) uses depth-shading and small features are enhanced further. For depth-shading, we compute the local depth of each point and use that to modify the amount of light reaching that point for shading. This assumes that on average deeper recesses receive less light.

[8] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. on Visualization and Computer Graphics*, 5(4), 1999.

[9] P. Besl and N. McKay. A method of registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[10] J. D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics*, 3(4):266–286, 1984.

[11] H. Edelsbrunner and E. Mucke. Three dimensional alpha shapes. *ACM Transactions on Graphics*, 13(1):43–72, 1994.

[12] M. Gopi, S. Krishnan, and C. Silva. Surface reconstruction using lower dimensional localized delaunay triangulation. *Computer Graphics Forum (Eurographics 2000)*, 19(3):467–478, 2000.

[13] B. Guo, J. Menon, and B. Willette. Surface reconstruction from alpha shapes. *Computer Graphics Forum*, 16(4):177–190, 1997.

[14] H. Hoppe, T. Derose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized point clouds. In *Proc. ACM Siggraph Annual Series*, pages 71–78, 1992.

[15] H. Hoppe, T. Derose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. In *Proc. ACM SIGGRAPH Annual Series*, pages 21–26, 1993.

[16] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America, A*, 4(4):629–642, 1987.

[17] M. Levoy and P. Hanrahan. Light field rendering. In *Proc. ACM SIGGRAPH Annual Series*, pages 31–42, 1996.

[18] T. Malzbender, B. Zuckerman, and W. Bodine. Image based relighting of cuneiform tablets.

[19] R. Mencl. A graph-based approach to surface reconstruction. *Computer Graphics Forum (Proc. Eurographics)*, 14(3):445–456, 1995.

[20] R. Mencl and H. Muller. Graph-based surface reconstruction using structures in scattered point sets. In *Proc. CGI 98*, 1998.

[21] R. Mencl and H. Muller. Interpolation and approximation of surfaces from three-dimensional scattered data points, 1998.

[22] M. Rioux. Three dimensional colour imaging, 1999. Canadian patent no. 2,043,336, issued 1999/07/27.

[23] M. Rioux, L. Cournoyer, et al. Tri-color laser triangulation.

[24] S. Rippa. Minimal roughness property of the delaunay triangulation. *CAGD*, 7(6):489–497, 1990.

[25] R. Szeliski and D. Tonnesen. Surface modeling with oriented particle systems. In *Proc. ACM SIGGRAPH Annual Series*, pages 185–194, 1992.

[26] M. Teichmann and M. Capps. Surface reconstruction with anisotropic density-scaled alpha shapes. In *Proc. IEEE Visualization*, pages 67–72, 1998.

[27] D. Terzopoulos, A. Witkin, and M. Kass. Constraints on deformable models: Recovering 3d shape and non-rigid motion. *Artificial Intelligence*, 36:91–123, 1988.

[28] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *Proc. ACM SIGGRAPH Annual Series*, pages 311–318, 1994.

[29] R. C. Veltkamp. Boundaries through scattered points of unknown density. *Graphical Models and Image Processing*, 57(6):441–452, 1995.

[30] L. Watkins and D. Snyder. The digital hammurabi project. In *Mueums and the Web*, 2003.