Automatic, Accurate Surface Model Inference for Dental CAD/CAM[1]

Chi-Keung Tang, Gérard Medioni
Inst. Rob. Intell. Sys., USC
Los Angeles, CA 90089-0273.

François Duret
School of Dentistry, USC
Los Angeles, CA 90089-0641.

Abstract

Dental CAD/CAM offers the prospects of drastically reducing the time to provide service to patients, with no compromise on quality. Given the state-of-the-art in sensing, design, and machining, an attractive approach is to have a technician generate a restorative design in wax, which will be milled by a machine in porcelain or titanium. The difficulty stems from the inherent outlier noise in the measurement phase. Traditional techniques remove noise at the cost of smoothing, degrading discontinuities such as anatomical lines which require accuracy up to 5 to 10 μm to avoid artifacts. This paper presents an efficient method for the automatic and accurate data validation and 3-D shape inference from noisy digital dental measurements. The input consists of 3-D points with spurious samples, as obtained from a variety of sources such as a laser scanner or a stylus probe. The system produces faithful smooth surface approximations while preserving critical curve features such as grooves and preparation lines. To this end, we introduce a voting technique, which efficiently ignores noise, infers smooth structures, and preserves underlying discontinuities. This method is non-iterative, does not require initial guess, and degrades gracefully with spurious noise, missing and erroneous data. We show results on real, complex, and noisy dental data.

Keywords

Dental CAD/CAM, patient-specific models, grouping and segmentation, representation, surface fitting, computer vision.

[1]This work was supported in part by the National Science Foundation under Grant No IRI-94231441.

Automatic, Accurate Surface Model Inference for Dental CAD/CAM

1 Introduction

Dental CAD/CAM has been revolutionizing dentistry for the past ten years, following Prof. Duret's pioneering work in the 1970s [4], and is ubiquitous in every major dental group and laboratory today. Its main components include data acquisition, modeling, and milling systems (Figure 1). The only system widely used commercially is

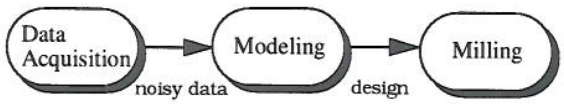

Figure 1 Dental CAD/CAM major components.

the CEREC system, produced by Siemens Inc. It is a self-contained unit with an imaging camera, a monitor, and an electrically controlled machine to mill inlay and onlay restorations from ceramic blocks. The accuracy is not good, and significant manual processing is required. Another system of interest, developed under Prof. Duret's leadership, was able to produce crowns with an average gap of 35 μm . The system is no longer commercially available, and suffered from lack of speed and a cumbersome interface.

An alternative approach, followed here, is to perform the restoration design manually, in a laboratory, by using a conventional medium such as wax; then to transfer this physical design into a digital model. This information can be directly used to control a CAM machine to mill the restoration from a ceramic (or other) block.

The major problem is the transfer from a physical, wax shape, to the computer model. Digital measurements are sampled from the wax model, using palpation or optical sensing [5]. Figure 2 shows the set up of one such commercial, multiple-view registration system, from which our data were obtained. As we shall see, the point sets obtained, though mostly accurate, also contain many *erroneous* outlier readings.

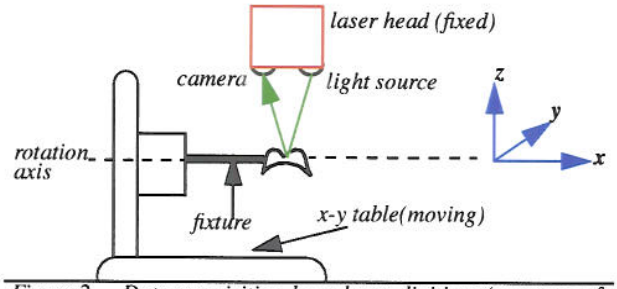

Figure 2 Data acquisition by a laser digitizer (courtesy of GC Coporation, Tokyo, Japan).

Given such noisy measurements, the challenge is to derive an *accurate* shape description *automatically*. Otherwise, even slight errors may result in artifacts which need to be corrected manually, or worse, may make the restoration unusable.

While filter-based techniques such as discrete Fourier transform [11] are effective in ignoring spurious samples, they often "oversmooth", degrading sharp discontinuities and distinct features that correspond to important anatomical (preparation) lines and features. At present, unless intensive human intervention is used, it is impossible to construct an accurate representation that respects both medical and dental criteria requiring accuracy of up to 5 to 10 μm [6].

Recently, much progress has been made in computer vision for robust surface inference from clouds of points. Their applications should result in better 3-D dental model descriptions. This paper presents an integrated approach for inferring dental models in the form of surfaces (for capturing smoothness) and 3-D curves (for preserving shape discontinuities) from noisy dental data.

Much work has been done in surface fitting to clouds of points. The majority of work use the *deformable model* approach (first proposed by Kass *et al.* in [10] and then in [19]) which attempt to deform an initial shape using energy minimization so that the deformed shape fits a set of points. The work by Boult and Kender [2], Poggio and Girosi [13], Blake and Zisserman [3], Fua and Sander [4], Szeliski *et al.* [14], and many others belong to this category. *Physics-based* approaches proposed by Terzopoulos *et al.* (in [17] and [18]) model the problem as a multi-particle system governed by (physics) laws and equations. The initial surface (or model), originally in equilibrium, is now subject to external forces exerted at each data point. Such forces make the system converge to another equilibrium. Hoppe *et al.* [9] and Boissonnat [1] use *computational geometry* techniques, which treat the data as vertices of a graph and construct edges based on local properties.

Most methods above are computationally expensive as an iterative process takes place. Also, they have limited potential in inferring faithful 3-D models from dental data: most do not work in the presence of spurious noise with surface discontinuities usually smoothed out.

Our method attacks these problems by applying a recent work on sparse 3-D data ([8] and [16]), where a non-linear *voting* process is used to achieve efficient feature segmentation and discontinuity detection. Our method is robust to outlier noises since its effect is reduced by accumulating a large number of vector votes. We give an overview, and outline our paper in the next section.

2 Overall approach

Our method is summarized in Figure 3. Each input point is first quantized in a 3-D voxel array. A preprocessing step is then applied to estimate the normal to the surface. This step, as well as the surface and curve inference processes, is implemented in a convolution-like operation with predefined vector fields. The design of these fields is described in section 3. Such fields, when aligned with an estimated input normal, associate to its neighborhood voxels a preferred orientation.

By such alignment, we accumulate, at each voxel, a collection of vector "votes". Two independent 3-D *saliency maps*, one for surface (SMap) and one for curve (CMap), are produced. Details on voting and saliency maps are explained in section 4, and its full version in [8].

Local extrema (yellow areas in Figure 3) in these maps are extracted, resulting in a triangulation mesh (surfaces) and a set of directed and connected 3-D poly-line segments (curves). Feature extraction is described in section 5 (also detailed in [15]). These features are then be refined to localize detected discontinuities and remove inconsistent surface patches (such as the spurious patches labelled in Figure 2). Consequently, the surfaces and curves inferred in the 3-D model are coherently integrated. Feature refinement is described in section 6 (also detailed in [16]). Finally, results on real dental data are shown in section 7 (also on the accompanying video tape).

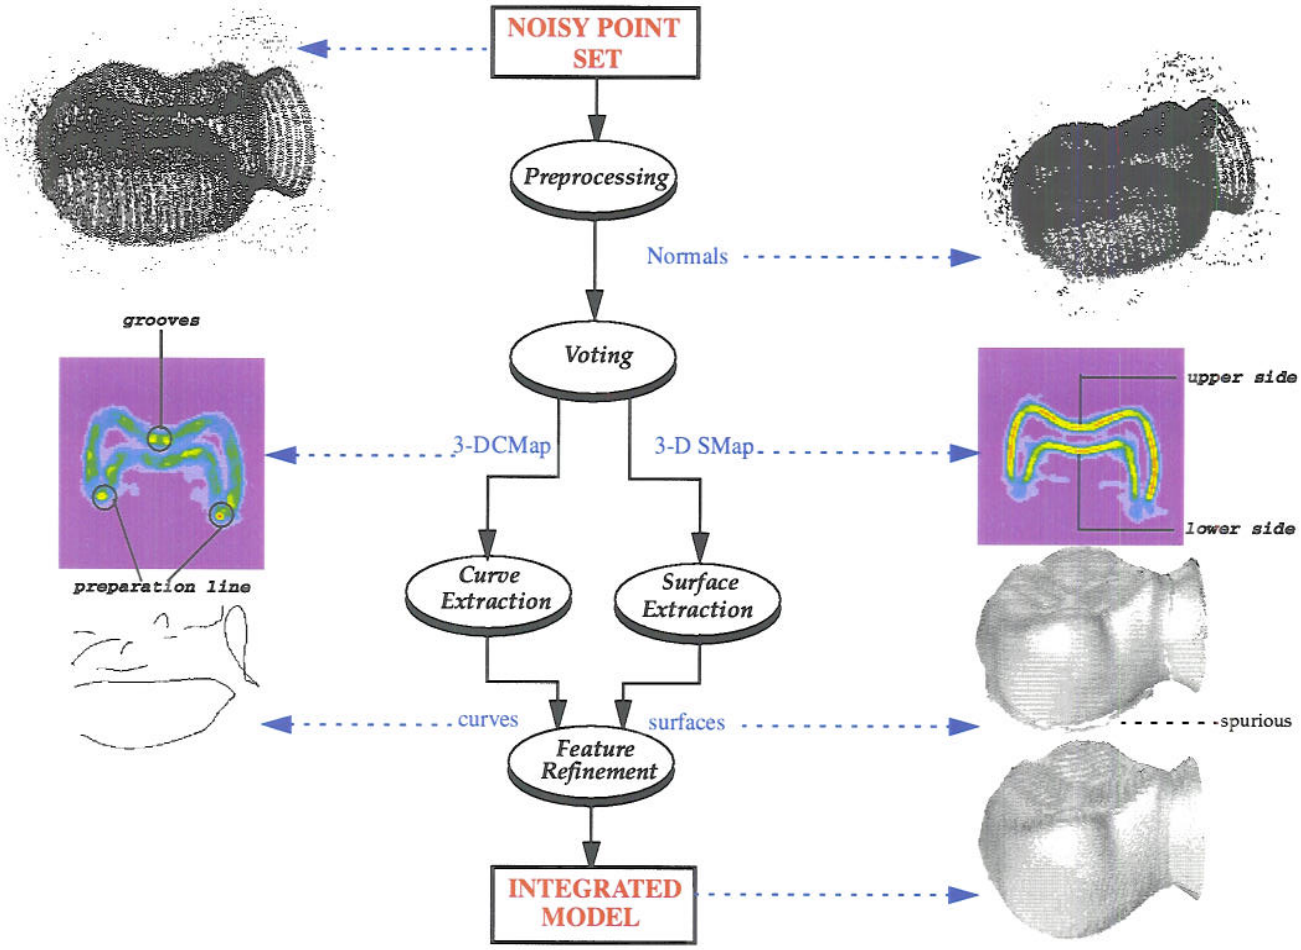

Figure 3 *Overall strategy of our method, using crown as a running example. (Only one slice of the 3-D SMap and CMap are shown)*

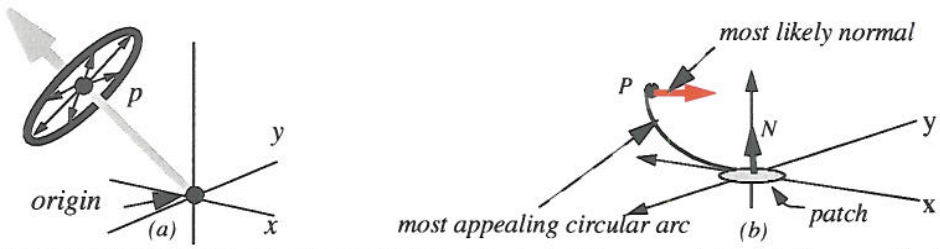

Figure 4 *The design of voting fields. (a) Given a point P, all normals (thin arrows) at point P are equally likely. We choose to represent all of them with one vector (thick arrow). (b) Given a patch and point P, the circular continuation is most appealing.*

3 The design of the fields

There are three types of 3-D vector voting fields in [8]: the *3-D point field* (*P-field*), *curve segment field* (*C-field*), and *Diabolo field* (*D-field*). P-field is used in preprocessing for normal estimation, and D-field is used for surface and curve inferences.

P-field. Without any *a priori* assumption, given a point P and an origin O, the most likely surface passing through

O and P is a family of planes containing these two points, represented by a single vector $\overrightarrow{OP}$ (Figure 4(a)).

D-field. Given a normal vector $\vec{N}$ at origin O and a point P (Figure 4(b)), without any *a priori* information, the most likely and "natural" surface through P is the circular continuation between O and P, because it keeps the curvature constant. The "most likely" normal is normal to that arc at P. The collection of such most likely normal vectors comprise the *D-field*. Its 3-D shape resembles a series of "concentric bowls" (Figure 5)

In all cases, the weight of each vector is inversely proportional to the distance from O and the curvature of the underlying circular arc connecting O and P. A Gaussian decay function is used for that purpose.

4 Vector voting and saliency maps

Suppose we already have normals at each point (preprocessing will be described shortly). The output of voting is two 3-D dense *saliency maps*, one for curve (CMap) and one for surface (SMap). Computing saliency maps is done by voting, which is realized by aligning each input normal vector (*voter*) with D-field (Figure 5). The resulting map is a collection of oriented fields. Each voter accumulates the 'votes' for its own preferred orientation and strength from every other input in the volume.

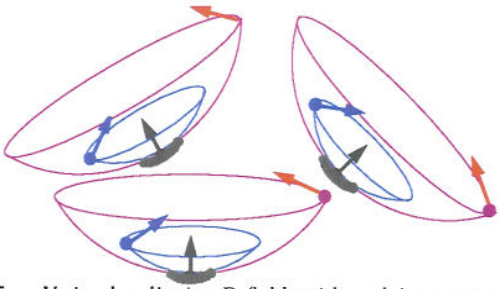

Figure 5 Voting by aligning D-fields with each input voter.

The contributions to a voxel are treated as vector weights, and we compute the central moments of the resulting system, which is equivalent to computing a 3-D ellipsoid having the same moments and principal axes. Such a physical model acts as an approximation to a majority vote which gives both the preferred direction and measure of the agreement.

Let $\overrightarrow{DF} = (DF^x, DF^y, DF^z)^T$ denote any vector constituting the D-field, and $A_{I_{ijk}}$ be the alignment operator for voter I_{ijk} at voxel (i, j, k). Then, for *each* voxel (x, y, z) in the entire 3-D array, we define the accumulated vote, O_{xyz}, as a 3x3 variance-covariance matrix where m^{xyz}_{uvw} is defined

$$O_{xyz} = \begin{bmatrix} m^{xyz}_{200} & m^{xyz}_{110} & m^{xyz}_{101} \\ m^{xyz}_{110} & m^{xyz}_{020} & m^{xyz}_{011} \\ m^{xyz}_{101} & m^{xyz}_{011} & m^{xyz}_{002} \end{bmatrix}$$

in the following, with $0 \le u, v, w \le 2$ and $u + v + w = 2$.

$$\sum_i \sum_j \sum_k \|I_{ijk}\|^2 \left[\left(A_{I_{ijk}} \cdot \overrightarrow{DF}\right)^x\right]^u \left[\left(A_{I_{ijk}} \cdot \overrightarrow{DF}\right)^y\right]^v \left[\left(A_{I_{ijk}} \cdot \overrightarrow{DF}\right)^z\right]^w$$

We decompose the above matrix of central moments O_{xyz} into the corresponding eigensystem, i.e.,

$$\begin{bmatrix} V_{min} & V_{mid} & V_{max} \end{bmatrix} \begin{bmatrix} \lambda_{min} & 0 & 0 \\ 0 & \lambda_{mid} & 0 \\ 0 & 0 & \lambda_{max} \end{bmatrix} \begin{bmatrix} V_{min} \\ V_{mid} \\ V_{max} \end{bmatrix}$$

where $\lambda_{min} \le \lambda_{mid} \le \lambda_{max}$ are the three eigenvalues, and the Vs denote the corresponding eigenvectors of the system. Such decomposition always yield real, non-negative eigenvalues as the matrix is positive and semi-definite.

The three eigenvectors correspond to the three principal directions of an ellipsoid in 3-D, while the eigenvalues describe the strength and agreement measures of the 3-D votes. On a smooth surface, the votes produce high agreement around one direction, so $\lambda_{max} \gg \lambda_{mid}, \lambda_{min}$ (Figure 6(a)). Along the curve bounding two surfaces, two of the eigenvalues are high, and one small, leading to $\lambda_{mid} \gg \lambda_{min}$ (Figure 6(b)).

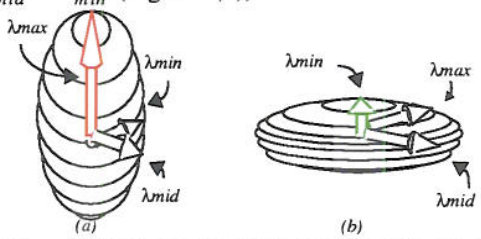

Figure 6 The two important voting ellipsoids. (a) λmax >> λmid ~ λmin, high agreement in exactly one direction (a surface). (b) λmax ~ λmid >> λmin, high agreement in exactly two orientations (an intersection, or 3-D curve).

Thus, after vote accumulation, two 3-D voxel maps defining the surface and curve *saliencies* are produced. Each voxel of these maps has a 2-tuple $(s, \hat{v})$, where s is a scalar saliency measure and $\hat{v}$ is a unit vector:

- surface map (*SMap*): $s = \lambda_{max} - \lambda_{mid}$ as **surface saliency**; $\hat{v} = V_{max}$ indicates the **normal** direction.
- curve map (*CMap*): $s = \lambda_{mid} - \lambda_{min}$ as **curve saliency**; $\hat{v} = V_{min}$ indicates the **tangent** direction.

Pre-processing for normal estimation

Normal inference is also achieved by voting: P-field is aligned with each input point and votes are collected exactly as described above, but only at the voters, making this step computationally inexpensive. As a result of this step, each voter now holds a normal, obtained as the eigenvector V_{max} corresponding to $\lambda_{max} - \lambda_{mid}$.

5 Surface and curve extraction

Local extrema in CMap (resp. SMap) are extracted by the extremal curve (resp. surface) algorithms to be outlined here. We refer the interested reader for more details to [15].

5.1 Extremal curve extraction

Each voxel in the CMap holds a 2-tuple $(s, \hat{t})$, where $s = \lambda_{mid} - \lambda_{min}$ is curve saliency and $\hat{t} = V_{min}$ indicates tangent direction. Suppose the CMap is continuous in

which $(s,\hat{t})$ is defined for every point p in 3-D space. A point p with $(s,\hat{t})$ is on an *extremal curve* if any displacement from p on the plane normal to $\hat{t}$ will result in a lower s value, i.e., $\frac{\partial s}{\partial u} = \frac{\partial s}{\partial v} = 0$, where u and v define the plane normal to $\hat{t}$ at the voxel center (Figure 7):

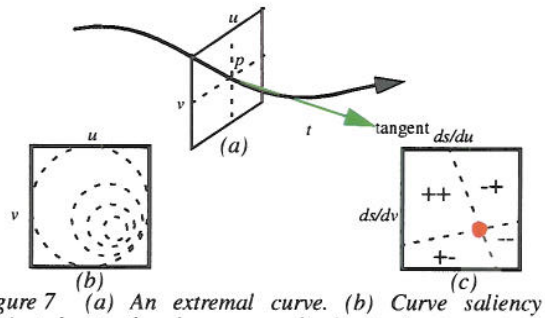

Figure 7 (a) An extremal curve. (b) Curve saliency is projected onto the plane perpendicular to the tangent to the curve. (c) A change in signs of derivatives in both u and v directions indicate a curve passes through the voxel.

This definition therefore involves the detection of zero crossing in the u-v plane normal to $\hat{t}$. To do this, we compute the saliency gradient $\hat{g}$ as,

$$\hat{g} = \left[\frac{\partial s}{\partial x}\ \frac{\partial s}{\partial y}\ \frac{\partial s}{\partial z}\right]^T = \nabla s$$

Define $\hat{q} = R \cdot (\hat{t} \times \hat{g})$ where R defines a frame aligning with the u-v plane. By construction, $\hat{q}$ is the projection of $\hat{g}$ onto the plane normal to $\hat{t}$. Therefore, an extremal curve is the locus of points for which $\hat{q} = \hat{0}$. The corresponding discrete $\hat{q}$ can be similarly defined, from which a tracing algorithm can be readily defined. The output is a set of connected and oriented poly-line segments representing the detected extremal curves.

5.2 Extremal surface extraction

Each voxel in the SMap holds a 2-tuple $(s,\hat{n})$ where $s = \lambda_{max} - \lambda_{min}$ indicating surface saliency and $\hat{n} = V_{max}$ denoting the normal direction. As before, suppose the SMap is continuous in which $(s,\hat{n})$ is defined for every point p in 3-D space. A point is on an *extremal surface* if its saliency s is locally extremal along the direction of the normal (Figure 8), i.e., $\frac{ds}{d\hat{n}} = 0$

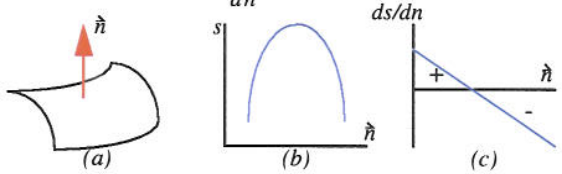

Figure 8 An extremal surface. (a) A surface patch with its normal $\hat{n} = V_{max}$, (b) the saliency along the normal, (c) the derivative of saliency.

This definition involves the detection of zero-crossing on the line aligned with $\hat{n}$, which is computed by defining a scalar $q = \hat{n} \cdot \hat{g}$, where $\hat{g}$ was defined earlier. Therefore, an extremal surface is the locus of points for which $q = 0$.

We define the corresponding discrete q, which can be processed directly by the Marching Cubes algorithm [12]: A polygonal mesh is thus produced.

6 Feature refinement by hybrid voting

The final phase integrates surfaces and curves as obtained to produce an unified 3-D model description. While the voting technique described in [8] and in section 4 provides good results on smooth structures, it only *detects* discontinuities but does not properly *localize* them. This is because SMap and CMap are *independently* interpreted without cooperation. For example, when SMap is used *alone* for surface inference, it may produce incorrect surface patches around surface discontinuities where they have low surface saliencies (e.g. the labelled spurious patches in Figure 2); or these discontinuities, though detected in CMap, may be smoothed out in SMap.

Thus, initial extracted curves and surfaces need to work together in this phase. To this end, hybrid voting, using different 3-D *excitatory* and *inhibitory* fields, which are derived from section 3, is used. In summary, a curve detected in CMap is first treated as *surface inhibitor* in SMap so that when a smooth surface is traced, it is forbidden to enter regions close to any detected discontinuities. With such "trimmed" surface, the same curve is then treated as *surface exciter* for computing precise and natural surface junction. A set of "extended" surfaces are produced. These surfaces will undergo subsequent refinement to produce a coherently integrated surface and curve descriptions.

Figure 9 summarizes the feature refinement process. We refer the interested readers to [16] for further details.

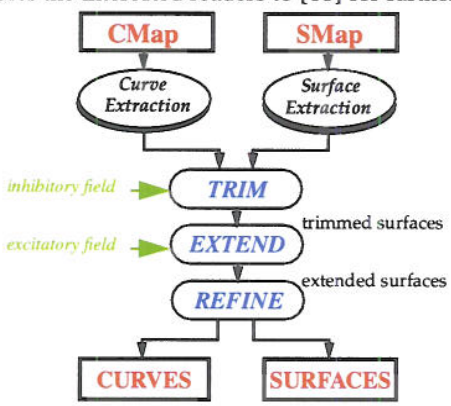

Figure 9 Summary of the cooperative feature refinement.

7 Results

We tested our working prototype on a variety of crowns and inlays. Data are acquired using the set-up shown in Figure 2. The wax shape is rotated about the x-axis in 15° (or 90°) increment so that 24 (or 4) successive views are visible to the sensor. In the following, we emphasize on: (a) input validation, (b) automatic surfaces, and (c) discontinuity curves extraction., with (b) and (c) already explained above. In fact, (a) is an immediate extension: After voting, we classify each point in the original noisy input as being "inlier" (resp. "outlier") if its associated surface or curve saliency is high (resp. low). This shows a simple application of spurious noise elimination when a faithful model is available.

With a fixed and small 5x5x5 vector field used in voting (the most time-consuming part), it has been shown in [16] that our algorithm runs in linear-time, or $O(n)$ time where n is total number of input points. The following tabulates all actual processing times (in min) on the example data, measured on a Sun Ultra 2. Owing to severe missing and erroneous data, feature refinement is not run except for *Crown-24*. Also, we shall see that 24 views are too many, while 4 are barely enough. Nonetheless, in all cases, our method produce accurate and faithful results and degrades gracefully with missing data and erroneous noise.

	no of points	Voting	Feature extraction	Feature refinement	Total
Crown-4	8844	6	2	-	8
Mod-4	5217	4	2	-	6
Inlay-4	2447	3	1	-	4
Crown-24	60095	35	5	5	45

7.1 Crown-4

We first demonstrate in this and the next two examples the graceful degradation of the voting method in the presence of spurious and missing samples. Because of this, the detected features are not further refined as outlined in section 6. Though, shape discontinuities are still faithfully and automatically marked. A sparse set of only 4 views of a *Crown* are digitized and quantized in a 100x100x100 voxel array, which contains 8844 points. Figure 10 depicts sample slices of the noisy and missing data. Despite that, we

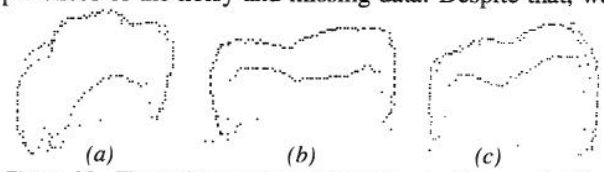

Figure 10 Three slices of the original input of Crown-4 at (a) x=30, (b) x=60, (c) x=79, showing missing and erroneous data.

can still recover the underlying surface model, which are automatically marked with curves representing grooves and preparation lines (though not as good as *Crown-24*).

With a faithfully inferred model, we can perform "inlayer" classification to filter out "outlier" noises present in the original data. Applications only requiring such filtered data set can skip feature extraction and refinement stages, thus speeding up the whole process. Figure 11 shows the result of inlier classification, and Figure 18 shows the resultant extremal surfaces and curves.

7.2 Mod-4

A set of 4 views of a *Mod* are digitized and quantized in 100x100x100 array, which contains 5217 points. This data set is more difficult because of its more complicated shape, and more missing and misleading data resulting from fewer registered views and self occlusion (Figure 12). Results of inlier classification are shown in Figure 13, and extremal surfaces and curves in Figure 19. While the final curve and surface descriptions are not as coherently integrated as that of *Crown-24*, the results are still reasonably faithful. Note that surfaces are only correct from the dis-

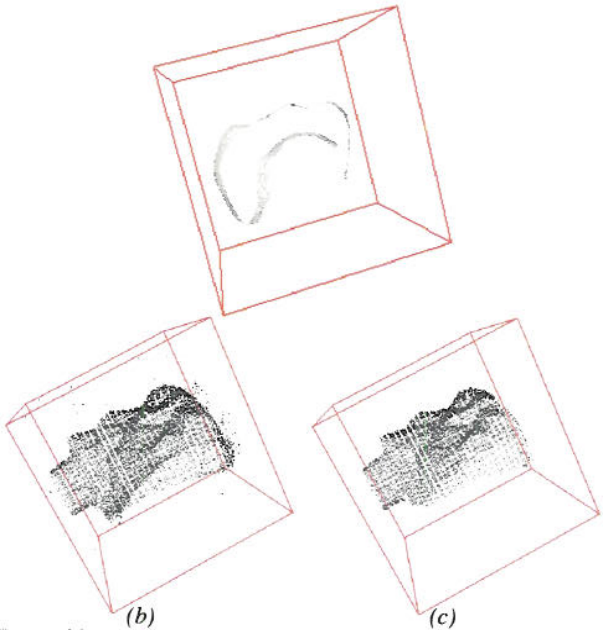

Figure 11 (a) A slice of the extremal surfaces of Crown-4. Part of the lower side is missing due to severe missing data (b) the original noisy data set and (c) result after inlier classification, with outliers removed.

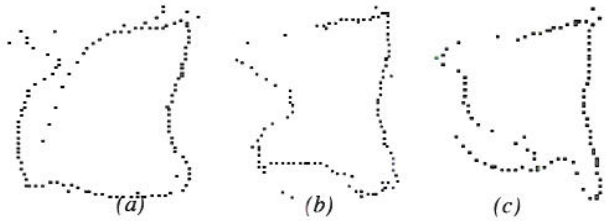

Figure 12 Three slices of the original input of Mod-4 cut at (a) x=40, (b) x=44, (c) x=48, showing missing and erroneous data.

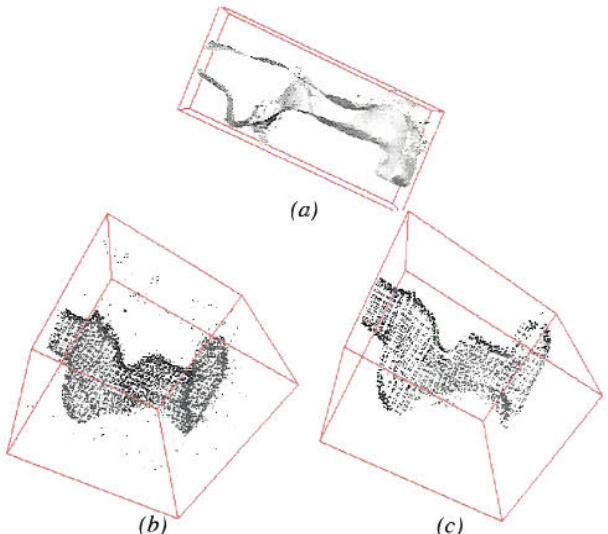

Figure 13 (a) A slice of the extremal surfaces of Mod-4, (b) the original noisy data set and (c) result after inlier classification, with outliers removed.

continuity curves, for which they have low surface saliencies. However, these regions are detected and marked by extremal curve algorithms as they are characterized by high curve saliencies. This allows further manual refinement or intervention to take place if necessary.

7.3 Inlay-4

A set of only 4 views of an *Inlay* are digitized, having 2447 data points quantized in a 70x50x70 array. This is a very complicated surface, and the data points are noisy, missing, and erroneous (Figure 14). Results of inlier classification is shown in Figure 15. The extremal surfaces and curves extracted are shown in Figure 20.

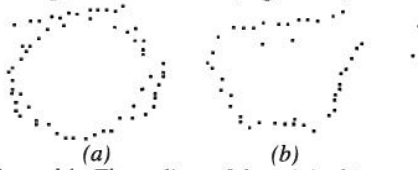

Figure 14 Three slices of the original input of Inlay-4 cut at (a) x=23, (b) x=36, (c) x=50, showing missing and erroneous data.

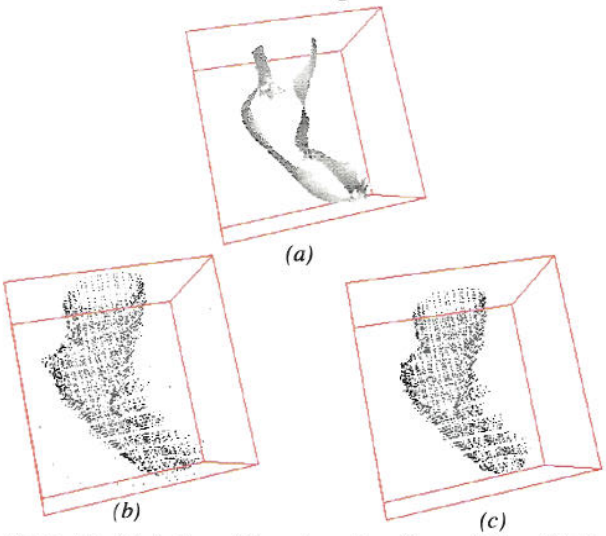

Figure 15 (a) A slice of the extremal surfaces of Inlay, (b) the original noisy data set in the same view, and (c) result after inlier classification, with outliers removed.

7.4 Crown-24

A set of 24 views of a *Crown* are registered, which contain 60095 points, quantized in a 100x100x100 array. Figure 16 depicts three slices of the data, showing mostly accurate but also redundant (thick), spurious and missing data. We can detect the upper and lower sides, the preparation line, and the grooves, which are in turn used to produce a coherently integrated surface and curve description. With an inferred model, we are able to perform inlier classification (Figure 17). Figure 21 shows the three views of the original input, the extremal discontinuity curves (i.e. grooves and preparation line) and surfaces (i.e. upper and lower sides) of the crown.

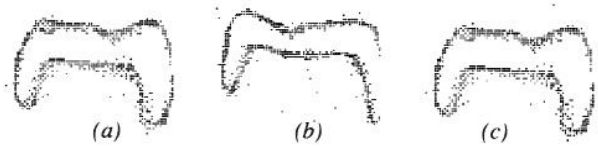

Figure 16 Three slices of the original input of Crown-24 at (a) x=42, (b) x=65, (c) x=80, with mostly accurate and erroneous data.

Note that the total processing time for *Crown-24* is much longer than other examples, which is due to the sheer size of data, many of which are indeed redundant. The optimal number of views may depend on various factors, and is the subject in need of further investigation.

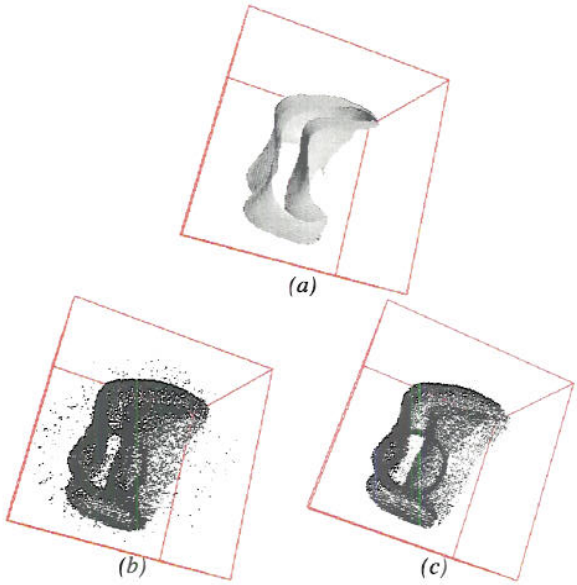

Figure 17 (a) A slice of the extremal surfaces of Crown-24 showing the upper and lower side with preserved and coherently integrated discontinuity, (b) the original noisy data set and (c) result after inlier classification, with outliers removed.

8 Conclusion

We have presented an efficient and robust "voting" method to automatically generate faithful dental models, consisting of surfaces and 3-D curves, from very noisy measurements. As shown by our results, while further improvement can still be made (or simply provide us with more views of data), our working prototype indeed demonstrates a very promising potential on significantly improving current dental CAD/CAM systems, as our system produces faithful results despite that some sample data are noisy, missing and confusing. We not only interpolate smooth structures, but also respect important anatomical lines, and filter out spurious outlier points as well. Our method is efficient, effective, and automatic which offers the prospects to reduce time to provide necessary service.

References

[1] J.D. Boissonnat, "Representation of Objects by Triangulating Points in 3-D space", in *Proc. Intl. Conf. Patt. Recogn.*, pp.830-832, 1982.

[2] T. E. Boult and J. R. Kender, "Visual Surface Reconstruction Using Sparse Depth Data", in *Proc. IEEE Comput. Vision Patt. Recogn.* (Miami Beach, FL), pp.68-76, 1986.

[3] A. Blake and A. Zisserman, "Invariant Surface Reconstruction Using Weak Continuity Constraints", in *Proc. IEEE Comput. Vision Patt. Recogn.* (Miami Beach, FL), pp.62-67, 1986.

[4] F. Duret, "Empreinte Optique", *Prosthodontic*, Lyon. Dental School: DDS thesis #FD73.

[5] F. Duret, The Dental Robotic: "State of the Science Dental CAD/CAM", *Clark's Clinic. Den. Ency*, Philadelphia, Lincholl ed.

[6] F. Duret, "Functionality and accuracy of Sopha CAD CAM system today", in *Proc. Int. Conf. Comput.*, vol. 1, p.500.
[7] P. Fua and P. Sander, "Segmenting Unstructured 3D Points into Surfaces",in *Proc. Euro. Conf. Comput. Vision* (Santa Margherita Ligure, Italy), pp. 676-680, 1992.
[8] G. Guy and G. Medioni, "Inference of Surfaces, 3-D Curves, and Junctions from Sparse, Noisy 3-D Data", *IEEE Trans. on Patt. Anal. and Machine Intell.*, vol.19, no.11, pp.1265-1277, 1997.
[9] H.Hoppe, T.DeRose, T.Duchamp, J.McDonald, W.Stuetzle, "Surface Reconstruction from Unorganized Points", *Computer Graphics*, 26, pp.71-78, 1992.
[10] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models", *Int. J. Comput. Vision*, pp.321-331, 1988.
[11] H.Kimura, H.Sohmura, et al., "Three Dimensional Shape Measurement of Teeth (part 3)", J Osaka Uni dent Sch.
[12] W.E. Lorensen and H.E. Cline, "Marching Cubes: A High Resolution 3D Surface Reconstruction Algorithm", *Computer Graphics*, 21(4), 1987.
[13] T. Poggio and F. Girosi, "A theory of networks for learning", *Science*, pp.978-982, 1990.
[14] R. Szeliski, D. Tonnesen, and D. Terzopoulos, "Modeling Surfaces of Arbitrary Topology with Dynamic Particles", in *Proc. IEEE Comput. Vision Patt. Recogn.* (New York City, NY), June 1993, pp.82-85.
[15] C.-K. Tang and G. Medioni, "Extremal Surface and Curve Algorithms", *IRIS-USC Technical Report.*
[16] C.-K. Tang and G. Medioni, "Integrated Surface, Curve and Junction Inference from Sparse 3-D Data Sets", in Proc. *IEEE Intl. Conf. Comput. Vision* (Bombay, India), Jan 1998, pp.818-824.
[17] D. Terzopoulos and D. Metaxas, "Dynamic 3D models with local and global deformations: deformable superquadratics", *IEEE Trans. on Patt. Anal. and Machine Intell.*, vol. 13, no. 7, pp.91-123, 1991.
[18] D. Terzopoulos and M. Vasilescu, "Sampling and reconstruction with adaptive meshes", in *Proc. IEEE Comput. Vision Patt. Recogn.*,(Lahaina, Maui, HI), 1991, pp.70-75.
[19] M. Vasilescu and D. Terzopoulos, "Adaptive Meshes and Shells: Irregular Triangulation, Discontinuities, and Heretical subdivision", in *Proc. IEEE Comput. Vision Patt. Recogn.*, pp.829-832, 1992.

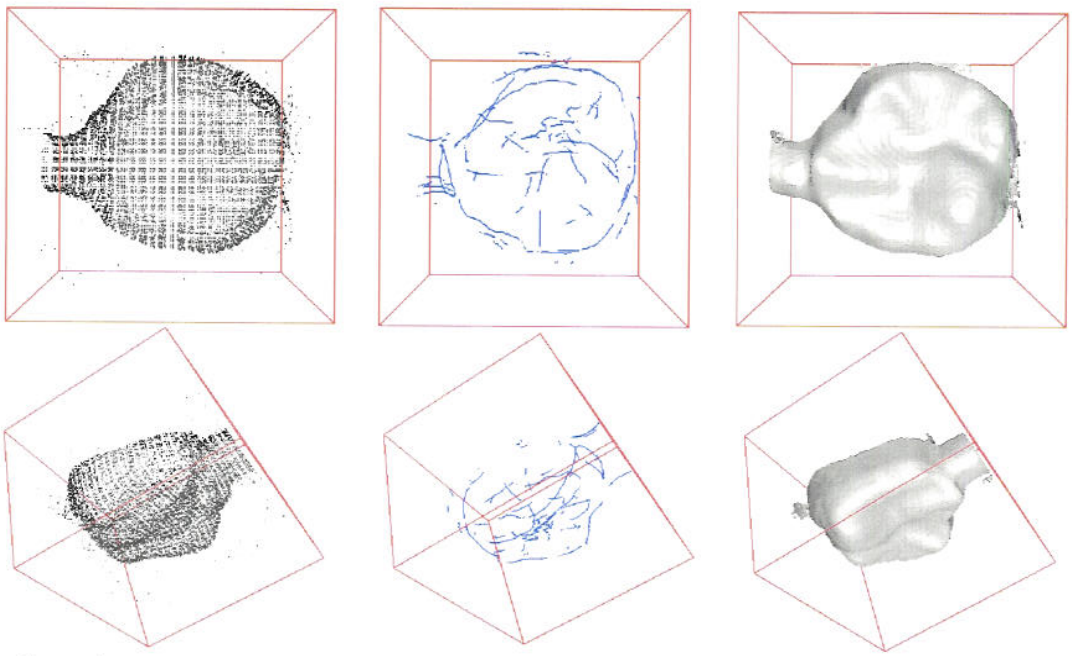

Figure 18 Crown-4. Two views of the original noisy data, the extremal discontinuity curves and surfaces inferred.

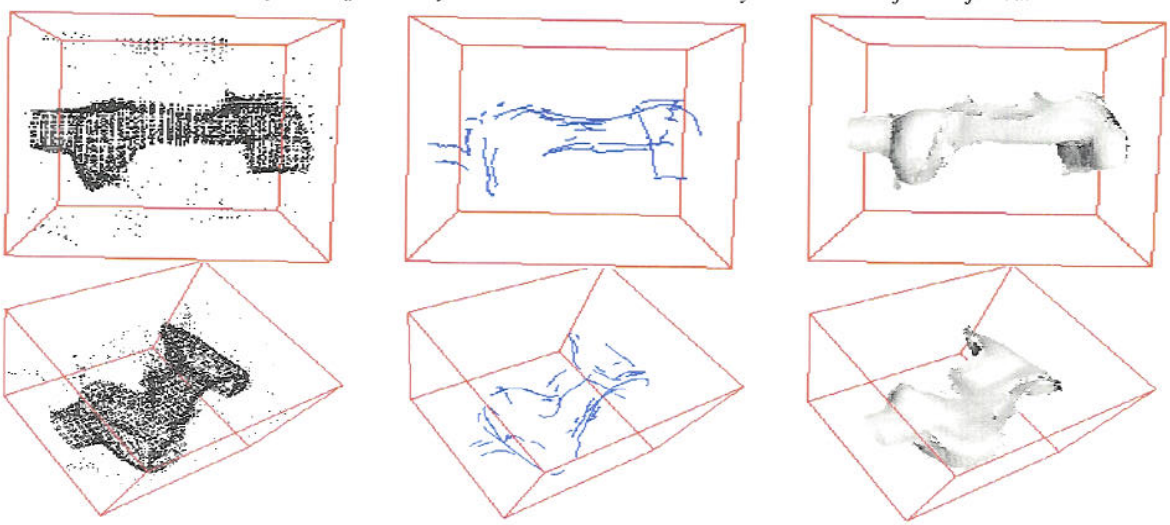

Figure 19 Mod-4. Two views of the original noisy data, the extremal discontinuity curves and surfaces inferred.

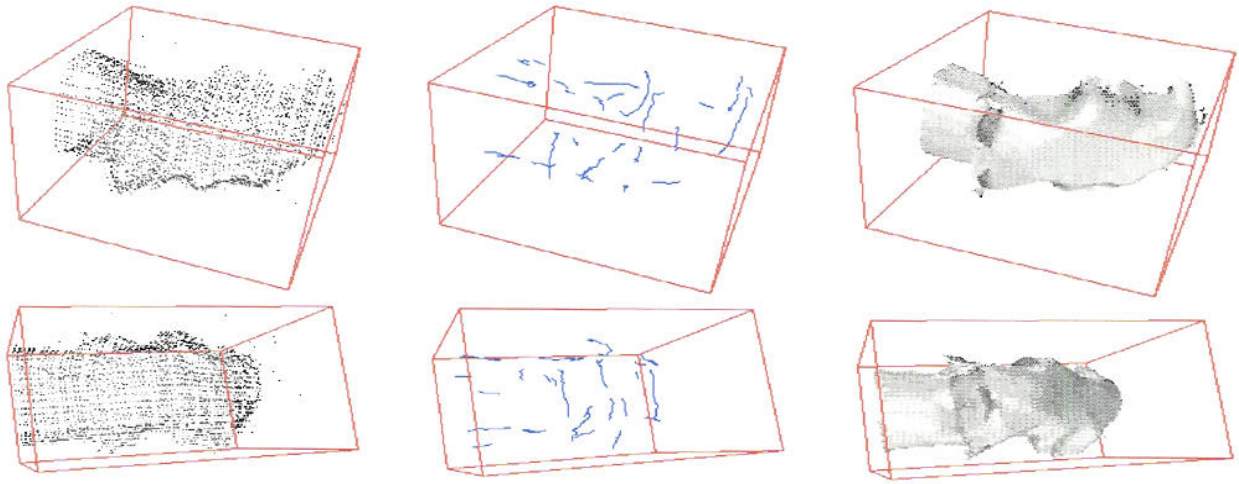

Figure 20 Inlay-4. Two views of the original noisy data, the extremal discontinuity curves and surfaces inferred.

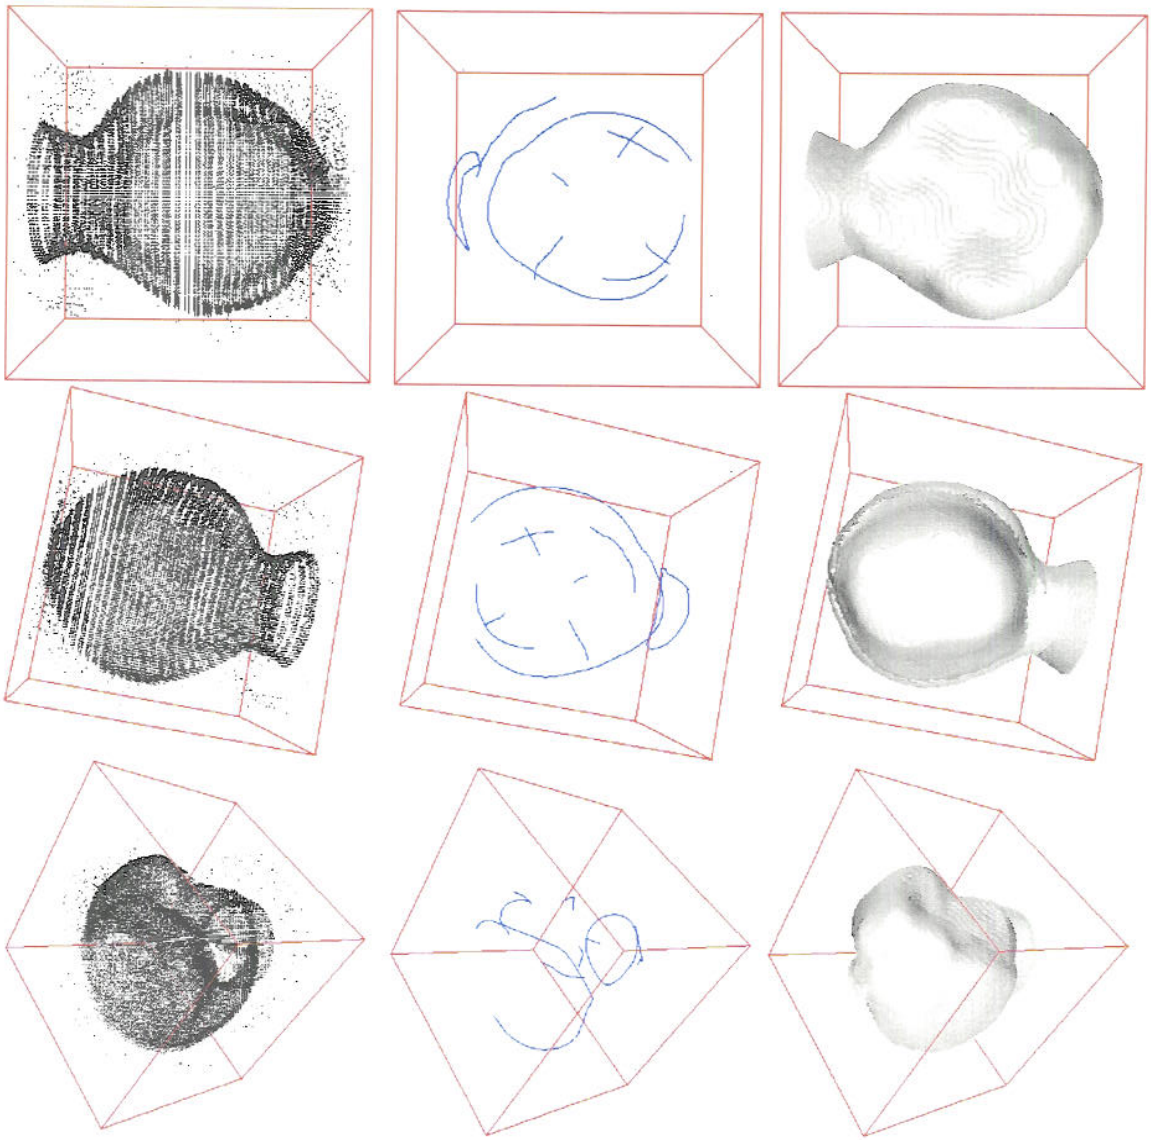

Figure 21 Crown-24. Top, bottom and side views of the original noisy data, the extremal curves (representing surface orientation discontinuities) and extremal surfaces extracted using voting.