Using Markov Random Fields and Algebraic Geometry

to Extract 3D Symmetry Properties

Yunfeng Sui
Dept. of Elec. and Comp. Engineering
University of North Carolina at Charlotte
Charlotte, NC 28223

Andrew R. Willis
Dept. of Elec. and Comp. Engineering
University of North Carolina at Charlotte
Charlotte, NC 28223

Abstract

In this paper, we present a new technique for solving the difficult problem of estimating the axis of symmetry for axially-symmetric surfaces. Accurate solutions to this problem are important in archaeology for systems that seek to reconstruct pottery vessels from measurements of their fragments. Our approach estimates quadratic surfaces at each measured surface point and uses a Markov Random Field superimposed on measured mesh surface to combine surface estimates into surface patches, each of which lies close to a single 3D quadratic surface. For each surface patch, we estimate a quadratic implicit polynomial whose coefficients directly provide an estimate of the unknown axis location and orientation. Using the bootstrap method, distributions for the axis parameters associated with each surface patch are computed. These competing estimates of the global axis are combined using Maximum Likelihood Estimation (MLE) to provide the final estimate. Our approach differs from past approaches by combining estimates derived from large surface regions that include many measurements instead of combining many local (often pointwise) estimates of the surface to determine the global estimate. Estimates from these large regions are more robust to noise and have sufficient data to provide statistical estimates of the uncertainty in the computed estimates. Major benefits of this technique include: (i) a pre-segmentation of the surface to identify noisy, i.e., a-symmetric, regions and (ii) a Bayesian method for merging estimates that makes the axis estimate robust to outliers and highly-variable estimates.

1 Introduction

In recent years there has been an increase in research pertaining to the automatic estimation of shape models from 3D surface data. However, in the estimation of axially symmetric shapes, the methods proposed have not been shown to be effective for small surface patches. But these cases are important in practice and challenging in concept. Some initial work on the subject was proposed in [1] where the authors apply concepts from algebraic geometry to develop a linear algorithm for estimating the axis of symmetry. This method has the benefit of providing a quick and reasonable estimate. Yet, in the interest of preserving linear computational complexity, the authors do not enforce the Plücker relation which guarantees that the solution is a valid Euclidean line. Additionally, by estimating the axis as a line intersected by locally estimated surface normals, the method cannot make use of the fact that the measured data lies on a continuous axially symmetric surface patch. Hence, the estimation accuracy proposed in [1] is not sufficient for our applications. In [2], a method of axis estimation is presented based on the fact that for a surface of revolution, maximal spheres tangent to the surface will have centers on the axis of symmetry. This method differs from ours since the authors are estimating osculating spheres for each data point/normal pair to obtain an estimate of the axis of symmetry. The centers of these spheres depend upon the principal curvature of the surface parallel which passes through each of the point/normal pairs. The authors add robustness to their estimator by detecting outliers in a

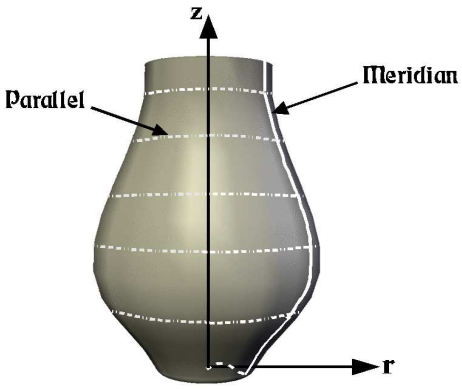

Figure 1: Geometry of a surface of revolution

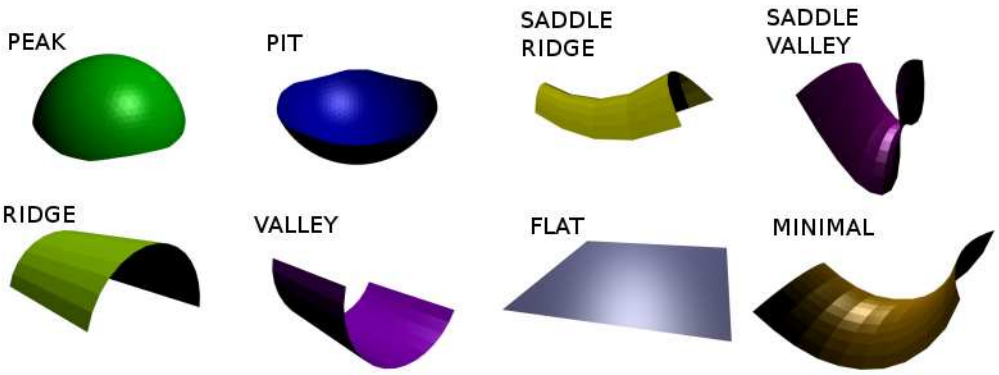

Figure 2: Eight different quadratic surface-types.

weighted iterative least-squares framework. Having computed their axis estimate, the authors then use this estimate to compute the profile curve using a cubic spline model fit to the sherd data. In [3], axially symmetric surfaces are fit to all surface data/normal pairs simultaneously using a weighted iterative least-squares fitting method. While this model incorporates information from both the meridians and the parallels of the surface of revolution and does not require use of any local operators, the algebraic curve fit tends to smooth out sharp curves and ridges which may provide important information regarding the central sherd axis. Sherd classification based on qualitative (e.g., global shape) of profiles with human-driven pre-processing is being developed at the Technical University of Vienna [4, 5].

The proposed approach is completely new and is not based on local estimates as in [1, 2] nor global estimates as in [3]. Instead, our approach seeks to find surface patches from the measured 3D data that lie very close to a common quadratic surface. To do so, the surface is segmented into distinct surface-type patches that are each associated with a single quadratic surface. We then proceed to fit algebraic surfaces to each of the quadratic patches and then use the coefficients of the fit to quickly and accurately compute the global fragment axis. The resulting estimates are more stable than those provided by local methods [1, 2] since we are fitting simple surfaces to large groups of measured 3D data points. Our model also provides a high-fidelity representation of the fragment shape by splitting it into simple quadratic-shaped patches, i.e., we are able to preserve ridges which are smoothed out in a global representation such as that in [3].

2 Archaeological Context

In antiquity, pots and vases were mass-produced by civilizations for use in a variety of contexts. Hence, pieces of pots and vases are prevalent artifacts uncovered in many archaeological excavation sites. By studying these fragments, archaeologists obtain great amounts of information about ancient civilizations. This paper provides a highly accurate solution to the difficult problem of extracting a geometric model of the unknown pot structure in the region associated with 3D measurements obtained from a sherd. The model extracted may be used in a variety of applications including shape-based searching of 3D sherd databases; sherd classification; and pot reconstruction [6, 7, 8].

	$\kappa_G>0$	$\kappa_G=0$	$\kappa_G<0$
$\kappa_H<0$	*Peak* $l=6$	*Ridge* $l=7$	*SaddleRidge* $l=8$
$\kappa_H=0$	impossible	*Flat* $l=4$	*MinimalSurface* $l=5$
$\kappa_H>0$	*Pit* $l=1$	*Valley* $l=2$	*SaddleValley* $l=3$

Table 1: Surface type labels as determined by the observed Gaussian and mean curvature.

3 Approach

Our approach to symmetry estimation involves three steps:

1. Segment the measured surface into quadratic surface patches. Small surface patches, i.e., patches that include less than 10% of the overall fragment measurement data, are removed (see §4).

2. A 3D quadratic implicit polynomial is fit to each of the remaining surface patches and polynomial coefficients are used to compute the axis of symmetry of the resulting algebraic surface (see §5).

3. Distributions for each axis estimate from (2) are obtained using the bootstrap method and the Maximum

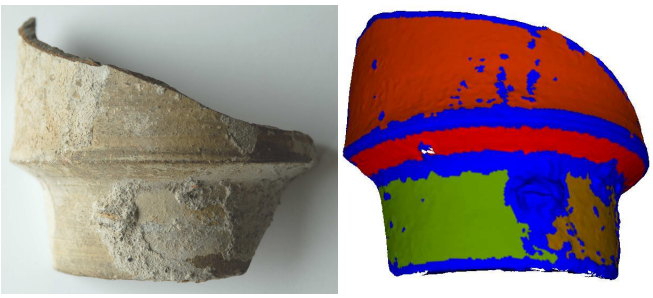

(a) (b)

Figure 3: (a) shows an archaeological pot fragment excavated in Petra, Jordan. (b) is the surface segmentation of the fragment in (a) using the proposed method (see §4). In (b) four disjoint quadratic surface patches are shown: orange, red, green, and brown. These regions are delimited by a blue surface patch which denotes surface points that are not used for axis estimation. Note that the method is able to identify and eliminate regions where calcified deposits have formed on the archaeological fragment over time. These surface regions, if not properly detected, can adversely impact estimates of the pot axis of symmetry.

Likelihood Estimate (MLE) of the fragment axis is computed (see §6).

Results using this approach and a comparison with axis estimation techniques from [1] are presented in §7.

4 MAP-MRF Surface Segmentation

Our goal is to estimate labels for each mesh vertex that indicates the local shape of the surface in the vicinity of the vertex as modeled by a quadratic surface. There are eight potential quadratic surface-types as shown in Figure (2) and the mathematical variables that we use to distinguish between these types is the surface mean curvature, κ_H, and Gaussian curvature, κ_G, measured at each mesh vertex. For computation, we assign each of these surface-types a label, $l \in \{1,2,\ldots,8\}$, as indicated in Table (1). Our MRF labeling problem then seeks to compute the best joint labeling for all of the MRF sites which provides a surface-type classification for each mesh vertex as having come from one of the eight surface-types in Figure (2). Contiguous collections of mesh vertices that share the same label are assumed to lie on similar (possibly the same) quadratic surfaces. In this way, the computed site labels provide a segmentation of the measured surface into a set of M distinct quadratic surface patches.

As others have in the past [9, 10], we impose a Markov Random Field (MRF) model on the measured surface mesh with N mesh vertices. Each vertex, $\mathbf{v}_i = (x,y,z)^t$, of the surface mesh is associated with a MRF site $\mathbf{x}_i$, a site label whose values we seek to estimate. We refer to the collection of all of the sites as $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ and the collection of all site *indices* as $S = \{1,2,\ldots,N\}$.

Cliques are collections of site indices and clique potential functions are defined over the cliques to describe stochastic dependencies between the clique labels. Our MRF cliques are defined in terms of a single-site *vertex clique*,$\mathscr{C}_1$, that models the probability of the data given a site labeling and a *1-ring clique*, $\mathscr{C}_{1-ring}$, that models the stochastic dependency between a site label and the labels of those sites in the *neighborhood* of the site. No edge clique is used since we care about the whole neighborhood instead of a particular one. For site i, the *neighborhood*, N_i, of that site is the collection of all mesh vertex *indices* that share a mesh edge with the site of interest. We can then refer to the set of sites within the neighborhood of site i as the set of sites with indices N_i which we specify notationally as $\mathbf{x}_{N_i}$, i.e., $\mathbf{x}_{N_i} = \{\mathbf{x}_j | j \in N_i\}$.

By the Markov property:

$$P(\mathbf{x}_i | \mathbf{x}_{S-\{i\}}) = P(\mathbf{x}_i | \mathbf{x}_{N_i})$$

where $S - \{i\}$ denotes the set of all site indices excluding index i. The joint probability of all the site variables may then be written as the Gibbs distribution (1)

$$P(\mathbf{X}) = \frac{1}{Z} e^{-U(\mathbf{X})} \tag{1}$$

where Z is referred to as the "partition function" and serves to normalize the exponential distribution such that it is a probability density function and $U(\mathbf{X})$ is referred to as the energy of the MRF. This energy function is specified as the sum of clique potential functions that are defined on each our two cliques as shown in (2).

$$U(\mathbf{X}|\alpha_1, \alpha_{1-ring}) = \alpha_1 \sum_{i \in \mathscr{C}_1} V_{\mathscr{C}_1}(\mathbf{x}_i) + \alpha_{1-ring} \sum_{i_k \in \mathscr{C}_{1-ring}} V_{\mathscr{C}_{1-ring}}(\mathbf{x}_i, \mathbf{x}_{N_i}) \tag{2}$$

For estimation, the only term which depends upon the label values in (1) is the energy function $U(\mathbf{X})$. Hence, estimation of the field labels is accomplished by finding the values of the site variables, $\mathbf{X}$, that maximize the probability, or equivalently, minimize the energy term.

We characterize the quadratic surface shape in the vicinity of each mesh vertex in terms of the mean and Gaussian curvatures. Estimates of these values are obtained at each mesh vertex using the technique specified in [11] and produce a pair of scalar values for each site, κ_H and κ_G, which are collected into a data vector. The data vector for site i is referred to as $\mathbf{d}_i = (\kappa_H, \kappa_G)^t$ and the collection of all the data is $\mathbf{D} = \{\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_N\}$.

	κ_G>0	κ_G=0	κ_G<0
κ_H<0	*Peak* $p_7(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\tau_G}{\kappa_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^6}\right)$	*Ridge* $p_8(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\kappa_G}{\tau_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^6}\right)$	*SaddleRidge* $p_9(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\tau_G}{\kappa_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^2}\right)$
κ_H=0	(*none*)	*Flat* $p_5(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\kappa_G}{\tau_G})^4}\right)\left(\frac{1}{1+(\frac{\kappa_H}{\tau_H})^4}\right)$	*MinimalSurface* $p_6(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\tau_G}{\kappa_G})^4}\right)\left(\frac{1}{1+(\frac{\kappa_H}{\tau_H})^4}\right)$
κ_H>0	*Pit* $p_1(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\tau_G}{\kappa_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^6}\right)$	*Valley* $p_2(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\kappa_G}{\tau_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^6}\right)$	*SaddleValley* $p_3(\mathbf{d}_i\|\tau)=\left(\frac{1}{1+(\frac{\tau_G}{\kappa_G})^4}\right)\left(\frac{1}{1+(\frac{\tau_H}{\kappa_H})^2}\right)$

Table 2: Design functions for $V_{\mathscr{C}_1}$, the vertex clique potential function. Note there is symmetry in the stated functions, i.e., $p_1 = p_7$, $p_2 = p_8$, and $p_3 = p_9$.

The interaction parameters of the MRF, $(\alpha_1, \alpha_{1-ring})$, determine the relative importance of each clique potential in the overall energy. To make our approach applicable to a general data input we normalize the observed curvatures. We choose a normalization which translates the object such that the centroid of the data lies at the origin (3) and subsequently scales the object such that vertex position variations, on average, lie in the vicinity of the 3D unit sphere centered on the origin (4).

$$\mathbf{v}_i^{'} = \mathbf{v}_i - \overline{\mathbf{v}}, \qquad \overline{\mathbf{v}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{v}_i \tag{3}$$

$$\tilde{\mathbf{v}}_i = \frac{1}{\frac{1}{3}\sqrt{trace(\Sigma_{\mathbf{v}})}}\mathbf{v}_i^{'}, \qquad \Sigma_{\mathbf{v}} = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{v}_i^{'})(\mathbf{v}_i^{'})^t \tag{4}$$

This normalization makes our approach invariant to homogeneous scaling. We also note that the approach is Euclidean invariant since both Gaussian and mean curvature are intrinsic surface parameters. MRF clique potentials are designed for this normalized space and are constructed such that the range of each potential function is bounded on the interval $[0,1]$. This allows predictable and relatively easy control of the segmentation behavior by manipulating the interaction parameters alone. It also facilitates learning and design of sufficiently universal potential functions.

Using the computed data values and the clique potentials we can then write down the MRF as a Bayesian estimation problem (5):

$$p(\mathbf{X}|\mathbf{D}) = \frac{p(\mathbf{D}|\mathbf{X})p(\mathbf{X})}{p(\mathbf{D})} \tag{5}$$

where $p(\mathbf{D})$ is a constant, $p(\mathbf{D}|\mathbf{X})$ is the likelihood of the observed data given a labeling, and $p(\mathbf{X})$ is the prior probability for that labeling. This implies that a factorization of the global distribution is possible that separates the energy function from (2) into two parts: (i) clique potentials that express the likelihood of the data given a labeling (functions that depend on both the data and the label values of sites in the clique), and (ii) clique potentials that encode our prior knowledge about the problem and serve constrain how labels within cliques are related (functions that depend only on the label values of sites in the clique). The sum of those clique potentials from (i) form the *likelihood energy*, $U(\mathbf{D}|\mathbf{X})$, and the sum of those clique potentials from (ii) form the *prior energy*, $U(\mathbf{X})$. The sum of these energies is then defined as the *posterior energy* as in (6).

$$U(\mathbf{X}|\mathbf{D}) = U(\mathbf{D}|\mathbf{X}) + U(\mathbf{X}) \tag{6}$$

The MRF Maximum A-Posteriori (MAP-MRF) estimate of the surface-type labels is then equivalent to minimizing the posterior energy function $U(\mathbf{X}|\mathbf{D})$ [12]. Explanation of the MRF-based segmentation algorithm is specified in the following two sections: §4.1 details the clique potential functions that make-up the likelihood and prior energies and §4.2 details our technique for minimizing the posterior energy to find the MAP-MRF solution to the posed labeling problem.

4.1 Clique Potential Functions

Clique potential functions define stochastic dependencies between the site variables and the data (the likelihood energy) and serve to constrain how labels may vary with respect to the labels of other sites within the clique (the prior energy). For segmentation we hand-specify clique potential functions for the vertex clique and the 1-ring clique. A surface segmentation then correspond to finding the MAP-MRF estimate for the surface-type labeling which maximizes the posterior distribution or equivalently, minimizes the total energy from (6). The labeling provides a segmentation of the measured surface into distinct surface-types as shown in Figure (2).

The definition of our clique potentials are similar in concept to that of the Ising model [13] which consists of two energy terms: (1) the data potential, i.e., energy due to those site labels that differ from the label suggested by the data, and (2) a homogeneity constraint that seeks to constrain neighboring values to have the same label. Here $V_{\mathscr{C}_1}$,

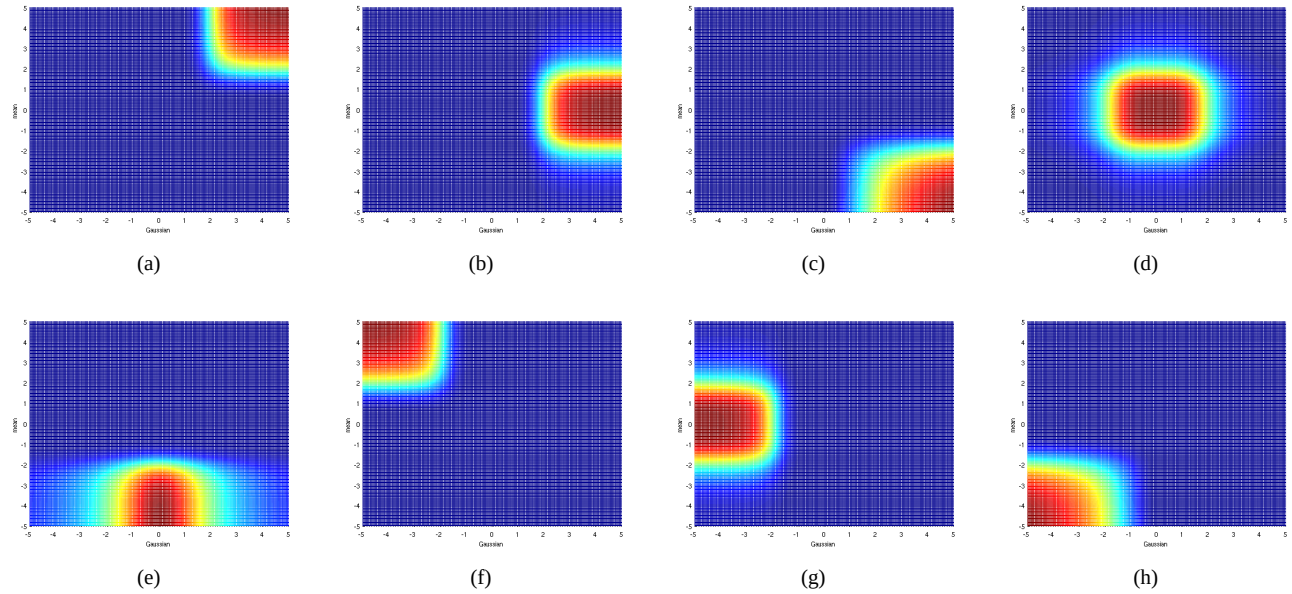

Figure 4: Graphical view of the design functions $p_1, p_2, \ldots, p_8$ from Table (2). Plots are ordered as follows (a) p_1 (b) p_2 (c) p_3 (d) p_5(e) p_6 (f) p_7 (g) p_8 (h) p_9. Blue denotes regions of the plot close to zero and red denotes region of the plot close to 1.

the vertex clique potential function, is the data potential and $V_{\mathscr{C}_{1-ring}}$, the 1-ring clique potential, enforces the homogeneity constraint. Our vertex clique potential is actually composed of eight *design functions* as shown in Table (2). Each function in Table (2) is a two-dimensional Butterworth filter defined on the (κ_G, κ_H) plane that may be loosely interpreted as a hand-specified class conditional density for each of the eight surface types of interest. As with any two-dimensional Butterworth filter, these functions depend upon two values $\tau = (\tau_G, \tau_H)^t$ that determine the cutoff locations of the filters. We use a common pair of constants, $\tau = (1,1)^t$, for all of the shown results and believe these to work well given the normalization steps suggested in equations (3,4). These vertex potential function is then simply:

$$V_{\mathscr{C}_1}(\mathbf{d}_i|\tau, \mathbf{x}_i = l) = 1 - p_l(\mathbf{d}_i|\tau)$$

Where $p_l(\mathbf{d}_i|\tau)$ is value of the l^{th} design function as specified in Table (2) and changes based on the assumed label value for the site, l. The vertex clique potential, $V_{\mathscr{C}_1}$, adds energy to the MRF when the site label differs from the label suggested by the estimated mesh surface curvatures [11].

The 1-ring clique potential seeks to enforce neighboring site values to be homogeneous. Hence, we introduce the following clique potential

$$V_{\mathscr{C}_{1-ring}}(\mathbf{x}_i, \mathbf{x}_j | j \in N_i) = \frac{1}{card(N_i)} \sum_{(i,j)\, pairs} (1 - \delta(\mathbf{x}_i - \mathbf{x}_j))$$

where $card(N_i)$ denotes the cardinality of the neighbor set of site i and $\delta(\mathbf{x})$ is the delta function, i.e., $\delta(\mathbf{x}) = 1$ iff $\mathbf{x} = 0$ (this occurs is the label values for the site pair are equal, i.e., $\mathbf{x}_i = \mathbf{x}_j = l_0$, otherwise $\delta(\mathbf{x}) = 0$.

Use of these clique potentials together for all sites, along with the applied interaction coefficients, $\alpha_1 = \alpha_{1-ring} = 1$, completes the definition of the MRF we apply to segment measured 3D surfaces. Segmentation results are shown in Figure (5a,f,k).

4.2 Optimization

The joint pdf (5) has an energy (6) defined on all possible surface-type labellings of the mesh vertices. Optimization corresponds to maximizing the joint pdf (5) by brute force evaluation. We proceed by computing the posterior energy for each site for each potential surface-type label. The site then takes on the value of the label with minimum energy. This proceeds until the optimization is interrupted or no label changes are observed, i.e., a local minima of the energy function is determined.

5 Extracting Symmetries using Algebraic Geometry

Our segmentation identifies collections of surface measurements which are assumed to lie on a common underlying

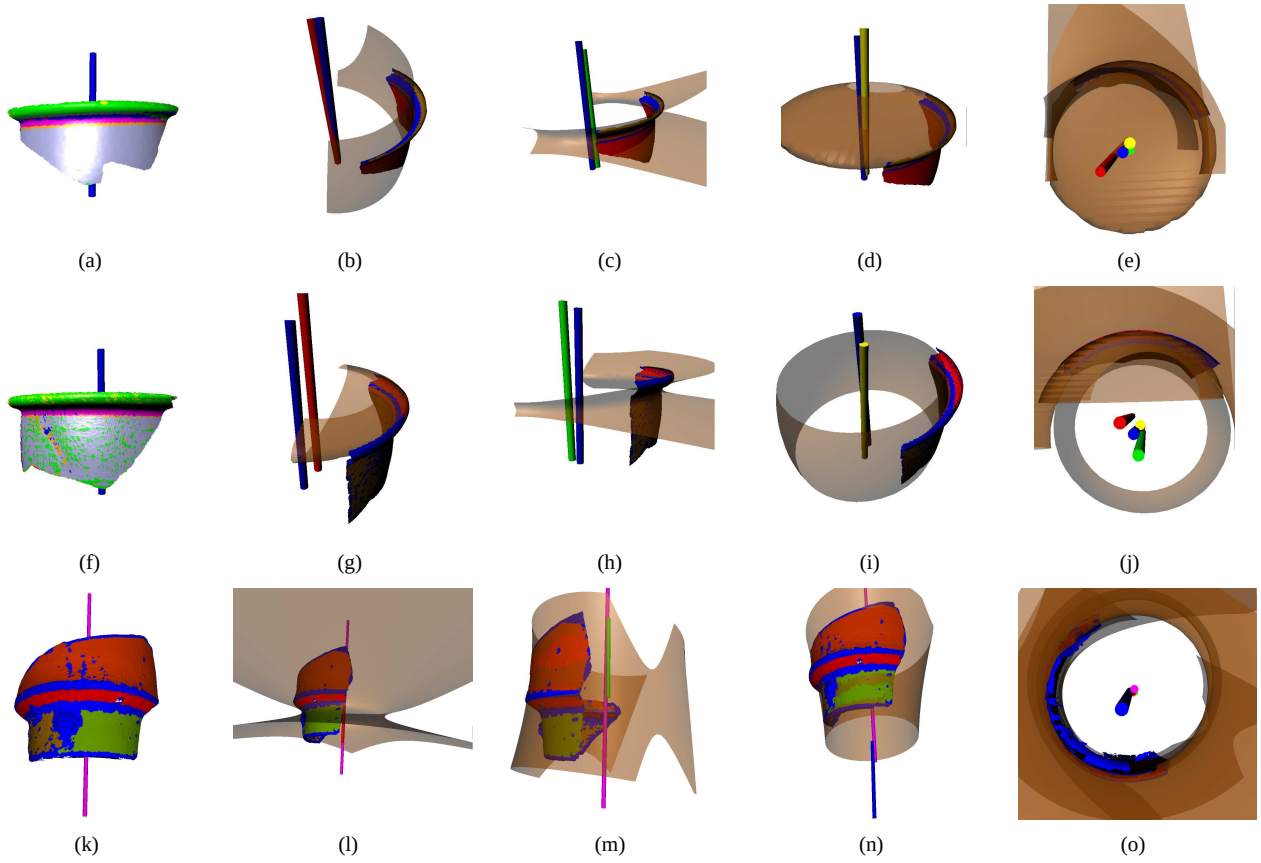

Figure 5: Estimation results for three fragments. The top and middle rows show results for fragments from a contemporary pot broken for the purpose of experimentation. The bottom row shows results for an archaeological pot fragment excavated from Petra, Jordan in 2002. (a,f,j) are the MRF segmentation results; different colors denote distinct quadratic surface patches. (b-d),(g-i),(k-m) show the three largest surface patches and the associated quadratic fits to each surface patch superimposed as a transparent surface and the estimated axis for each quadratic surface as a colored cylinder. The right-hand column, (e,j,n), shows the final axis estimates in blue, the quadratic surfaces (transparent brown), and individual and global fragment axis estimates together as red, green, and yellow cylinders.

quadratic surface. We now fit a quadratic implicit polynomial to the collection of 3D (x,y,z) points within each surface patch using the technique described in [14].

Having successfully estimated the central axis of each surface patch we now wish to fit quadratic functions to each of the surface patches. Symmetry properties may be extracted from each of these patches by expressing the quadratic in matrix-vector form.

Let a 3D quadratic function be specified as follows :

$$f(x,y,z) = ax^2 + bxy + cxz + dy^2 + eyz + fz^2 + gx + hy + kz + m = 0$$

The same equation may be rewritten as a matrix-vector equation where the vector $\mathbf{p} = (x,y,z)^t$

$$\mathbf{p}^t\mathbf{Q}\mathbf{p} + \mathbf{L}\mathbf{p} + m = 0, \quad \mathbf{Q} = \begin{bmatrix} a & \frac{b}{2} & \frac{c}{2} \\ \frac{b}{2} & d & \frac{e}{2} \\ \frac{c}{2} & \frac{e}{2} & f \end{bmatrix} \quad \mathbf{L} = \begin{bmatrix} g & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & k \end{bmatrix}$$

Extracting the symmetry properties of the fit quadratic surface patches requires the following steps:

1. Compute the rotation which will diagonalize the matrix of quadratic terms of $\mathbf{Q}$. This provides the rotation matrix, $\mathbf{R}_i$, that aligns the axes of symmetry of the quadratic with the coordinate system axes.

2. Complete the square of the resulting quadratic form. This provides the location of the quadratic form, $\mathbf{m}_i$, which, for closed surfaces such as an ellipse, is the center of the ellipse.

Completion of these steps provides us with a full 6-parameter rigid transformation that defines the location and direction of the axes for each of the fit quadratic surfaces. Each quadratic surface will have 3 axes and the curvature of

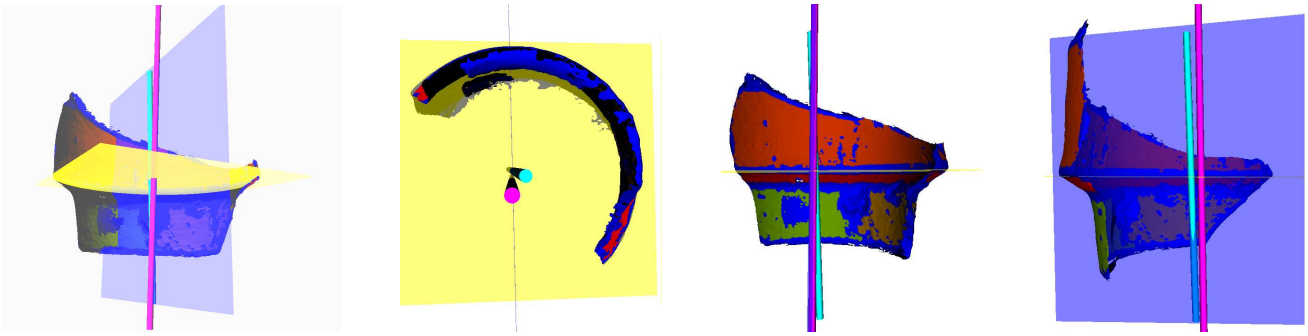

(a) (b) (c) (d)

Figure 6: (a-d) shows four views of the final axis estimates for the drinking cup fragment shown in Figure 3(a). The estimate obtained from the proposed method is shown in pink and the axis obtained using method [1] is shown in cyan. As the true axis of the fragment is unknown to us, one cannot say which is more accurate to ground truth. However, a yellow plane visible in (a) and (b) has been manually positioned in the scene. It approximates the plane upon which the cup would sit if placed on a table. The blue plane is perpendicular to the yellow plane and serves to highlight the subtle angular differences between the two estimates. The top-down view in (b) shows the subtle (yet important) differences between the two estimates location and the side views in (c) and (d) highlight the angular differences between the two axis estimates. While these subtle differences may not appear significant, such discrepancies have very significant impacts when attempting to match and classify surfaces.

the surface in the direction of these axes is easily obtained from the quadratic equation in standard form.

For a surface of revolution, the cross section of the surface at any height is a circle. Hence, the curvature values in the direction of two coordinate axes should be nearly equal. To determine which direction is the central axis, we detect the coordinate plane perpendicular to this axis which is taken as the pair of curvature values whose ratio is closest to 1. The remaining coordinate axis, referred to as $\mathbf{e}_i^{'}$, is then considered to be pointing the direction of the transformed surface meridian. We obtain the direction of the measured surface meridian by applying the inverse transform on $\mathbf{e}_i^{'}$, i.e., $\mathbf{e}_i = \mathbf{R}_i^t \mathbf{e}_i^{'}$. We can now define the axis of axial symmetry for the estimated implicit surface as the 3D line passing through the location, $\mathbf{m}_i$, having direction $\mathbf{e}_i$. For each quadratic surface patch, the axis is estimated as a six-parameter vector $\mathbf{l} = (\mathbf{m}_i, \mathbf{e}_i)$.

6 Combining Axis Estimates

For a given fragment there will typically be 3-6 quadratic surface patches. Surface measurements associated with each of these surface patches provide competing estimates of the global fragment axis. The matter is further complicated by the fact that noise, i.e., perturbations of the surface that deviate from symmetry about the true axis, is spatially varying and arises from numerous phenomenon such as the calcified deposits shown in Figure 3(a) and other deleterious effects on the objects such as chipping, cracking, and erosion. As a result, the variability of each axis estimate is unknown to us and simply averaging the estimates can provide poor results since some surface patches may provide highly unstable estimates.

Fortunately, extracted surface patches will include at least 10% of the overall measurement data, which ensures that each patch will include a large number of point measurements. For our experiments on fragments ~10cm in height, surface patches typically included ~8k-60k when scanned by a Konica-Minolta Vivid 910 scanner. This abundance of measurements within each surface patch enables us to use non-parametric techniques to estimate the variability in the estimate, specifically the bootstrap method.

The bootstrap method is applied by resampling (with replacement) the data points within each surface patch. For each resampling, a quadratic polynomial is fit providing a new axis estimate. Since the surface fitting procedure and axis estimation may be computed explicitly, we can quickly compute many different estimates. The resulting group of estimates are assumed to have a Gaussian distribution and we proceed by estimating the mean and covariance of the bootstrapped parameter estimates to provide a probability distribution on the axis parameters for each estimated quadratic surface. We denote the bootstrap distribution for the i^{th} surface patch as $p(\mathbf{l}_i)$. The global pot axis estimate is then taken as the MLE mean vector of the axis parameter distribution obtained by merging the estimated distributions, $p(\mathbf{l}_i)$, into a single distribution as defined in many pattern recognition texts [15]. This Bayesian approach for computing the global axis addresses the issues discussed at the beginning of this section and weights highly-variable estimates much less than those which exhibit stability.

7 Results

Results obtained from the proposed axis estimation technique are shown for three different ceramic fragments in Figure 5 (see caption for details). Figure 6 compares axis estimates obtained for the fragment from Figure 3 using our technique and that from [1]. Unfortunately, ground truth is not known for this fragment and nothing can be concretely said regarding the relative accuracies of the two methods. Yet the subtle differences in the pose of the axes is clear and a high degree of precision is required for such estimates in the previously mentioned archaeological contexts.

8 Conclusions

This article has introduced a new algorithm for the estimating the surface and axis of symmetry of an axially symmetric object from 3D scan data. The estimation procedure fits algebraic surfaces to large regions of measured the measured surface that have been identified to have come from a single quadratic surface. Using these large regions for estimation produces stable and accurate quadratic surface estimates. The coefficients of each fit algebraic surface directly provide the axes of symmetry for the object and are combined in an MLE framework to provide a global axis estimate for axially symmetric objects. The results section shows estimates provided by algorithm for axially symmetric ceramic fragments which demonstrate these abilities. The major innovation of this approach is the ability to incorporate large numbers of data points that often span large regions of the surface to generate a single axis estimate. Current investigations seek to extend this technique for estimating symmetry properties in more general surfaces such as the medial axis of a generalized cylinder or perhaps used for quickly estimating the medial axis of complex objects.

9 Acknowledgements

This work is supported by the National Science Foundation under NSF grant #IIS-0205477 and by the University of North Carolina at Charlotte (UNCC) Faculty Research Grant (FRG) program.

References

[1] H. Pottmann, M. Peternell, and B. Ravani, "An introduction to line geometry with applications," *Computer-Aided Design*, vol. 31, pp. 3–16, 1999.

[2] D. Mumford and Y. Cao, "Geometric structure estimation of axially symmetric pots from small fragments," in *Proc. of Int. Conf. on Signal Processing, Pattern Recognition, and Applications*, pp. 92–97, 2002.

[3] A. Willis, D. Cooper, *et al.*, "Accurately Estimating Sherd 3D Surface Geometry with Application to Pot Reconstruction," in *CVPR Workshop : ACVA*, June 2003.

[4] S. Tosovic and R. Sablatnig, "3D modeling of archaeological vessels using shape from silhouette," in *Third International Conference on 3D Digital Imaging and Modeling*, (Quebec, Canada), June 2001.

[5] M. Kampel and R. Sablatnig, "Automated segmentation of archaeological profiles for classification," in *ICPR*, vol. I, pp. pp. 57–60, 2002.

[6] A. Gilboa, A. Karasik, I. Sharon, and U. Smilansky, "Towards computerized typology and classification of ceramics," *J. Archaeological Sci.*, vol. 31, no. 6, pp. 681–694, 2004.

[7] I. Marie and H. Qasrawi, "Virtual assembly of pottery fragments using moire surface profile measurements," *J. Archaeological Sci.*, vol. 32, no. 10, pp. 1527–1533, 2005.

[8] M. Kampel and R. Sablatnig, "Virtual reconstruction of broken and unbroken pottery," in *3D Digital image and Modeling*, pp. 318–325, 2003.

[9] A. Willis, J. Speicher, and D. Cooper, "Rapid prototyping 3d objects from scanned measurement data," *Journal of Image and Vision Computing*, vol. 25, no. 7, pp. 1174–1184, 2007.

[10] S. Li, "Towards 3d vision from range images: An optimization framework and parallel networks," *CVGIP: Image Understanding*, vol. 55, no. 3, pp. 231–260, 1992.

[11] A. McIvor and A. Valkenburg, "A comparison of local surface geometry estimation method," *Machine Vision and Applications*, vol. 10, pp. 17–26, 1997.

[12] S. Z. Li, *Markov Random Field Modeling in Image Analysis*. Springer, 2001.

[13] E. Ising, "Beitrag zur theorie des ferromagnetismus," *Zeitschr.. f. Physik*, pp. 253–258, 1925.

[14] G. Taubin, "Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 11, pp. 1115–1138, 1991.

[15] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley-Interscience Publication, 2000.