

Curve Skeleton Extraction from Incomplete Point Cloud

Andrea Tagliasacchi
School of Computing Science
Simon Fraser University

Hao Zhang
School of Computing Science
Simon Fraser University

Daniel Cohen-Or
School of Computer Science
Tel-Aviv University




Figure 1: From a raw scan with significant missing data, our algorithm extracts a complete curve skeleton, shown in the middle. RBF surface reconstruction infers the wrong surface topology from the original input, while succeeding with the assistance of the curve skeleton.

Abstract

We present an algorithm for curve skeleton extraction from imperfect point clouds where large portions of the data may be missing. Our construction is primarily based on a novel notion of generalized *rotational symmetry axis* (ROSA) of an oriented point set. Specifically, given a subset S of oriented points, we introduce a variational definition for an oriented point that is most rotationally symmetric with respect to S . Our formulation effectively utilizes normal information to compensate for the missing data and leads to robust curve skeleton computation over regions of a shape that are generally cylindrical. We present an iterative algorithm via planar cuts to compute the ROSA of a point cloud. This is complemented by special handling of non-cylindrical joint regions to obtain a centered, topologically clean, and complete 1D skeleton. We demonstrate that quality curve skeletons can be extracted from a variety of shapes captured by incomplete point clouds. Finally, we show how our algorithm assists in shape completion under these challenges by developing a skeleton-driven point cloud completion scheme.

Keywords: curve skeleton, incomplete data, rotational symmetry

1 Introduction

The skeleton of a shape, especially an articulated shape such as a human or animal, provides an intuitive and effective abstraction which facilitates shape understanding and manipulation. The best known skeletal representation is Blum’s medial axis [Blum 1967],

which, along with its variants collectively referred to as medial representations [Siddiqi and Pizer 2009], is designed to capture *reflectional* symmetries in a shape [Bouix et al. 2006]. The medial axis of a 3D model is generally a non-manifold containing 2D sheets that are hard to store and manipulate. A 1D curve skeleton, on the other hand, is more useful in practice due to its topological simplicity, leading to computational efficiency and ease of manipulation.

Many algorithms for curve skeleton extraction exist [Cornea et al. 2007], dealing mostly with shapes specified by closed polygonal meshes [Hilaga et al. 2001; Katz and Tal 2003; Chuang et al. 2004; Dey and Sun 2006; Au et al. 2008]. Computing a curve skeleton from point cloud data is possible, e.g., via a deformable blob grown from the “inside” of the input cloud [Sharf et al. 2007a] or relying on the Voronoi diagram of the point set [Ogniewicz et al. 1992]. However, by design, none of the proposed methods so far handle point clouds with large portions of the data missing.

Incomplete data are common as the result of acquisition via laser scanning, due to self occlusion or less than ideal conditions of the surface material. Such problems are more prevalent during real-time capture of moving objects [Mitra et al. 2007; Wand et al. 2007; Pekelny et al. 2008; Sharf et al. 2008]. As the number of camera views is typically limited, the raw point clouds contain large gaping holes and severe under-sampling.

Overview of ROSA In this paper, we present an algorithm for curve skeleton extraction directly from incomplete point cloud data. Our problem setting is unique in that we aim to cope with significant missing data without the aid of a data sequence to complement each other and possibly fill in the gaps. One of the key insights is that instead of considering reflectional symmetry, as for medial axes, a curve skeleton is most appropriately thought of as a generalized *rotational* symmetry axis (ROSA) of a shape. Indeed, a ROSA ought to be a 1D structure. This seemingly obvious observation proves to be an essential step towards robust handling of incomplete point data. The second key is to effectively exploit orientation information to compute ROSA so as to compensate for the missing data.

To handle significant missing data with only a single point cloud




Figure 2: ROSA definition for a set of oriented points in 3D. Left: Optimal direction (red arrow) minimizes sum of angular variations with the surrounding normals. Right: Optimal point (red dot) minimizes sum of projected distances to the normal extensions.




Figure 3: Stability of ROSA point position and orientation with more than half of the points missing.

available, an appropriate shape prior is necessary. A general premise of our approach is that the shapes of interest should be covered by generally *cylindrical* regions except at their joints. This is a reasonable assumption as only such shapes would admit meaningful curve skeletons; a serving plate, baseball cap, or bowling ball, for example, do not belong to this category and they possess no natural curve skeletons. ROSA is designed to skeletonize generally cylindrical regions, even with significant missing data.

To compute ROSA, we take advantage of available point normals and introduce a variational formulation which works on a local subset S of oriented samples in the input point cloud. We define an oriented point $p = (x_p, \mathbf{v}_p)$, called a *ROSA point*, with position x_p and normal \mathbf{v}_p , that is most rotationally symmetric about S . Our definition applies to 3D quantities and it requires that

1. the orientation \mathbf{v}_p minimizes the variance of the angles between \mathbf{v}_p and the normals in S . In other words, \mathbf{v}_p is to make the *same angle* with these normals as much as possible, consistent with the notion of rotational symmetry;
2. the position x_p minimizes the sum of squared distances to the *line extensions of the point normals* in S .

Figure 2 illustrates our definition and Figure 3 shows it at work with significant missing data. One can immediately recognize the difference between our formulations and the use of average normal or centroid — both are sensitive to missing data.

Figure 4 reinforces the point that orientation information can effectively compensate for missing data in curve skeleton extraction. Relying on point orientation for shape inference has been a common practice in several other contexts, such as neighborhood identification via Mahalanobis distance [Amenta and Kil 2004; Lehtinen et al. 2008] and surface reconstruction with the assistance of point normals [Carr et al. 2001; Kazhdan et al. 2006].

Building on the local ROSA formulation, we present an iterative procedure via planar cuts to locate the best subset S for deriving the rotational symmetry axis of the whole point cloud. It is interesting to note that if such procedures were to be applied to a 2D shape, we would arrive at a point on the *symmetry set* of the shape boundary [Giblin and Brassett 1985], of which the medial axis is a subset. This connection is discussed further in Section 3.




Figure 4: Orientation information compensates for missing data. Left: With missing samples from two touch circles and no normals, the computed centroid (green dot) leads to a wrong interpretation of the data. Right: Sample normals reveal two clusters, leading to two ROSA points (red dots). The correct shape can be inferred.

Contributions The main contributions of our work include a novel definition of generalized rotational symmetry axis for a point set that is: (i) designed to model generally cylindrical regions of a shape, and (ii) robust to missing data. We develop a curve skeleton extraction algorithm from incomplete point clouds based on recursive planar cuts and local ROSA construction, as well as a scheme to robustly detect and skeletonize the non-cylindrical joint regions of a shape. The computed curve skeleton is complete, despite a potentially incomplete data source, and it is guaranteed to be a 1D structure with associated correspondence to points in the input cloud. We demonstrate experimentally that quality curve skeletons can be constructed from a variety of shapes under imperfect data conditions. Finally, we show an application of our algorithm for shape completion under significant missing data; see Figure 1.

2 Related work

Skeletons are effective shape abstractions that have been used in various applications including mesh segmentation [Li et al. 2001; Au et al. 2008], animation [Lewis et al. 2000], and shape matching [Hilaga et al. 2001]. We focus on 1D curve skeletons embedded in 3D. For a coverage on medial axis and other higher dimensional medial representations, we point out the recent book by Siddiqi and Pizer [2009]. There are many existing algorithms for curve skeleton extraction and we only mention a subset. For the rest, we refer the reader to the recent survey by Cornea and Min [2007].

The field-based approach to curve skeleton extraction relies on an Euclidean distance field [Malandain and Fernández-Vidal 1998] or an implicit potential field [Chuang et al. 2004] corresponding to the input shape, resulting in a voxelized representation of the internal volume. The skeleton is then computed via volumetric thinning, ridge extraction, or force following along the ridges of a potential field. These methods generally require clear knowledge about the interior of the input shape, making them inapplicable to incomplete data. The recent work of Au et al. [2008] shares some flavors with field-based methods; it skeletonizes a shape by shrinking it using constrained Laplacian smoothing. Excellent results are obtained, but they can only come from watertight meshes, since Laplacian smoothing requires mesh connectivity and a full model is needed to balance the shrinking process so as to obtain a centered skeleton.

Other works which require complete meshes as input include Li et al. [2001] which uses mesh decimation, the segmentation-based approach, e.g., Katz and Tal [2003], and topology-driven methods, e.g., via Reeb graphs [Hilaga et al. 2001; Patane et al. 2008], which extract and link critical points of a function over the mesh surface. The latter rely on geodesic distances, which are difficult to obtain on a point cloud with significant missing data. Geodesic distances also play a key role in the work of Dey and Sun [2006] where they define a medial geodesic function over the medial axis of a shape to skeletonize the medial axis into a 1D structure.




Figure 5: Overview of our algorithm. (a) Input point cloud with a joint (blue). (b) Optimal cutting plane and relevant neighborhood points (blue) anchored at a point cloud sample (red particle). (c) Skeletal cloud after ROSA and joint recovery. (d) After thinning with branch (green) and joint (blue) identification. (e) After re-centering, still a skeletal cloud. (f) The final 1D curve skeleton.

Recently, skeleton extraction from a sequence of deforming meshes [James and Twigg 2005; de Aguiar et al. 2008] has been of interest. These methods utilize vertex correspondence across frames and assumptions on motion, e.g., piecewise rigidity, to extract the skeleton or bones from complete data. In our problem setting, the input is a single point cloud with significant missing data. Shape completion prior to skeletonization, e.g., via RBF [Carr et al. 2001], can be a means to handle incomplete point clouds. However, these surface reconstruction schemes are not designed to cope with significant missing data. In such cases, user intervention is often necessary to clarify object topology [Sharf et al. 2007b].

Curve skeleton extraction directly from point clouds is possible via Voronoi skeletons [Ogniewicz et al. 1992]. However, the robustness of this class of methods relies on specific sampling conditions which are generally un-attainable in practice and certainly far from being fulfilled in our setting. Sharf et al. [2007a] grows a smooth blob from the inside of a point cloud and the growing fronts trace out a curve skeleton. This approach can handle moderate missing data, where properly set tension parameters can prevent the blob from “leaking outside” over incomplete data regions. However, this is difficult to achieve for significant missing data.

There has been a great deal of recent research on symmetry detection and symmetry-aware geometry processing, e.g., [Kazhdan et al. 2004; Mitra et al. 2006; Podolak et al. 2006; Ovsjanikov et al. 2008], to name a few. Most methods deal with reflectional symmetry and when general symmetries, including rotational ones, are considered, the focus has so far been on intrinsic analysis over a shape using geodesic distances. Hence these methods do not apply to our problem. More relevant is the work of Thrun and Wegbreit [2005] which utilizes symmetries for shape completion from 3D range images, but the types of symmetries considered are only elementary ones. To the best of our knowledge, our work is the first to introduce a generalized, local rotational symmetry axis definition to oriented point clouds and apply it to curve skeleton extraction.

We focus on shapes composed of generally cylindrical regions except at their joints. There exist works on fitting closed polygonal meshes using cylinders [Raab et al. 2004], ellipsoids [Simari and Singh 2005; Lu et al. 2007], or general swept volumes [Kim et al. 2003]. Fitting generalized cylinders parametrically over a point cloud seems difficult without first extracting a skeleton [Chuang et al. 2004]. With missing data, this is exactly the problem we address, to which we take a local, non-parametric approach.

3 Overview

Given a single point cloud with normals, we extract a curve skeleton, a compact 1D abstraction of the sampled shape. Point normals can be acquired via photometric stereo [Nehab et al. 2005] or computed from unorganized points via a normal estimation and orientation scheme, e.g., [Hoppe et al. 1992]. Denoising or outlier removal from a raw point set is not a focus of our work. When necessary, we apply the parameterization-free projection operator LOP by Lipman et al [2007] to obtain a clean and evenly distributed point set. Note that LOP makes no attempt at completing missing data.

Our method does not rely on any ambient field, volumetric discretization, sampling assumptions, or an intermediate surface representation. We assume that the input point cloud samples a shape which is composed of generally cylindrical regions, which we call *branch regions*, except at their joints, as shown in Figure 5(a). The branch regions are well described by rotational symmetry and ROSA is designed to extract a generalized, local rotational symmetry axis from such a region to form a skeleton. Joint regions are typically non-cylindrical and require special handling. Note that for illustration purposes only, in this section and the next, we adopt a point cloud example without missing data.

Cutting planes We observe that each point on the rotational symmetry axis of a generally cylindrical shape should correspond to a narrow “band” on the shape which is approximately planar; see Figure 5(b). This motivates the use of planar cuts over the input point cloud to *localize* the search of a ROSA point on the shape’s skeleton. Obviously, not all cutting planes imply desirable rotational symmetries. We search for a best cutting plane and *anchor* the search at each sample point in the input cloud. There are three advantages to anchoring the search. First, the anchor point can seed the search for the relevant set of samples near a cutting plane for ROSA construction; see Section 4.1 for details. Secondly, anchoring of the cutting plane implies a natural correspondence between the point cloud and the computed skeleton. Finally, this approach leads to a simplified search for the best cutting plane.

ROSA construction Instead of optimizing for both orientation and position of a ROSA point at the same time, leading to a higher dimensional search, we decouple the two components and optimize for orientation and then position, each of which is a linear problem




Figure 6: 2D illustration of iterative ROSA construction and connection to medial axis. (a-c) Cutting plane anchored at blue sample goes from p_0 to p_2 and converges. During iteration, the normal of the next cutting plane (red) makes the same angle with those normals (black) at the boundary, corresponding to the current cutting plane. (d) When converging, ROSA point (green) is the intersection between two boundary normals and lies on the medial axis.

and can be solved in closed form. Specifically, through each sample in the point cloud, we find a best cutting plane whose normal minimizes the variance of angles with the normals at a set of relevant points close to the cutting plane; see Figure 5(b) for an illustration of this relevant set. The optimal orientation is found iteratively, as illustrated in 2D in Figure 6(a-c). Once the cutting plane is found, we compute the optimal position of a ROSA point based on that relevant set of oriented points; see Section 4.1.

Connection to medial axis Figure 6(d) reveals that there is a close connection between ROSA construction in 2D and the medial axis. Indeed, if the iteration converges, simple geometric arguments show that the optimal ROSA point is the center of a bi-tangent circle of the shape boundary. This point lies precisely on the medial axis if the bi-tangent circle is inside the shape. Since we do not constrain the circle to be inside, the ROSA points generally belong to the symmetry set of the boundary curve, which is the loci of centers of all bi-tangent circles [Giblin and Brassett 1985]. Note however that since we incorporate point orientations, the set of ROSA points is more restricted than the full symmetry set.

Joint handling and curve skeleton extraction A joint region is generally non-cylindrical and lacks a simple rotational symmetry axis. We exploit *spatial coherence* between the point cloud and skeleton to ensure that points on the skeletal structure provide a smooth connection between the ROSAs of the branches, as shown in Figure 5(c). Since this step does not constrain the structure near joints to be 1D or well centered, we apply thinning and centering in post-processing. The thinning process uses 1D moving least squares (MLS) construction [Lee 2000] which also allows us to differentiate between joints and branches; see Figure 5(d). Points on the resulting skeletal structure are centered according to ROSA within a branch and collapsed to a single center within a joint and then connected to nearby branches. The resulting structure, shown in Figure 5(e), is sufficiently close to being 1D and can be easily converted into a set of curve segments, as shown in Figure 5(f).

4 Curve skeleton extraction via ROSA

We first present ROSA construction in Section 4.1. Section 4.2 describes our joint handling procedure. This is followed by thinning and re-centering of the resulting skeletal cloud and the extraction of the complete 1D skeleton. In the following, we refer to a point from the input cloud as a point cloud sample, a point computed by the ROSA formulation as a ROSA point, and a point which lies on the currently computed skeletal structure as a skeletal sample.

4.1 ROSA and skeletal cloud construction

Cutting planes and relevant neighborhoods Let p_i be a point cloud sample. Let us consider a cutting plane π_i through p_i , with orientation \mathbf{v}_i , and identify a narrow band of point cloud samples within a distance less than δ from π_i . The thickness value δ is applied globally and it is a free parameter set to be 2.5% of the bounding box diagonal of the input point cloud in all examples. Determination of the plane orientation \mathbf{v}_i will be described later.

For a complex shape, the cutting plane may encompass multiple shape parts. Thus we first need to further identify from within the narrow band of points near π_i , a *relevant neighborhood* N_i of point cloud samples, for ROSA construction. In general, the configuration of points in the entire band can be complex. However, we avoid having to solve a full-fledged clustering problem since N_i is anchored at p_i , i.e., $p_i \in N_i$. Note also that while having point positions may only lead to ambiguities when we group around p_i under missing data, point normals can effectively compensate for the missing data in identifying the relevant neighborhood N_i .

Therefore, we utilize Mahalanobis distance, which combines Euclidean and orientation-space information, to derive the relevant neighborhood N_i . We adopt the formulation of Mahalanobis distance from Lehtinen et al. [2008] to compute an orientation-aware distance d_{Mah} between point cloud samples with normals. We use this distance measure and an appropriately chosen threshold ϵ_{Mah} to construct a graph on all the point cloud samples, where there is an edge between p_j and p_k if and only if $d_{\text{Mah}}(p_j, p_k) < \epsilon_{\text{Mah}}$.

The relevant neighborhood N_i at p_i is then extracted by computing the connected component in the Mahalanobis distance based graph defined above: we execute a breadth first search rooted at p_i and recursively add point cloud samples from within the narrow band around π_i . Note that while the set N_i changes with the threshold ϵ_{Mah} , the final computed ROSA point remains quite stable due to the robustness of ROSA definition to missing data. Thus, we choose a relatively aggressive threshold in our experiments.

Optimal cutting plane At sample p_i , we wish to find an optimal cutting plane π_i^* through p_i which best models its local rotational symmetry. Specifically, the normal of π_i^* should be most rotationally symmetric about the point normals in the relevant neighborhood N_i . The corresponding optimization problem is difficult and non-linear, thus we solve it by an iterative approach. We start with an initial orientation \mathbf{v}_i^0 and iteratively update the orientation by solving the following variational problem involving the variance of angles between the plane normal and point normals from the relevant neighborhood associated with the current cutting plane:

$$\mathbf{v}_i^{t+1} = \underset{\mathbf{v} \in \mathbb{R}^3, \|\mathbf{v}\|=1}{\operatorname{argmin}} \operatorname{var} \{ \langle \mathbf{v}, \mathbf{n}(p_j) \rangle : p_j \in N_i^{(t)} \}, t \geq 0, \quad (1)$$

where $N_i^{(t)}$ is the relevant neighborhood for the cutting plane at the t -th iteration and $\mathbf{n}(p_j)$ is the unit normal at point p_j . Problem (1) has a closed form solution as it can be re-written as one which minimizes the quadratic form $\mathbf{v}^T M \mathbf{v}$, subject to $\|\mathbf{v}\| = 1$, where

$$M = \begin{bmatrix} \overline{\mathcal{X}^2} - \overline{\mathcal{X}}^2 & 2\overline{\mathcal{X}\mathcal{Y}} - 2\overline{\mathcal{X}}\overline{\mathcal{Y}} & 2\overline{\mathcal{X}\mathcal{Z}} - 2\overline{\mathcal{X}}\overline{\mathcal{Z}} \\ 2\overline{\mathcal{X}\mathcal{Y}} - 2\overline{\mathcal{X}}\overline{\mathcal{Y}} & \overline{\mathcal{Y}^2} - \overline{\mathcal{Y}}^2 & 2\overline{\mathcal{Y}\mathcal{Z}} - 2\overline{\mathcal{Y}}\overline{\mathcal{Z}} \\ 2\overline{\mathcal{X}\mathcal{Z}} - 2\overline{\mathcal{X}}\overline{\mathcal{Z}} & 2\overline{\mathcal{Y}\mathcal{Z}} - 2\overline{\mathcal{Y}}\overline{\mathcal{Z}} & \overline{\mathcal{Z}^2} - \overline{\mathcal{Z}}^2 \end{bmatrix}.$$

Here \mathcal{X} denotes a random variable for the x -component of the point normals in $N_i^{(t)}$, similarly for \mathcal{Y} and \mathcal{Z} , and $\overline{\mathcal{X}}$ denotes the sample mean, which in our case, is simply an average over the set $N_i^{(t)}$. The quadratic problem can be solved analytically via singular value




Figure 7: Optimal cutting plane orientations at branches (well-behaved) and joints (noisy due to lack of rotational symmetry).

decomposition. In Figure 7, we zoom in and show optimal orientations computed near a branch region; they are well-behaved. However, near a joint, the lack of local rotational symmetry makes the orientations noisy. This issue is addressed in Section 4.2.

The initial direction \mathbf{v}_i^0 is selected randomly among those perpendicular to the normal at p_i , as motivated by the cylindrical shape prior. Experimentally, we have observed fairly fast convergence, with no more than ten iterations required before the plane orientation stabilizes. However, we currently do not have a convergence proof. Problematic local minima are rare and such occurrences can be corrected by enforcing spatial coherence during joint handling.

Skeletal point computation Given an optimal cutting plane π_i^* at p_i , we next compute the corresponding ROSA point r_i^* , the center of local rotational symmetry. The computed ROSA points collectively form the initial skeletal cloud. We again utilize orientation information at the points and solve the following problem to minimize the sum of squared distances from p_i to the normal lines,

$$r_i^* = \underset{r \in \mathbb{R}^3}{\operatorname{argmin}} \sum_{p_j \in N_i^*} \|(r - p_j) \times \mathbf{n}(p_j)\|^2, \quad (2)$$

where N_i^* is the relevant neighborhood for the optimal cutting plane. Problem (2) is again a standard quadratic minimization and has a closed form solution by straightforward differentiation.

4.2 Joint handling and curve skeleton extraction

Enforcing spatial coherence While in branch regions, the computed skeletal cloud resembles a 1D structure, the same does not hold for joints, as shown in Figure 8(a). Indeed, a joint region is generally non-cylindrical and does not possess a meaningful optimal cutting plane. As a result, ROSA points within a joint can be rather scattered. To remedy this problem, we exploit spatial coherence: close-by samples over the underlying shape of the point cloud should correspond to close-by skeletal samples.

To enforce spatial coherence, we apply Laplacian smoothing to the ROSA points where the point connectivity is implied from Mahalanobis proximities between the corresponding point cloud samples. Intuitively, this operation “tightens” the skeletal cloud within joint regions and reveals more clearly the branch connections; see Figure 8(b). In addition, skeletal cloud noise caused by imperfect cutting planes, e.g., due to local minima, can be removed.

Thinning and branch/joint identification While being spatially coherent, the resulting skeletal cloud requires further thinning to converge to a 1D structure. This is especially necessary within branch regions near a joint, as the smoothing step above may distort the linearity of the skeletal cloud therein. We apply 1D MLS construction on the skeletal cloud for thinning. MLS techniques are well known for curve or surface reconstruction and in the 1D case, we repeatedly project skeletal samples onto their corresponding locally best-fitting lines, via principal component analysis (PCA).




Figure 8: Joint handling, a zoomed-in view. (a) After ROSA construction. (b) After joint recovery via smoothing. (c) After thinning. (d) After re-centering and joint collapsing. (e) After re-distribution of skeletal samples. (f) Final 1D curve segments.

Local PCA and well-thinned branch regions give us a simple and robust way to distinguish between branch and joint skeletal samples. Specifically, we examine a standard linearity measure

$$\psi(r_i) = \lambda_i^{(1)} / (\lambda_i^{(1)} + \lambda_i^{(2)} + \lambda_i^{(3)})$$

at skeletal sample r_i , where $\lambda_i^{(j)}$ is the j -th largest eigenvalue from local PCA at r_i . We only apply 1D MLS to r_i when $\psi(r_i) < \epsilon_{\text{MLS}}$, typically indicating that r_i is in a branch. Applying 1D MLS to skeletal samples within a joint is not so meaningful, as the data there do not possess a line-like structure. The tolerance ϵ_{MLS} is set conservatively to 0.4 throughout our experiments. Figures 8(b-c) show the effect of thinning.

Centering and 1D curve extraction Steps taken so far for joint handling may distort the centered-ness of the skeletal cloud. Thus we re-center the skeletal cloud, according to the type of skeletal samples, either within a joint or a branch. The scalar field given by ψ , defined above, on the skeletal samples provides exactly the means to achieve the latter. Re-centering within a branch follows Equation (2). Specifically, a branch skeletal sample r_i is moved to the ROSA point computed for a small neighborhood of point cloud samples which correspond to a small neighborhood of r_i in the skeletal cloud. Within a joint, skeletal samples collapse to a unique *joint center* again according to (2), using all the point cloud samples corresponding to the joint; see Figure 8(d).

In final post-processing, we re-distribute the skeletal samples by smoothing constrained by spatial coherence as before. The resulting skeletal cloud still maintains a line-like structure within the joints with only the joint skeletal samples “pulled” out of the collapsed center and distributed along line-like connections to the adjacent branches; see Figure 8(e). From the final skeletal cloud, that is sufficiently close to being 1D, we apply subsampling and connect the samples with short curve segments [Lee 2000] to obtain a classical 1D curve skeleton representation; see Figure 8(f).

5 Results and applications

We demonstrate results of curve skeleton extraction from incomplete point clouds and also present a skeleton-driven point cloud completion scheme to assist conventional surface reconstruction algorithms, such as RBF and Poisson reconstruction, in proper recovery of surface topology when the missing data is significant. For visualization, we render the back-facing points in black and front-




Figure 9: Curve skeletons extracted as virtual scans are removed. Here two scans are sufficient to extract a complete skeleton, where about 40% of the points for the fertility model and about 55% of the points for the camel model are missing.

facing points using transparent colored splats. The final skeleton is shown by segments of thin red cylinders.

5.1 Curve skeleton from incomplete point clouds

To demonstrate the robustness of our method, we first use a virtual scanner to collect point data from a set of views around a complete surface model and then progressively remove points captured from the different views to simulate incomplete data. Specifically, we place k viewpoints on a bounding sphere of the object centered at the object centroid O . From each viewpoint V_i , we generate a set of samples via orthographic ray casting from a uniform grid on the plane through V_i and perpendicular to the vector V_iO .

In Figure 9, we show the curve skeletons extracted as points captured by the k views are removed. While one scan does not result in complete skeletons, two are sufficient for the camel and fertility models. Difficulties with the camel involve the hump and close-by points near the knees. The skeletons extracted around these regions show the robustness of our schemes for finding optimal cutting planes and relevant neighborhoods.

The fertility model shows our algorithm at work on a sampled high-genus shape whose generally cylindrical parts may not have a circular profile; it also demonstrates robustness in joint handling. As we can see, beyond two scans, the extracted curve skeletons remain quite consistent. However, certain topological changes can occur within regions where there are significant differences in the point samples, e.g., compare (c-3) with (c-2) and (f-6) with (f-4).

Figure 10 presents a small gallery of curve skeleton results on a variety of 3D objects. The mannequin sequences were captured from few views, almost a single view for the standing pose, and this leads to a great deal of missing data. However, as the mannequin is an articulated model whose body parts are mostly perfectly cylindri-

cal, our shape prior works strongly to allow the recovery of quality curve skeletons. These results suggest the potential applicability of our method to the analysis of deformable point cloud sequences when the views are limited and frames of capture are sparse.

For each raw input cloud, we remove outliers, denoise, and subsample the raw data to 10,000 points using the LOP operator [Lipman et al. 2007]. The required point normals in all tests are estimated via local PCA followed by the normal orientation algorithm of Hoppe et al. [1992]. Note that normals obtained this way are not guaranteed to be accurate or even possess the correct orientations, especially near regions of missing data; these potential errors are quite well tolerated by our algorithm, as the experiments show.

The entire curve skeleton extraction algorithm has been developed within the MATLAB environment, hence the speed of execution is not optimized. The most time-consuming step is the construction of the Mahalanobis neighborhoods during ROSA construction. Typically, on a point cloud with 10K points, ROSA computation takes about 3 minutes, while joint recovery and other post-processing steps are accomplished at interactive rates.

5.2 Skeleton-driven point cloud completion

Missing data are common during 3D shape acquisition. When such cases are only caused by shallow concavity or occasional surface material artifacts, the missing regions are quite small and interpolation schemes, such as RBF [Carr et al. 2001] or Poisson [Kazhdan et al. 2006] construction, have proven to be effective in completing the shape. However, due to deep cavities in a shape, severe occlusion, or limited views generated by the scanner, large portions of the data may be missing and shape completion over these regions can be next to impossible without a reliable shape prior or user assistance [Sharf et al. 2007b]. Data completion over large gaps




Figure 10: A gallery of curve skeleton extraction results. Except for the synthetic “fork” model (green) and the horse, whose missing data are simulated via the virtual scanner, all input come from raw scanned data.

is possible using the smoothness prior and via energy minimization in the Laplacian processing framework [Sorkine and Cohen-Or 2004]. However, a prerequisite to these methods is mesh connectivity, which is difficult to obtain in our case in the first place.

Indeed, the recovery of correct surface topology from imperfect point clouds is where the main challenge lies. Most problems occur between near-by surfaces, where missing data typically occur and reconstruction schemes such as RBF may infer the wrong topology; see Figure 1. This is where a complete and topologically correct curve skeleton can be of great assistance. Our method effectively exploits an additional shape prior and allows shape completion in highly challenging situations, as demonstrated in Figure 1 and 11.

Overview Based on the curve skeletons extracted in this work, we develop a *mesh-less* approach to shape completion where we fill

the missing data regions using points only. The curve skeleton and associated correspondence to the point cloud samples provide the necessary reference frame to allow a “space-time” reconstruction over large gaps using both the cylindrical shape and smoothness priors. Here the time axis is the direction along the skeleton and the local parameterization is a radial one within the cutting planes, assuming that locally the 2D profile is star-shaped.

The more complete point cloud can then be subjected to RBF or Poisson reconstruction to obtain the final surface. In Figure 11, we show Victoria reconstructed via Poisson [Kazhdan et al. 2006] based on the skeleton shown in Figure 10. Similar to Figure 1, the assistance given by the curve skeleton effectively steers the reconstruction towards a better interpretation of the object topology.

Algorithm First, we traverse branches of the obtained curve skeleton, gathering point cloud samples lying near the corresponding cutting planes. For every cutting plane, we construct a cylindrical parameterization of the collected samples. Specifically, at every radial angle θ , the distance from a point cloud sample and the current ROSA center on the skeleton can be estimated.

Given t_i , which discretely parameterizes the skeletal curve, and a set of discrete radial directions θ_i , we then construct a distance matrix $D(t_i, \theta_i)$. Under the smoothness prior, the distances parameterized in D should be well represented by a smooth function. Thus, over the matrix D , we solve an implicit function optimization, minimizing the membrane energy within the common Laplacian processing framework. While in these works the mesh connectivity is available, in our mesh-less solution, adjacency information is naturally provided by the cylindrical parameterization: elements with close-by parameters t_i and θ_i are connected.

We solve the optimization problem iteratively, which is quite efficient. The solution provides estimated distance values over missing data regions, which are filled with new point samples reflecting the distances. Finally, we employ RBF or Poisson reconstruction on the resulting point cloud to obtain a smooth surface. We should stress here that the local cylindrical parameterizations utilized in our approach are meaningful only within branch regions. Hence we do not insert new samples over joint regions. However, it is interesting to note that by its nature, a joint is typically well constrained by the meeting branches. Thus smooth implicit schemes such as RBF or Poisson typically perform quite well in surface completion over the joint regions, as shown in Figure 1 and 11.

5.3 Limitations

The obvious conceptual limitation to our approach is the shape prior used. Furthermore, ROSA is not guaranteed to be inside the perceived shape, e.g., when the cut profile is highly concave. On the theory front, we lack a convergence proof for the iterative planar cut procedure, though we have not experienced any non-converging behavior of ROSA construction experimentally. The main practical issue is parameter tuning. Throughout our experiment, we aim to select a robust set of parameters, δ , ϵ_{Mah} , and ϵ_{MLS} , and in most cases, we have succeeded in doing so. However, as we observe from Figure 9 for the 2-scan camel and 6-scan fertility models, using the same threshold values can result in topological inconsistency of the extracted curve skeletons. More careful parameter tuning can fix the problem, however a more principled approach is to find an automatic parameter-setting scheme.

Lastly, there are always extreme cases, e.g., highly inaccurate normals or several under sampling such as the 1-scan cases in Figure 9, which can cause our method to fail. However, it is worth iterating that our algorithm is statistical in nature; it benefits from the data ensemble so that particular issues encountered in parts of the data may not heavily influence the final result negatively due to more reliable data or more suitable processing from other parts. Thus, occasional extreme data artifacts are typically well compensated within such a statistical framework.

6 Conclusion and future work

Curve skeletons are intuitive and powerful abstractions for a large variety of shapes, in particular articulated shapes and more generally those formed by generally cylindrical components. There has been a great deal of work on curve skeleton extraction from surface representations while the method presented in this paper applies, for the first time, to point clouds where large portions of the data may be missing. Our approach builds upon the notion of gener-




Figure 11: Skeleton-assisted Poisson reconstruction (right) more faithfully recovers the topology of the Victoria model, while the same reconstruction (left) on the original input does not, due to missing data; see Figure 10 for locations of missing data.

alized rotational symmetry axis for curve skeleton extraction and with proper joint handling, it leads to complete, characteristic curve skeletons even with significant missing data. We demonstrate the effectiveness of our algorithm using numerous examples and also show how the extracted curve skeletons can assist in RBF or Poisson shape completion in challenging situations.

For future work, we would like to realize the potential of curve skeletons for registration and reconstruction of time-varying point clouds. We believe curve skeletons will serve as effective reduced models for such tasks. In particular, the ability to extract complete skeletons from incomplete point clouds should allow less scans per model and less captures across time for a more efficient solution to this important and difficult problem. Finally, shape inference from incomplete data is an ill-posed problem and there are always challenging scenarios where user assistance is needed. Curve skeletons are not only intuitive shape abstractions, but also easy to manipulate compared to the point cloud itself. We wish to exploit this characteristic of curve skeletons and the skeleton-shape correspondences obtained from our method for interactive shape repair.

Acknowledgments

The authors would like to thank the anonymous reviewers for their helpful comments. This work is supported in part by grants from NSERC (No. 611370 and No. 613326), an MITACS Research Grant (No. 699127), the Israeli Ministry of Science, and the Israel Science Foundation. The camel and fertility models are from the AIM@SHAPE shape repository and the horse model from the Stanford 3D Scanning Repository. The Victoria and fork models are courtesy of Dr. Andrei Sharf. The “Y”-shaped model is courtesy of Francois Xavier Nhieu. The remaining point cloud models were acquired by a Polhemus FastScan 3D scanner.

References

- AMENTA, N., AND KIL, Y. J. 2004. Defining point-set surfaces. *ACM Trans. on Graph.* 23, 3, 264–270.
- AU, O. K.-C., TAI, C.-L., CHU, H.-K., COHEN-OR, D., AND LEE, T.-Y. 2008. Skeleton extraction by mesh contraction. *ACM Trans. on Graph.* 27, 3, 44:1–44:10.
- BLUM, H. 1967. *A transformation for extracting new descriptors of shape*. Models for the perception of speech and visual form. MIT Press, 362–380.

- BOUIX, S., SIDDIQI, K., TANNENBAUM, A., AND ZUCKER, S. 2006. Medial axis computation and evolution. *Statistics and Analysis of Shapes*.
- CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *Proc. of SIGGRAPH*, 67–76.
- CHUANG, J.-H., AHUJA, N., LIN, C.-C., TSAI, C.-H., AND CHEN, C.-H. 2004. A potential-based generalized cylinder representation. *Computers & Graphics* 28, 6, 907–918.
- CORNEA, N. D., MIN, P., AND SILVER, D. 2007. Curve-skeleton properties, applications, and algorithms. *IEEE Trans. Vis. & Comp. Graphics* 13, 3, 530–548.
- DE AGUIAR, E., THEOBALT, C., THRUN, S., AND SEIDEL, H.-P. 2008. Automatic conversion of mesh animations into skeleton-based animations. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2 (4), 389–397.
- DEY, T., AND SUN, J. 2006. Defining and computing curve-skeletons with medial geodesic function. In *Symp. on Geom. Proc.*, 143–152.
- GIBLIN, P. J., AND BRASSETT, S. A. 1985. Local symmetry of plane curves. *American Mathematical Monthly*, 689–707.
- HILAGA, M., SHINAGAWA, Y., KOHMURA, T., AND KUNII, T. L. 2001. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proc. of SIGGRAPH*, 203–212.
- HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *Proc. of SIGGRAPH*, 71–78.
- JAMES, D., AND TWIGG, C. 2005. Skinning mesh animations. *ACM Trans. on Graph.* 24, 3, 399–407.
- KATZ, S., AND TAL, A. 2003. Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Trans. on Graph.* 22, 3, 954–961.
- KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2004. Symmetry descriptors and 3D shape matching. In *Symp. on Geom. Proc.*, 115–123.
- KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Symp. on Geom. Proc.*, 61–70.
- KIM, Y. J., VARADHAN, G., LIN, M. C., AND MANOCHA, D. 2003. Fast swept volume approximation of complex polyhedral models. In *Proc. of Symp. on Solid Modeling and App.*, 11–22.
- LEE, I. 2000. Curve reconstruction from unorganized points. *Computer Aided Geometric Design* 17, 2, 161–177.
- LEHTINEN, J., ZWICKER, M., TURQUIN, E., KONTKANEN, J., DURAND, F., SILLION, F., AND AILA, T. 2008. A meshless hierarchical representation for light transport. *ACM Trans. on Graph.* 27, 3, 37:1–37:10.
- LEWIS, J. P., CORDNER, M., AND FONG, N. 2000. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *Proc. of SIGGRAPH*, 165–172.
- LI, X., WOON, T., TAN, T., AND HUANG, Z. 2001. Decomposing polygon meshes for interactive applications. In *Proc. of Symposium on Interactive 3D graphics*, 35–42.
- LIPMAN, Y., COHEN-OR, D., LEVIN, D., AND TAL-EZER, H. 2007. Parameterization-free projection for geometry reconstruction. *ACM Trans. on Graph.* 26(3), 22:1–22:6.
- LU, L., CHOI, Y.-K., WANG, W., AND KIM, M.-S. 2007. Variational 3D shape segmentation for bounding volume computation. *Computer Graphics Forum (Proc. of Eurographics)* 26, 3, 329–338.
- MALANDAIN, G., AND FERNÁNDEZ-VIDAL, S. 1998. Euclidean skeletons. *Image and Vision Computing* 16, 5, 317–327.
- MITRA, N., GUIBAS, L., AND PAULY, M. 2006. Partial and approximate symmetry detection for 3D geometry. *ACM Trans. on Graph.* 25, 3, 560–568.
- MITRA, N. J., FLÖRY, S., OVSIANIKOV, M., GELFAND, N., GUIBAS, L., AND POTTMANN, H. 2007. Dynamic geometry registration. In *Symp. on Geom. Proc.*, 173–182.
- NEHAB, D., RUSINKIEWICZ, S., DAVIS, J., AND RAMAMOORTHY, R. 2005. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. on Graph.* 24, 3, 536–543.
- OGNIEWICZ, R., ILG, M., AND ZURICH, E. 1992. Voronoi skeletons: theory and applications. In *Proc. IEEE Conf. on Comp. Vis. and Pat. Rec.*, 63–69.
- OVSIANIKOV, M., SUN, J., AND GUIBAS, L. 2008. Global intrinsic symmetries of shapes. In *Computer Graphics Forum (Proc. of Symp. on Geom. Proc.)*, vol. 27, 1341–1348.
- PATANE, G., SPAGNUOLO, M., AND FALCIDIENO, B. 2008. Reeb graph computation based on a minimal contouring. In *Proc. IEEE Conf. on Shape Modeling and App.*, 73–82.
- PEKELNY, Y., , AND GOTSMAN, C. 2008. Articulated object reconstruction and motion capture from depth video. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2, 399–408.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Trans. on Graph.* 25, 3, 549–559.
- RAAB, R., GOTSMAN, C., AND SHEFFER, A. 2004. Virtual woodwork: Making toys from geometric models. *Int. J. of Shape Modeling* 10, 1, 1–30.
- SHARF, A., LEWINER, T., SHAMIR, A., AND KOBELT, L. 2007. On-the-fly curve-skeleton computation for 3D shapes. *Computer Graphics Forum (Proc. of Eurographics)* 26, 3, 323–328.
- SHARF, A., LEWINER, T., SHKLARSKI, G., TOLEDO, S., AND COHEN-OR, D. 2007. Interactive topology-aware surface reconstruction. *ACM Trans. on Graph.* 26, 3, 43:1–43:10.
- SHARF, A., ALCANTARA, D. A., LEWINER, T., GREIF, C., SHEFFER, A., AMENTA, N., AND COHEN-OR, D. 2008. Space-time surface reconstruction using incompressible flow. *ACM Trans. on Graph.* 27, 5, 110:1–110:10.
- SIDDIQI, K., AND PIZER, S. 2009. *Medial Representations: Mathematics, Algorithms and Applications*. Springer.
- SIMARI, P. D., AND SINGH, K. 2005. Extraction and remeshing of ellipsoidal representations from mesh data. In *Proc. of Graphics Interface*, 161–168.
- SORKINE, O., AND COHEN-OR, D. 2004. Least-squares meshes. In *Proc. IEEE Conf. on Shape Modeling and App.*, 191–199.
- THRUN, S., AND WEGBREIT, B. 2005. Shape from symmetry. In *Proc. of Int. Conf. on Comp. Vis.*, vol. 2, 1824–1831.
- WAND, M., JENKE, P., HUANG, Q., BOKELOH, M., GUIBAS, L., AND SCHILLING, A. 2007. Reconstruction of deforming geometry from time-varying point clouds. In *Symp. on Geom. Proc.*, 49–58.