

Fast and robust curve skeletonization for real-world elongated objects

Amy Tabb

USDA-ARS-AFRS
Kearneysville, West Virginia, USA

amy.tabb@ars.usda.gov

Henry Medeiros

Marquette University, Electrical and Computer Engineering
Milwaukee, Wisconsin, USA

henry.medeiros@marquette.edu

Abstract

We consider the problem of extracting curve skeletons of three-dimensional, elongated objects given a noisy surface, which has applications in agricultural contexts such as extracting the branching structure of plants. We describe an efficient and robust method based on breadth-first search that can determine curve skeletons in these contexts. Our approach is capable of automatically detecting junction points as well as spurious segments and loops. All of that is accomplished with only one user-adjustable parameter. The run time of our method ranges from hundreds of milliseconds to less than four seconds on large, challenging datasets, which makes it appropriate for situations where real-time decision making is needed. Experiments on synthetic models as well as on data from real world objects, some of which were collected in challenging field conditions, show that our approach compares favorably to classical thinning algorithms as well as to recent contributions to the field.¹²

1. Introduction

The three-dimensional reconstruction of complex objects under realistic data acquisition conditions results in noisy surfaces. We describe a method to extract the curve skeleton of such noisy, discrete surfaces for the eventual purpose of making decisions based on the curve skeleton. Much work has been done in the computer graphics community on the problem of skeletonization. In general, curve skeletonization converts a 3D model to a simpler represen-

tation, which facilitates editing or visualization [7] as well as shape searching and structure understanding [3, 11, 16]. However, the work in the computer graphics community usually assumes noise-less surfaces or surfaces with negligible noise levels. In this work, we intend to use curve skeletons as an intermediate step between surface reconstruction and computing measurements of branches for automation applications in which robustness to noise and fast execution are important, such as in the automatic modeling of fruit trees in an orchard [28].

There are two commonly-used types of skeletons, the medial axis transform (MAT) skeleton, and the curve skeleton. Skeletons using MAT are curves in 2D while in 3D they are locally planar. They allow for the original model to be reconstructed but are very sensitive to local perturbations [19]. Curve skeletons consist of one-dimensional curves for surfaces in 3D, which provides a simpler representation than MAT-type skeletons (see [7] for a comprehensive review). However, because there are different definitions of curve skeletons, there is an abundance of methods, with different advantages and disadvantages.

The problem we explore in this paper is to compute curve skeletons of discrete 3D models represented by voxels, which may be sparse, noisy, and are characterized by elongated shapes. The curve skeleton must be thin and one-dimensional except in the case of junction points, which should be detected during the computation of the curve skeleton. The skeleton must also be centered, but because of noise and the use of voxels we use the relaxed centeredness assumption (see [7] for more details). To make our approach robust to noise, we identify spurious curve skeleton segments in the course of the algorithm, which removes the need for a separate pruning step [4, 21, 31]. Finally, since our work is mainly concerned with real trees, branch crossings occur and support poles may be attached to the trees via ties. As a result, our datasets include loops and cycles and breaking a loop is not desirable. Hence, our method is able to determine that curve skeleton segments which do not terminate at a surface voxel are part of a loop.

We propose a path-based algorithm for the real-time

¹The citation information for this paper is: A. Tabb and H. Medeiros, “Fast and robust curve skeletonization for real-world elongated objects”, 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Tahoe, NV/CA. DOI 10.1109/WACV.2018.00214

²Mention of trade names or commercial products in this publication is solely for the purpose of providing specific information and does not imply recommendation or endorsement by the U.S. Department of Agriculture. USDA is an equal opportunity provider and employer. A. Tabb acknowledges the support of US National Science Foundation grant number IOS-1339211.

computation of curve skeletons of elongated objects with noisy surfaces that takes into account all of the criteria above. Specifically, our contributions are: 1) our method is robust to noise and there is no requirement for additional pruning, 2) it has time complexity $\mathcal{O}(n^{\frac{7}{3}})$ – where n is the number of occupied voxels in 3D space – and as a result is suitable for automation contexts, 3) the method can handle loops, 4) we provide an extensive evaluation on synthetic models as well as on real-world objects, and 5) we provide source code that is publicly available [27].

2. Related work

Using laser data, there has been some related work on the problem of extracting cylinders from point cloud data [6, 15, 17, 24, 29]. Such methods cannot be used for the datasets we consider because these works assume that the underlying shapes are cylinders. While elongated shapes may be locally cylindrical, they may have many curves and cylinder fitting may not be appropriate for these shapes.

In [2], the authors compute MAT-style skeletons and combine those ideas with those of traditional thinning, for the purposes of object recognition and classification with an emphasis on reconstruction. Their algorithm is efficient, but like many other algorithms it depends on a pruning step which requires parameter setting. Other works which combine the ideas of thinning with computing skeletons are [10], where a bisector function is used to compute a surface skeleton, and [5] where new kernels are used for thinning.

In [30], the authors provide an algorithm for computing a curve skeleton from a discrete 3D model, assuming that some noise is present. This is accomplished by a shrinking step that preserves topology, as well as a thinning step to create 1D structures, and finally a pruning step. While that method is able to deal with some noise, the shrinking step involved in the algorithm would result in missing some branches with small scale.

A recent approach that is most similar to the method we present was proposed by Jin *et al.* in [12, 13]. In these works, curve skeletons are extracted from medical data that contains noise, and once a seed voxel has been identified, new curve skeleton segments are found via searches based on the geodesic distance. While our method shares a similar overall structure in that paths are iteratively discovered, we do not make assumptions concerning the thickness and lengths of neighboring branches. In addition, that method was not designed to handle loops. Finally, in that work, the authors note the problems with computational speed in their approach because of the use of geodesic path computations. Our method was conceived to be executed in real-time and is hence computationally inexpensive.

3. Method description

The proposed method to compute a curve skeleton is composed of four main steps:

1. Determine the seed voxel for the search for curve skeleton segments (Section 3.2);
2. Determine potential endpoints for curve skeleton segments (Section 3.3);
3. Determine prospective curve segments (Section 3.4);
4. Identify and discard spurious segments and detect loops (Section 3.5).

Step 1 is executed once at initialization, whereas the remaining steps are executed iteratively until all curve segments have been identified. These steps are described in detail below and the entire process is illustrated in Figure 1. In this description, we assume that there is only one connected component, but if there are additional connected components, all four steps are performed for each component.

3.1. Preliminaries

The set of occupied voxels in 3D space is \mathbb{V} , and $n = |\mathbb{V}|$ is the number of occupied voxels. We note here that voxels are defined with respect to a uniform three-dimensional grid, and we let the number of voxels (occupied and empty) in such a grid be N . However, in this work, we only operate on occupied voxels, and since the objects we treat are elongated, $n \ll N$. In the remainder of the document, *voxels* will mean occupied voxels. We interpret the voxels as nodes in a graph and assume that voxel labels are binary: occupied or empty. The edges of the graph are defined by a neighborhood relationship on the voxels. We represent the set of occupied neighbors for a voxel v_i as \mathcal{N}_i . In our implementation we use a 26-connected neighborhood. A surface voxel has $|\mathcal{N}_i| < 26$ and the set of surface voxels is $\mathbb{S} \subseteq \mathbb{V}$.

Each segment of the curve skeleton is represented by a set of voxels. At each iteration m of the skeletonization algorithm, a new curve skeleton segment $\mathbb{C}^{(m)}$ is discovered. Then, the overall skeleton is represented by the set \mathbb{C} which is the set of skeleton segment sets:

$$\mathbb{C} = \bigcup_m \{\mathbb{C}^{(m)}\}. \quad (1)$$

3.1.1 Modified breadth-first search algorithm

Our skeletonization method is heavily based on a proposed modification of breadth-first search (BFS) which allows one to alter the rate at which nodes are discovered according to a weighting function. We now discuss this BFS modification generally and then show its application to our approach to curve skeletonization in future sections.

In classic BFS, there are three sets of nodes: non-frontier, frontier, and undiscovered nodes and the result of the BFS is a label for each node in the graph. The starting

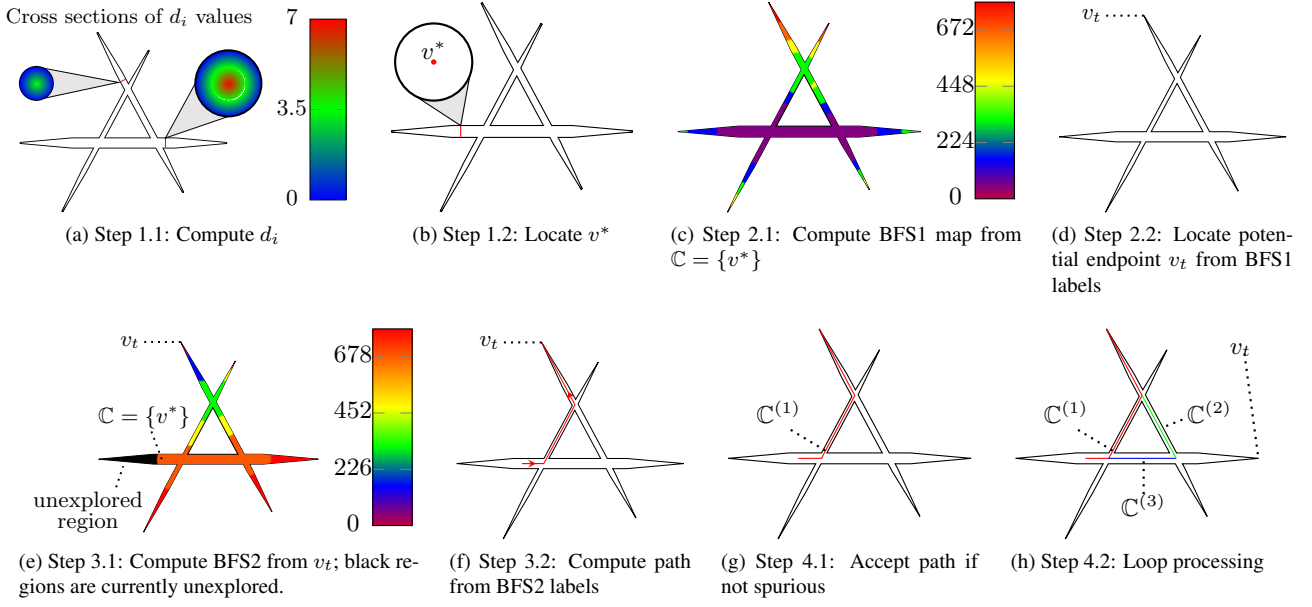


Figure 1. **Best viewed in color.** Illustration of the computation of curve skeletons with our proposed method on an artificially created three-dimensional object consisting of three intersecting segments of varying diameters. Legends for colormaps are indicated to the right of the figures. In step 1.1 shown in 1a, the distance labels d_i are computed; the cross sections in 1a show the pattern of labels in the object’s interior. Then 1b shows step 1.2, where v^* is selected from the local maxima of d_i ; any voxel in the local maxima may be selected. In Step 2.1, the BFS1 map (Section 3.3.1) from v^* given d_i is computed (Figure 1c). In 1d, step 2.2, the maximal label from step 2.1 is selected as a proposed endpoint v_t . Figure 1e shows step 3.1 in which we compute the BFS2 labels (Section 3.4.1) from v_t to the current curve skeleton, which is $C = \{v^*\}$. Step 3.2 in Figure 1f consists of tracing the path through the BFS2 labels from v^* to v_t . In step 4.1, we accept the path as part of the curve skeleton if it passes the spurious path test (Figure 1g). Finally, in Figure 1h we show the loop handling procedure; v_t is found on the right-hand side, and the loop is incorporated into the curve skeleton. This completes the steps for iteratively adding a curve segment. Steps 2.1 through 4.2 are then repeated until there are no more proposed endpoints that pass the spurious path test.

node has a label of 0 and the labels of other nodes are the number of edges that need to be traversed on a shortest path from any node to the starting node. In our modified BFS algorithm, the labels represent the sum of pairwise distances along the shortest path to a given node.

Each voxel v_i also has a weight assigned to it, w_i . Voxels with smaller values of w_i are incorporated into the frontier before neighboring voxels with higher weight values. The speed of discovery of nodes can thereby be altered to favor paths that go through voxels with characteristics which are desirable for a specific purpose as explained in detail in Section 3.3.

Our modified BFS is shown in Algorithm 5. The frontier for a particular iteration k is $\mathbb{F}^{(k)}$, which is composed of two subsets, $\mathbb{F}_A^{(k)}$ and $\mathbb{F}_B^{(k)}$. The initialization of $\mathbb{F}_A^{(0)}$ depends on the intended use of the algorithm (details are given in Sections 3.3.1 and 3.4.1), and $\mathbb{F}_B^{(0)}$ is always initially empty. The label of voxels is given by:

$$l_i = \begin{cases} \infty & \text{if } v_i \notin \mathbb{F}_A^{(0)} \\ 0 & \text{if } v_i \in \mathbb{F}_A^{(0)} \end{cases} \quad (2)$$

The algorithm progresses as follows. A voxel $v_i \in \mathbb{F}_A^{(k)}$ has neighbors v_j which had been discovered previously as well as neighbors that were discovered later than itself as determined by the labels of v_i and its neighbors. In line 3, only neighbors discovered later than the voxels in $\mathbb{F}_A^{(k)}$ are updated based on the label of v_i , the weight w_i , and the distance between the neighboring voxels. The set $\mathbb{N}^{(k)}$ in line 6 is the set of frontier candidates beyond the current frontier at iteration k . $\mathbb{N}^{(k)}$ is used to select a label threshold, l_{min} . If the voxels in $\mathbb{F}^{(k)}$ have a label greater than this threshold, they remain in the frontier (specifically $\mathbb{F}_B^{(k+1)}$) for the next iteration. If a voxel v_i in $\mathbb{F}^{(k)}$ has a label smaller than this threshold, then those neighbors of v_i which were discovered later than v_i are placed in $\mathbb{F}_A^{(k+1)}$ and v_i is removed from the frontier for the next iteration. The distinction between $\mathbb{F}_A^{(\cdot)}$ and $\mathbb{F}_B^{(\cdot)}$ allows for a more efficient implementation because only labels in $\mathbb{F}_A^{(\cdot)}$ must be updated in line 4.

3.2. Determination of the seed voxel v^*

As mentioned above, the first step of our skeletonization algorithm is the determination of the seed voxel. This step

Algorithm 1 Modified BFS Algorithm

Input: Set of occupied voxels \mathbb{V} , initial frontier voxels $\mathbb{F}_A^{(0)}$, voxel weights w_i , initial voxel labels l_i

Output: Updated voxel labels l_i

```

1:  $k = 0, \mathbb{F}_B^{(0)} = \emptyset$ 
2: while  $|\mathbb{F}_A^{(k)}| > 0$  do
3:   for each voxel  $v_i \in \mathbb{F}_A^{(k)}$  do
4:     for each voxel  $v_j \in \mathcal{N}_i$  such that  $(l_j > l_i)$  do
5:        $l_j = \min(l_j, l_i + w_j + ||v_j - v_i||)$ 
6:    $\mathbb{F}^{(k)} = \mathbb{F}_A^{(k)} \cup \mathbb{F}_B^{(k)}$ 
7:    $\mathbb{N}^{(k)} = \{v_j | l_j > l_i, \forall v_j \in \mathcal{N}_i, \forall v_j \notin \mathbb{F}^{(k)}, \forall v_i \in \mathbb{F}^{(k)}\}$ 
8:    $l_{min} = \min_{v_j \in \mathbb{N}^{(k)}} l_j$ 
9:    $\mathbb{F}_A^{(k+1)} = \{v_j | l_j < l_{min}, \forall v_i \in \mathbb{F}^{(k)}, \forall v_j \in \mathbb{N}^{(k)}\}$ 
10:   $\mathbb{F}_B^{(k+1)} = \{v_i | l_i \geq l_{min}, \forall v_i \in \mathbb{F}^{(k)}, |\mathcal{N}_i \cap \mathbb{N}^{(k)}| > 0\}$ 
11:   $k = k + 1$ 

```

consists of two sub-steps: computation of distance labels, and localization of an extreme point as explained below.

3.2.1 Distance label computation

To compute the distance labels $d_i, i = 0, \dots, n - 1$, we compute the distance transform using Euclidean distances so that d_i represents the distance from v_i to the closest voxel in \mathbb{S} , i.e., a surface voxel. To compute the distance labels efficiently, we use the linear-time algorithm of [18] on the occupied voxels in our graph representation (note that the pseudo-code in [18] considers regular grids instead). In addition, in our implementation, the three scans of the algorithm are executed in parallel.

3.2.2 Finding the seed voxel v^*

Once the distances d_i are computed, we find the voxels with maximum distance label, d_{max} , which ensures that the curve skeleton goes through the thickest part of the object. There may be several voxels with $d_i = d_{max}$, and we arbitrarily select one of them to be v^* . If a point of the curve skeleton is known to be a desirable seed point for a specific application, that point may be selected to serve as v^* without affecting the subsequent steps we describe in this paper.

3.3. Determination of endpoint candidates

This section describes the first step to determine the curve skeleton segments $\mathbb{C}^{(m)}$: the identification of the endpoints of prospective segments that are connected to existing segments in the skeleton. This is done in two sub-steps. First, we compute the breadth-first search distances from

the existing curve segments to potential endpoints. We call this step BFS1. Then, a candidate endpoint is identified from the extreme points in this set. These sub-steps are described in detail below.

3.3.1 BFS1 Step

At the first iteration of our algorithm, the curve skeleton consists of a single voxel, v^* . We initialize $\mathbb{F}_A^{(0)} = \{v^*\}$ in Algorithm 5, and initialize the labels as in Equation 2. We then perform Algorithm 5 using weights $w_i = d_{max} - d_i$, where d_i and d_{max} are the distance labels and the maximum distance label, respectively (shown in Figure 1a). These weights increase linearly according to a voxel's Euclidean distance to a surface voxel, i.e., surface voxels have $w_i = d_{max}$. The overall effect of weighting the search in such a way is that paths which pass through the center of the object are explored first. This procedure finds the distances from each voxel to the existing curve skeleton along a centered path. As explained in detail below, points with maximal distance are endpoint candidates.

At each subsequent iteration of the algorithm, new curve skeleton segments (which are identified as described in Sections 3.4 and 3.5 below) are added to $\mathbb{F}_A^{(0)}$ and the BFS1 labels are updated accordingly. For improved efficiency, on subsequent iterations, BFS1 labels are simply updated instead of computed from scratch. Suppose that on iteration m the set of approved curve skeleton segment voxels is $\mathbb{C}^{(m-1)}$, so that $\mathbb{F}_A^{(0)} = \mathbb{C}^{(m-1)}$. We leave the existing BFS1 labels from iteration $m - 1$ unchanged, except for those in $\mathbb{F}_A^{(0)}$:

$$l_i^{(m)} = \begin{cases} l_i^{(m-1)} & \text{if } v_i \notin \mathbb{F}_A^{(0)} \\ 0 & \text{if } v_i \in \mathbb{F}_A^{(0)} \end{cases} \quad (3)$$

Then Algorithm 5 progresses as usual given these labels.

3.3.2 Identification of an endpoint candidate v_t from BFS1

We next identify candidate endpoint voxels of the curve skeleton. A candidate endpoint voxel is a surface voxel which is not yet connected to the curve skeleton. An endpoint candidate is given by the voxel $v_t \in \mathbb{S}$ such that the label of v_t is greater than or equal to any other BFS1 label for any other surface voxels, i.e.,

$$v_t = \arg \max_{v_i \in \mathbb{S}} (l_i), \quad (4)$$

where l_i is the BFS1 label of v_i . There may be multiple voxels with the same maximum label value. As in the seed voxel selection step in Section 3.2.2, v_t may be chosen arbitrarily from the set of voxels with the maximum label value.

3.4. Determination of prospective curve segments

The existing curve skeleton might be reachable from a proposed endpoint v_t via more than one path (see Figure 1h, for example). We use the breadth-first search distance from the prospective endpoint to the existing curve skeleton to identify those branches and junction points. This is also done in two sub-steps. First, we compute the BFS distances from v_t , which we call the BFS2 step. Then we determine the curve skeleton segments by analyzing connected components. These sub-steps are explained in detail below.

3.4.1 BFS2 Step

In order to determine where a proposed curve skeleton segment intersects with the existing curve skeleton, we use Algorithm 5 with weights $w_i = d_{max} - d_i$, as in the BFS1 step. Unlike the BFS1 step, however, for each iteration of the algorithm, the labels are now initialized using Equation 2 with $\mathbb{F}_A^{(0)} = \{v_t\}$. When an existing curve skeleton section is encountered, the search is stopped for that region. The output of this step are the BFS2 labels.

Once the BFS2 labels are computed, the frontier arcs are then analyzed and grouped by connected components. The number of connected components in the frontier voxel set is the number of curve skeleton paths from v_t to the existing curve skeleton.

3.4.2 Identification of curve skeleton segments from BFS2

Let a frontier connected component (FCC) from BFS2 be the set of voxels \mathbb{F}_{CC} . We determine the voxels in \mathbb{F}_{CC} that are neighbors of the existing curve skeleton \mathbb{C} and denote these voxels as \mathbb{F}_{CSN} :

$$\mathbb{F}_{CSN} = \{v_i | (v_i \in \mathbb{F}_{CC}) \wedge (\exists v_j \in \mathcal{N}_i) \wedge (v_j \in \mathbb{C})\}. \quad (5)$$

From \mathbb{F}_{CSN} we determine the voxel with the smallest BFS2 label in the set and denote this voxel as $v_{s,1}$, i.e.,

$$v_{s,1} = \arg \min_{v_i \in \mathbb{F}_{CSN}} (l_i), \quad (6)$$

where l_i is the BFS2 label. As before, there may be many voxels with the same smallest label, and one may be chosen arbitrarily. The next step is to determine the path from $v_{s,1}$ to v_t such that the path is centered. We accomplish this with the BFS2 labels as well as the distance transform labels d_i . This combination improves centeredness on curved portions as compared to only using BFS2 labels.

The process of determining a new curve skeleton segment is sketched in Algorithm 6. The sequence of current voxels v_c creates a new curve skeleton segment. We start from the neighbor of the existing curve skeleton, $v_{s,1}$, and set v_c equal to $v_{s,1}$. We determine d^* by examining v_c 's

Algorithm 2 Determination of curve skeleton segment from BFS2 and d_i

Input: Set of occupied voxels \mathbb{V} , BFS2 voxel labels l_i , voxel distance transforms d_i , proposed endpoint v_t , voxel in the existing curve skeleton $v_{s,1}$

Output: Curve skeleton segment $\mathbb{C}^{(m)}$

- 1: $v_c = v_{s,1}$
 - 2: $\mathbb{C}^{(m)} = \{v_{s,1}\}$
 - 3: **while** $(v_c \neq v_t) \wedge (v_c \notin \mathbb{C})$ **do**
 - 4: Determine $d^* = \max_{v_j \in \mathcal{N}_c \wedge l_c > l_j} (d_i)$ where d_i is the distance transform of v_i
 - 5: Compute $v_n = \arg \min_{v_j \in \mathcal{N}_c^*} (l_j)$ where l_j is the BFS2 label of v_j and $\mathcal{N}_c^* = \{v_j | v_j \in \mathcal{N}_c, d_j = d^*\}$
 - 6: $\mathbb{C}^{(m)} = \mathbb{C}^{(m)} \cup \{v_n\}$
 - 7: $v_c = v_n$
-

neighbors and the BFS2 labels. Let the BFS2 label of v_c be l_c . d^* is the maximum distance label of v_c 's neighbors whose distance labels l_i are less than l_c . Then v_n is determined as the voxel, out of v_c 's neighbors, with the smallest BFS2 label among those voxels with distance transform label equal to d^* . The practical ramifications for these choices are as follows. By choosing the next voxel v_n as a voxel with a smaller BFS2 label value than v_c , we are guaranteed to be moving towards v_t within the voxels that are occupied. Secondly, by requiring that the next distance $d_n = d^*$, we are choosing the voxel most in the center of all of the neighboring voxels that are on a path towards v_t in the occupied voxels. Algorithm 6 is performed for all FCCs.

3.5. Identification of spurious segments and loops

Due to the noisy nature of our surfaces, it is necessary to reject some of the curve skeleton segments identified in the previous step. In the next section, we describe our approach to classify curve skeleton segments as spurious or non-spurious using the frontier voxels from the BFS2 step. When the number of FCCs is one, the proposed curve segment $\mathbb{C}^{(m)}$ undergoes the spurious curve segment classification described in the next section. When the number of FCCs is greater than one, a loop is present, and this case is handled using the approach described in Section 3.5.2.

3.5.1 Spurious curve segment classification

Our classification approach assumes that the surface voxels are disturbed by additive three-dimensional Gaussian noise $\eta \sim (\mu, \Sigma)$ and checks whether the endpoint of the proposed curve segment v_t belongs to this distribution.³ If

³Although we used a Gaussian model for mathematical convenience, our experiments showed that the real noise distribution has little impact on the performance of our approach. That is the case for real-world data or

it does not, the segment is considered spurious. To compute this distribution, note that \mathbb{F}_{CC} is composed of interior and surface voxels. Let the set of surface voxels from \mathbb{F}_{CC} be \mathbb{F}_S , and let $v_{s,0}$ be the closest voxel from the existing skeleton to the proposed segment. For each voxel $v_j \in \mathbb{F}_S$, we determine the difference vector $v'_j = v_j - v_{s,0}$. Then, the sample mean and the sample covariance of η are given by $\mu = \frac{1}{N} \sum_{i=1}^N v'_i$ and $\Sigma = \frac{1}{N} \sum_{i=1}^N (v'_i - \mu)(v'_i - \mu)^T$ where $N = |\mathbb{F}_S|$.

The squared difference vectors $\|v'_j\|^2$ are χ^2 -distributed with three degrees of freedom. In order to classify a curve segment $\mathbb{C}^{(m)}$, we compute the probability that the shifted segment endpoint $v'_t = v_t - v_{s,0}$ belongs to this χ^2 -distribution. Let $x = (v'_t - \mu)^T \Sigma^{-1} (v'_t - \mu)$, its probability density function is given by

$$f(x) = \frac{x^{1/2} e^{-x/2}}{2^{3/2} \Gamma(3/2)}. \quad (7)$$

If $f(x) > t$, where t is a user-supplied acceptance probability, then the curve segment's tip v_t is considered part of the surface voxel's distribution and consequently discarded as a spurious segment. Otherwise, $\mathbb{C}^{(m)}$ is incorporated into the existing curve skeleton. This approach allows curve segments to be classified as spurious or not depending on local conditions and not on absolute parameters.

3.5.2 Loop handling for multiple frontier connected components

The presence of multiple FCCs indicates that one or more tunnels are present in the surface, which corresponds to one or more loops in the curve skeleton. In this case, we do not pursue the spurious curve classification step and instead handle the loop first (Algorithm 7); once the loop has been located, the algorithm returns to spurious curve segment classification (§3.5.1). Let the j -th FCC be represented by the set of voxels $\mathbb{F}_{CC,j}$, where $j = 0, 1, \dots, a-1$ and a is the number of FCCs. Then for each $\mathbb{F}_{CC,j}$, we compute a proposed curve skeleton segment using Algorithm 6, and let this proposed curve skeleton segment be denoted \mathbb{C}'_j . The proposed curve skeleton segments \mathbb{C}'_j s have some voxels in common. In particular, all of the proposed segments travel through the region near the tip v_t , which is a surface voxel. We remove this common region near the tip from all of the \mathbb{C}'_j s (lines 1, 3). Then, the \mathbb{C}'_j s are processed such that there are no redundancies between \mathbb{C}'_j s (line 5) and added to the set of curve skeleton segments \mathbb{C} (line 6). A figure illustrating this process is in the supplemental materials.

Algorithm 8 summarizes the complete proposed approach. Line 1 corresponds to the seed localization step performed at initialization as described in Section 3.2. Lines 3-

even when a substantial amount of shot-like noise is introduced into the models (See Figures 2, 3, and the figures in the supplementary materials).

Algorithm 3 Loop handling

Input: Proposed skeleton segments with common voxels \mathbb{C}'_j , number of FCCs a

Output: Set of skeleton segments \mathbb{C} with disjoint loop segments

- 1: $\mathbb{C}_t = \cap_{j \in [0, a-1]} \mathbb{C}'_j$
 - 2: **for** $j = 0$ to $a - 1$ **do**
 - 3: $\mathbb{C}'_j = \mathbb{C}'_j \setminus \mathbb{C}_t$
 - 4: **for** $k = j + 1$ to $a - 1$ **do**
 - 5: $\mathbb{C}'_j = \mathbb{C}'_j \setminus \mathbb{C}_k$
 - 6: $\mathbb{C} = \mathbb{C} \cup \{\mathbb{C}'_j\}$
-

Algorithm 4 Proposed skeletonization algorithm

Input: Set of occupied voxels \mathbb{V} representing the object of interest, user-supplied acceptance probability t

Output: Object skeleton \mathbb{C}

- 1: Determine seed voxel v^* and make initial curve skeleton $\mathbb{C}^{(0)} = v^*$
 - 2: **repeat**
 - 3: Update BFS1 labels using Alg. 5 with $\mathbb{F}_A^{(0)} = \mathbb{C}^{(m-1)}$
 - 4: Locate a proposed endpoint v_t from BFS1
 - 5: Create BFS2 labels using Alg. 5 with $\mathbb{F}_A^{(0)} = \{v_t\}$
 - 6: Create $\mathbb{C}^{(m)}$ by tracing paths in BFS2 labels from existing curve skeleton v_t according to Alg. 6
 - 7: **if** (Number of FCCs == 1) **then**
 - 8: Accept or decline curve skeleton segments using the method of Section 3.5.1 according to the acceptance probability parameter t . If accepted $\mathbb{C} = \mathbb{C} \cup \{\mathbb{C}^{(m)}\}$ and $m = m + 1$
 - 9: **else**
 - 10: Check for loops using Alg. 7.
 - 11: **until** No more endpoint hypotheses are found
-

4 show the iterative endpoint localization method presented in Section 3.3. The determination of prospective segments of Section 3.4 is carried out by Lines 5-6. Finally, lines 7-10 perform the spurious segment classification and loop handling routines described in Section 3.5. The computational complexity of the entire method is $\mathcal{O}(n^{\frac{7}{3}})$. As a matter of fact, the method runs in $\mathcal{O}(\|V_t\| \times d_{max} \times n)$, where d_{max} is the largest voxel to surface distance, and $\|V_t\|$ is the number of proposed endpoints v_t . Since $\|V_t\|$ and d_{max} tend to be orders of magnitude smaller than n for elongated objects, the method runs extremely fast in practice. A detailed analysis of the computational complexity as well as a discussion of the topological stability of the proposed approach are included in the supplementary materials.

Table 1. Characteristics of the nine datasets used in our evaluation. n is the number of occupied voxels/nodes, N is the number of voxels in the grid, d_{max} is the largest voxel to surface distance, and $||\mathbb{V}_t||$ is the number of proposed endpoints v_t .

ID	n	N	d_{max}	$ \mathbb{V}_t $
A	55,156	27,744,000	16	15
B	88,407	56,832,000	5	174
C	88,798	45,240,000	7	145
D	92,892	80,640,000	6	154
E	98,228	58,464,000	9	50
F	136,497	80,640,000	7	134
G	158,686	64,512,000	9	128
H	176,820	80,640,000	10	97
I	246,654	80,640,000	10	83

4. Experiments

We evaluated our method and four comparison approaches that reflect the state of the art on curve skeletonization on real datasets consisting of nine different trees, denoted as trees A - I. These trees are real-world objects with an elongated shape. Most of the trees are three meters or taller, and the surfaces are noisy. The data for six out of the nine trees was acquired outdoors, and the reconstructions were generated using the method of [26], but other reconstruction algorithms could be applied (e.g., [23]). Table 1 lists the main characteristics of the nine datasets. All five methods were evaluated with respect to their accuracy and robustness to noise as well as run times. We additionally performed a qualitative evaluation of the performance of our method in non-elongated synthetic models commonly used in the evaluation of skeletonization algorithms.

The first comparison method is the classical medial-axis thinning algorithm of [14], which maintains the Euler characteristics of the object during its execution. The second and third comparison methods, denoted PINK skel and PINK filter3d, are also medial-axis type thinning approaches based on the discrete bisector function [10] and critical kernels [5]. The fourth comparison method is the approach of Jin *et al.* [12, 13], discussed in Section 2.

We implemented our method in C/C++ on a machine with a 12 core Intel Xeon(R) 2.7 GHz processor and 256 GB RAM.⁴ For all the results shown in this section, the spurious branch probability, the only parameter of our algorithm, was set to $t = 1e^{-12}$. The implementation of the thinning algorithm from [14] is provided through Fiji/ImageJ2 in the Skeletonize3D plugin, authored by Ignacio Arganda-Carreras. The implementations of [10] and [5] were provided by the scripts ‘skel’ and ‘skelfilter3d’, respectively, from the PINK library [9]. The implementation of the

method of Jin *et al.* was kindly provided by the authors. We did try to evaluate our datasets using the curve skeleton algorithm and implementation of [8], but that approach failed to return a skeleton. We hypothesize that this failure is a result of the thinness of some of the structures in our datasets, which are sometimes only one voxel wide, since the discussion in [8] specifically mentions that the algorithm may fail for thin structures.

4.1. Accuracy and robustness to noise

In order to illustrate the accuracy of our method in comparison with the state-of-the-art approaches, Figure 2 shows the original surface and curve skeletons computed with all five methods for Dataset B. As expected, the thinning algorithm, PINK skel, and PINK filter3d methods were not able to deal adequately with the noise in our datasets, and created many extra, small branches. In addition, the PINK filter3d method removes some branches. Jin *et al.*’s method performed better than the thinning algorithms with respect to noise, although it still presented some small spurious branches. In addition, this method is unable to deal with loops or cycles in the original structure. Our method is robust to the noise in our datasets and also was able to deal with loops in the curve skeleton. High-resolution images and results for the other datasets are also available in the supplementary materials.




Figure 2. **Best viewed in color.** Detailed view of the results from Dataset B: the surface reconstruction with noise of a real tree with a supporting metal pole, and curve skeletons computed with the thinning algorithm, PINK skel script, PINK filter3d script, Jin *et al.* method, and our proposed method. The different colors in 2f represent the curve skeleton segments identified during the course of the algorithm. Figures for Datasets A, the complete view of B, and C-I are given in the supplementary materials.

We also assessed the performance of our method under increasingly noisy conditions on a synthetic 3D model, which serves as a ground truth. We iteratively add noise to the ground truth model by randomly choosing, with uniform probability, $(p/2) \times n$ surface voxels which have non-surface neighbors to be deleted and another set of the same size to which a new neighboring voxel is added. The parameter p represents the proportion of voxels to be altered, and in our experiments $p = 0.05$. For subsequent iterations, we repeat the process using the voxel occupancy map from the previous iteration such that, after n_i iterations, the model has either noisy protrusions or depressions of at most

⁴The source code is available at [27].

n_l voxels. A closeup view of the model without noise and with a noise level of $n_l = 14$ as well as the corresponding curve skeletons computed using our method can be found in the supplementary materials.

To quantify the effect of noise on each of the skeletonization methods, we compute the root mean squared error (RMSE) of the skeletons generated by each approach in comparison with the ground truth skeleton. That is, for each voxel in the curve skeleton of the ground truth model, we find the closest voxel in the curve skeleton of the noisy model and use the sum of the squared closest distances to compute the RMSE. Figure 3 shows the corresponding results for our method and the comparison approaches. The x-axis in the graph represents the maximum voxel noise according to the process described in the preceding paragraph, and the y-axis shows the RMSE in terms of voxel distances. As the figure indicates, our method outperforms all the other approaches by a significant margin. The second best approach is given by the method of Jin *et al.*, which has an average RMSE 20% higher than our method.




Figure 3. **Best viewed in color.** Root mean squared error of the curve skeletons computed with the comparison methods and our method, as compared to the ground truth curve skeleton. The PINK filter3d method is not reported, since its minimum error value is 3356.

4.2. Computational efficiency

Figure 4 summarizes the time performance of the methods on the nine trees. A general ordering with respect to increasing run time is: 1) our method, 2) thinning, 3) PINK skel, 4) PINK filter3d, and 4) the Jin *et al.* method. The method proposed by Jin *et al.*, which has as one of its components geodesic path computation, has the longest run time of the methods. Note that the run times of the comparison methods include loading and saving the results, whereas ours does not. The loading and saving portions required to run the thinning algorithm for dataset E, which has $N = 58$ million voxels, is 1.2 seconds. Since we do not have access to the source code of Jin *et al.*, assessing the time spent loading and saving is difficult, but we assume that it is of the same order of magnitude. For the PINK scripts, intermediate results are loaded and saved in temporary locations,

which affects run time. Nevertheless, our curve skeletonization method is able to compute curve skeletons one to three orders of magnitude faster than the other methods. It executes in less than four seconds even for very large models.




Figure 4. **Best viewed in color.** Run times of our curve skeletonization method on the nine datasets in comparison with the existing approaches. The vertical axis is shown on a logarithmic scale due to the dramatic differences between our method and the existing approaches.

4.3. Results on traditional skeletonization datasets and influence of t on the results

While we are interested in real-world, elongated, noisy objects, we also evaluated our method on some commonly-used smooth models. The proposed curve skeleton method used in the context of smooth models produced the general structure of those objects and was also able to detect loops when they were present, for instance for the camel and seahorse examples. We also performed experiments with the only user-supplied threshold in our method. When $t = 1$, no segments are discarded. The resulting curve skeleton resembles a more dense version of the thinning result. When $t = 0.0001$, the result resembles the Jin *et al.* result except that loops are preserved. Results for both of these items are found in the supplementary materials.

5. Conclusions

Understanding the structure of complex elongated branching objects in the presence of noise is a challenging problem with important real-world applications. In this paper, we presented a fast and robust algorithm to compute curve skeletons of such real-world objects. These curve skeletons provide most of the information necessary to represent the structure of these objects. A large portion of the paper centered on how the ideas of BFS could be exploited to create an efficient curve skeletonization procedure. Our approach is able to detect connected segments and performs pruning in the course of the algorithm, so those steps do not

need to be performed separately. The small run times of less than a few seconds make this method suitable for automation tasks where real-time decisions are required.

References

- [1] AIM@Shape. Visualization virtual services. <http://visionair.ge.imati.cnr.it/>, 2011. last accessed 20 January, 2016. **34**
- [2] C. Arcelli, G. S. di Baja, and L. Serino. Distance-driven skeletonization in voxel images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(4):709, 2011. **2**
- [3] C. Aslan, A. Erdem, E. Erdem, and S. Tari. Disconnected skeleton: Shape at its absolute scale. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(12):2188–2203, 2008. **1**
- [4] X. Bai, L. Latecki, and W. yu Liu. Skeleton pruning by contour partitioning with discrete curve evolution. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(3):449–462, March 2007. **1**
- [5] G. Bertrand and M. Couprie. Two-dimensional parallel thinning algorithms based on critical kernels. *Journal of Mathematical Imaging and Vision*, 31(1):35–56, 2008. **2, 7**
- [6] T. Chaperon, F. Goulette, and C. Lurgeau. Extracting cylinders in full 3D data using a random sampling method and the Gaussian image. In *Proceedings of the Vision Modeling and Visualization Conference 2001, VMV '01*, pages 35–42, 2001. **2**
- [7] N. D. Cornea, D. Silver, and P. Min. Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 13(3):530–548, May 2007. **1**
- [8] N. D. Cornea, D. Silver, X. Yuan, and R. Balasubramanian. Computing hierarchical curve-skeletons of 3D objects. *The Visual Computer*, 21(11):945–955, 2005. **7**
- [9] M. Couprie. Pink image processing library. <http://pinkhq.com/>, December 2013. last accessed 20 January, 2016. **7**
- [10] M. Couprie, D. Coeurjolly, and R. Zour. Discrete bisector function and euclidean skeleton in 2D and 3D. *Image and Vision Computing*, 25(10):1543 – 1556, 2007. **2, 7**
- [11] W.-B. Goh. Strategies for shape matching using skeletons. *Computer Vision and Image Understanding*, 110(3):326 – 345, 2008. *Similarity Matching in Computer Vision and Multimedia*. **1**
- [12] D. Jin, K. Iyer, E. Hoffman, and P. Saha. A new approach of arc skeletonization for tree-like objects using minimum cost path. In *Pattern Recognition (ICPR), 2014 22nd International Conference on*, pages 942–947, Aug 2014. **2, 7**
- [13] D. Jin, K. S. Iyer, C. Chen, E. A. Hoffman, and P. K. Saha. A robust and efficient curve skeletonization algorithm for tree-like objects using minimum cost paths. *Pattern recognition letters*, 76:32–40, 2016. **2, 7**
- [14] T.-C. Lee, R. L. Kashyap, and C.-N. Chu. Building skeleton models via 3-d medial surface/axis thinning algorithms. *CVGIP: Graph. Models Image Process.*, 56(6):462–478, Nov. 1994. **7**
- [15] Y.-J. Liu, J.-B. Zhang, J.-C. Hou, J.-C. Ren, and W.-Q. Tang. Cylinder detection in large-scale point cloud of pipeline plant. *Visualization and Computer Graphics, IEEE Transactions on*, 19(10):1700–1707, Oct 2013. **2**
- [16] D. Macrini, K. Siddiqi, and S. Dickinson. From skeletons to bone graphs: Medial abstraction for object recognition. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, June 2008. **1**
- [17] H. Medeiros, D. Kim, J. Sun, H. Seshadri, S. A. Akbar, N. M. Elfiky, and J. Park. Modeling dormant fruit trees for agricultural automation. *Journal of Field Robotics*, pages n/a–n/a, 2016. **2**
- [18] A. Meijster, J. B. Roerdink, and W. H. Hesselink. A general algorithm for computing distance transforms in linear time. In *Mathematical Morphology and its applications to image and signal processing*, pages 331–340. Springer, 2002. **4, 12**
- [19] B. Miklos, J. Giesen, and M. Pauly. Discrete scale axis representations for 3D geometry. In *ACM SIGGRAPH 2010 Papers, SIGGRAPH '10*, pages 101:1–101:10, New York, NY, USA, 2010. ACM. **1**
- [20] P. Min. Binvex. <http://www.patrickmin.com/>, 2015. last accessed 20 January, 2016. **34**
- [21] A. S. Montero and J. Lang. Skeleton pruning by contour approximation and the integer medial axis transform. *Computers & Graphics*, 36(5):477 – 487, 2012. *Shape Modeling International (SMI) Conference 2012*. **1**
- [22] F. Nooruddin and G. Turk. Simplification and repair of polygonal models using volumetric techniques. *Visualization and Computer Graphics, IEEE Transactions on*, 9(2):191–205, April 2003. **34**
- [23] M. P. Pound, A. P. French, J. A. Fozard, E. H. Murchie, and T. P. Pridmore. A patch-based approach to 3d plant shoot phenotyping. *Machine Vision and Applications*, 27(5):767–779, Jul 2016. **7**
- [24] T. Rabbani and F. Van Den Heuvel. Efficient hough transform for automatic detection of cylinders in point clouds. *ISPRS WG III/3, III/4*, 3:60–65, 2005. **2**
- [25] J. Silevela and J. Portillo. Breadth-first search and its application to image processing problems. *Image Processing, IEEE Transactions on*, 10(8):1194–1199, Aug 2001. **12**
- [26] A. Tabb. Shape from silhouette probability maps: Reconstruction of thin objects in the presence of silhouette extraction and calibration error. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 161–168, June 2013. **7**
- [27] A. Tabb. Code from: Fast and robust curve skeletonization for real-world elongated objects. <http://dx.doi.org/10.15482/USDA.ADC/1399689>, 2017. **2, 7**
- [28] A. Tabb and H. Medeiros. A robotic vision system to measure tree traits. In *IEEE RSJ International Conference on Intelligent Robots and Systems*, 2017. **1**
- [29] T.-T. Tran, V.-T. Cao, and D. Laurendeau. Extraction of cylinders and estimation of their parameters from point clouds. *Computers & Graphics*, 46(0):345 – 357, 2015. *Shape Modeling International 2014*. **2**

- [30] Y.-S. Wang and T.-Y. Lee. Curve-skeleton extraction using iterative least squares optimization. Visualization and Computer Graphics, IEEE Transactions on, 14(4):926–936, July 2008. [2](#)
- [31] A. Ward and G. Hamarneh. The groupwise medial axis transform for fuzzy skeletonization and pruning. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 32(6):1084–1096, June 2010. [1](#)

Fast and robust curve skeletonization for real-world elongated objects: Supplementary materials

6. Computational complexity

In this section, we analyze the computational complexity of the algorithms in this paper. For reference ease, we include all of the algorithms from the main paper.

Algorithm 5 Modified BFS Algorithm

```

1: while  $|\mathbb{F}_A^{(k)}| > 0$  do
2:   for each voxel  $v_i \in \mathbb{F}_A^{(k)}$  do
3:     for each voxel  $v_j \in \mathcal{N}_i$  such that  $(l_j > l_i)$  do
4:        $l_j = \min(l_j, l_i + w_j + ||v_j - v_i||)$ 
5:    $\mathbb{F}_A^{(k)} = \mathbb{F}_A^{(k)} \cup \mathbb{F}_B^{(k)}$ 
6:    $\mathbb{N}^{(k)} = \{v_j | l_j > l_i, \forall v_j \in \mathcal{N}_i, \forall v_j \notin \mathbb{F}^{(k)}, \forall v_i \in \mathbb{F}^{(k)}\}$ 
7:    $l_{min} = \min_{v_j \in \mathbb{N}^{(k)}} l_j$ 
8:    $\mathbb{F}_A^{(k+1)} = \{v_j | l_i < l_{min}, \forall v_i \in \mathbb{F}^{(k)}, \forall v_j \in \mathbb{N}^{(k)}\}$ 
9:    $\mathbb{F}_B^{(k+1)} = \{v_i | l_i \geq l_{min}, \forall v_i \in \mathbb{F}^{(k)}, |\mathcal{N}_i \cap \mathbb{N}^{(k)}| > 0\}$ 
10:   $k = k + 1$ 

```

6.1. Complexity of the modified BFS algorithm

We now analyze the computational complexity of Algorithm 5. Note that a voxel is a member of $\mathbb{F}_A^{(k)}$ for any k exactly once. On line 8, voxels in $\mathbb{F}_A^{(k+1)} \subseteq \mathbb{N}^{(k)}$ and a voxel in $\mathbb{N}^{(k)}$ cannot be in $\mathbb{F}^{(k)}$: $\mathbb{F}^{(k)} \cap \mathbb{N}^{(k)} = \emptyset$ (from line 6). In addition, the condition $l_j > l_i$ ensures that a voxel in $\mathbb{F}_A^{(k+1)}$ has not already been discovered, meaning that voxels in $\mathbb{F}_A^{(k+1)}$ cannot be members of the frontier $\mathbb{F}^{(m)}$ at any previous iteration m , $m = 0, 1, \dots, k - 1$. Hence $\mathbb{F}_A^{(k)} \cap \mathbb{F}_A^{(j)} = \emptyset$ for $k \neq j$. Finally, since all voxels are eventually discovered, the sets $\mathbb{F}_A^{(k)}$ form a partition of \mathbb{V} , implying

$$\bigcup_k \mathbb{F}_A^{(k)} = \mathbb{V} \quad (\text{S1})$$

and

$$\sum_k |\mathbb{F}_A^{(k)}| = |\mathbb{V}| = n \quad (\text{S2})$$

Consequently, for all iterations of Algorithm 5, line 2 will be performed n times.

On the other hand, a voxel v_i may be a member of $\mathbb{F}_B^{(k)}$ multiple times, waiting for the condition $l_i < l_{min}$ to be true so that its neighbors are moved into $\mathbb{F}_A^{(k)}$ (line 8). Because the conditions for computing set $\mathbb{N}^{(k)}$, and therefore l_{min} (lines 6 - 7), for the members of $\mathbb{F}_B^{(k)}$ are dependent on the

Algorithm 6 Determination of curve skeleton segment from BFS2 and d_i

```

1:  $v_c = v_{s,1}$ 
2:  $\mathbb{C}^{(m)} = \{v_{s,1}\}$ 
3: while  $(v_c \neq v_t) \wedge (v_c \notin \mathbb{C})$  do
4:   Determine  $d^* = \max_{v_i \in \mathcal{N}_c \wedge l_c > l_i} (d_i)$  where  $d_i$  is the distance transform of  $v_i$ 
5:   Compute  $v_n = \arg \min_{v_j \in \mathcal{N}_c^*} (l_j)$  where  $l_j$  is the BFS2 label of  $v_j$  and  $\mathcal{N}_c^* = \{v_j | \exists v_j \in \mathcal{N}_c, d_j = d^*\}$ 
6:    $\mathbb{C}^{(m)} = \mathbb{C}^{(m)} \cup \{v_n\}$ 
7:    $v_c = v_n$ 

```

Algorithm 7 Loop handling

```

1:  $\mathbb{C}_t = \bigcap_{j \in [0, a-1]} \mathbb{C}'_j$ 
2: for  $j = 0$  to  $a - 1$  do
3:    $\mathbb{C}'_j = \mathbb{C}'_j \setminus \mathbb{C}_t$ 
4:   for  $k = j + 1$  to  $a - 1$  do
5:      $\mathbb{C}'_j = \mathbb{C}'_j \setminus \mathbb{C}_k$ 
6:    $\mathbb{C} = \mathbb{C} \cup \{\mathbb{C}'_j\}$ 

```

Algorithm 8 Proposed skeletonization algorithm

```

1: Determine seed voxel  $v^*$  and make initial curve skeleton  $\mathbb{C}^{(0)} = v^*$ 
2: repeat
3:   Update BFS1 labels using Alg. 5 with  $\mathbb{F}_A^{(0)} = \mathbb{C}^{(m-1)}$ 
4:   Locate a proposed endpoint  $v_t$  from BFS1
5:   Create BFS2 labels using Alg. 5 with  $\mathbb{F}_A^{(0)} = \{v_t\}$ 
6:   Create  $\mathbb{C}^{(m)}$  by tracing paths in BFS2 labels from existing curve skeleton  $v_t$  according to Alg. 6
7:   if (Number of FCCs == 1) then
8:     Accept or decline curve skeleton segments using classification method. If accepted  $\mathbb{C} = \mathbb{C} \cup \{\mathbb{C}^{(m)}\}$  and  $m = m + 1$ 
9:   else
10:    Check for loops using Alg. 7.
11: until No more endpoint hypotheses are found

```

composition of $\mathbb{F}_A^{(k)}$, l_{min} cannot be stored from previous iterations and still be relevant.

Finding set $\mathbb{N}^{(k)}$ in Algorithm 5 takes time $|\mathbb{F}_A^{(k)}| + |\mathbb{F}_B^{(k)}|$ per iteration k , and line 7 can be computed while \mathbb{N} is found. Lines 8 and 9 also take time $|\mathbb{F}_A^{(k)}| + |\mathbb{F}_B^{(k)}|$ per iteration k .

Since we know from Equation S2 that $\sum_k |\mathbb{F}_A^{(k)}| = n$, the asymptotic lower bound is $\Omega(n)$, and it is a strict bound. For instance, consider $\mathbb{F}_A^{(0)}$ contains one voxel, which is an endpoint of a 1-dimensional line in 3D space. In this case, $|\mathbb{F}_A^{(k)}| = 1$ and $|\mathbb{F}_B^{(k)}| = 0$ for all k , and the maximum value

of k is n .

Determining the value of $\sum_k |\mathbb{F}_B^{(k)}|$ in general is problematic as it depends strongly on the shape involved, on the composition of $\mathbb{F}_A^{(0)}$, and weights w_i .

6.2. Complexity of modified BFS algorithm with zero weights

The Euclidean distance between neighboring voxels is $d \in \{1, \sqrt{2}, \sqrt{3}\}$. We claim that a voxel can only be in one of the frontier sets at most three times; first, it enters the frontier through $\mathbb{F}_A^{(k)}$, and if the voxel remains in the frontier, the only other sets it may belong to are $\mathbb{F}_B^{(k+1)}$ and $\mathbb{F}_B^{(k+2)}$, at which point it exits the frontier sets. Hence

$$\sum_k (|\mathbb{F}_A^{(k)}| + |\mathbb{F}_B^{(k)}|) \leq 3n \quad (\text{S3})$$

The asymptotic upper bound of Algorithm 5 is then $\mathcal{O}(n)$.

Algorithm 5 using zero weights has many of the same characteristics as the breadth-first distance transform described in Section V of [25]. One major difference is that in [25] neighbors are assumed to have the same distance from each other, which is not an assumption we share when working with 26-connected voxels in 3D. However, the algorithm with zero weights retains the asymptotic upper bound of $\mathcal{O}(n)$ as the method in [25].

6.3. Complexity of modified BFS with non-zero weights

The complexity of the modified BFS algorithm with non-zero weights is $\mathcal{O}(nd_{max})$, assuming that the weights are computed based on a distance transform $w_i = d_{max} - d_i$.

Our discussion of Algorithm 5 for zero weights provided an asymptotic upper bound $\mathcal{O}(n)$, and determining this bound came down to determining how many times a voxel could possibly be a member of the frontier. For the analysis of Algorithm 5 when weights are non-zero, we return to similar questions. From Algorithm 5, voxels are in \mathbb{F}_A once, and are in \mathbb{F}_B multiple times, until the condition $l_i < l_{min}^{(k)}$ is satisfied.

Let us determine the maximum number of times a voxel resides in the frontier sets. The maximum distance label d_{max} is the maximum difference between any two distance labels, and the L2 norms between any two neighboring voxels belong to the set $\{1, \sqrt{2}, \sqrt{3}\}$. Then, the maximum number of times a voxel can possibly be in the frontier sets is $3d_{max}$, leading to asymptotic upper bound $\mathcal{O}(nd_{max})$. We consider shapes that produce the largest values of $\frac{d_{max}}{n}$. The shape that maximizes $\frac{d_{max}}{n}$ is a sphere. In that scenario, d_{max} is the radius of the sphere, and the relationship between n and d_{max} is $n = \frac{4}{3}\pi(d_{max})^3$. Then, performing the relevant substitutions we have the asymptotic upper bound of this step as $\mathcal{O}(n^{\frac{4}{3}})$.




Figure S1. Example of a shape with a high proportion of proposed endpoints t relative to n .

6.4. Computational complexity of the curve skeleton method

We now consider the complexity of the curve skeleton method of Algorithm 8, excluding the complexity of the distance transform step, which is $\mathcal{O}(n)$ using the approach of [18]. Let the number of proposed endpoints, or iterations of Algorithm 8, be $||\mathbb{V}_t||$. Computing the BFS1 labels (step 2.1) and BFS2 labels (step 3.1) per iteration has asymptotic upper bound $\mathcal{O}(d_{max}n)$. The computation of the BFS1 and BFS2 labels is repeated $||\mathbb{V}_t||$ times, giving $\mathcal{O}(||\mathbb{V}_t||d_{max}n)$ as an upper bound. As in Section 6.3, we try to determine shapes that maximize $\frac{||\mathbb{V}_t||}{n}$ to find an asymptotic upper bound on $||\mathbb{V}_t||$. We have given an example of a particular shape in Figure S1 where $||\mathbb{V}_t|| = \frac{2}{3}(n+2)$. More generally, the number of extremities can be no larger than the number of voxels, hence we can assume that $||\mathbb{V}_t|| = \mathcal{O}(n)$. In summary, the asymptotic upper bound of the complete algorithm is $\mathcal{O}(n^{\frac{7}{3}})$.

We note that in our experiments involving elongated objects, both $||\mathbb{V}_t||$ and d_{max} are extremely small relative to n as shown in Table 1. In those datasets, the maximum value of $||\mathbb{V}_t||$ relative to n is $||\mathbb{V}_t|| \leq 0.002n$ (Dataset B), and the greatest value of d_{max} is 16, or in terms of n , $d_{max} \leq 0.0003n$ (Dataset A). Consequently, while the asymptotic upper bound is greater than quadratic, if $||\mathbb{V}_t||$ and d_{max} can be assumed to be small constants as in the case of elongated objects, in practice the algorithm runs quickly. This is highlighted in Figure S2, which shows the runtime of our method as a function of the number of occupied voxels n for each of the datasets in Table 1. As the figure indicates, despite a four times increase in the number of occupied voxels between the smallest and the largest dataset, the run time of our algorithm increases slowly.

7. Topological Stability

There are three steps in the algorithm where decisions may be made deterministically or randomly in the case of equal labels: (1) the selection of v^* in Section 3.2, (2) endpoint candidates v_t in 3.3.2, and (3) frontier connecting voxel $v_{s,1}$ in 3.4.2. If topological stability is desired, the




Figure S2. Execution time of our algorithm as a function of the number of occupied pixels in each dataset.

Table 1. Characteristics of the nine datasets used in our evaluation. n is the number of occupied voxels/nodes, N is the number of voxels in the grid, d_{max} is the largest voxel to surface distance, and $||\mathbb{V}_t||$ is the number of proposed endpoints v_t .

ID	n	N	d_{max}	$ \mathbb{V}_t $
A	55,156	27,744,000	16	15
B	88,407	56,832,000	5	174
C	88,798	45,240,000	7	145
D	92,892	80,640,000	6	154
E	98,228	58,464,000	9	50
F	136,497	80,640,000	7	134
G	158,686	64,512,000	9	128
H	176,820	80,640,000	10	97
I	246,654	80,640,000	10	83

following protocol could be employed to deterministically select a voxel when there are multiple voxels with the same label. The voxels of equal label are placed into a vector, and then the x, y, z coordinates would be sorted as specified by the user. One such ordering would be to sort the coordinates by x value, and then in case of ties by y value, and then in case of ties by z . This kind of repeatable ordering would produce reproducible results that would be more topologically stable than a random ordering.

8. Additional figures for loop handling step (section 3.5.2 in the main paper)

Figure S3 shows the loop handling procedure. As shown in Figure S3c, only regions with loops are recovered.




Figure S3. **Best viewed in color.** An example of the loop handling procedure. S3a shows the surface in blue, and S3b is the existing curve skeleton in red. Fig. S3c shows the result after Algorithm 7 is performed.

9. Additional Results

9.1. Simulated noise experiment

Figure S4 shows the models used for the simulated noise experiment, and computed curve skeletons using our proposed method.

As in the experiments with the real datasets, the comparison methods are characterized by greater numbers of spurious voxels than our method as the noise level increases. We represent this in Figure S6, which shows the number of voxels in the curve skeletons for each of the methods as a function of the noise level. The figure shows that, as the noise level rises, the number of voxels also increases for all the methods. For the thinning and PINK methods, this increase is dramatic, in some cases up to 9 times the initial value. The Jin *et al.* method also shows some additional voxels as the noise rises, particularly for noise levels higher than 6, showing a total increase of approximately 50% from zero noise to noise level 14. Our method has the least additional voxels, showing a total increase of only 4%.

9.2. Tree models acquired in field conditions

Figures S7 to S24 show high resolution images of the nine real trees used to evaluate our method as described in Table 1 of the main paper. See the discussion in the main paper in the Experiments section for a more detailed description of the figures below.




Figure S4. **Best viewed in color.** Detail of a synthetic model of a tree without noise S4a and with noise S4c at iteration 14. The curve skeleton of the ground truth is given in the top row, while curve skeleton of the noisy object is given in the second row.

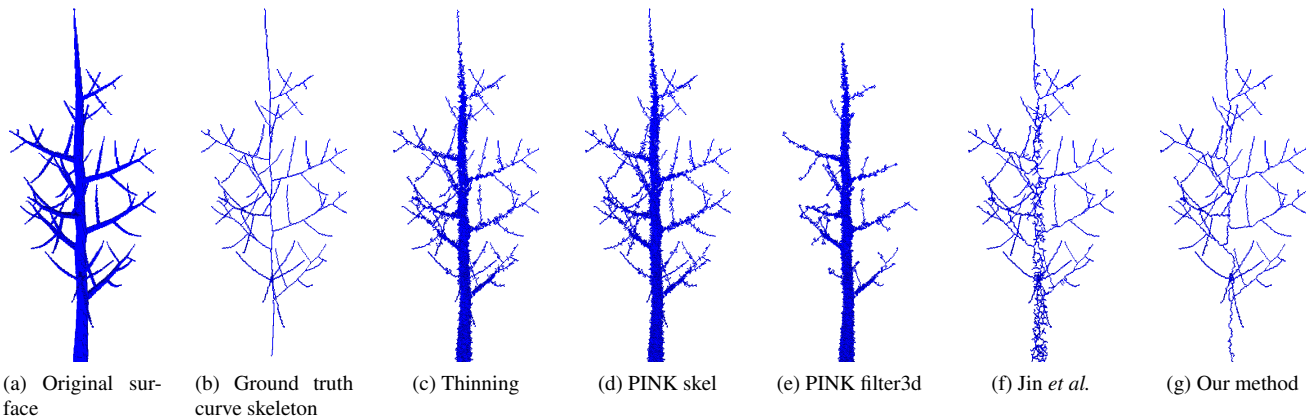


Figure S5. Synthetic model used in the evaluation of the robustness to noise and the corresponding outputs of each algorithm for a noise level of 14. Note that in the presence of noise the skeletons generated by the thinning algorithm as well as the two PINK methods contain a substantially higher number of voxels and no longer consist of thin one-dimensional segments.






Figure S6. **Best viewed in color.** Number of voxels in the curve skeletons of the synthetic model as a function of the noise level. The vertical axis is shown on a logarithmic scale due to the dramatic differences between our method and most of the comparison approaches.



(a) Surface (b) Thinning (c) PINK skel

Figure S7. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset A (part 1 of 2).



(a) PINK filter3d (b) Jin *et al.* (c) Our method
Figure S8. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset A (part 2 of 2).


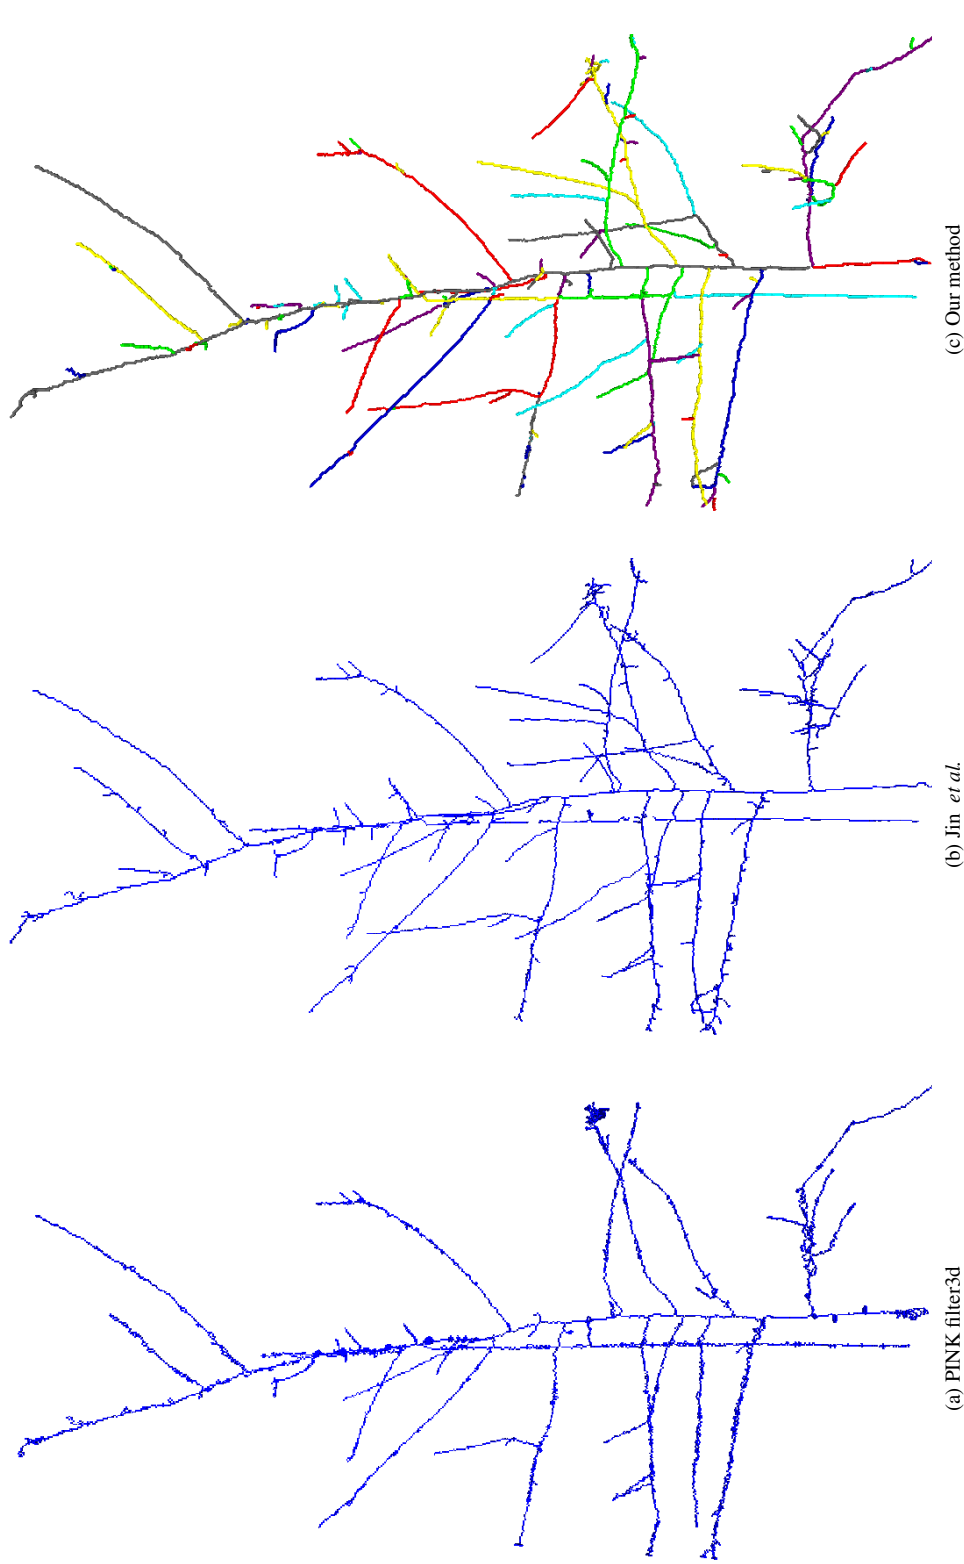


Figure S9. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset B (part 1 of 2).



(a) PINK filter3d (b) Jin *et al.* (c) Our method

Figure S10. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset B (part 2 of 2).





Figure S11. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset C (part 1 of 2).



(a) PINK filter3d
 Figure S12. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset C (part 2 of 2).
 (b) Jin *et al.*
 (c) Our method






Figure S13. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset D (part 1 of 2).




(a) PINK filter3d
 Figure S14. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset D (part 2 of 2).
 (b) Jin *et al.*
 (c) Our method




(a) Surface (b) Thinning (c) PINK skel

Figure S15. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset E (part 1 of 2).




(a) PINK filter3d (b) Jin *et al.* (c) Our method
 Figure S16. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset E (part 2 of 2).




(a) Surface (b) Thinning (c) PINK skel

Figure S17. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset F (part 1 of 2).




(a) PINK filter3d (b) Jin *et al.* (c) Our method

Figure S18. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset F (part 2 of 2).



(a) Surface (b) Thinning (c) PINK skel

Figure S19. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset G (part 1 of 2).




(a) PINK filter3d

(b) Jin *et al.*


(c) Our method

Figure S20. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset G (part 2 of 2).



(a) Surface
(b) Thinning
(c) PINK skel

Figure S21. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset H (part 1 of 2).



(a) PINK filter3d

(b) Jin *et al.*

(c) Our method

Figure S22. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset H (part 2 of 2).





Figure S23. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset I (part 1 of 2).



(a) PINK filter3d (b) Jin *et al.* (c) Our method

Figure S24. **Best viewed in color.** Original surface, comparison curve skeletons, and curve skeleton computed with our method, for Dataset I (part 2 of 2).

9.3. Computer graphics models


The curve skeleton algorithm is demonstrated on commonly-used computer graphics models in Figures S25-S25. Surfaces shown in S25a to S28a provided courtesy of INRIA, owner of S28c unknown, all via AIM@SHAPE-VISIONAIR Shape Repository [1]. All models were converted from mesh to voxels using the algorithm of [22] as implemented in [20].

9.4. Influence of parameter t on results


We performed some experiments with the one user-supplied threshold, t in our method, and show results in Figure S29. When $t = 1$, no segments are discarded. The resulting curve skeleton resembles a more dense version of the thinning result. When $t = 0.0001$, the result resembles the Jin *et al.* result except that loops are preserved. All results generated in this paper for the proposed method, except in this section, were generated with $t = 1e^{-12}$. Consequently, different values of t may be chosen depending on the application.

10. Additional results without comparisons


To further illustrate the performance of our approach, Figures S30 to S39 show additional high resolution images of results obtained using our method.




(a) Surface



(b) Skeleton




(c) Surface




(d) Skeleton


Figure S25. Original surfaces and skeletons computed with our method.




(a) Surface



(b) Skeleton




(c) Surface




(d) Skeleton


Figure S26. Original surfaces and skeletons computed with our method.




(a) Surface



(b) Skeleton




(c) Surface



(d) Skeleton


Figure S27. Original surfaces and skeletons computed with our method.




(a) Surface



(b) Skeleton



(c) Surface



(d) Skeleton

Figure S28. Original surfaces and skeletons computed with our method.

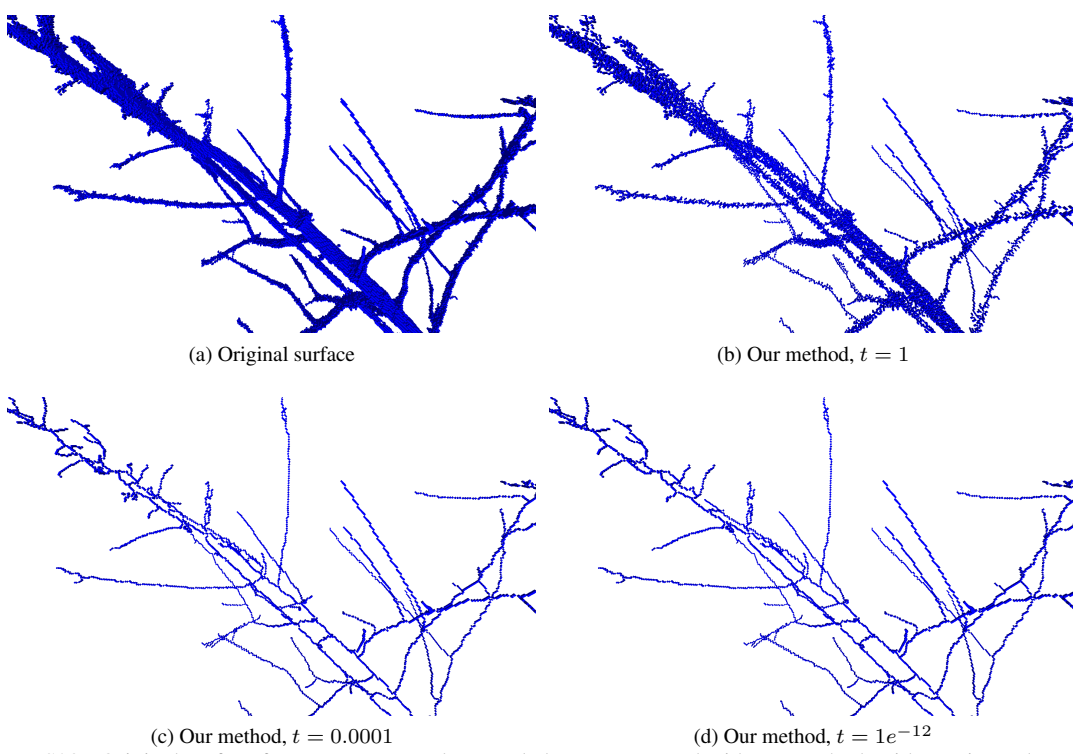




Figure S29. Original surface from Dataset B and curve skeletons computed with our method, with varying values of t .




(a) Surface

(b) Our method

Figure S30. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface




(b) Our method

Figure S31. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface




(b) Our method

Figure S32. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface



(b) Our method

Figure S33. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface



(b) Our method

Figure S34. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface



(b) Our method

Figure S35. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface



(b) Our method

Figure S36. **Best viewed in color.** Original surface and curve skeleton computed with our method.




(a) Surface



(b) Our method

Figure S37. **Best viewed in color.** Original surface and curve skeleton computed with our method.



(a) Surface

(b) Our method

Figure S38. **Best viewed in color.** Original surface and curve skeleton computed with our method.



(a) Surface

(b) Our method

Figure S39. **Best viewed in color.** Original surface and curve skeleton computed with our method.