Robust Extraction of 1D Skeletons from Grayscale 3D Images

Emilio Antúnez
Dept. of Electrical Engineering
Stanford University
eantunez@stanford.edu

Leonidas Guibas
Dept. of Computer Science
Stanford University
guibas@cs.stanford.edu

Abstract

We propose novel thinning conditions for grayscale 3D images that do not rely on image thresholding or segmentation. The resulting grayscale skeletonization routine extracts a topologically-consistent skeleton of curve-like features at all scales and intensities of the image. Our thinning conditions yield rich structural data that can be used to identify and delete low-contrast shape features created by blurring and noise.

1. Introduction

Medical imaging techniques such as MRI, CT, and EM generate grayscale 3D images of a scanned volume. Often, this representation is unnecessarily dense; for imaged subjects like blood vessels or proteins, relevant structural information is more effectively conveyed by a *curve skeleton* – a thin, largely one-dimensional subset of voxels centered on high-intensity regions.

For binary images, well-known morphological thinning techniques extract a skeleton that is topologically equivalent to the foreground set. These methods are not easily applied to grayscale images, where fuzzy boundaries and global intensity variations hinder the segmentation of desired structures. Extracting a skeleton from one or more binary versions of the image can be unnecessarily costly and yield fragile results.

We propose a version of morphological thinning that is performed directly on the grayscale volume. The resulting skeletonization routine is not only more efficient than previous work, but more informative; by inspecting all gray levels at once, our algorithm detects local structure not visible at one intensity. These local cues are useful for filtering out low-contrast features and noise.

2. Background and Previous Work

2.1. Basic Concepts

Conventional 3D images are defined over the 3D integer space, $\mathbb{Z}^3$. Each point in an image is called a *voxel*.

Every voxel v has 26 *neighbors* whose coordinate values each differ by at most one unit from its own; this *neighborhood* is labeled $N_{26}[v]$. Of these neighbors, 18 will differ from v in at most two coordinates, and 6 will differ from v in at most one coordinate; these neighborhoods are labeled $N_{18}[v]$ and $N_6[v]$, respectively. The set of voxels in $N_k[v]$ are said to be *k-adjacent* to v.

A *k-connected path* is a voxel sequence where every voxel is k-adjacent to the one before it.

2.2. Discrete 3D Images

Discrete images are defined by their *intensity function*, $I : \mathbb{Z}^3 \to \mathbb{R}$. For binary images, a bilevel I divides voxels into a high-valued *foreground set* and a low-valued *background set*. For grayscale images, the multilevel I implies a one parameter-family of foreground sets $I_\lambda = \{x \in \mathbb{Z}^3 : I[x] \geq \lambda\}$ called λ*-sections*.

To interpret shapes in a binary image or λ-section, we must decide which type of k-adjacency defines connected components of the foreground and background. Our methods can be applied to images of any *connectivity class*, but we assume a 6-connected foreground and 26-connected background to shorten our discussion.

2.3. Binary 3D Image Thinning

Binary morphological thinning iteratively deletes as many points as possible from the foreground without changing the image's topology, leaving a thin "skeleton" of the original foreground set. These deletable voxels are called *simple points*; they are well-studied in both 2D and 3D, and many tests exist to detect them.

Our algorithm builds on this version of the test:

Theorem 1 ([2]). *A voxel x is simple if and only if:*

1. *There is at least one foreground point in $N_6[x]$.*
2. *There is at least one background point in $N_{26}[x]$.*
3. *Every two foreground points in $N_6[x]$ are connected by a foreground path in $N_{18}[x]$.*
4. *Every two background points in $N_{26}[x]$ are connected by a background path in $N_{26}[x]$.*

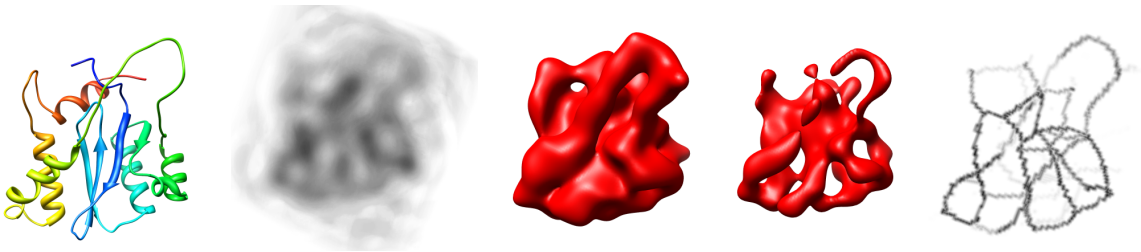

(a) (b) (c) (d) (e)

Figure 1. An electron density map at 8Å resolution (b) is generated from the atomic structure (PDB ID 2ITG) shown in (a). As illustrated by isosurfaces (c) and (d), no single threshold captures all desired features. The result of our skeletonization technique appears in (e).

2.4. Grayscale 3D Image Thinning

True grayscale thinning algorithms are rare; most applications assume that an initial segmentation step will extract an accurate foreground set for binary thinning. A few thinning algorithms, however, *do* consider grayscale information. One approach is to sample a limited number of λ-sections, then add thinning constraints to merge them into a single skeleton [1]. Another method constructs a rough foreground set, then deletes simple points in order of increasing intensity to carve out the skeleton [4][6].

An altogether different approach to grayscale thinning – one closely related to our own – uses no explicit foreground set. Instead, voxel values are iteratively reduced while preserving the topology of each binary section [3]. A skeletal voxel's final value equals the intensity of the grayscale structure it represents. To compute a voxel's new value on every thinning pass, though, the authors still rely on explicit simple point tests on up to 24 binary image sections. In the discussion to follow, we show that this binary approach is needlessly costly and ignores a great deal of useful grayscale structure.

3. Thinning Using Least-Extremal Paths

Our core contribution is an algorithm to identify all λ-sections for which a voxel is simple, without brute-force simplicity testing at each section. More importantly, in sections where a voxel is nonsimple, our method provides a local structure descriptor that indicates *why* it is nonsimple, and how the image must change locally to make the voxel deletable. This will be the key to detecting and filtering skeletal noise.

3.1. Least-Extremal Paths

Theorem 1 identified simple points based on foreground and background paths. We now introduce an analogue of these paths for grayscale images.

For any connected image path p, we define $\max(p)$ and $\min(p)$ to be the maximum and minimum intensities of any voxel in p, respectively.

A k-connected path between voxels a and b, restricted to set $X \subset \mathbb{Z}^3$, is called a *least-negative path* ($LNP[a, b, k, X]$) if there is no other such path with greater $\min(p)$ and a *least-positive path* ($LPP[a, b, k, X]$) if there is no other such path with lower $\max(p)$. Collectively, LNPs and LPPs are referred to as *least-extremal paths* (LEPs).

Though our LEP definition is general, our thinning conditions will only consider paths in a fixed domain, with fixed connectivity, and one common endpoint. For brevity, we will use the shortened notations $LNP[a] = LNP[a, n_{max}, 6, N_{18}]$ and $LPP[a] = LPP[a, n_{min}, 26, N_{26}]$, where n_{max} is a max-valued voxel in N_6 and n_{min} is a min-valued voxel in N_{26}.

The set of LNPs (or LPPs) from a single source to all other voxels can be computed easily using a modified version of Dijkstra's algorithm. This algorithm can return not only the optimal path intensities but also a tree of optimal paths rooted at the source.

3.2. LEP-Based Simplicity and Destructibility Tests

We now define simple point conditions, analogous to those of Theorem 1, for sections of a grayscale image.

Theorem 2. *In the λ-section I_λ of a (6,26)-connected grayscale image with intensity function I, a voxel x is simple if and only if:*

1. $\lambda \leq I[n_{max}]$.
2. $\lambda > I[n_{min}]$.
3. $\lambda \notin (\min(LNP[n]), I[n]], \forall n \in N_6$.
4. $\lambda \notin (I[n], \max(LPP[n])], \forall n \in N_{26}$.

To find the 32 *nonsimple ranges* above, we need only to compute $n_{max} \in N_6$, $n_{min} \in N_{26}$, an LNP tree rooted at n_{max}, and an LPP tree rooted at n_{min}.

As in [3], we define a *destructibility test* to check how far a voxel's value can be reduced without changing the topology of any λ-section. Unlike [3], we do not need to perform simple point tests on multiple sections; we simply compute the max value in any nonsimple range that falls at or below the voxel value.

3.3. The Extended Nonbranching Range

An extra condition is needed to preserve the endpoints of open curves. To be compatible with noise-filtering techniques discussed later, this test also must specify whether a voxel is the end of a *high-contrast* curve. Commonly-used endpoint conditions fail to do so if a grayscale curve fades out gradually. Instead, we use a condition that will preserve a voxel if it appears in a locally-1D curve for a wide range of binary sections.

A foreground voxel in a binary image is a *branching point* if it is 6-adjacent to three or more foreground voxels. For each voxel x in a grayscale image, we define the *extended nonbranching range* $(\lambda_k, I[x]]$, where I_{λ_k} is the highest section in which x is connected to a branching point by a foreground path of length k or less. Our thinning algorithm preserves x's value if this range is nonempty. (We found $k = 2$ gives good results.)

3.4. Grayscale Thinning Algorithm

With a destructibility condition and the nonbranching test defined, the final grayscale thinning algorithm visits each voxel in the image and lowers its value as far as possible without crossing into one of the nonsimple or nonbranching ranges, known collectively as *protected ranges*. The thinning is complete when no voxel values change after a given thinning pass.

3.5. Curve Skeleton Extraction

The grayscale image returned by our thinning routine will be optimally "thin"; any further reduction of voxel values would change the topology of some binary section. Extraction of the desired curve skeleton, however, involves more than simply accepting all voxels above a fixed background intensity.

Voxel surfaces/volumes may keep arbitrarily high intensities if bounded by higher-intensity curves/surfaces; these structures are known as *tunnels/cavities*.

We can identify voxels belonging to curvelike structures using the rules from Theorem 2: most of the desired skeletal voxels will have nontrivial protected ranges for condition 1, condition 3, or the nonbranching condition, which is not true for voxels belonging only to high-intensity surfaces or volumes.

4. Local Filtering of Topological Noise

The algorithm described in the previous section was designed to preserve skeletal curves for "foreground" shapes appearing at any intensity. This behavior is not always desirable, however. Image blurring and noise can create spurious, low-constrast shape features – appearing in only a narrow range of binary sections – that are not desirable in the final skeleton. Our grayscale simplicity test generates local cues that can be used to identify and eliminate these *shallow* structures.

4.1. Naive Contrast Filtering

The naive approach to this problem is to ignore all nonsimple ranges below some minimum length. This behaves as a contrast filter; only grayscale features with sharp, high-contrast boundaries are preserved.

4.2. Deleting Shallow Curves

Lack of a sharp boundary does not necessarily indicate a shallow feature; a high-intensity shape may have a wide, smoothly-varying boundary region. To preserve such shapes while deleting shallow curves, one should only require a minimum length for protected regions resulting from condition 2.1 or the branching condition. The resulting thinning algorithm erodes a low-contrast curve from its endpoints, and deletes the curve once it is reduced to a single voxel. A deeper curve's endpoints eventually achieve high-enough contrast during the thinning to halt the erosion and preserve the curve.

4.3. Collapsing Shallow Cavities

During thinning, a shallow cavity's interior is shielded from erosion by a thin, slightly higher-intensity bounding surface. For reasons mentioned above, contrast filtering is not a safe solution to this problem. We can, however, modify the thinning behavior so that a voxel surface barely separating two regions of vastly-different intensity is incrementally pushed toward the higher-intensity side. This will shrink and collapse shallow cavities as low intensities reach their exterior.

We can detect these deformable surface points during thinning by examining each $LPP[n]$: the intensity of n will be much higher than n_{min} and barely lower than $\max(LPP[n])$. If so, n is on the higher-intensity side of the surface; by setting $I[n] \leftarrow \max(LPP[n])$, the boundary is fattened inward (see Fig. 2). We do this for each of a boundary voxel's LPPs then repeat its destructibility test. In most cases, the voxel's value can safely be lowered, thus completing a surface deformation step.

4.4. Collapsing Shallow Tunnels

During thinning, a shallow tunnel becomes a voxel surface shielded from further erosion by a slightly higher-intensity boundary curve. Again, we would like to push this boundary inward until the tunnel collapses.

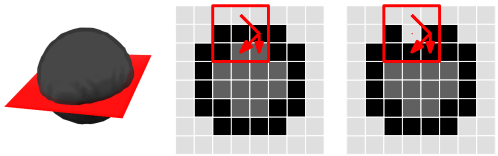

(a) (b) (c)

Figure 2. Cross-sections (b,c) of a cavity (a) show the collapsing step for one voxel.

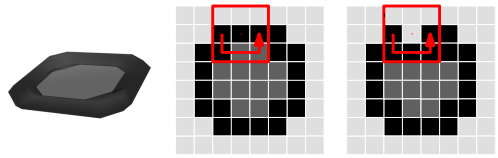

(a) (b) (c)

Figure 3. Planar views (b,c) of a tunnel (a) show the collapsing step for one voxel.

For cavities, we used $I[n_{min}]$ as a measure of "erosive pressure", to ensure that the surface was only deformed if a major intensity difference formed across it. The equivalent value here is λ_{se}, where $I_{\lambda_{se}}$ is the lowest section for which the voxel is a surface endpoint. As suggested by [5] for binary images, we label x a surface endpoint for those sections I_λ for which there exists no 6-connected cycle, c, in $N_{18}[x]$ with $\min(c) \geq \lambda$.

If a voxel x is a deformable curve point, then for some $n \in N_6[x]$ the "potential boundary intensity" $\lambda_{pb} = \min(I[x], I[n])$ will be much higher than λ_{se} and barely higher than $\min(LNP[n])$. If true, then n_{max}, x, and n are all boundary voxels, and $LNP[n]$ is a path bypasses x through the slightly lower-intensity voxel surface. By setting $I[y] \leftarrow \max(I[y], \lambda_b)$ for every $y \in LNP[n]$, the boundary curve is thickened inward near x (see Fig. 3). If we repeat the destructibility test on x, the voxel's value can safely be reduced in most cases, thus completing the boundary curve contraction step.

5. Results

An example of our algorithm's output can be seen in Figure 1, performed on an 8Å-resolution synthetic electron density map, generated from an atomic structure using the CNS software package. A similar dataset appears in [1]; note that our skeleton does not include the small, spurious loops seen in their output.

Each point in our skeleton also has an associated intensity. Curves resulting from blurring often have lower intensity than nearby curves; this may provide further opportunities for nonlocal noise filtering.

(a) (b) (c)

Figure 4. A shallow cavity (a) is thinned with cavity contraction off (b) and on (c).

Figure 4 shows our thinning result with and without cavity contraction. Note that the off-center cavity does not bias the skeleton in Figure 4(c); the result is a straight line, as expected for that solid without its cavity.

6. Conclusion

We presented a new thinning/skeletonization algorithm for grayscale 3D images that does not require segmentation of high-intensity regions. By default, it generates skeletal points for curve-like structures at any scale and intensity of the image. Using the LEPs constructed during the destructibility test, we can extend the default thinning behavior to filter out shallow grayscale shape features and noise.

In the future, we will expand this algorithm to extract grayscale surface skeletons. We will also capitalize on the intensity associated with each voxel to perform more robust global pruning of the skeleton.

7. Acknowledgements

This work was supported by an NIH-NIGMS grant (T32 GM63495-07). We thank Julie Bernauer for the synthetic electron density maps we used for testing.

References

[1] S. Abeysinghe, M. Baker, W. Chiu, and T. Ju. Segmentation-free skeletonization of grayscale volumes for shape understanding. *SMI*, pages 63–71, June 2008.

[2] G. Bertrand and G. Malandain. A new characterization of three-dimensional simple points. *Pattern Recognition Letters*, 15(2):169–175, 1994.

[3] M. Couprie, F. Bezerra, and G. Bertrand. Topological operators for grayscale image processing. *Journal of Electronic Imaging*, 10(4):1003–1015, 2001.

[4] P. Dokladal, C. Lohou, L. Perroton, and G. Bertrand. A new thinning algorithm and its application to extraction of blood vessels. *Proc. of Biomedsim*, pages 32–37, 1999.

[5] T. Ju, M. Baker, and W. Chiu. Computing a family of skeletons of volumetric models for shape description. *Computer-Aided Design*, 39(5):352–360, 2007.

[6] S. Svensson, I. Nystrom, C. Arcelli, and G. S. di Baja. Using grey-level and distance information for medial surface representation of volume images. *ICPR*, 2, 2002.