Robust Gap Removal from Binary Volumes

A. Sobiecki [1] and A. C. Jalba[2] and A. C. Telea [1]

[1]University of Groningen, The Netherlands
[2]Eindhoven University of Technology, The Netherlands

Abstract

We present a method for the robust detection and removal of cracks and holes from binary voxel shapes based on the shapes' surface and curve skeletons. For this, we first classify gaps or indentations in the input shape by their position with respect to the shape's curve skeleton, into details (which should be preserved) and defects (which should be removed). Next, we remove defects, and preserve details, by using a local reconstruction process that uses the shape's surface skeleton. We illustrate our method by comparing it against classical morphological solutions on a wide collection of real-world shapes.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

Digital 3D shapes can be affected by various types of defects. Superficial defects, *e.g.* small-scale noise, are easy to remove while preserving the shape's salient features (*e.g.*, edges) by classical fairing and filtering [Tau95, TWBO02]. Profound defects, *e.g.* gaps and cracks that deeply penetrate inwards from the shape's surface, can be caused by faults in the original scanned shape, such as in the case of archaeological artifacts being scanned prior to restoration [BSK05], and are harder to remove while preserving details.

Defect detection and removal is typically done for 2D shapes by inpainting methods [SC05, BCHS06]. For 3D volumes, morphological filters, *e.g.* closing, find and remove (close) gaps whose size is under a user-given threshold [Ser82]. While simple and fast, such methods must be set up with great care so that they do not either remove relevant details (false positives) or leave defects (false negatives). More advanced 3D shape restoration methods exist, *e.g.* [BVG11, BSK05, KSY09, ZGL07]. Yet, such methods detect and remove *surface*, rather than *volumetric*, defects, and also require delicate manual defect-delineation and parameter setting.

We propose a method to detect and remove gap-like defects from binary volumetric shapes, with the next main advantages: (1) detection-and-removal of complex (deep, large, noisy) gaps and cracks with full surface detail preservation; (2) full automatic working; and (3) simple implementation and linear complexity *vs* shape size. For this, we use a novel combination of the input shape's curve skeleton (to detect defects) and surface skeleton (to remove gaps), described in Sec. 2. The method demonstrably obtains better results than classical morphological gap-removal on a variety of real-world complex shapes, as shown in Sec. 3. Section 4 concludes our paper.

2. Proposed method

We first outline the scope of our method: (1) Our input shapes are uniformly sampled 3D binary voxel volumes $\Omega \subset \mathbb{Z}^3$. (2) Defects to be found and removed are thin-and-elongated structures that enter deeply into the shape Ω from its boundary $\partial\Omega$. For 3D shapes, such defects can be 2D structures, such as surface-like cuts and cracks. No restriction is put on the topology or geometry of these defects. (3) Small-scale shape features, such as detail on $\partial\Omega$, should be altered as little as possible by the restoration.

Our method adapts to 3D the 2D gap-filling method in [SJB*14], which is outlined next: Given a binary pixel shape $\Omega \subset \mathbb{Z}^2$, one first computes the medial axis, or Euclidean skeleton, $S_{\Omega_{oc}}$ of the shape Ω_{oc} obtained by morphologically opening, next closing, Ω. $DT_{\partial\Omega} : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ is the Euclidean distance transform [MRH00] of Ω Using a disk structuring element of radius ρ fills all gaps in Ω whose local thickness is smaller than ρ. Next, the set of skeletal fragments $F = S_{\Omega_{oc}} \setminus \Omega$, *i.e.*, points of the skeleton of the defect-free shape Ω_{oc} which are outside of the original shape Ω , is computed. Such fragments correspond to gaps that cut *deeply* in the input shape Ω. Finally, gaps in Ω are filled by convolving the skeleton-fragment-set F with 2D disks whose radii equal the distance transform $DT_{\Omega_{co}}$ of the shape Ω_{co} obtained by first closing, then opening, Ω.

We can easily generalize this method to 3D shapes $\Omega \in \mathbb{Z}^3$ by using their 3D surface skeletons and distance transforms. Figure 1 (bottom path) shows this for a frog sliced by a thick planar cut. Voxels $\mathbf{x} \in \Omega$ are red, and voxels added by gap-filling are green. Using the surface skeleton fills the cut well, but also fills many detail gaps – see *e.g.* green details between the fingers. This is not surprising: The shape Ω_{oc} whose surface-skeleton $S_{\Omega_{oc}}$ we use, 'inflates' Ω so as to close gaps in Ω before the skeleton computation. However,

DOI: 10.2312/egsh.20161004

this also closes *detail* gaps, like the ones mentioned above, which creates surface-skeleton fragments outside Ω. Restoring the shape from such fragments yields the undesired fill-ins of detail gaps.

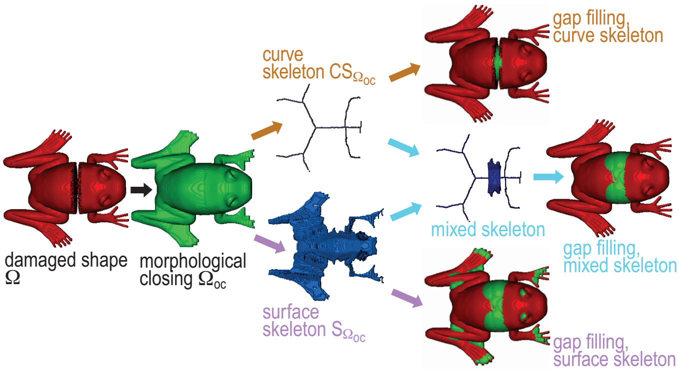

Figure 1: *3D gap restoration by surface and curve skeletons.*

The above problem is due to the higher sensitivity of 3D surface skeletons to inflation-induced changes as compared to the 2D medial axes in [SJB*14]. To alleviate this, one could simplify $S_{\Omega_{oc}}$ prior to its use in computing the gap-set F. However, this yields an undesired locally non-smooth gap filling. Indeed: a gap in a 3D shape is, usually, far from circular. To fill this gap in a plausible way, one must use the (almost) full surface skeleton, and not the simplified one. This issue does not occur in 2D: Skeletal fragments present in 2D gaps are 1D curve segments (internal skeleton branches) which perfectly capture the simpler configuration of a 2D gap, and which are not affected by skeleton simplification. In 3D, a gap contains a *mix* of both surface-skeleton terminal manifolds (which capture shape details) and internal surface-skeleton manifolds (which capture the shape topology) [SP09].

We could fix the above problem by using the curve skeleton $CS_{\Omega_{oc}}$ in the 3D detection-and-reconstruction process, instead of the surface skeleton $S_{\Omega_{oc}}$. Curve skeletons are 1D structures locally centered within surface skeletons [CSM07]. Curve skeletons have the same simple 1D structure as 2D medial axes, so robustly finding their fragments outside the shape, *i.e.* $CS \setminus \Omega$, is easy. Yet, curve skeletons do not capture the local shape *geometry* well [CSM07,SP09], so using them to reconstruct Ω will fill the detected gaps by locally *tubular* structures (Fig. 1, top orange path). Summarizing, we now avoid filling detail gaps, but we cannot fully fill the large cut (which has a non-circular cross-section).

Summarizing the above, we conclude that, for a 3D shape, curve skeletons are good to detect, but not to remove, gaps; and surface skeletons are good to remove, but not to detect, gaps. Hence, we combine the two skeleton types to solve our problem:

1. Compute the curve skeleton $CS_{\Omega_{oc}}$ and use it to create the gap-set F^{CS} of all curve-skeleton voxels outside the input shape Ω;
2. Compute the set F^S of surface-skeleton voxels outside Ω;
3. From F^S, we remove all voxels not connected (within F^S) to at least a voxel in F^{CS}, by a simple flood-fill from F^{CS} onto F^S. This removes from F^S all spurious fragments which filled detail gaps, yielding the final fragment-set F^{final};
4. Restore Ω by convolving voxels $\mathbf{x} \in F^{final}$ by balls of radii $DT_{\partial\Omega_{co}}(\mathbf{x})$, where $DT_{\partial\Omega_{co}}$ is the distance transform of the input shape Ω after morphological closing followed by opening.

Figure 1 (middle cyan path) shows the result of our mixed curve-and-surface skeleton method on our test shape: The large central cut is restored just as nicely as when using the surface-skeleton only (Fig. 1, bottom purple path), *and* no spurious detail is filled in, just as when using the curve-skeleton only (Fig. 1, top orange path).

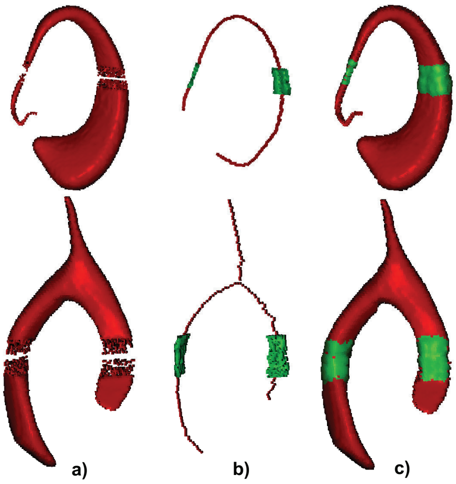

Figure 2: *Reconstruction of tubular shapes. (a) Input shape; (b) Curve skeleton (red) with selected surface skeleton parts (green); (c) Reconstruction result.*

3. Results

We implemented our method as follows: Surface-skeleton-only reconstruction uses the integer medial axis method (IMA) [HR08], which is fast, simple, and yields high-quality (unsimplified) surface skeletons and 3D distance transforms [SJT14]. Reconstruction by combined curve-and-surface skeletons uses the method in [JST16], which is fast, yields both (simplified) curve and surface skeletons, and guarantees that curve skeletons are contained in the corresponding surface skeletons. This containment is *crucial*, as step 3 of our method (Sec. 2) does a flood-fill from the curve skeleton onto the surface skeleton.

We ran our method, implemented in C++, on a Core i7 PC at 3.40 GHz with 16 GB RAM, using about 30 shapes voxelized from polygonal models at resolutions up to 512^3 voxels with *binvox* [NT03]. Figure 2 shows the reconstruction of a simple tubular shape. Our combined curve-and-skeleton method successfully detects and fills even the large jagged-edge gap shown. Figure 3 shows additional examples. Here, jagged gaps were created procedurally, by cutting the input shapes at various places with planes of various orientations and spheres of various radii, respectively. Noise was added on the cut surfaces by removing a small random set of voxels thereof. We compared our curve-and-surface reconstruction to four other methods: simple top-hat morphological gap-closing by using three structuring elements (ball, axis-aligned cube, and axis-aligned cross, see Fig. 3 top); and the surface-skeleton 3D extension of [SJB*14]. In Fig. 3, red shows points in Ω, and green

Figure 3: *Comparison of our gap-filling method (right column) with four other methods.*

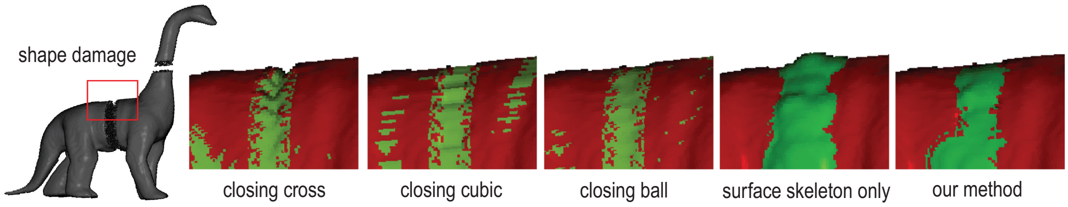

Figure 4: *Reconstruction smoothness. Our method (right) yields the locally smoothest results as compared to four other variants.*

shows points added by the reconstruction. As visible, our combined curve-and-surface restoration method yields the best results in filling deep gaps *and* preserving surface detail.

We next discuss several relevant aspects of our method:

Smoothness: As noted in most inpainting-related works, restoration should produce a 'plausible' reconstruction of damaged areas [BCHS06], which is reflected by the *smoothness* of the reconstructed signal over, and along the boundaries of, the reconstructed area. We adopt this requirement for our 3D context too. As seen in Figs. 3 and 4, morphological closing produce noisier results than our skeleton-based reconstruction. We also see that the combined curve-and-surface reconstruction produces the smoothest results – the reconstructed surface (green) closely follows the curvature of the surrounding original surface (red).

Locality: As outlined in [SJB*14] (for the 2D case), gap-filling reconstruction should detect and remove deep gaps and leave shallow gaps (detail indentations) untouched. Figs. 1 and 3 show that morphological closing indiscriminately fills all gaps smaller than the structuring-element size. This is visible in the amount of green in the reconstructed images which does not correspond to cut-locations visible in the shapes in the leftmost column, *e.g.*, filling the small gaps between the details of the dragon surface (top row), spikes of the trident (neptune model, fourth row from top), or frog's fingers (third row from bottom). The combined curve-and-surface method suffers far less from such issues.

Simplicity and scalability: Our method uses trivial morphological opening and closing; computing 3D curve and surface skeletons; a flood fill on voxel volumes; and reconstructing a shape from a selected set of skeleton points by its medial axis transform (MAT). All these operations can be computed in linear time in the size $\|\Omega\|$ of the input shape [JST16]. This makes our method work in seconds on volumes up to 512^3 voxels on a commodity PC.

4. Conclusions

In this paper, we have presented a new method for automatic detection and restoration of gaps and cracks present in 3D volumetric shapes. For this, we extend the 2D gap detection and removal method in [SJB*14] to 3D, by using a novel combination of curve skeletons (for gap detection) and surface skeletons (for gap repairing). Such skeletons can be readily and efficiently provided by the 3D skeletonization method in [JST16]. Our proposed 3D restoration method was shown to produce better results, in terms of detail preservation and gap removal, as compared to classical morphological closing methods, while having an identical cost – linear in terms of the voxel count of the input shape.

Acknowledgments

We acknowledge the financial support of CNPq, Brazil, through the grant 202535/2011-8.

References

[BCHS06] Bertalmio M., Caselles V., Haro G., Sapiro G.: PDE-based image and surface inpainting. In *Handbook of Mathematical Models in Computer Vision* (2006), Springer, pp. 33–61. 1, 4

[BSK05] Bendels G., Schnabel R., Klein R.: Detail-preserving surface inpainting. In *Proc. EG Symp. on Virtual Reality, Archaeology and Cultural Heritage* (2005), pp. 41–48. 1

[BVG11] Bermano A., Vaxman A., Gotsman C.: Online reconstruction of 3D objects from arbitrary cross-sections. *ACM TOG 30*, 4 (2011), 348–359. 1

[CSM07] Cornea N., Silver D., Min P.: Curve-skeleton properties, applications, and algorithms. *IEEE TVCG 13*, 3 (2007), 87–95. 2

[HR08] Hesselink W., Roerdink J.: Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform. *IEEE TPAMI 30*, 12 (2008), 2204–2217. 2

[JST16] Jalba A., Sobiecki A., Telea A.: An unified multiscale framework for planar, surface, and curve skeletonization. *IEEE TPAMI 38*, 1 (2016), 30–45. 2, 4

[KSY09] Kawai N., Sato T., Yokoya N.: Efficient surface completion using principal curvature and its evaluation. In *Proc. IEEE ICIP* (2009), pp. 521–524. 1

[MRH00] Meijster A., Roerdink J., Hesselink W.: A general algorithm for computing distance transforms in linear time. *Math. Morphology and Its Applications to Image and Signal Processing* (2000), 331–340. 1

[NT03] Nooruddin F., Turk G.: Simplification and repair of polygonal models using volumetric techniques. *IEEE TVCG 9*, 2 (2003), 191–205. see also www.cs.princeton.edu/~min/binvox. 2

[SC05] Shih T., Chang R.: Digital inpainting – survey and multilayer image inpainting algorithms. In *Proc. IEEE ICITA* (2005), pp. 15–24. 1

[Ser82] Serra J.: *Image Analysis and Mathematical Morphology*. Academic Press, London, 1982. 1

[SJB*14] Sobiecki A., Jalba A., Boda D., Diaconeasa A., Telea A.: Gap-sensitive segmentation and restoration of digital images. In *Proc. EG GVC* (2014), pp. 136–144. 1, 2, 4

[SJT14] Sobiecki A., Jalba A. C., Telea A. C.: Comparison of curve and surface skeletonization methods for voxel shapes. *Pattern Recognition Letters 47* (2014), 147–156. 2

[SP09] Siddiqi K., Pizer S.: *Medial Representations: Mathematics, Algorithms and Applications*. Springer, 2009. 2

[Tau95] Taubin G.: Estimating the tensor of curvature of a surface from a polyhedral approximation. In *Proc. ICCV* (1995), pp. 902–907. 1

[TWBO02] Tasdizen T., Whitaker R., Burchard P., Osher S.: Geometric surface smoothing via anisotropic diffusion of normals. In *Proc. IEEE Visualization* (2002), pp. 125–132. 1

[ZGL07] Zhao W., Gao L., Lin S.: A robust hole-filling algorithm for triangular mesh. *Visual Comput.* (2007). 1