

A Medial Point Cloud based Algorithm for Dental Cast Segmentation

(Invited Paper)

Jacek Kustra
and Marko de Jager
Philips Research
Email: jacek.kustra@philips.com,
marko.de.jager@philips.com

Andrei Jalba
Eindhoven University of Technology
Email: a.jalba@tue.nl

Alexandru Telea
University of Groningen
Email: a.c.telea@rug.nl

Abstract—We present a novel approach towards the analysis and segmentation of dental surfaces (such as dental cast scans) for the application in consumer oral healthcare devices. Although these can be acquired in great detail using techniques such as Computer Tomography, the development of mass consumer products is dependent of non-invasive safe acquisition techniques such as the bite imprint, commonly used in orthodontic treatments. This provides a limited level of detail on the global teeth and gum shape, therefore understanding the shape local and global properties is important. In this paper we present a segmentation algorithm based on the surface’s medial point cloud, and we present it’s application for different dental shape analysis requirements.

I. INTRODUCTION

Oral Healthcare devices targeted at the consumer market require a good understanding of teeth shape properties. Although several imaging techniques, such as Computer Tomography can provide very detailed shapes, including the teeth roots. However, due to their invasive and costly nature, they are not suitable to be used in personalized consumer devices. To find the shape properties of a personal oral healthcare device, one has to use simpler, however less detailed imaging techniques. Bite imprints provide such a possibility as they take the subject’s teeth surface and several commercially available scanners exist which can digitize this surface. Advances in range imaging and 3D scanning technologies opened the possibilities to capture dental scans directly from a patient [1], and next to digitize such shapes into 3D surface meshes of the teeth-and-gum structure. All these techniques provide a single manifold surface: They do not directly provide any interdental space characteristics and very limited information is available on the tooth-gum interface characteristics. Also, any further assessment of any of these properties requires the identification of the individual teeth elements in the first place. Therefore, shape analysis techniques for the dental surface are important.

In this paper, we present a novel method of interpreting teeth shape. Rather than using local information such as curvature, we take a *global* approach, based on the surface skeleton of the input 3D surface. Key to our method is the observation that surface skeletons capture *all* input surface creases, regardless of their sharpness. To compute an accurate and high-resolution

surface skeleton, we use a recent GPU-based method which delivers point-cloud skeleton representations for models of hundreds of thousands of polygons in a few seconds. Next, we regularize the surface skeleton, thereby eliminating small manifolds that do not correspond to segments large enough to be of interest. Next, we use the resulting skeleton cloud to infer the teeth and gum separation, visualize the arch and finally segment the individual teeth elements using a mean-shift approach. Finally, we project back the identified skeletal segments onto the input surface, and use a motivated nearest-neighbor approach to produce a segmentation that entirely covers the input shape.

The structure of this paper is organized as follows. Section II reviews the related work on skeletonization and current approaches towards dental cast segmentation. Section III details our method. Section IV presents our results. Section V discusses our method. Section VI concludes the paper.

II. RELATED WORK

Dental segmentation: Segmentation of dental surfaces from dental cast scans has been widely explored in the last years due to the increasing popularity and availability of digital models. Most of the work available focuses on orthodontic workflow, where manual segmentation takes a prohibitively long time. Therefore, several methods have been proposed to automate this process, ranging from fully automated methods [2], approaches requiring minimal user interaction [3], [4], [5]. [2] avoids the 3D mesh processing complexity by first transforming the 3D acquired data into a plan-view range image on which the segmentation is performed. Recently, a snake-based approach for teeth segmentation has been proposed, using the surface curvature properties to guide the snake iterations [3]. However, all of these methods have problems when the local curvature of the input shape is low, *i.e.*, the creases separating teeth are shallow. Also, all these methods assume a way to locate the potential positions and amount of teeth prior to the segmentation, *e.g.* such as the specification of the arch.

Skeletonization: Since their introduction by Blum [6], skeletons, or medial axes, have been widely used for different shape operations. The medial axis consists of the loci of

maximally inscribed circles (in 2D) or spheres (in 3D). The 2D medial axis is a collection of curves. 3D shapes admit two types of skeletons: *Surface* skeletons are a collection of curved manifolds containing the loci of maximally inscribed spheres. They fully describe the shape, *i.e.*, the shape can be reconstructed from the skeletal information. *Curve* skeletons are a collection of 1D curves locally centered in the shape, according to various heuristics. They cannot fully capture the geometry of the input shape, and are most appropriate for tubular objects [7]. Given these, we next focus on surface skeletons solely. Several methods exist for computing surface skeletons either from a meshed surface or a volumetric (voxel) model. Volumetric methods typically have significantly higher memory and speed costs and lower accuracy. Skeletonization methods can be divided into thinning [8], [9], geometric [10], [11], and field-based [13], [14], [17]. For a complete survey, we refer to [18]. Only very few 3D segmentation methods have employed skeletons in their design. Exceptions are [19], which uses curve skeletons to produce part-type segmentations of smooth, tubular, shapes, and [20], which uses surface skeletons to produce part-type segmentations of shapes exhibiting sharp edges. One of the main blockers in using surface skeletons for surface segmentation is the difficulty to efficiently and robustly compute accurate skeletons of complex models. Recently, this has been overcome by [21] and [22] who compute point-cloud skeletal representations for mesh models of up to millions of polygons in seconds on the GPU. However, since these methods produce an unstructured point cloud rather than a compact voxel model (as in *e.g.* [19]), they cannot be easily used to segment surfaces following the ideas in [20].

III. METHOD

Our proposed segmentation method first transforms the surface into its medial domain. We next exploit skeletal point density properties to perform several operations in this domain. Since teeth and gums present distinct density properties in the medial domain, a separation can be achieved by thresholding the medial surface based on these properties (Fig. 3a). For the separation of individual teeth, we cluster the medial cloud points, also based on their local density. Finally, the medial segmentation is projected back to the surface. We describe all these steps below.

A. Skeletonization preliminaries

Given a shape $\Omega \subset \mathbb{R}^3$ with boundary $\partial\Omega$, we first define its distance transform $DT_{\partial\Omega} : \mathbb{R}^3 \rightarrow \mathbb{R}^+$

$$DT_{\partial\Omega}(\mathbf{x} \in \Omega) = \operatorname{argmin}_{\mathbf{y} \in \partial\Omega} \|\mathbf{x} - \mathbf{y}\|. \quad (1)$$

The skeleton of Ω is next defined as

$$S_{\Omega} = \{\mathbf{x} \in \Omega \mid \exists \mathbf{f}_1, \mathbf{f}_2 \in \partial\Omega, \mathbf{f}_1 \neq \mathbf{f}_2, \|\mathbf{x} - \mathbf{f}_1\| = \|\mathbf{x} - \mathbf{f}_2\| = DT_{\partial\Omega}(\mathbf{x})\}, \quad (2)$$

where \mathbf{f}_1 and \mathbf{f}_2 are two of the contact points with $\partial\Omega$ of the maximally inscribed ball in Ω centered at \mathbf{x} , also called *feature transform* (FT) points [23], [24]. The vectors $\mathbf{f}_1 - \mathbf{x}$ and $\mathbf{f}_2 - \mathbf{x}$

linking the skeleton points \mathbf{x} with their feature points are called *feature vectors* or *spoke vectors* [11].

Given a densely sampled input surface $\partial\Omega$, represented as a polygonal mesh, we extract the corresponding densely sampled surface skeleton following the method in [22], [21]. To eliminate small-scale skeleton details corresponding to equally small-scale convexities (bumps) on $\partial\Omega$, we next *regularize* the skeleton. For this, we compute an importance metric $\rho : S \rightarrow \mathbb{R}^+$ for each skeleton point, such that $\rho(\mathbf{s} \in S)$ equals the shortest-path distance on $\partial\Omega$ between the feature points \mathbf{f}_1 and \mathbf{f}_2 of \mathbf{s} . As demonstrated in [19], [22], this metric monotonically increases from the skeleton boundaries inwards so that thresholding with a small value τ delivers a simplified skeleton

$$S_{\tau} = \{\mathbf{s} \in S_{\Omega} \mid \rho(\mathbf{s}) > \tau\} \quad (3)$$

which captures all salient branches of S_{Ω} but eliminates those which correspond to surface details shorter than τ .

B. Surface curvature vs skeleton density

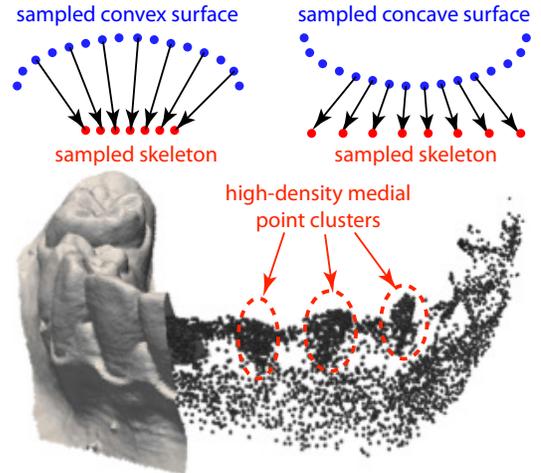


Fig. 1. Relationship between local surface curvature and medial cloud density. **Top:** Concept sketched in 2D. **Bottom:** High-density point clusters are formed inside positive-curvature 3D surface areas (front teeth).

To further segment our skeleton point-cloud, we use the following observations linking the curvature of $\partial\Omega$ and point density on S_{Ω} . Consider a densely sampled surface $\partial\Omega$. For a region of $\partial\Omega$ having positive curvature (a bump, or convex, region), the ball-shrinking directions $r\mathbf{n}$, identical to the feature vector directions, will point inwards in a converging fashion. Hence, the density of the skeletal points for this region will be higher than the surface density. Conversely, for a region with negative curvature (a crease, or concave, region), the ball-shrinking directions will point inwards in a diverging fashion, so the density of the skeletal points for this region will be lower than the surface density. Figure 1 (top) illustrates this in 2D. Figure 1 (bottom) shows an actual example for a 3D skeletal cloud computed from a dental cast. We observe that skeletal parts *enclosed* in the front teeth present a high point density, since these teeth are indeed convex shape parts.

C. Point Cloud Density and Mean shift clustering

We now show how to use the density-related observations from Sec. III-B to segment the skeleton point cloud. Since the skeletal cloud exhibits strong density variations, it should be possible to segment it into point clusters representing the dense regions based on a method which exploits such density variations. An ideal such method is *mean shift clustering* [26], which we extend to our segmentation needs, as follows. We start by selecting a set of seed points $P \subset S_\tau$ from the simplified skeleton. Each seed point $\mathbf{x} \in P$ is assigned a unique ‘segment id’. For each seed point $\mathbf{x} \in P$, we aim to find its so-called convergence point $\mathbf{c}(\mathbf{x}) \in \mathbb{R}^3$. For this, we first find all neighbors $N_{\mathbf{x}}^\varepsilon \subset S_\tau$ of \mathbf{x} within a small fixed radius ε and determine the centroid of $N_{\mathbf{x}}^\varepsilon$

$$\mathbf{m}(\mathbf{x}) = \frac{\sum_{\mathbf{y} \in N_{\mathbf{x}}^\varepsilon} K(\|\mathbf{x} - \mathbf{y}\|)\mathbf{y}}{\sum_{\mathbf{y} \in N_{\mathbf{x}}^\varepsilon} K(\|\mathbf{x} - \mathbf{y}\|)} \quad (4)$$

where K is a Gaussian kernel

$$K(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \quad (5)$$

following the kernel density estimation idea in [26]. ε is set to a small fraction (about 5%) of the model size. We next iteratively shift the seed points \mathbf{x} to their centroids $\mathbf{m}(\mathbf{x})$ following Eqn. 4 (see also Fig. 2 a) until these stabilize, *i.e.*, move at one iteration less than a small threshold $\lambda = \|\mathbf{m}(\mathbf{x}) - \mathbf{x}\|$, set in practice to 10^{-4} . Also, for each non-seed point (which is *not* shifted), we define a voting weight $v(\mathbf{y})$, initialized to zero at the beginning of the algorithm. At every mean-shift iteration, we add a value $K(\|\mathbf{y} - \mathbf{m}(\mathbf{x})\|)$ to $v(\mathbf{y})$ for each non-seed point $\mathbf{y} \in N_{\mathbf{x}}^\varepsilon$, and also add a pointer from \mathbf{y} to $\mathbf{m}(\mathbf{x})$, to indicate that \mathbf{y} was in the neighborhood of $\mathbf{m}(\mathbf{x})$. When $\mathbf{m}(\mathbf{x})$ has converged, we search its neighborhood for other existing convergence points \mathbf{c}' than itself. If one exists, we merge the ids of $\mathbf{m}(\mathbf{x})$ and \mathbf{c}' . Otherwise, we create a new convergence point $\mathbf{c} = \mathbf{m}(\mathbf{x})$. At the end of the mean shift, all seed points have thus converged to a set C of convergence points (see Fig. 2 b). The ids of the points $\mathbf{c} \in C$ give us the final segments. Finally, to assign each non-seed point $\mathbf{y} \in S_\tau$ to a segment, is done by assigning to \mathbf{y} the id of the convergence point that it is linked to and which has the highest amount of votes within the k last iterations (Fig. 2 c). Different segmentation levels can be achieved by considering the voting of only the last k iterations of the mean shift process. This way, only the areas around the skeleton-cloud density *peaks* are considered. This is illustrated by the dental cast models, where the gum areas remain mostly unsegmented (Fig. 3 a-e), for which we used a value $k = 20$. In contrast, for the other shapes in Fig. 3 f-k, the full mean-shift path has been considered for the voting, leading to the full surface being segmented into patches.

D. Seed point detection

We find the initial seed points P used to initialize the mean-shift iterations by taking the points along the highest point on the medial cloud which have a low distance $DT_{\partial\Omega}(\mathbf{s})$

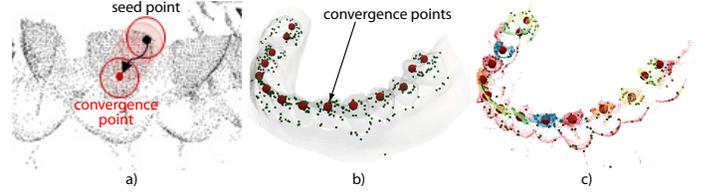


Fig. 2. Mean shift clustering: (a) Principle: A seed point (black dot) is shifted to the centroid of its skeleton-cloud neighborhood density until convergence (red dot). (b) Final convergence points for dental cast scan. (c) Segment ids assigned to all skeleton points.

to the original surface $\partial\Omega$ and (b) exhibit a high curvature, determined by the angle between the feature vectors $\mathbf{f}_1 - \mathbf{s}$, $\mathbf{f}_2 - \mathbf{s}$.

E. Segmentation transfer to surface

As a last step of our method, we transfer the segmentation from the skeleton to the surface, as follows. For each point \mathbf{s} on the simplified skeleton S_τ , we copy the segment id of \mathbf{s} to its two feature points \mathbf{f}_1 and \mathbf{f}_2 . However, this does not assign a segment id to *all* points on $\partial\Omega$, since we segmented the simplified skeleton S_τ rather than the full skeleton S_Ω . For all surface points $\mathbf{p} \in \partial\Omega$ which are not assigned a segment id, we search the closest surface point $\mathbf{p}' \in \partial\Omega$ which has a segment id assigned, and assign to \mathbf{p} the same segment id as \mathbf{p}' . This effectively fills the gaps between segments on $\partial\Omega$ in distance order, yielding a full, non-overlapping, segmentation of the input surface.

IV. RESULTS

Figure 3 presents several results obtained with our medial point cloud method. On Fig. 3a) we observe the resulting absolute density values on the medial surface. We observe that the teeth exhibit very different properties from the gums and can be clearly distinguished in the medial domain. On Fig. 3c) these density values are transferred to the surface using the feature vectors. A clear separation between the teeth and gums is visible. The medial cloud density properties can also be used for determining the dental arch, which is visualized in Fig. 3b) showing a vertical projection of (alpha-blended) points. Finally, we demonstrate the teeth segmentation on images d-f). We see that the method can separate very well the incisors, canines, molars and pre-molars, both from each other, and also from the gum area. However, we also see several problems. The molars are over-segmented (Fig. 3 a,b), and occasionally the gum area is also oversegmented (Fig. 3 c,e). The oversegmentation of the gum area is *not* problematic since for most applications the user is interested in teeth geometrical properties and the gums closely surrounding the teeth. The oversegmentation of the molars can be explained by the fact that they have a more complex geometry than the other teeth (more internal creases). However, their inter-separation is much clearer on the surface *i.e.*, on the dental cast their separation from each other is usually clear, not requiring the medial domain to separate them.

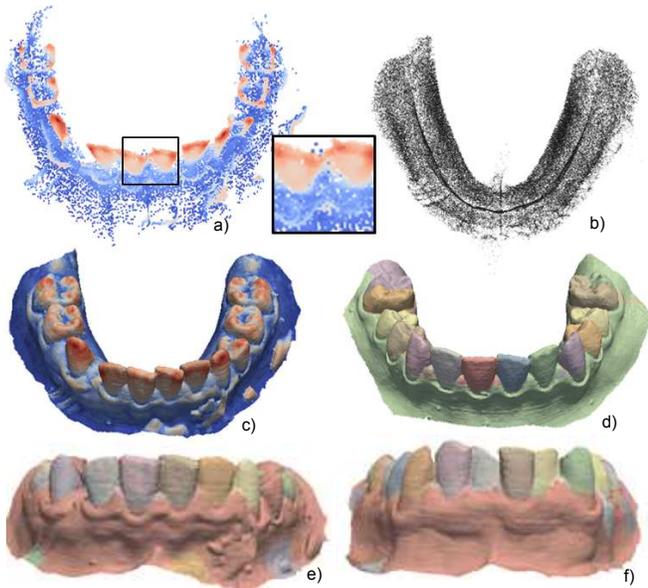


Fig. 3. a) Regularized medial cloud colored with point density values: The absolute point density values clearly distinguish between teeth and the gums. b) Medial point cloud top projection: The dental arch is clearly visible and can be extracted using the projected point density. c) Point density transferred to the surface: Teeth and gums have a clear medial density contrast. d-f) Final segmentation on different dental casts

V. DISCUSSION

A key asset of the proposed method is its algorithmic simplicity. We make use of existing algorithms with proven accuracy, convergence, and complexity properties [26], [19], [22]. From a novelty perspective, our method is the first we are aware of which uses surface skeletons computed on mesh models to segment input surfaces, its only other competitor, using the more expensive and lower-resolution voxel-based models being [20]. Also, this is the first application to our knowledge to use medial point cloud density properties for surface segmentation. Using medial surface properties to access low resolution scans, opens new possibilities to extract relevant properties from these datasets. An example is the interdental space properties which are difficult to observe on the surface due to its limited resolution, however these can be accessed by observing the data in the medial domain. Also, the medial point cloud reveals itself as a powerful tool for extracting global shape properties such as the arch.

These results are important in the context of personal oral healthcare devices, where the user can only be subjected to simple scanning techniques for accessing its unique oral shape features, which can be used towards development of personalized devices or to fine tune a device to a specific user.

We implemented our method in C++ using ANN [27] to find nearest neighbors and the GPU-based skeletonization method in [22]. The latter also provides us with the necessary distance transform, feature points, and importance metric (Sec. III-A). On a Windows PC at 2.66 GHz with an Nvidia 690 GTX, the segmentation takes roughly 10s.

Future work is needed to study the most effective way these properties can be translated into consumer oral healthcare device properties and settings. Also, more work is required to study extensions of the algorithm to more general surfaces.

VI. CONCLUSIONS

We have demonstrated the application of medial surfaces for the analysis of (limited resolution) dental scans. For this, we use the medial points density properties, which can be used for separating teeth and gums, segmenting individual teeth elements or inferring the arch. To our knowledge, our method is the second existing technique able to use surface skeletons to segment 3D surfaces. In contrast to the first published technique in this area [20], we can directly handle meshed models without a costly voxelization step; we do not require the complex and sensitive detection of skeletal boundaries; and we can treat significantly more complex shapes than the earlier cited method in this class.

REFERENCES

- [1] Atron, Inc., "3D intraoral scanner," 2013, www.a-tron3d.com/en/products/id-3d-intraoral-scanner.html.
- [2] T. Kondo, S. Ong, and K. W. C. Foong, "Tooth segmentation of dental study models using range images," *Medical Imaging, IEEE Transactions on*, vol. 23, no. 3, pp. 350–362, 2004.
- [3] T. Kronfeld, D. Brunner, and G. Brunnert, "Snake-based segmentation of teeth from virtual dental casts," *CAGD*, vol. 7, no. 2, pp. 221–233, 2010.
- [4] M. Zhao, L. Ma, W. Tan, and D. Nie, "Interactive tooth segmentation of dental models," in *Proc. EMBS*, 2005, pp. 654–657.
- [5] Y. Hirogaki, T. Sohmura, H. Satoh, J. Takahashi, and K. Takada, "Complete 3-d reconstruction of dental cast shape using perceptual grouping," *Medical Imaging, IEEE Transactions on*, vol. 20, no. 10, pp. 1093–1101, 2001.
- [6] H. Blum, "A Transformation for Extracting New Descriptors of Shape," in *Models for the Perception of Speech and Visual Form*, W. Wathen-Dunn, Ed. Cambridge: MIT Press, 1967, pp. 362–380.
- [7] N. Cornea, D. Silver, and P. Min, "Curve-skeleton properties, applications, and algorithms," *IEEE TVCG*, vol. 13, no. 3, pp. 87–95, 2007.
- [8] C. Pudney, "Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images," *CVIU*, vol. 72, no. 3, pp. 404–413, 1998.
- [9] K. Palagyi and A. Kuba, "Directional 3D thinning using 8 subiterations," in *Proc. DGCI*, vol. 1568. Springer LNCS, 1999, pp. 325–336.
- [10] N. Amenta, S. Choi, and R. Kolluri, "The power crust," in *Proc. SMA. ACM*, 2001, pp. 65–73.
- [11] S. Stolpner, S. Whitesides, and K. Siddiqi, "Sampled medial loci and boundary differential geometry," in *Proc. IEEE 3DIM*, 2009, pp. 87–95.
- [12] —, "Sampled medial loci for 3D shape representation," *CVIU*, vol. 115, no. 5, pp. 695–706, 2011.
- [13] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. Zucker, "Hamilton-Jacobi skeletons," *IJCV*, vol. 48, no. 3, pp. 215–231, 2002.
- [14] S. Pizer, K. Siddiqi, G. Szekeley, J. Damon, and S. Zucker, "Multiscale medial loci and their properties," *IJCV*, vol. 55, no. 2-3, pp. 155–179, 2003.
- [15] S. Bouix, K. Siddiqi, and A. Tannenbaum, "Flux driven automatic centerline extraction," *Medical Image Analysis*, vol. 9, no. 3, pp. 209–221, 2005.
- [16] S. Bouix, K. Siddiqi, A. Tannenbaum, and S. Zucker, "Medial axis computation and evolution," in *Statistics and analysis of shape*. Springer LNCS, 2006, ch. 1, pp. 1–28.
- [17] D. Reniers, J. J. van Wijk, and A. Telea, "Computing multiscale skeletons of genus 0 objects using a global importance measure," *IEEE TVCG*, vol. 14, no. 2, pp. 355–368, 2008.
- [18] K. Siddiqi and S. Pizer, *Medial Representations: Mathematics, Algorithms and Applications*. Springer, 2009.

- [19] D. Reniers and A. Telea, "Part-type segmentation of articulated voxel-shapes using the junction rule," *CGF*, vol. 27, no. 7, pp. 1837–1844, 2008.
- [20] —, "Patch-type segmentation of voxel shapes using simplified surface skeletons," *CGF*, vol. 27, no. 7, pp. 1954–1962, 2008.
- [21] J. Ma, S. Bae, and S. Choi, "3D medial axis point approximation using nearest neighbors and the normal field," *Vis. Comput.*, vol. 28, no. 1, pp. 7–19, 2012.
- [22] A. Jalba, J. Kustra, and A. Telea, "Surface and curve skeletonization of large 3D models on the GPU," *IEEE TPAMI*, vol. 35, no. 6, pp. 1495–1508, 2013.
- [23] R. Strzodka and A. Telea, "Generalized distance transforms and skeletons in graphics hardware," in *Proc. VisSym*, 2004, pp. 221–230.
- [24] D. Reniers and A. Telea, "Tolerance-based feature transforms," in *Advances in Comp. Graphics and Comp. Vision (eds. J. Jorge et al.)*. Springer, 2007, pp. 187–200.
- [25] A. Telea and A. Jalba, "Computing curve skeletons from medial surfaces of 3D shapes," in *Proc. TPCG (Eurographics UK)*, 2012, pp. 273–280.
- [26] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE TPAMI*, vol. 24, no. 5, pp. 603–619, 2002.
- [27] D. Mount and S. Arya, "Approximate nearest neighbor search software," 2011, www.cs.umd.edu/~mount/ANN.