Eurographics Symposium on Geometry Processing 2024
R. Hu and S. Lefebvre
(Guest Editors)

Volume 43 (2024), Number 5

Coverage Axis++: Efficient Inner Point Selection for 3D Shape Skeletonization
Supplementary Materials

Zimeng Wang*[1], Zhiyang Dou*[1,2], Rui Xu[3], Cheng Lin[1], Yuan Liu[1], Xiaoxiao Long[1], Shiqing Xin[3], Taku Komura[1], Xiaoming Yuan†[1], Wenping Wang†[4]

[1]The University of Hong Kong [2]TransGP [3]Shandong University [4]Texas A&M University

Equal Contribution; †Equal Advising.

This supplementary material covers: more experimental results (Sec. A), more detailed running time statistics of existing methods (Sec. B), coverage rate of Coverage Axis++ (Sec. C), computation complexity analysis on Coverage Axis (Sec. D), and an additional discussion on the skeletal point connectivity of MAT (Sec. E).

Appendix A: More Experimental Results

As discussed in Sec. 4.7.1, the proposed method has polynomial complexity with respect to the number of surface samples, denoted as $|S|$. In contrast, the Coverage Axis algorithm showcases exponential complexity in $|S|$. This is evidenced by Tab. 1 of the main paper, where we report the running times of our method and Coverage Axis for various models when $|S|$ is set to 1500. Furthermore, it can be expected that when increasing the number of surface samples, the performance gap between our method and the Coverage Axis will widen further. We conducted additional experiments by increasing the number of surface samples to 5000 and comparing the performance of our method with the Coverage Axis across different shapes. The results are summarized in Tab. D4.

Note that the proportion of timeout models (i.e., models with a running time exceeding 1000s) of Coverage Axis increases from 21.2% to 50% as $|S|$ increases from 1500 to 5000. Furthermore, comparing the results obtained with $|S|$ set as 1500 and 5000, we find that Coverage Axis++ exhibits a more significant advantage over Coverage Axis in terms of running time. On average, our method achieves a reconstruction error of 4.429% in 9.7s, whereas Coverage Axis achieves a reconstruction error of 5.448% in over 516.3s.

Appendix B: More Detailed Running Time Statistics.

We provide more statistics of the running time of Neural Skeleton [CD23], MATFP [WWWG22]. All evaluations are performed on the same machine to ensure a fair comparison.

Table B1: *Detailed running time statistics of MATFP [WWWG22] using mesh inputs.*

Model	Stage1 (s)	Stage2 (s)	Stage3 (s)	Stage4 (s)	Total Time (s)
Vase	1.62	3.30	0.66	2.69	8.27
Hand	4.15	5.42	1.57	3.54	14.68
Dog	0.59	0.92	1.51	7.53	10.55
Fertility	2.32	4.12	2.19	8.33	16.96
Crab	2.49	4.90	1.48	6.25	15.12
Average	2.23	3.73	1.48	5.67	13.11

Table B2: *Detailed running time statistics of Neural Skeleton [CD23] using point cloud inputs.*

Model	Network Optimization Time (s)	Coverage Skeleton Time (s)	Total Time (s)
Ant-2	9.95	10.53	20.48
Bottle	9.88	11.61	21.49
Chair-2	9.86	121.52	131.38
Dog	9.96	15.42	25.38
Dolphin	10.07	24.88	34.95
Fertility	9.92	108.15	118.07
Guitar	9.86	107.32	117.18
Hand-1	10.01	21.36	31.37
Kitten	9.89	46.88	56.77
Snake	9.85	25.74	35.59
Average	9.92	49.39	59.27

Appendix C: Coverage Rate

In the main paper, we conducted extensive experiments for Coverage Axis++ on various models to demonstrate that our method is capable of extracting compact underlying shape representation of MAT while achieving low reconstruction error, as summarized in Tab. 1. Since the selected skeletal points form a good representation of the medial axis, the corresponding dilated balls should cover a significant portion of the surface samples to build the correspondence between the medial balls and the shape volume. In other words, the coverage rate of our method should be high. The idea has been verified by Coverage Axis, as it directly requires a 100% coverage rate by imposing a hard constraint during the optimization. We record the coverage rates of each model for Coverage Axis++, which are summarized in Tab. C3. The result shows that

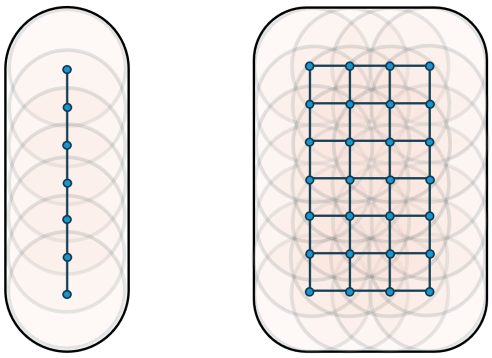

(a) (e)

Figure D1: *Distribution of randomly sampled candidate skeletal points in two different shapes. (a) tubular structure. (b) planar structure.*

our method consistently achieves very high coverage rates, with an average of 96.134%.

Table C3: *Coverage rate of Coverage Axis++. The inputs are meshes. $|V|$: The number of skeletal points. $\overleftrightarrow{\epsilon}$: Two-sided HD between original surface and reconstruction.*

Model	$\lvert V\rvert$	$\overleftrightarrow{\epsilon}$	Coverage rate
Bear-1	68	4.047%	98.662%
Bear-2	49	4.572%	98.461%
Bird	83	2.560%	95.914%
Bug	137	4.511%	98.274%
Bunny	73	5.771%	97.650%
Camel	93	5.556%	96.919%
Chair	115	2.947%	96.082%
Crab-1	93	2.744%	97.874%
Crab-2	107	2.825%	98.608%
Cup	243	9.247%	95.411%
Dinosaur	66	2.672%	99.331%
Elephant	94	3.874%	98.079%
Fandisk	127	4.067%	92.029%
Femur	26	3.095%	97.536%
Fish	43	4.809%	88.587%
Giraffe	71	4.208%	97.335%
Guitar	71	2.497%	96.804%
Hand	47	1.975%	95.143%
Human	46	2.105%	98.869%
Lifebuoy	33	3.324%	96.151%
Octopus	74	3.005%	97.989%
Plane	44	4.306%	83.389%
Pot	90	5.278%	95.786%
Rocker	112	3.652%	96.944%
Seahorse	60	5.168%	96.144%
Spectacles	97	3.080%	96.075%
Spider	54	4.493%	95.987%
Vase	118	3.875%	95.710%
Average	-	3.938%	96.134%

Appendix D: Complexity Analysis on Coverage Axis

Next, we provide more analysis of the complexity of the Coverage Axis. From Tab. 1 of the main paper, we observe that the running times of the Coverage Axis for certain models (such as Bird, Chair, Guitar, and Spectacles) are extremely long, indicating that solving SCP for these models is particularly challenging. Note that a common characteristic of these shapes is that they all have planar shapes. We illustrate why solving SCP for models with planar shapes is difficult.

Given an arbitrary shape, under the meaning of medial axis transform, it can be divided into two types of shapes: *tubular* and *planar* structures, as depicted in Fig. D1. Without loss of generality, we assume that the size and sample density of the two shapes are the same. In such cases, if the tubular shape contains q points, the planar shape would have q^2 points. Recall that the SCP amounts to the following $0-1$ integer programming:

$$\begin{aligned} &\min \|\mathbf{v}\|_2, \\ &s.t. \quad \mathbf{D}\mathbf{v} \geq \mathbb{1}, \end{aligned} \tag{1}$$

where $\mathbb{1}$ is the vector of all ones, $\geq$ is applied element-wise, and $\mathbf{v} \in \{0,1\}^{n\times 1}$ is a decision vector, with its i-th element $v_i \in \{0,1\}$ indicating if the candidate skeletal point p_i is selected. Consequently, for the tubular shape, there are 2^q possible combinations of v, while for the planar shape, there are 2^{q^2} possible combinations. As a result, solving the SCP becomes significantly more challenging when the model involves planar shapes due to the exponential increase in the number of possible combinations. In comparison, Coverage Axis++ sidesteps solving an optimization problem with high computation complexity by designing a simple yet effective heuristic algorithm, taking coverage and uniformity into consideration, which thus achieves significant acceleration.

Appendix E: Skeletal Point Connectivity

In this Section, we discuss the necessity of mesh connectivity in the medial axis for various downstream applications. 1) Reconstruction: The connectivity of meshes and edges in the Medial Axis Transform (MAT) is important for shape reconstruction. This is because the connectivity allows for interpolation between medial axis balls, a.k.a., slab meshes, thereby achieving relatively higher accuracy in shape reconstruction. In this process, the connectivity between medial axis points is key to establishing this interpolation. If we disregard this connectivity, the results will exhibit obvious errors; See Fig. E2. 2) Shape Segmentation: As MAT contains a medial surface, previous unsupervised structural decomposition/shape segmentation [LLL*20, LLL*21] has shown that detecting dimensional changes and non-manifold branches on the medial mesh can be effective for the shape segmentation, e.g., use the information from manifold branch and dimensional change for segmentation; See Fig. E3 (a). 3) Hexmeshing: Based on the shape description ability of medial mesh, one could produce all-hex meshing using topological and geometrical information of the medial mesh, which is more general than previous methods based on curve-skeleton as pointed out by [ZXW*24]; See Fig. E3 (b).

Table D4: *Quantitative comparison on run time and shape approximation error between Coverage Axis and Coverage Axis++ when the number of surface samples is set as* 5000*. The inputs are meshes.* ***Time****: Runtime measured in seconds.* $|V|$*: The number of skeletal points. Note we set the number of selected skeletal points of our method the same as the Coverage Axis [DLX*22] for a fair comparison.*

		Coverage Axis				Coverage Axis++			
Model	$\|V\|$	**Time**	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	**Time**	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$
Bear-1	70	115.2	3.651%	4.958%	4.958%	**5.1**	2.752%	2.795%	**2.795%**
Bear-2	38	7.5	7.412%	7.385%	7.412%	**2.5**	4.552%	5.997%	**5.997%**
Bird	101	>1000	1.959%	1.857%	**1.959%**	**9.8**	2.177%	2.438%	2.438%
Bug	145	>1000	5.701%	5.636%	5.701%	**12.2**	3.074%	3.005%	**3.074%**
Camel	116	25.6	2.488%	2.238%	**2.488%**	**9.5**	2.513%	2.936%	2.936%
Chair	137	>1000	1.920%	1.871%	**1.920%**	**13.6**	1.909%	2.105%	2.105%
Crab-1	92	43.9	3.738%	3.293%	3.738%	**9.5**	2.577%	2.738%	**2.738%**
Crab-2	83	210.5	7.705%	6.223%	7.705%	**8.9**	3.168%	3.212%	**3.212%**
Cup	286	>1000	9.176%	5.026%	9.176%	**27.8**	9.174%	6.390%	**9.174%**
Dog	65	7.6	4.557%	4.386%	4.557%	**5.0**	2.642%	2.876%	**2.876%**
Elephant	89	147.1	6.015%	7.221%	7.221%	**8.5**	3.168%	3.925%	**3.925%**
Fandisk	210	>1000	2.925%	2.778%	**2.925%**	**23.7**	2.613%	3.057%	3.057%
Femur	27	13.6	2.842%	2.782%	2.842%	**2.0**	2.035%	2.035%	**2.035%**
Fish	72	>1000	1.769%	2.271%	**2.271%**	**5.7**	1.994%	2.722%	2.722%
Giraffe	50	25.9	5.841%	5.853%	5.853%	**5.1**	2.627%	5.267%	**5.267%**
Guitar	63	>1000	2.040%	4.312%	4.312%	**7.4**	2.232%	2.383%	**2.383%**
Hand	46	131.4	2.874%	2.607%	**2.874%**	**4.2**	3.315%	2.884%	3.315%
Horse	56	200.2	5.174%	6.042%	6.042%	**4.6**	3.608%	5.246%	**5.246%**
Lifebuoy	38	96.7	4.667%	4.196%	4.667%	**2.8**	2.52%	3.307%	**3.307%**
Neptune	120	206.8	2.115%	5.751%	5.751%	**11.4**	3.128%	5.542%	**5.542%**
Pig	51	16.6	8.535%	8.575%	8.575%	**3.5**	6.038%	8.327%	**8.327%**
Plane	84	>1000	1.708%	2.275%	**2.275%**	**8.7**	1.965%	2.533%	2.533%
Pliers	32	13.5	2.072%	1.925%	**2.072%**	**2.8**	2.081%	2.055%	2.081%
Pot	152	>1000	3.570%	3.100%	**3.570%**	**14.4**	3.504%	3.690%	3.690%
Rapter	74	>1000	3.322%	6.545%	**6.545%**	**6.2**	3.322%	6.634%	6.634%
Rocker	130	>1000	3.632%	3.496%	**3.632%**	**14.4**	3.777%	2.790%	3.777%
Seahorse	53	>1000	2.763%	11.156%	11.156%	**4.8**	2.288%	5.515%	**5.515%**
Spectacles	112	>1000	1.982%	1.952%	1.982%	**14.4**	1.660%	1.653%	**1.660%**
Vase	109	226.1	5.260%	6.587%	6.587%	**9.2**	3.132%	3.192%	**3.192%**
Wine-glass	335	>1000	22.680%	7.904%	22.680%	**34.6**	21.310%	8.549%	**21.310%**
Average	-	516.3	4.670%	4.673%	5.448%	**9.7**	3.695%	3.860%	**4.429%**

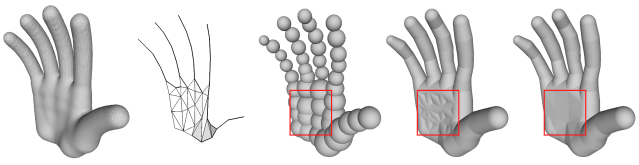

Figure E2: *The medial mesh part among the skeletal points in the MAT enables interpolation between medial axis balls, resulting in higher accuracy shape reconstruction (e). In contrast, reconstruction using only skeletal points (c) or without a medial mesh connection (d) typically leads to larger approximation errors.*

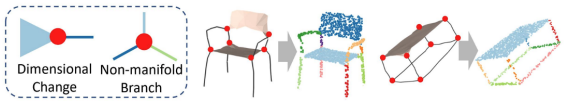

(a) [LLL*21]

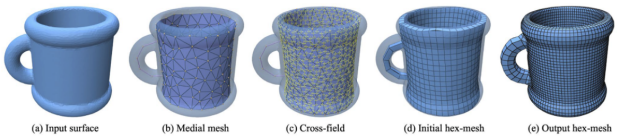

(b) [ZXW*24]

Figure E3: *(a) The medial axis transform (MAT) includes a medial surface. Unsupervised shape segmentation methods show that detecting dimensional changes and non-manifold branches on the medial mesh effectively aids segmentation [LLL*20]. (b) In the application of all-hex meshing [ZXW*24], the medial axis helps with computing the cross-field to guide the generation of the hexahedral mesh during the optimization.*

References

[CD23] CLÉMOT M., DIGNE J.: Neural skeleton: Implicit neural representation away from the surface. *Computers & Graphics 114* (2023), 368–378. 1

[DLX*22] DOU Z., LIN C., XU R., YANG L., XIN S., KOMURA T., WANG W.: Coverage axis: Inner point selection for 3d shape skeletonization. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 419–432. 3

[LLL*20] LIN C., LIU L., LI C., KOBBELT L., WANG B., XIN S., WANG W.: Seg-mat: 3d shape segmentation using medial axis transform. *IEEE Transactions on Visualization and Computer Graphics* (2020). 2, 3

[LLL*21] LIN C., LI C., LIU Y., CHEN N., CHOI Y.-K., WANG W.: Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 4277–4286. 2, 3

[WWWG22] WANG N., WANG B., WANG W., GUO X.: Computing medial axis transform with feature preservation via restricted power diagram. *ACM Transactions on Graphics (TOG) 41*, 6 (2022), 1–18. 1

[ZXW*24] ZHANG S., XU G., WU H., GU R., QI L., PANG Y.: Medial hex-meshing: high-quality all-hexahedral mesh generation based on medial mesh. *Engineering with Computers* (2024), 1–21. 2, 3