EG23

SIC Saarland Informatics Campus

Parameter-Free and Improved Connectivity for Point Clouds

DIANA MARIN*, STEFAN OHRHALLINGER, MICHAEL WIMMER

Institute of Visual Computing & Human-Centered Technology, TU Wien

*dmarin@cg.tuwien.ac.at

Introduction

Retrieving connectivity from points without a known surface has been a long-standing research problem in computer graphics. It is a necessary precondition in surface reconstruction since some type of connectivity is required in order to be able to create a triangulated surface from points. Usually, this problem is solved by computing a graph and filtering some of the edges to retain the best encoding connectivity.

We propose the 3D version of SIGDT [MOW22], which is a restriction of the spheres-of-influence proximity graph to the 3D Delaunay graph, and name it SIGDT3D. This graph recovers the original connectivity of the surface better than the commonly used kNN graph and is parameter free.

Related Work

For determining the neighborhood of unstructured points in 3D, the kNN method is widely used, e.g., when computing the tangent plane for connectivity or normals [HDD∗92]. Many subsequent surface reconstruction methods rely on kNN, up to recent deep learning methods that require it for training purposes, e.g., [EGO∗20].

The spheres-of-influence graph (SIG) has been introduced in [Tou88] as a clustering method. Two vertices are connected in SIG if the distance between them is less or equal to the sum of the distances to their respective nearest neighbours. The rarely used SIG has recently gotten attention for its improved connectivity when reconstructing curves from unstructured points [MOW22].

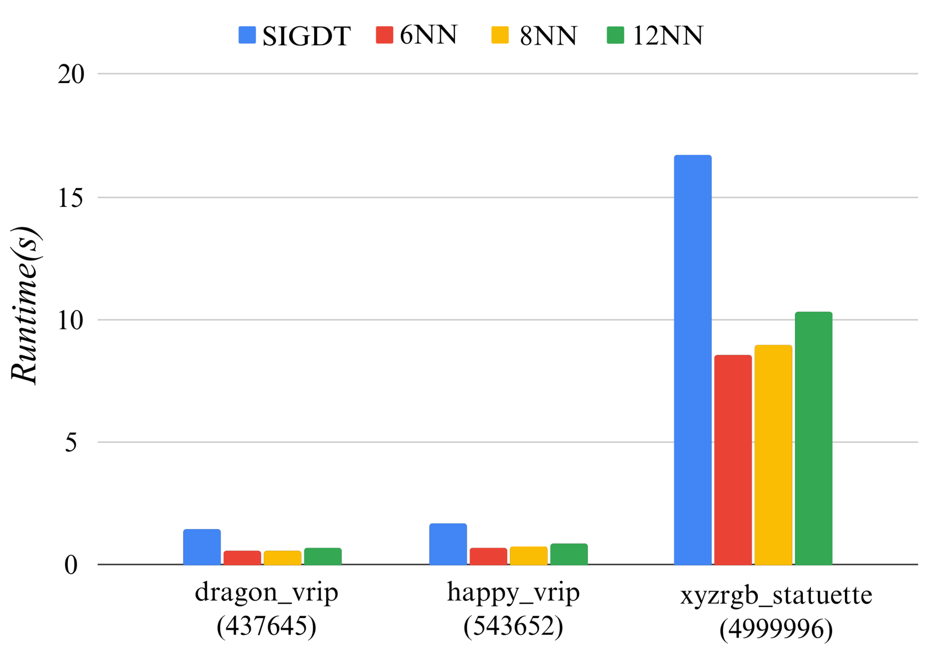

Figure 1: Runtime in seconds for datasets with various number of points.

Method

Our algorithm first computes the 3D Delaunay graph on the input points. We then filter the Delaunay edges that are part of the SIG. The only information required for this computation is the distance to the nearest neighbour, and the Delaunay triangulation already contains these edges. Hence, we pass over all edges in order to save the nearest neighbour distances at the vertices, without increasing time or space complexity over the Delaunay graph computation. Then, we apply the SIG criterion to only keep the Delaunay edges that satisfy the condition.

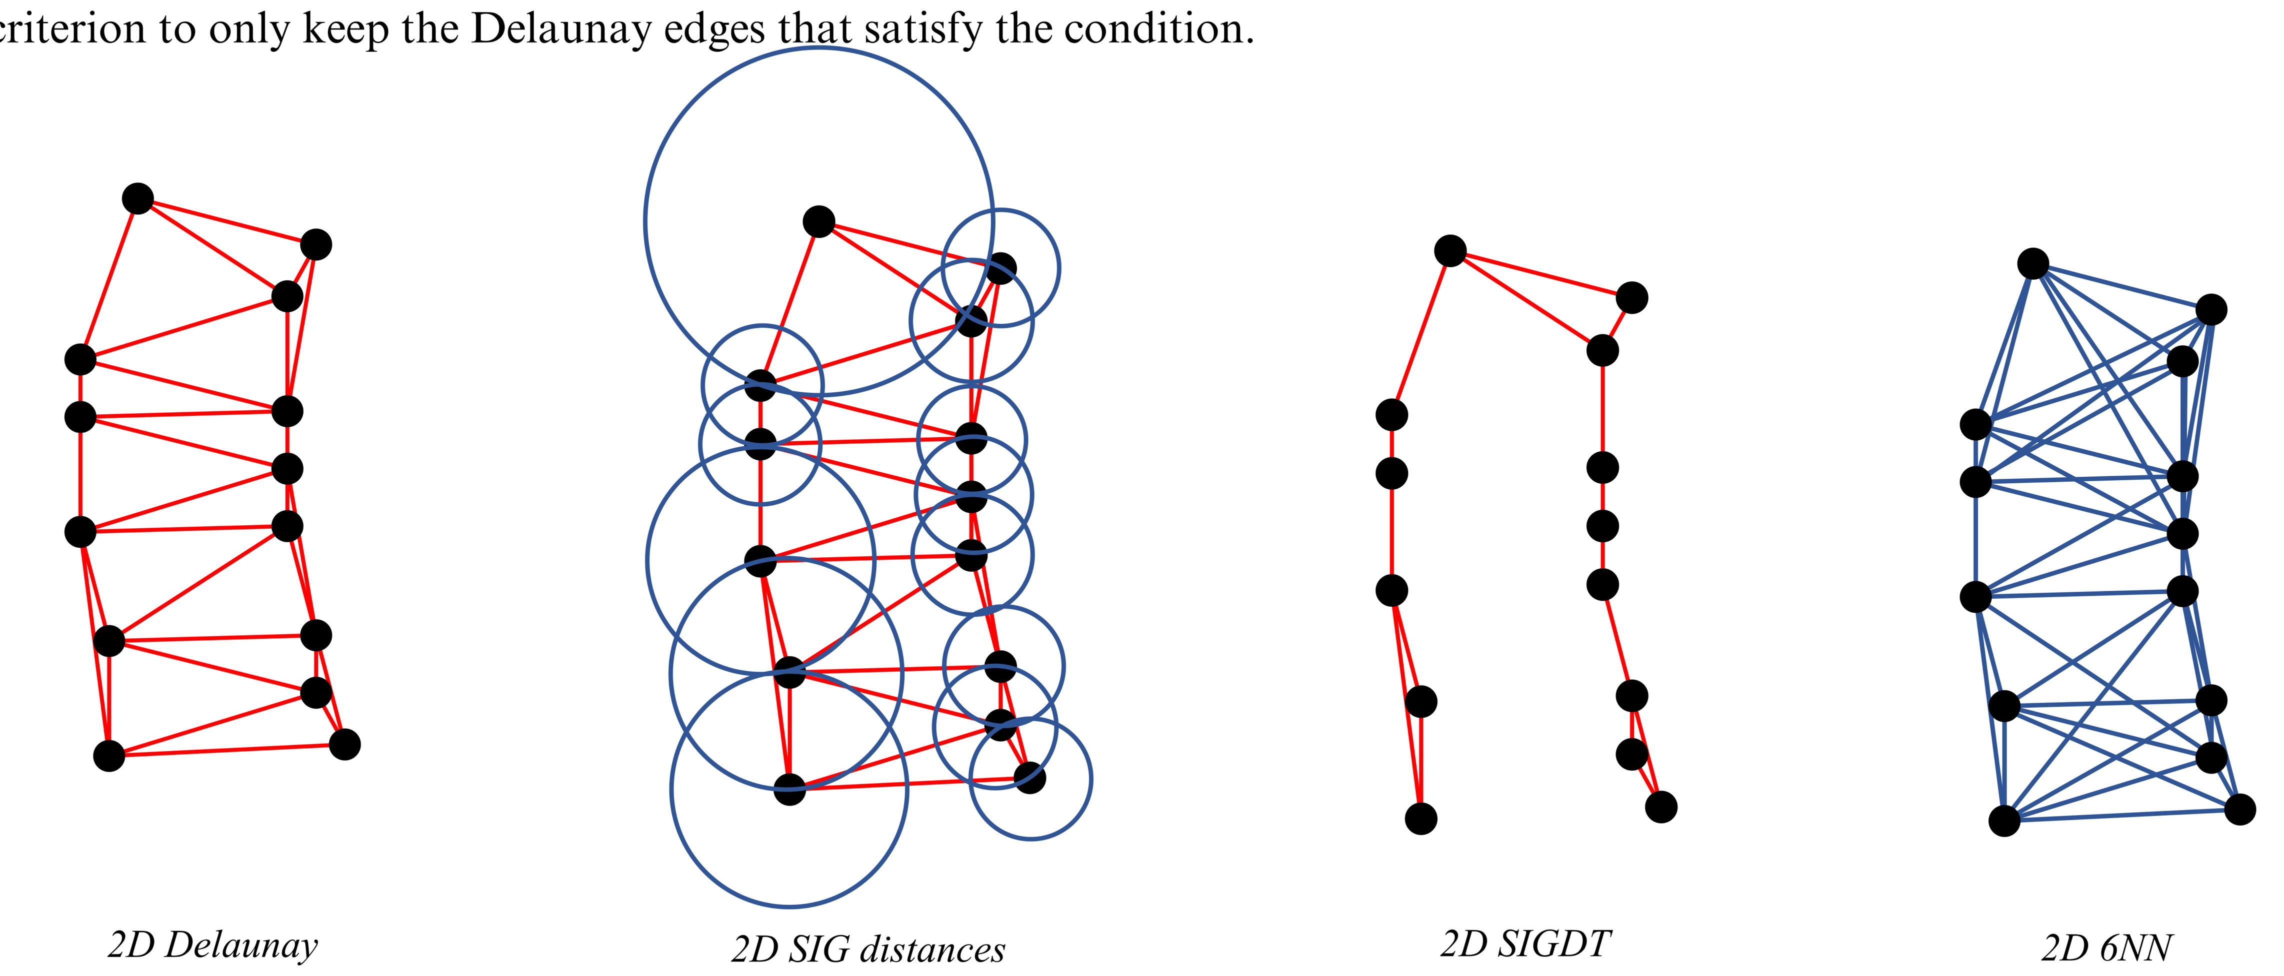

Results

We have tested SIGDT3D on various models at different resolutions from the Stanford Scanning Repository. We have also computed the kNN graphs for typical values used in points' connectivity, k=6,8,12, and extracted the edges of the original reconstruction. In order to assess the quality of the connectivity encoding, we computed the intersection over union (IoU) between the obtained and the original edge set. Using this metric, a score of 1 means a perfect match while 0 means no overlap at all. A visual representation is presented in Figure 2.

For almost all 3D models, SIGDT3D had a higher IoU value than any of the kNN graphs (Figure 3). This value decreases significantly when k is increased, due to the number of extra edges that are not contained in the ground truth surface.

However, our method proves to be slower due to the need to compute the Delaunay triangulation. We have tested our method against kNN, with k=6,8,12, on the models provided in the Stanford repository, and our method is, on average, 3 times slower than the kNN neighbourhood (Figure 1). Both implementations took advantage of CGAL's parallel of Delaunay triangulation and the nearest neighbourhood search respectively. However, the Delaunay triangulation accounts, on average, for 60% of the computation, which indicates the possibility of improvement in our code, but clearly not above the performance of kNN.

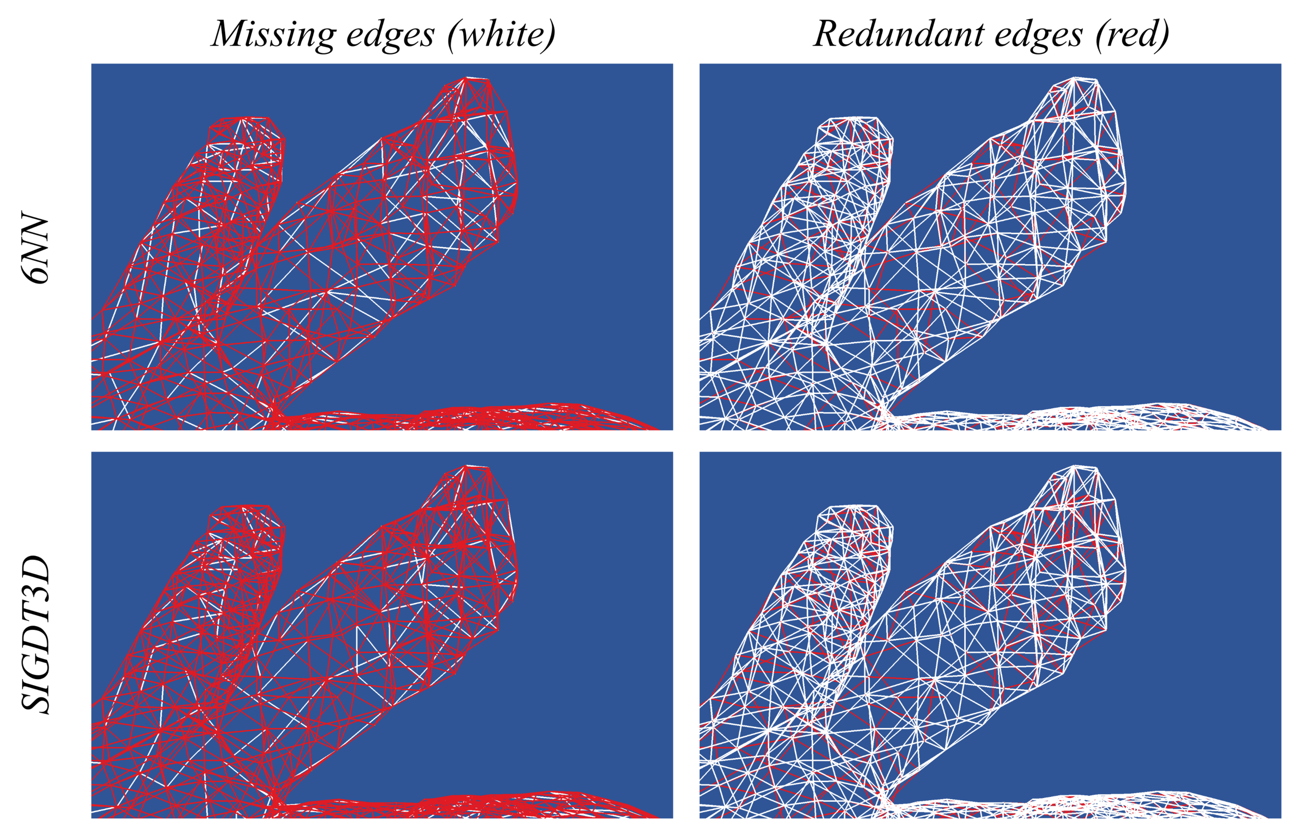

Figure 2: Visual results of IoU for a sparsely sampled ears region: original edges are in white, while reconstructed ones are in red. We overlay them to show the number of redundant/missing edges compared to the original.

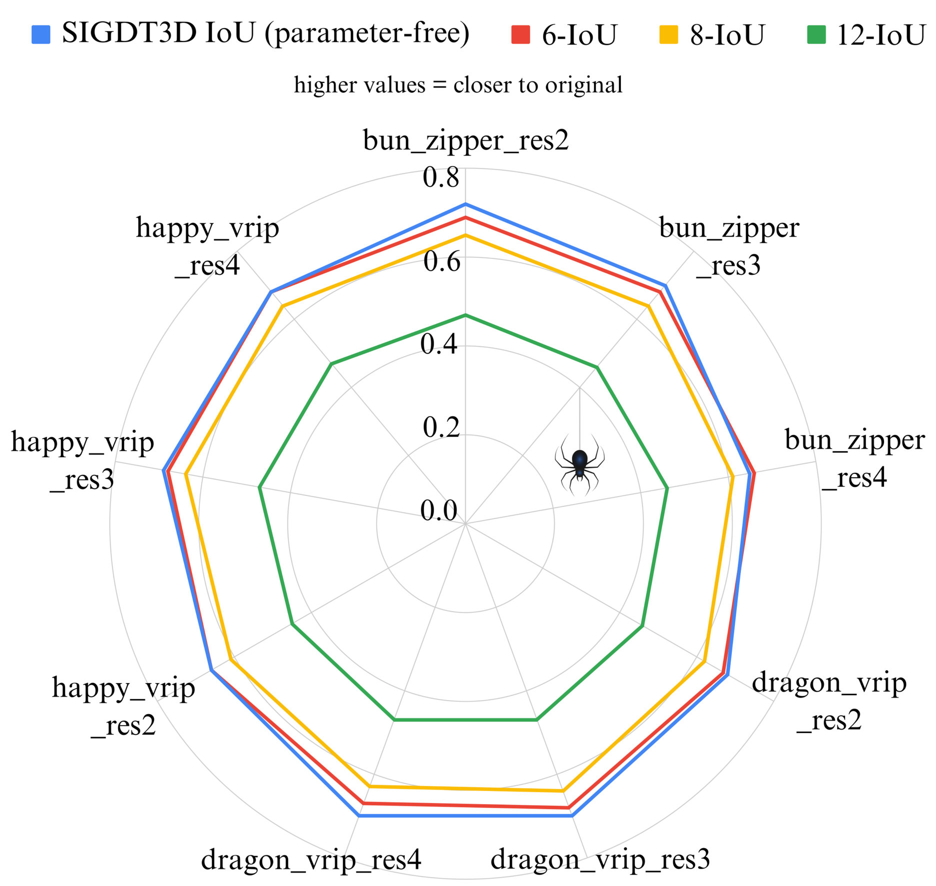

Figure 3: Spider chart for IoU over multiple datasets. Our method in blue performs best (highest score) for almost all tested meshes.

Conclusion and Future Work

We propose an extension of the SIGDT proximity graph to 3D. This graph shows significantly better proximity encoding, while keeping a low number of edges and not requiring a parameter. Thus, it offers two important advantages over kNN. The current results also encourage more experiments with the many other applications of kNN in improvements in normal estimation, surface reconstruction, motion planning, learning from 3D data, simulations, and many more.

References

[EGO∗20] ERLER P., GUERRERO P., OHRHALLINGER S., MITRA N. J., WIMMER M.: Points2surf learning implicit surfaces from point clouds. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V (2020), Springer, pp. 108–124.

[HDD∗92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In Proceedings of the 19th annual conference on computer graphics and interactive techniques (1992), pp. 71–78.

[MOW22] MARIN D., OHRHALLINGER S., WIMMER M.: Sigdt: 2d curve reconstruction. Computer Graphics Forum 41, 7 (Oct. 2022), 25–36. doi:10.1111/cgf.14654.

[Tou88] TOUSSAINT G. T.: A graph-theoretical primal sketch. In Machine Intelligence and Pattern Recognition, vol. 6. Elsevier, 1988, pp. 229–260.

TU WIEN

EG23