

Approximating the Medial Axis Transform of LiDAR point clouds

Ravi Peters
r.y.peters@tudelft.nl


1. Introduction

LiDAR point clouds are highly detailed and cover large areas. This brings great advantages for applications such as flood modeling, crisis management and 3D city modeling. Unfortunately, and despite recent developments on this subject, current methods from practice are unable to fully take advantage of modern LiDAR datasets. First, because of their huge data volume they do not fit in a computer's internal memory. As a result, many of the conventional software tools have become very inefficient. And second, many existing methods use only 2.5D data-structures and algorithms. While this alleviates memory requirements and simplifies computation, it comes at the price of a significant loss of information, because valuable 3D information that is present in LiDAR point clouds is ignored.



LiDAR dataset from the municipality of Rotterdam

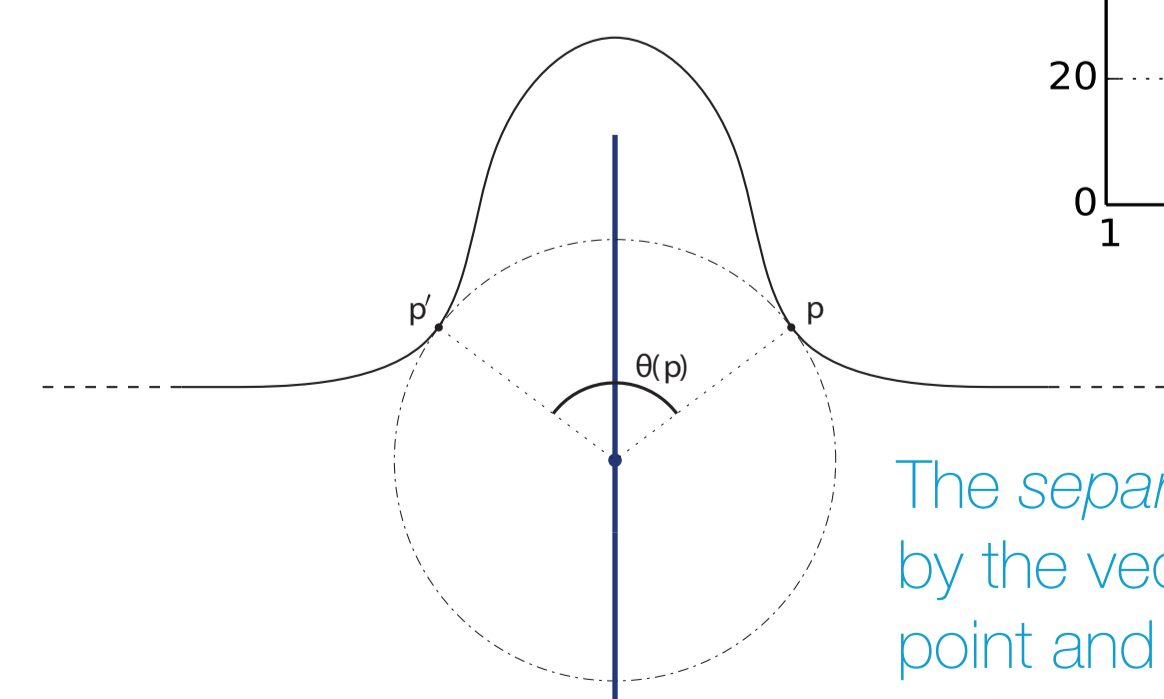

My research aims to develop methods for efficient simplification and shape matching in LiDAR point clouds that use truly 3D data-structures and algorithms based on the Medial Axis Transform (MAT).



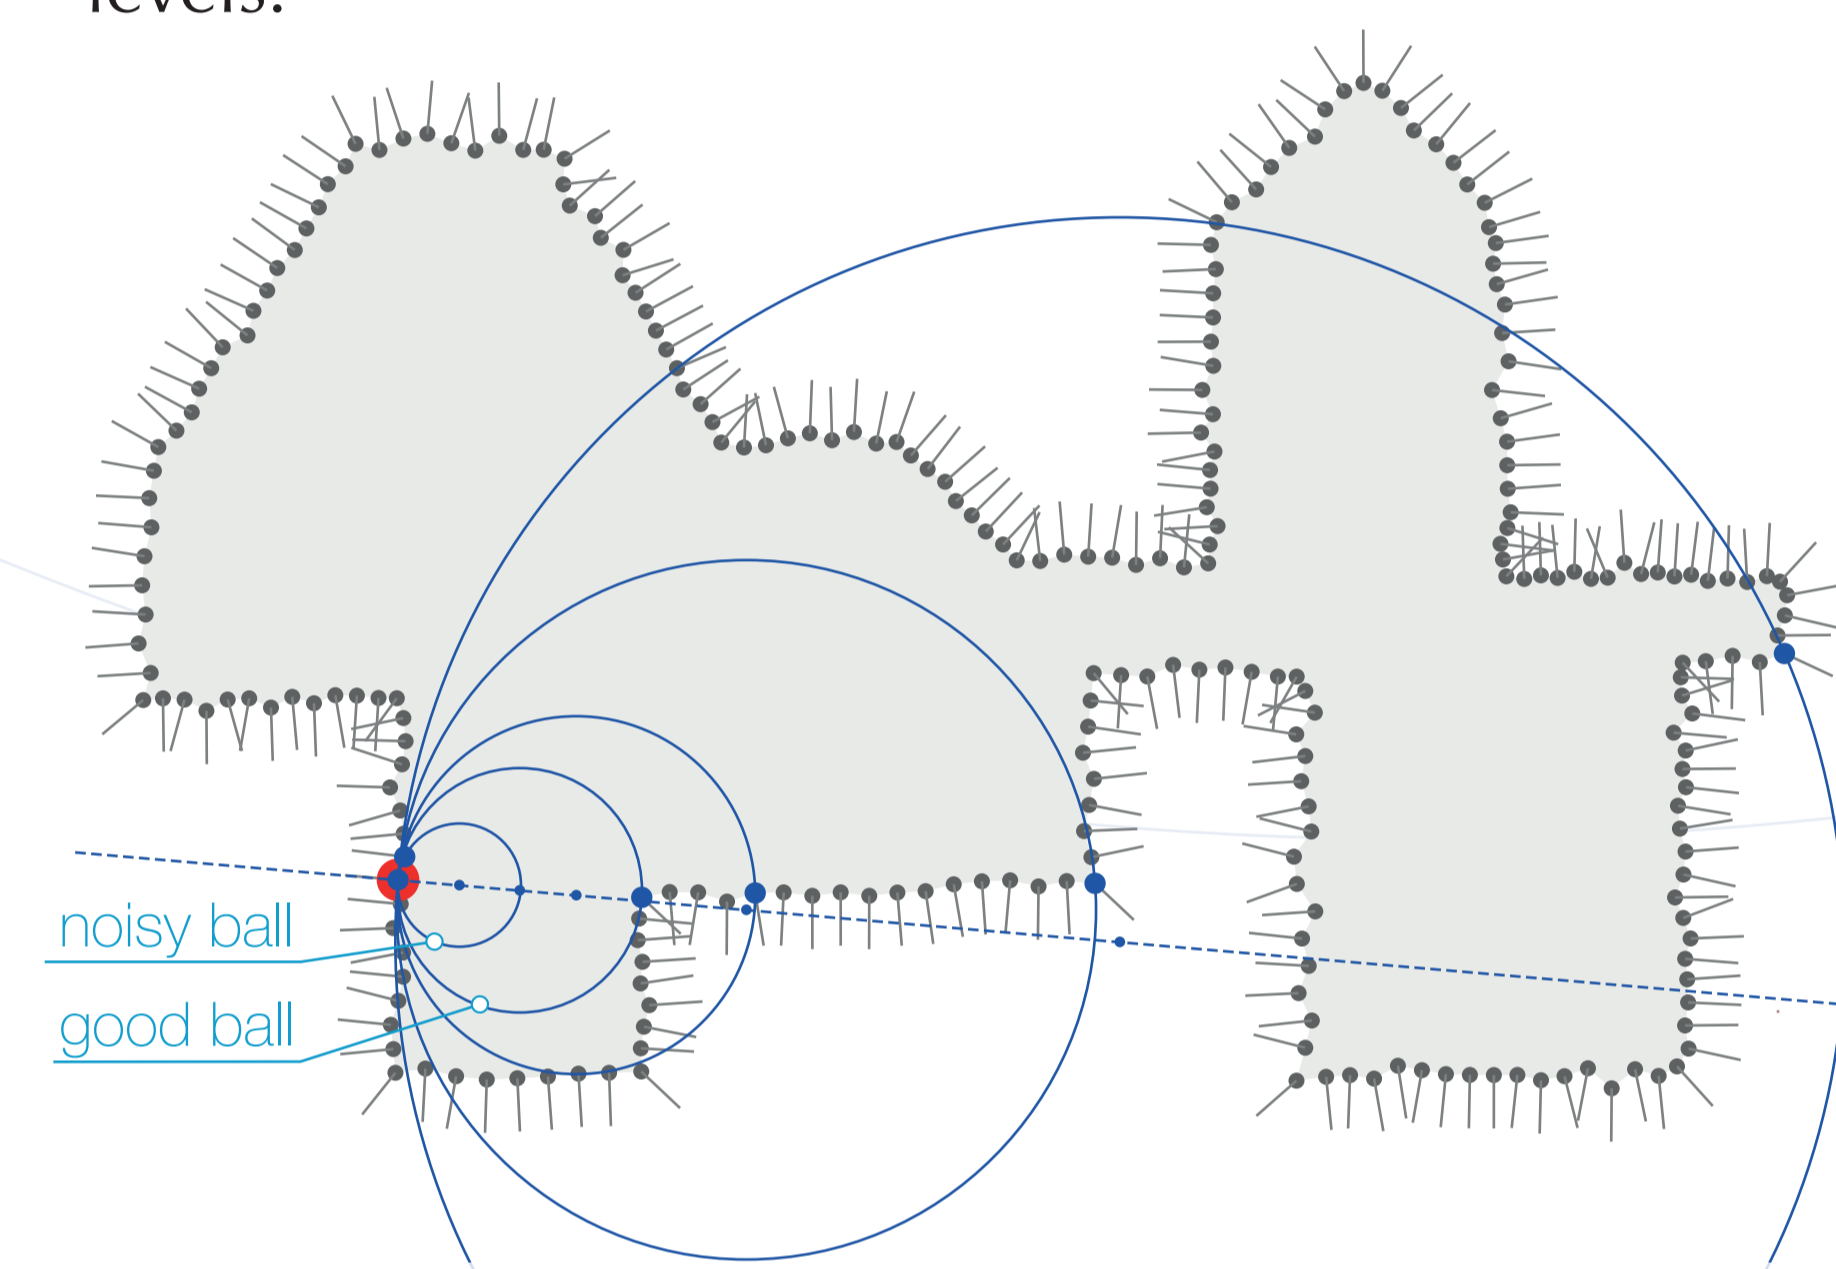
2. Shrinking balls

Ma et al. (2012) introduced the shrinking ball algorithm to approximate a point approximation of the MAT from an oriented input point cloud. For each sample point an empty maximal tangent ball is found by iteratively reducing its radius using nearest neighbor queries. Compared to earlier (Voronoi-based) algorithms (e.g. Amenta et al., 2001) it is simple, fast, robust in practice, and easy to parallelize. This makes it a good choice for approximating the MAT of large LiDAR point clouds.

The *Medial Axis Transform (MAT)* is defined as the set of maximal balls tangent to the object surface at two or more points. The centers of these balls form the Medial Axis, a medial skeletal structure.

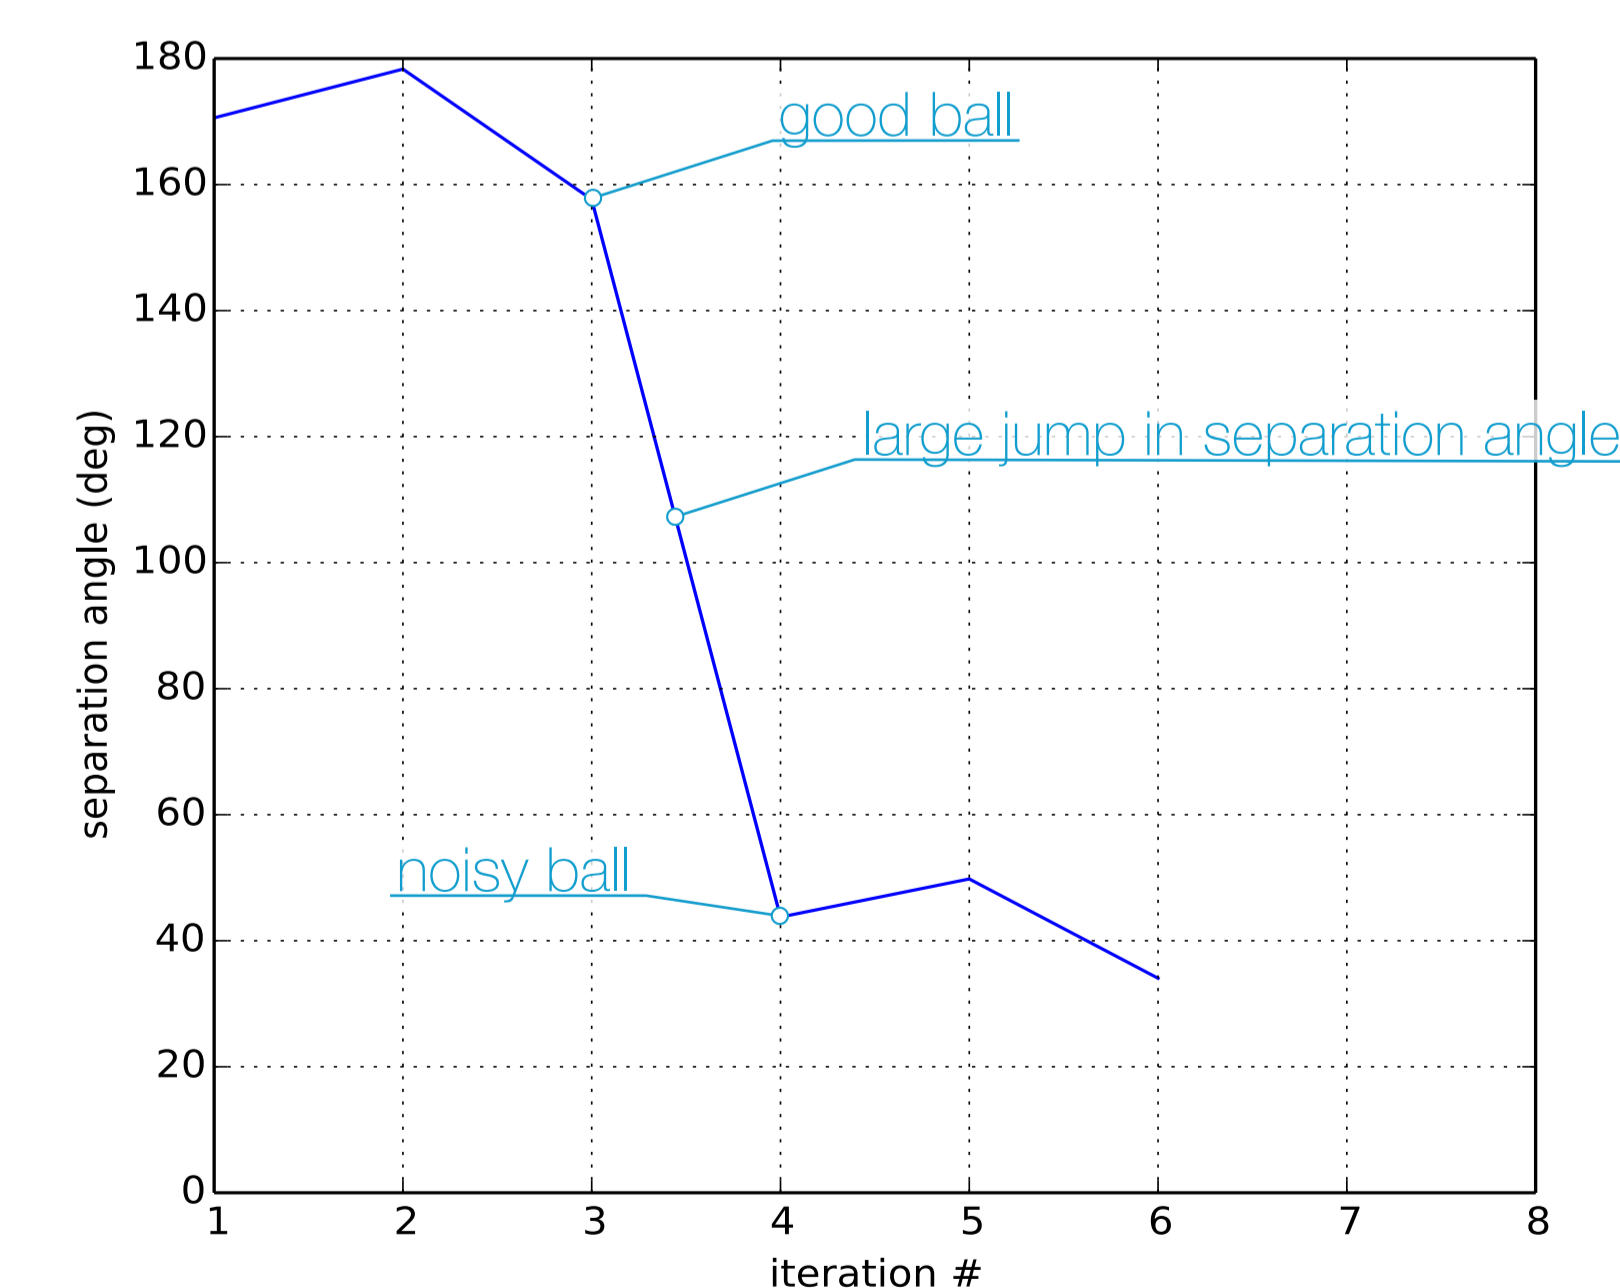


The *separation angle* θ is formed by the vectors between a medial point and its two corresponding surface points.



Therefore, I have modified the shrinking ball algorithm to consider the continuity in *separation angle* of a shrinking ball. Whenever a large decrease in separation is about to happen, the ball is shrunk no further (and kept as a medial ball).


Compared to simpler thresholding methods on the final separation angle of a medial ball, this method delivers less false positives in practice, and delivers a much denser MAT approximations for noisy datasets.




3. Handling noise

The MAT is notorious for its sensitivity to small perturbations in the object surface. Since LiDAR point clouds typically contain significant levels of noise, it is essential to deal with this problem.

Existing methods to deal with noisy input points suffer from mainly two flaws: 1) they detect a lot of false positives, and 2) they *remove* noisy MAT points. The resulting MAT approximation can thus become very sparse for higher input noise levels.



Original shrinking ball algorithm
Modified shrinking ball algorithm



Top view of a building extracted from the LiDAR dataset of the municipality of Rotterdam. On top, the MAT approximation obtained using the original shrinking ball algorithm. The bottom figure show the MAT approximation obtained using the modified noise-aware version of the shrinking ball algorithm.

Case study

The Dutch Kadaster is currently working on a national topographic map in 3D using the LiDAR-based national elevation model (AHN2). To keep processing times on a reasonable level, the AHN2 dataset must be thinned in a preprocessing phase. Currently, they employ a simple n^{th} point filter. Unfortunately, due to non-adaptiveness this easily results in the destruction of significant surface characteristics.

For this case study I investigate an adaptive thinning filter that considers local point configuration based on the MAT-based local feature size metric, similar to the work of Dey et al. (2001).

References

- Jaehwan Ma, Sang Won Bae, and Sunghee Choi. 3D medial axis point approximation using nearest neighbors and the normal field. *The Visual Computer*, 28(1):7–19, 2012
- Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, pages 249–266, 2001
- Tamal Dey, Joachim Giesen, and James Hudson. Decimating samples for mesh simplification. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, pages 85–88, 2001