

The Medial Axis Transform

Lesson 09*

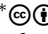
1	Defining the MAT	3
1.1	Medial geometry	3
1.2	Exterior MAT and the MAT of a DTM	4
2	Computing the MAT	5
2.1	MAT approximation	5
2.1.1	Methods based on the Voronoi diagram	5
2.1.2	The shrinking-ball algorithm	5
2.2	Pruning	8
3	Applications of the MAT for DTMs	8
4	Notes and comments	11

Until now, we have described the shape of a terrain using the boundary representation. The boundary representation, or *B-rep*, is the most commonly used shape representation for DTMs, ie we model a terrain by describing explicitly its boundary surface, eg using points, grid cells, triangles, or contour lines. The medial axis transform (MAT) gives us another way to represent shape. Contrary to the B-rep, the MAT describes a shape by its skeleton instead of its boundary (compare Figures 1b and 1c).

The MAT can be considered a *dual* representation to the B-rep, similar to how the Voronoi diagram is dual to the Delaunay triangulation. Both the MAT and the B-rep contain exactly the same information and it is possible to convert one to the other without loss of information. However, the two different representations expose different properties of a shape which means they both have their own unique use cases.

For example, the MAT allows us to split a shape into parts simply by looking at the branches of the skeleton. The resulting shape parts often turn out to be meaningful in practice. Observe for instance that for the gingerbread man in Figure 1, its arms, legs, torso and head each have one corresponding branch in its medial axis (compare Figures 1a and 1c). For DTMs, equally meaningful decompositions into parts can be made, eg the MAT allows us to decompose a DTM into separate hills, watercourses and other objects on top the DTM (see Figure 2).

In the context of DTMs both the 2D and the 3D MAT are used. The 2D MAT is used for operations on sets of 2D curves, such as isolines or river networks. In these cases no elevation information is used to construct the MAT. With the 3D MAT on the other hand, the whole surface of the terrain, including the elevation, is used during construction.

* Ravi Peters, Hugo Ledoux, Ken Arroyo Ohori. This work is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>) (last update: December 10, 2018)




Figure 1: Different ways to represent the shape of gingerbread man




Figure 2: For a terrain the MAT is typically subdivided into open clusters that correspond to features such as hills, buildings and watercourses in the terrain. Shown here is a vertical cross section of a DTM. Exterior MAT in light blue, interior MAT in dark blue.

1 Defining the MAT

There are two equivalent definitions of the MAT¹. One is based on the distance transform, and one is based on medial balls. Both definitions describe how to obtain the MAT from the boundary, denoted \mathcal{B} of an object (Figure 1b). And both can be applied to both 2D and 3D shapes.

Grassfire analogy Imagine that everything is made of grass and that all the points on \mathcal{B} are simultaneously set on fire at time $t = 0$. The fire spreads evenly to all directions at constant speed. Now, the MAT is defined as the set of points where the fire front meets itself. This concept is illustrated in Figure 1e, where each contour can be seen as a fire front at some constant time t . The medial axis is drawn where the fire front meets itself.

Medial balls A *medial ball* is a ball that fits completely inside \mathcal{B} and does not contain any other ball that would fit inside \mathcal{B} . The MAT is defined as the set of points that are the centres of all medial balls of \mathcal{B} (see Figure 1d). Notice that each medial ball touches \mathcal{B} in at least two points, called its *feature points*.

As illustrated in Figure 1c, the MAT can be subdivided into *medial branches* and *junctions*. Junctions are locations where three or more medial branches coincide. The points of the MAT are called *medial atoms*, or simply *atoms*. Observe that if an atom has exactly two feature points, it is part of a medial branch, and if it has more than two feature points it lies on a junction or on the tip of a medial branch. The medial branches, its junctions and how those are are connected define the *medial structure*.

For a 2D object, such as in Figure 1, the medial branches are curves and the the junctions are points. However, for a 3D object, the medial branches can also be surfaces (see Figure 3), and the junctions can also be curves. The branches of the 3D MAT are therefore also called *medial sheets*.


1.1 Medial geometry

The medial geometry describes how atoms are related to the object boundary \mathcal{B} . It is defined for each medial atom that is part of a medial sheet. Figure 4 illustrates the complete medial

¹Sometimes the MAT is referred to as medial axis function, stick figure, skeleton or surface skeleton. Inventor Harry Blum finally settled on symmetry axis, as he considered symmetry to be the crucial role of the MAT (Blum, 1973).



Figure 3: An example of the 3D MAT (Amenta et al., 2001).



Symbol	Description
$B(\mathbf{c}, r)$	medial ball
\mathbf{c}	medial atom
r	radius
\mathbf{p}, \mathbf{q}	feature points
$\vec{\mathbf{s}}_{\mathbf{p}}, \vec{\mathbf{s}}_{\mathbf{q}}$	spoke vectors
θ	separation angle
$\vec{\mathbf{b}}$	medial bisector

Figure 4: The geometry of a medial atom.

geometry of an atom. The medial ball B has the atom \mathbf{c} at its center and has a radius r , ie the shortest distance from \mathbf{c} to \mathcal{B} . The medial ball B touches the boundary \mathcal{B} at the feature points \mathbf{p} and \mathbf{q} . The vectors from \mathbf{c} to \mathbf{p} and \mathbf{q} are called the *spoke vectors*, denoted $\vec{\mathbf{s}}_{\mathbf{p}}$ and $\vec{\mathbf{s}}_{\mathbf{q}}$. The angle between the spoke vectors is called the *separation angle*, denoted θ and the bisector of the spoke vectors is called the *medial bisector*, denoted $\vec{\mathbf{b}}$.

Using the medial geometry we can describe a number of interesting properties of the MAT.

1. Any atom \mathbf{c} is always *medial* to \mathcal{B} , ie it is equidistant to the feature points of \mathbf{c} (hence the name of the MAT).
2. The medial ball B is always tangential to \mathcal{B} at the feature points.
3. The radius r can be used to define the ‘thickness’ of an object, since it measured the distance to the ‘middle’ of the object where the MAT is located.

1.2 Exterior MAT and the MAT of a DTM

The MAT can be divided into an interior part and an exterior part. So far we have only looked at the *interior MAT*, which consists of medial balls that reside entirely on the inside of an object. However, in many cases it is also possible to define medial balls that reside entirely on the outside of an object. That part of the MAT is called the *exterior MAT*. An object can only have an exterior MAT if the shape of that object is non-convex, since for convex object it is not possible to find exterior medial balls with a finite radius.

The separation between inside and outside is very clear and unambiguous for a an object with a closed boundary such as the gingerbread man of Figure 1. But, a DTM typically only covers part of the surface of the planet. This means it is not a closed boundary, unless we have a DTM of the entire planet’s surface. We then follow the convention that the ‘ground side’ of the DTM is the interior, and the ‘sky side’ is the exterior, as follows from Figure 5. Following this convention, we can still define the interior and exterior MAT of a DTM, see for instance Figure 2.

A second peculiarity for DTMs is that both its interior and its exterior MAT can consist of multiple disjoint parts (also visible in Figure 2), whereas for closed objects, the interior MAT always consists of one connected part. This is again a side effect of the open boundary. These disjoint parts are called *medial clusters*. In most cases, one object on the terrain corresponds to one medial cluster.

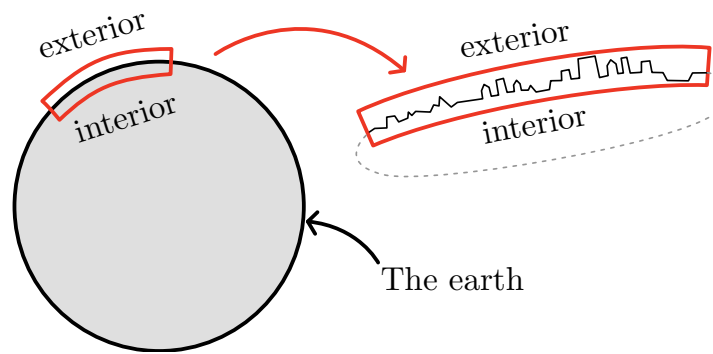


Figure 5: The interior and exterior of a DTM.

2 Computing the MAT

Computing the MAT from the boundary \mathcal{B} of an object is typically done in two steps. During the first step, ie *MAT approximation*, a noisy approximation of the MAT is obtained, and during the second step, ie *pruning*, the noise is removed.

2.1 MAT approximation

For DTMs the MAT is typically approximated using the Voronoi diagram in case the 2D MAT is needed, or using the shrinking-ball algorithm in case the 3D MAT is needed.

2.1.1 Methods based on the Voronoi diagram

If you compute the Voronoi diagram (VD) of a set of points on the boundary of an object, the MAT can be found by taking a subset of the edges of that VD. This can be seen in Figure 6a. Notice that the MAT approximation becomes more accurate if more boundary points are used to compute the VD. Voronoi-based methods revolve around selecting the correct edges from the VD that approximate the MAT.

A simple way to check if a Voronoi edge is part of the MAT is based on so-called ‘duality circles’. A duality circle is defined for each pair of dual edges, ie one edge from the VD and one from the DT that are each others dual. The duality circle is drawn through the two endpoints of the Delaunay edge and one point of the Voronoi edge (see the two circles in Figure 6b). If and only if the second point of the Voronoi edge lies inside the duality circle, the Voronoi edge is part of the MAT.

2.1.2 The shrinking-ball algorithm

The shrinking-ball algorithm works well for robustly approximating the MAT of 3D objects that are represented with boundary points, eg a noisy aerial lidar point cloud of a terrain would be a good candidate. It is a simple and fast algorithm that can be made robust to noise in the boundary points. This means that pruning the MAT after approximation is not necessary with the shrinking-ball algorithm. The shrinking-ball algorithm takes an oriented point cloud as input, ie a point cloud that includes a normal vector for each point. It outputs a disjoint set of medial atoms.

The algorithm is based on the observation that the medial atom corresponding to a boundary point \mathbf{p} must be positioned somewhere on the line L through the normal $\vec{\mathbf{n}}$ of \mathbf{p} . This

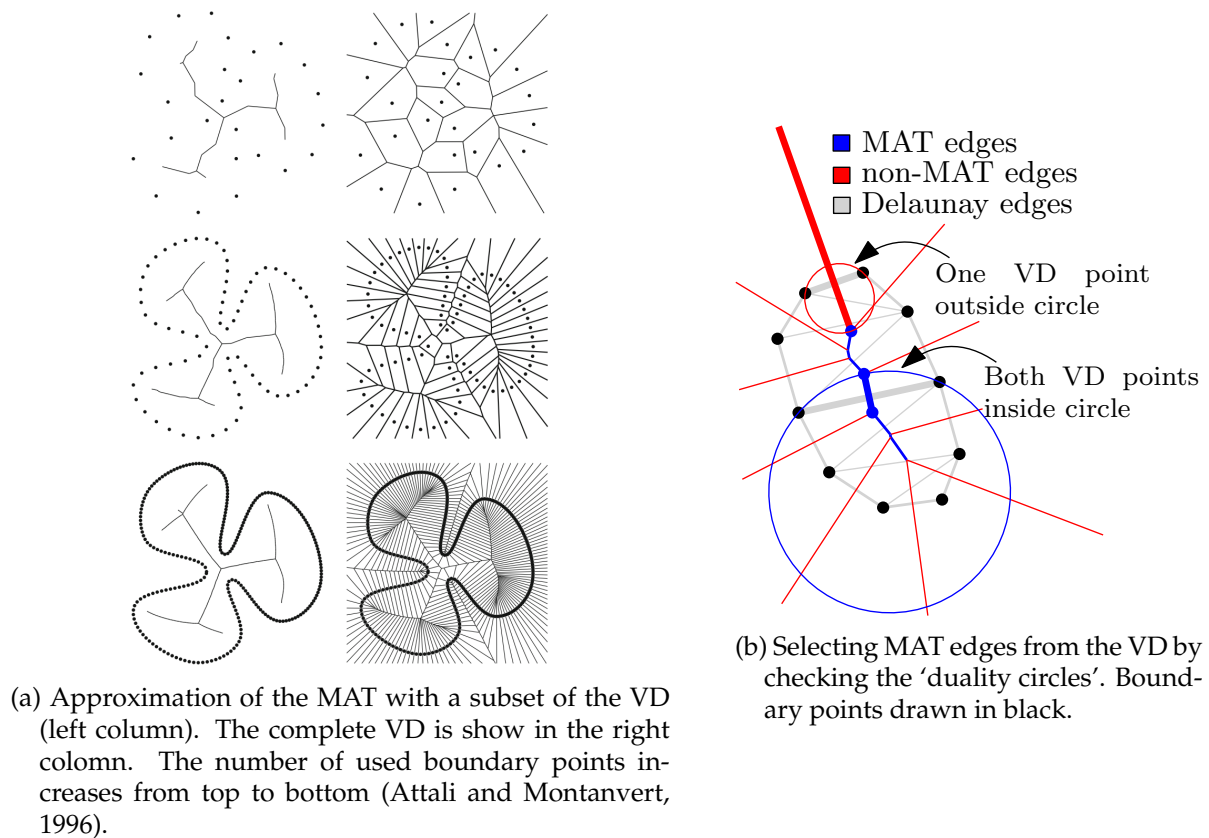


Figure 6: MAT approximation with the Voronoi Diagram

observation is used to restrict the search space for the medial ball of \mathbf{p} to the line L . As illustrated by Figure 7, the algorithm begins with a very large candidate ball for \mathbf{p} that is centered on L . At each consecutive iteration, a new candidate ball is constructed that is smaller than the previous one and closer to the final *medial* ball. Every ball is constructed so that it touches \mathbf{p} and is centred on L . Only the candidate feature point \mathbf{q} changes with each iteration. A new \mathbf{q} , denoted \mathbf{q}_{next} , is found by selecting the closest point from the centre \mathbf{c} of the current ball. Using \mathbf{p} , \mathbf{n} , \mathbf{q}_{next} we can compute the centre of the next ball \mathbf{c}_{next} , at which point we move on to the next iteration. The algorithm terminates when an empty ball is found, which is the case when the radius no longer shrinks. Algorithm 1 gives the pseudo-code for the shrinking-ball algorithm.

The algorithm is run for each point in the input point cloud. If the normal vectors point away from the interior, the interior MAT is computed. And by flipping the orientation of the normal vector, the exterior MAT can also be computed. If normals are not available for a point cloud, these can be estimated using local plane fitting, ie by fitting a plane to the k nearest neighbours of each boundary point. The vector perpendicular to that plane then becomes the estimated normal vector. A KD-tree is typically used to speed up the nearest neighbour searches, both for normal estimation and the shrinking-ball algorithm. Figure 8 gives an example result of the shrinking-ball algorithm for a terrain point cloud. Observe how the MAT describes the valleys (exterior MAT) and ridges (interior MAT) in the terrain with its skeletal structure.

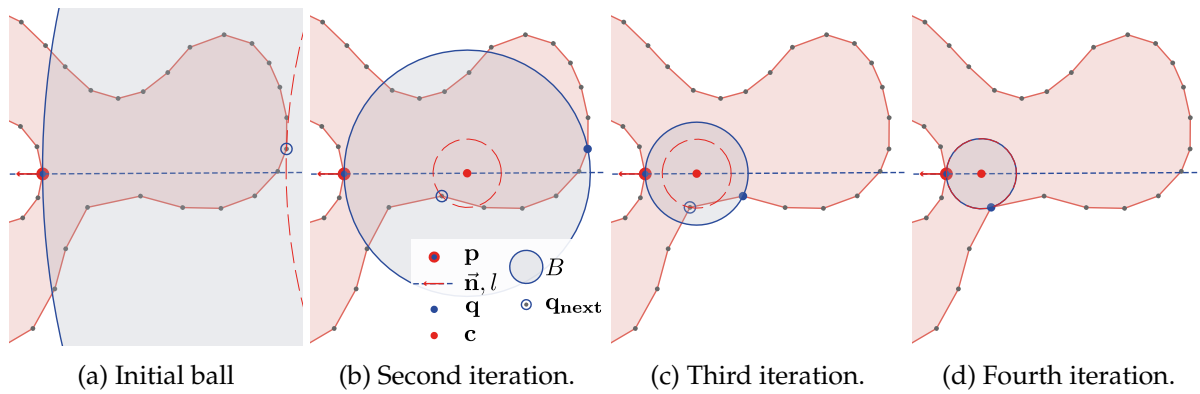


Figure 7: Ball shrinking iterations with the shrinking-ball algorithm. The final iteration yields a *medial ball*. A legend is given in b.

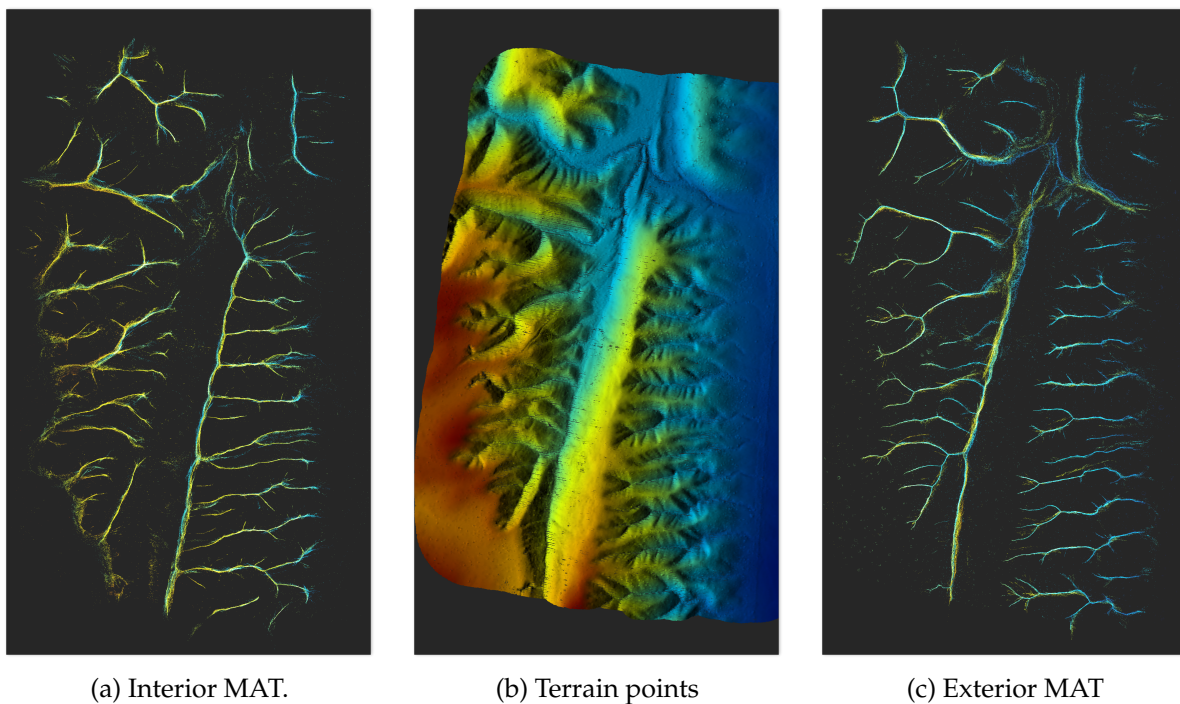


Figure 8: MAT approximation (a,c) of a lidar terrain point cloud (b) obtained with the shrinking-ball algorithm. Top view.

Algorithm 1: The shrinking-ball algorithm.

Input : a KD-tree of the surface point cloud T ,
a surface point \mathbf{p}
it's normal vector $\vec{\mathbf{n}}$, and
the initial ball radius r_{init}

Output: the medial ball centre \mathbf{c} ,
the medial ball radius r

- 1 $r \leftarrow r_{init}$
- 2 $\mathbf{c} \leftarrow$ the centre of the ball that touches \mathbf{p} is centered on $\vec{\mathbf{n}}$ with a radius r
- 3 **repeat**
- 4 $q_{next} \leftarrow$ the nearest point to \mathbf{c} , obtained quickly using T
- 5 $r_{next} \leftarrow$ radius of the next ball that touches \mathbf{p} and \mathbf{q}_{next} and is centred on $\vec{\mathbf{n}}$
- 6 $\mathbf{c}_{next} \leftarrow$ centre of the next ball, can be computed with \mathbf{p} , $\vec{\mathbf{n}}$, and r_{next}
- 7 **if** $r_{next} = r$ **then**
- 8 **break**
- 9 $\mathbf{c} \leftarrow \mathbf{c}_{next}$
- 10 $r \leftarrow r_{next}$
- 11 **until** a *break* statement is executed


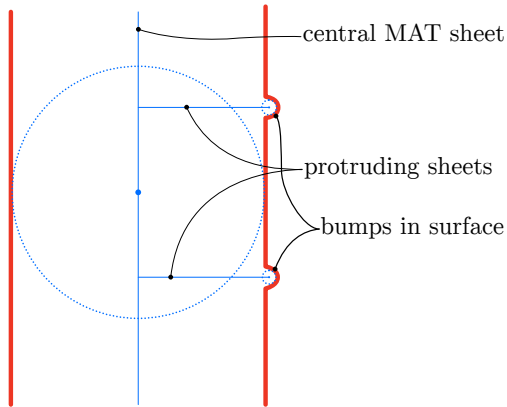
2.2 Pruning

Pruning is the process of retracting or removing unimportant branches from the MAT. It is often necessary because the MAT is unstable, ie it is extremely sensitive to small bumps in \mathcal{B} . As illustrated in Figures 9 and 10a, tiny deviations in \mathcal{B} can lead to big spurious branches in the structure of the MAT. This is especially problematic when \mathcal{B} has some noise as is the case with the typical DTM. The resulting MAT can become so distorted by the spurious branches that it becomes hard to distinguish its main medial structure. The main aim of pruning is to remove these spurious branches.

Most pruning methods are based on the properties of the medial geometry. Based on these properties, an importance measure for each medial atom is defined, which is then used as a threshold to filter medial atoms. The resulting (pruned) MAT is usually a subset of the original MAT. Some methods preserve topology, others do not or do so only up to a certain level. The main challenge is often selecting the optimal threshold value—a compromise between removing noisy MAT parts and not removing fine detail, ie often the endpoints of good MAT branches are also affected by pruning. Examples of importance measures for pruning are the separation angle θ (recall this is the angle between the spoke vectors) and the separation distance λ , ie the distance between the two feature points of a medial atom. These values are typically low for noisy parts of the MAT. Figure 10 gives an example of pruning with the separation distance.

3 Applications of the MAT for DTMs


There are several applications of the 2D MAT. Often it is applied on sets of lines such as isolines or river networks to perform operations that require the preservation of topology, ie there should not be any self-intersections or intersection between neighbouring lines. This is often a problem with methods that treat the lines separately without considering the complete system of lines. However, by working on the MAT, and keeping the boundary in sync during the MAT processing, this requirement can be satisfied rather easily.



(a) Small bumps in the boundary can cause big spurious branches to appear. Object boundary drawn in red, MAT in blue.

(b) Spurious branches can appear in the MAT due to a small amount of noise in the boundary.


Figure 9: Instability of the MAT.



(a) MAT without pruning ($\lambda = 0$)




(b) Reconstructed object for $\lambda = 0$.



(c) MAT with medium pruning ($\lambda = 6$)



(d) Reconstructed object for $\lambda = 6$. Despite the pruned MAT, the corresponding boundary remains almost unchanged.




(e) MAT with strong pruning ($\lambda = 10$).




(f) Reconstructed object for $\lambda = 10$. Notice how the tree is separated from the the ground and that edges have become rounder.

Figure 10: The effect of different levels of pruning on the MAT.



(a) Perspective view of exterior MAT and ground points of a terrain with several watercourse.



(b) Plan view of exterior skeleton sheets. The medial sheets reflect the structure of the watercourses.

Figure 11: Water course detection with the 3D MAT. Each distinct medial sheet was assigned a random colour. The surface points are coloured by elevation (yellow = low; blue = high).

For example, Gold et al. (2001) generalise the isolines, i.e. to simplify their appearance for (small scale) map production. This is achieved by pruning spurious branches from the MAT and then reconstructing the isolines from the pruned MAT. The method also works for river networks.

Karimipour et al. (2013) describe a way to estimate catchment areas based solely on the MAT of an existing river network. A catchment area is a hydrological unit where precipitations that fall into this area, eventually end up in the same river. Ideally, these areas are computed based on the elevation of a DEM by simulating the natural flow of water. However, this can be a very time consuming process and the elevation data may not always be available, and with the MAT, a quick and reasonable approximation of catchment areas can be obtained. The catchment areas are polygons that are constructed by connecting several simplified MAT branches.

Finally, Dakowicz and Gold (2003) demonstrate how the MAT can be used to solve the wedding cake effect for TINs that was described in Lesson 08. Recall that this effect happens because sometimes flat triangles are created between the vertices of one isoline. The problem is solved by adding points with a plausible elevation along the MAT between the isolines prior to triangulation.

The idea of applying the 3D MAT to digital terrain modelling is still relatively new. As mentioned in the introduction of this lesson, one interesting application is object detection in a terrain. Broersen et al. (2017) show how the 3D MAT, computed with the shrinking-ball algorithm, can be used to detect watercourses in the Dutch landscape (see Figure 11). Here the bottom edges of the medial sheets are projected to the xy plane to obtain 2D lines that delineate the watercourses. Peters (2018) gives some examples of how the 3D MAT can be used to perform building detection in lidar point clouds.

A second application of the 3D MAT is visibility analysis in point clouds (Peters et al., 2015). Here the difficulty is that point do not have an area and can thus not be used to obstruct visibility rays. However, the medial balls of the interior MAT do represent the interior volume of objects and can thus be used to obstruct the visibility rays.

4 Notes and comments

The Medial Axis Transform was originally introduced in 1967 by Harry Blum, a biologist (Blum, 1967).

The algorithm to approximate the 2D MAT from the VD is described in detail by Gold (1999) and Gold and Snoeyink (2001).

Ma et al. (2012) introduced the shrinking ball algorithm. Peters (2018) explains how to make the algorithm robust so that it can be successfully applied to lidar point cloud inputs and how to obtain a sheet segmentation using a region-growing approach.

References & further reading

Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications, SMA '01*, pages 249–266, New York, NY, USA, 2001. ACM.

Dominique Attali and Annick Montanvert. Modeling noise for a better simplification of skeletons. In *Image Processing, 1996. Proceedings., International Conference on*, volume 3, pages 13–16. IEEE, 1996.

Harry Blum. A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form*, 19(5):362–380, 1967.

Harry Blum. Biological shape and visual science (part i). *Journal of Theoretical Biology*, 38(2): 205 – 287, 1973.

Tom Broersen, Ravi Peters, and Hugo Ledoux. Automatic identification of watercourses in flat and engineered landscapes by computing the skeleton of a lidar point cloud. *Computers & Geosciences*, 106:171–180, September 2017.

Maciej Dakowicz and Christopher Gold. Extracting meaningful slopes from terrain contours. *International Journal of Computational Geometry & Applications*, 13(04):339–357, 2003.

C. Gold and J. Snoeyink. A one-step crust and skeleton extraction algorithm. *Algorithmica*, 30(2):144–163, 2001.

Christopher Gold. Crust and anti-crust: a one-step boundary and skeleton extraction algorithm. In *Proceedings of the fifteenth annual symposium on Computational geometry, SCG '99*, pages 189–196, New York, NY, USA, 1999. ACM.

C.M. Gold, D. Thibault, and Z. Liu. Map generalization by skeleton retraction. In *ICA Workshop on Map Generalization, Ottawa (August 1999)*, 2001.

Farid Karimipour, Mehran Ghandehari, and Hugo Ledoux. Medial axis approximation of river networks for catchment area delineation. In *Developments in Multidimensional Spatial Data Models*, pages 1–15. Springer, 2013.

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.

Ravi Peters. *Geographical point cloud modelling with the 3D medial axis transform*. PhD thesis, Delft University of Technology, 2018.

Ravi Peters, Hugo Ledoux, and Filip Biljecki. Visibility analysis in a point cloud based on the medial axis transform. In Filip Biljecki and Vincent Tourre, editors, *Eurographics Workshop on Urban Data Modelling and Visualisation*. The Eurographics Association, 2015.