

The Medial Axis Transform

Lesson 4.2*

Contents

1	Defining the MAT	1
1.1	Medial geometry	3
1.2	Exterior MAT and the MAT of a DTM	4
1.3	Medial clusters	4
2	Computing the MAT	6
2.1	MAT approximation	6
2.1.1	The shrinking-ball algorithm	6
2.2	Pruning	7
3	Notes and comments	7
4	Exercises	9

The Medial Axis Transform (MAT) is yet another way to represent a 3D model. It can be considered a *dual* representation to the b-rep, similar to how the Voronoi diagram is dual to the Delaunay triangulation. Contrary to the b-rep, that represents a model by describing explicitly its boundary surface, the MAT describes a model by its *skeleton* (compare Figures 1b and 1c). 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.

Compared to other shape representations, this skeleton structure makes different properties of the model explicit. 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 for example, 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 4b).

1 Defining the MAT

The MAT can be computed both for 2D and 3D objects (compare Figures 2a and 2b). In both cases

*©️🇩🇪 Ravi Peters. This work is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>) (last update: February 5, 2021)






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



(a) The MAT for a 2D box consists of medial balls (blue) and the medial axis (red). (b) 3D medial balls of a 3D box shape. (c) 3D medial axis of a 3D box shape.

Figure 2: The MAT in 2D and in 3D for a box shape.



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

Figure 3: The geometry of a medial atom.

there are two equivalent definitions of the MAT¹ that apply. 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 objects .

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 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 2c), 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 3 illustrates the complete medial geometry of an atom. The medial ball B has the atom c at its center and has a radius r , ie the shortest distance from c to \mathcal{B} . The medial ball B touches the boundary \mathcal{B} at the feature points p and q . The vectors from c to p and q are called the *spoke vectors*, denoted \vec{s}_p and \vec{s}_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{b} .

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

1. Any atom c is always *medial* to \mathcal{B} , ie it is equidistant to the feature points of c (hence the name of the MAT).

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 (?).

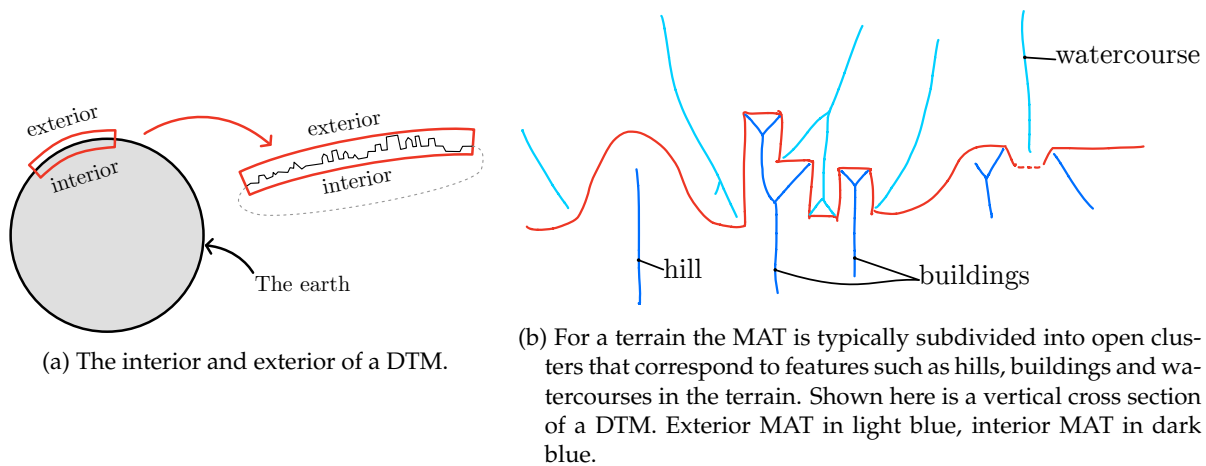


Figure 4: Defining interior and exterior for an open surface such as a terrain.

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 or for any perfectly manifold boundary. However, for objects that are not completely closed this separation is less clear, as there could be MAT sheets that connect the interior and exterior parts though holes in the boundary surface. In some cases with an open boundary a reasonable distinction can still be made. For example for a DTM we can follow the convention that the 'ground side' of the DTM is the interior, and the 'sky side' is the exterior, as follows from Figure 4a. Following this convention, we can still define the interior and exterior MAT of a DTM, see for instance Figure 4b.

1.3 Medial clusters

The interior and exterior MAT can consist of multiple disjoint parts (eg in Figure 4b). For closed objects the interior MAT is always one part, whereas the exterior MAT can be multiple parts depending on the number of concavities in the object boundary. The disjoint parts are called *medial clusters*. Each medial cluster is in fact a set of adjacent sheet where each adjacency is also a junction between medial sheets.

For objects with open boundaries like DTMs, there can also be multiple interior medial clusters. In this case one object on the terrain typically corresponds to one medial cluster (Figure 4b). Figure 5 also illustrates how the MAT can thus be used to meaningfully subdivide an object into parts. For an input that is simply a surface point cloud that happens to contain several object, we can detect easily these objects by looking at the medial clusters of its MAT. This effectively decomposes the object into meaningful sub-objects.




Figure 5: Interior MAT part in.

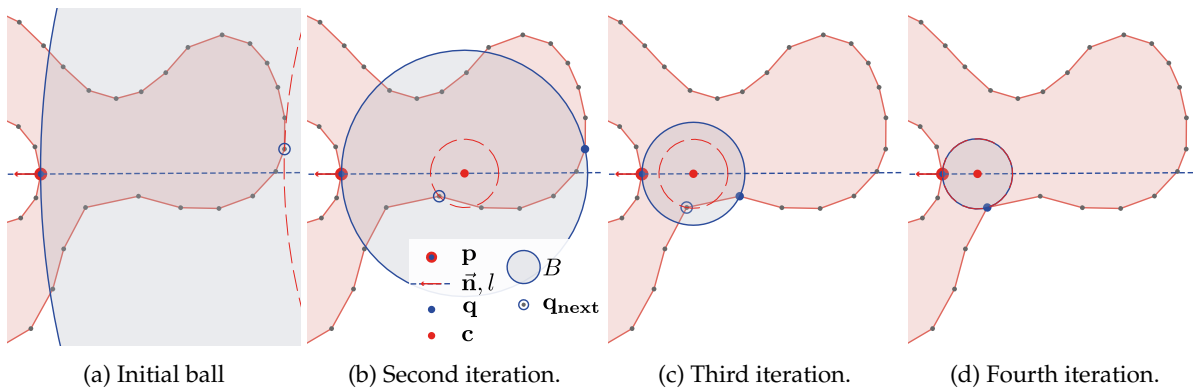


Figure 6: Ball shrinking iterations with the shrinking-ball algorithm. The final iteration yields a *medial* ball. A legend is given in b.

2 Computing the MAT

Computing the MAT from the boundary 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

The MAT can be approximated in various ways, eg by using voxels and distance transforms or as a subset of the Voronoi diagram. In this lesson, however, we will focus on the so-called shrinking-ball algorithm.

2.1.1 The shrinking-ball algorithm

The shrinking-ball algorithm works well for robustly approximating the MAT of 3D objects that are represented using boundary points, ie point clouds. It is a simple and fast algorithm that can be made robust to noise in the boundary points. 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 p must be positioned somewhere on the line L through the normal \tilde{n} of p . This observation is used to restrict the search space for the medial ball of p to the line L . As illustrated by Figure 6, the algorithm begins with a very large candidate ball for 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 p and is centred on L . Only the candidate feature point q changes with each iteration. A new q , denoted q_{next} , is found by selecting the closest point from the centre c of the current ball. Using $p, \tilde{n}, q_{\text{next}}$ we can compute the centre of the next ball 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 7 gives an example result of the shrinking-ball algorithm for a terrain point cloud.

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  $q_{next}$  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
```

Observe how the MAT describes the valleys (exterior MAT) and ridges (interior MAT) in the terrain with its skeletal structure.

Notice that the shrinking ball algorithm subdivides the output only in an interior and exterior part based on the point normals in the input. It is possible to further segment the MAT into for example medial sheets and medial clusters. This can be achieved for example with a region-growing segmentation algorithm that uses properties of the medial geometry (eg the medial bisector). Figures 5c and 5d show the result of such a segmentation.

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 8 and 9a, 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 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 , iethe distance between the two feature points of a medial atom. These values are typically low for noisy parts of the MAT. Figure 9 gives an example of pruning with the separation distance.

3 Notes and comments

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




Figure 7: MAT approximation (a,c) of a lidar terrain point cloud (b) obtained with the shrinking-ball algorithm. Top view.




Figure 8: Instability of the MAT.

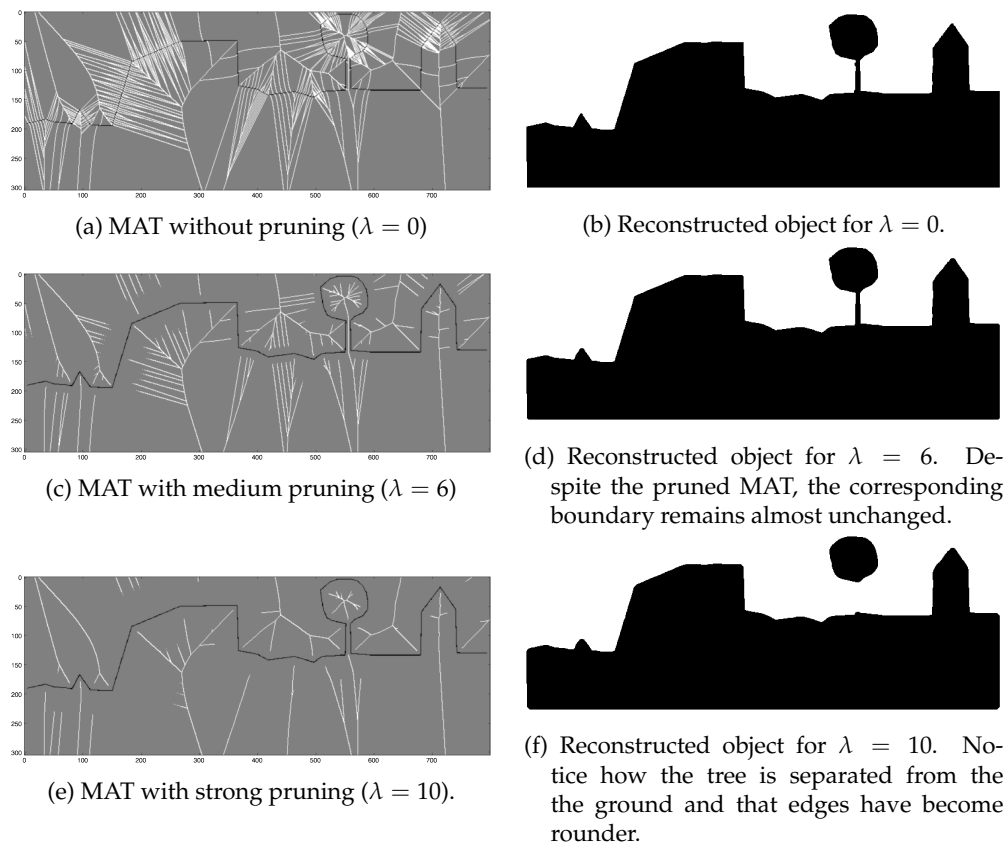


Figure 9: The effect of different levels of pruning on the MAT.

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 and cluster segmentation using a region-growing approach.

4 Exercises

1. Draw the medial axis of a 2D box. Then draw the medial bisectors. How could the medial bisector be used to distinguish between the different sheets?
2. For what closed object the MAT is a single point?
3. Do think the MAT could help us to detect thick and thin parts of an object? If so, how?

References & further reading

- Blum H (1967). A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form*, 19(5):362–380.
- Ma J, Bae SW, and Choi S (2012). 3D medial axis point approximation using nearest neighbors and the normal field. *The Visual Computer*, 28(1):7–19.
- Peters R (2018). *Geographical point cloud modelling with the 3D medial axis transform*. Ph.D. thesis, Delft University of Technology. ISBN: 978-94-6186-899-2.