

Robust skeletonization using the discrete lambda-medial axis

John Chaussard^a, Michel Couprie^a, Hugues Talbot^a

^a*Université Paris-Est, LIGM, Équipe A3SI, ESIEE Paris, France*

Abstract

Medial axes and skeletons are notoriously sensitive to contour irregularities. This lack of stability is a serious problem for applications in *e.g.* shape analysis and recognition. In 2005, Chazal and Lieutier introduced the λ -medial axis as a new concept for computing the medial axis of a shape subject to single parameter filtering. The λ -medial axis is stable under small shape perturbations, as proved by these authors. In this article, a discrete λ -medial axis (DLMA) is introduced and compared with the recently introduced integer medial axis (GIMA). We show that DLMA provides measurably better results than GIMA, with regard to stability and sensibility to rotations. We give efficient algorithms to compute the DLMA, and we also introduce a variant of the DLMA which may be computed in linear-time.

Introduction

The notion of medial axis has been introduced by Blum in the 60s [7, 8]. It has proved its usefulness in many practical applications, and numerous works were devoted to its properties and implementations. The original definition of the medial axis by Blum was based on a fire propagation analogy. However, its simplest definitions only require elementary geometry. In the continuous Euclidean space, the two following definitions can be used to formalise this notion: let X be a bounded subset of \mathbb{R}^n ;

- Interpretation (a) of the medial axis of X consists of the centers of the n -dimensional balls that are included in X but that are not included in any other n -dimensional ball included in X .
- Interpretation (b) of the medial axis of X consists of the points $x \in X$ that have more than one nearest points on the boundary of X .

Email addresses: j.chaussard@esiee.fr (John Chaussard), m.couprie@esiee.fr (Michel Couprie), h.talbot@esiee.fr (Hugues Talbot)

This work has been partially supported by the “ANR BLAN07-2_184378 MicroFiss” project.

These two definitions differ only by a negligible set of points (see [28]), in general interpretation (a) of the medial axis is a strict subset of interpretation (b). Notice that in some works, the term “skeleton” is used to refer to both interpretations, especially in the continuous framework. In this paper, we shall restrict the use of this term to the cases where the skeleton is topologically equivalent to the original shape.

To compute the medial axis approximately or exactly, different methods have been proposed, relying on different frameworks: discrete geometry [9, 21, 27, 32, 22], digital topology [18, 40, 39, 31], mathematical morphology [34, 37], computational geometry [2, 30, 3], partial differential equations [36], and level-sets [24]. In this paper, we focus on medial axes in the discrete grid \mathbb{Z}^2 or \mathbb{Z}^3 , which are centered in the shape with respect to the Euclidean distance.

A major difficulty when using the medial axis in applications (*e.g.*, shape recognition), is its sensitivity to small contour perturbations, in other words, its lack of stability. A recent survey [1] summarises selected relevant studies dealing with this topic. This difficulty can be expressed mathematically: the transformation which associates a shape to its medial axis is only semi-continuous. This fact, among others, explains why it is usually necessary to add a filtering step (or pruning step) to any method that aims at computing the medial axis.

Hence, there is a rich literature devoted to medial axis pruning, in which different criteria were proposed in order to discard “spurious” medial axis points or branches: see Attali et al. [4], Ogniewicz and Kübler [30], Attali and Montanvert [3], Malandain and Fernández-Vidal [27], Attali and Lachaud [2], Svensson and Sanniti di Baja [38], Bai et al. [5], Couprie et al. [17], Hesselink and Roerdink [22], to cite only a few. However, we lack theoretical justification, that is, a formalised argument that would help to understand why a filtering criterion is better than another.

In 2005, Chazal and Lieutier introduced the λ -medial axis and studied its properties, in particular those related to stability [12]. Consider a bounded subset X of \mathbb{R}^n , as for example, for $n = 2$, the region enclosed by the solid curve depicted in Fig. 1 (left). Let x be a point in X , we denote by $\Pi(x)$ the set of points of the boundary of X that are closest to x . For example in Fig. 1, we have $\Pi(x) = \{a, b\}$, $\Pi(x') = \{a', b'\}$ and $\Pi(x'') = \{a''\}$. Let λ be a non-negative real number, the λ -medial axis of X is the set of points x of X such that the smallest n -dimensional ball¹ including $\Pi(x)$ has a radius greater than or equal to λ . Notice that the 0-medial axis of X is equal to X , and that any λ -medial axis with $\lambda > 0$ is included in the medial axis according to definition b). We show in Fig. 1 (right) two λ -medial axes with different values of λ .

A major outcome of Chazal and Lieutier [12] is the following property: informally, for “regular” values of λ , the λ -medial axis remains stable under perturbations of \overline{X} that are small with respect to the Hausdorff distance. Typical non-regular values are radii of locally largest maximal n -dimensional balls.

This property is a strong argument in favor of the λ -medial axis, especially

¹The center of this ball is also the projection of x onto the convex hull of $\Pi(x)$ (see [11]).

in the absence of such result for other proposed criteria.

In the field of computational geometry, the λ -medial axis has been exploited in particular by Samozino et al. [33] to propose a robust method for reconstructing surfaces from point clouds. Also, notions closely related to the λ -medial axis have led Chazal et al. to propose stable approximations of tangent planes and normal cones from noisy samples [11].




Figure 1: Illustration of the λ -medial axis. Left: Points x, x' and x'' and their respective closest boundary points. Top right: λ -medial axis with $\lambda = \epsilon$, a very small positive real number. Bottom right: λ -medial axis with $\lambda = d(a', b')/2 + \epsilon$.

In the discrete grids, namely \mathbb{Z}^2 and \mathbb{Z}^3 , a similar filtering criterion has been considered in independent works [27, 22]. It consists of selecting for each medial axis point two of its closest boundary points, and using the distance between these two points as filtering criterion. The work of Hesselink and Roerdink [22] provides a linear-time algorithm to compute a filtered medial axis based on this criterion, which exhibits good noise robustness properties in practice.

In this article (which extends Chaussard et al. [10], a preliminary version published in the DGCI conference proceedings), we introduce the definition of a discrete λ -medial axis (DLMA) in \mathbb{Z}^n . We evaluate experimentally its stability and rotation invariance in 2D and 3D. In this experimental study, we compare it with the previously introduced integer medial axis (GIMA) [23, 22] and show that the DLMA provides measurably better results. Furthermore, we introduce a variant of the DLMA which may be computed in linear time, for which the results are very close to those of the DLMA, and which is only slightly slower than the one proposed by Hesselink and Roerdink [22].

1. The λ -medial axis

Let us first recall the original definition of the λ -medial axis given by Chazal and Lieutier.

Let $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \mathbb{R}^n$, the Euclidean distance between x and y is denoted by $d(x, y)$, in other terms, $d(x, y) = (\sum_{k=1}^n (y_k - x_k)^2)^{\frac{1}{2}}$. Let X be a finite subset of \mathbb{R}^n or \mathbb{Z}^n . We set $d(y, X) = \min_{x \in X} \{d(y, x)\}$. We denote by $|X|$ the number of elements of X , and by \overline{X} the complement of X .

Let E be either \mathbb{R}^n or \mathbb{Z}^n . Let $x \in E, r \in \mathbb{R}^+$, we denote by $B_r(x)$ the n -dimensional ball (or ball for simplicity) of radius r centered on x , defined by $B_r(x) = \{y \in E \mid d(x, y) \leq r\}$. We also define $\mathring{B}_r(x) = \{y \in E \mid d(x, y) < r\}$.

Let $X \subseteq E$. A ball $B_r(x) \subseteq X$ is *maximal for X* if it is not strictly included in any other ball included in X . The *Euclidean medial axis of X* , denoted by $\text{EMA}(X)$, is the set of the centers of all the maximal balls for X .

Let S be a non-empty subset of E , and let $x \in E$. The *projection of x on S* , denoted by $\Pi_S(x)$, is the set of points y of S which are at minimal distance from x ; more precisely,

$$\Pi_S(x) = \{y \in S \mid \forall z \in S, d(y, x) \leq d(z, x)\}.$$

If X is a subset of E , the *projection of X on S* is defined by $\Pi_S(X) = \bigcup_{x \in X} \Pi_S(x)$.

Let $S \subset \mathbb{R}^n$, we denote by $R(S)$ the radius of the smallest ball enclosing S , that is, $R(S) = \min\{r \in \mathbb{R} \mid \exists y \in \mathbb{R}^n, B_r(y) \supseteq S\}$.

The λ -medial axis may now be defined based on these notions.

Definition 1 (Chazal and Lieutier [12]). *Let X be an open bounded subset of \mathbb{R}^n , and let $\lambda \in \mathbb{R}^+$. The λ -medial axis of X is the set of points x in X such that $R(\Pi_{\overline{X}}(x)) \geq \lambda$.*

2. A discrete λ -medial axis

Transposing the original definition of the λ -medial axis directly to \mathbb{Z}^n would lead to an unsatisfactory result. For instance, consider a horizontal ribbon in \mathbb{Z}^2 with constant, even width and infinite length. Clearly, the projection of any point of this set on its complementary set is reduced to a singleton. This is why, if we keep the same definition, any λ -medial axis of this object with $\lambda > 0$ would be empty.

In order to avoid such unwanted behaviour, we replace the projection by the so-called extended projection [17]. The extended projection was originally introduced in order to propose a discrete definition of the bisector function, another indicator used to filter skeletons.

For each point $x \in \mathbb{Z}^n$, we define the *direct neighborhood of x* as $N(x) = \{y \in \mathbb{Z}^n \mid d(x, y) \leq 1\}$. The direct neighborhood comprises $2n + 1$ points.

The set $\Pi_{\overline{X}}(X)$ is composed of the points of \overline{X} that are in the direct neighborhood of a point of X , that is, $\Pi_{\overline{X}}(X) = \{y \in \overline{X} \mid N(y) \cap X \neq \emptyset\}$. We call $\Pi_{\overline{X}}(X)$ the (*external*) *boundary of X* .

Let $X \subseteq \mathbb{Z}^n$, and let $x \in X$. The *extended projection of x on \overline{X}* , denoted by $\Pi_{\overline{X}}^e(x)$, is the union of the sets $\Pi_{\overline{X}}(y)$, for all y in $N(x)$ such that

$$d(y, \overline{X}) \leq d(x, \overline{X}). \quad (\text{condition 1})$$

Using the extended projection allows us to cope with the problem pointed out previously. Considering the example of an horizontal ribbon in \mathbb{Z}^2 with constant even width and infinite length, points close to the “middle” of the ribbon will each have an extended projection consisting of a pair of points. Informally, the

extended projection permits to “see” both sides of the object from a point x close to the “middle” of the object, whereas only one side is taken into account when considering the projection.

As illustrated in Fig. 2, condition 1 avoids producing multiple medial axis points when only one is sufficient, in other words, it yields a thinner axis according to definition 2.

Fig. 2 helps us to illustrate this definition and the following one. Consider the object X (white pixels) depicted in Fig. 2a, and the two pixels x, y in X . Fig. 2b shows the squared Euclidean distance map of X , that is, the value in each pixel is the square Euclidean distance between this pixel and the nearest background pixel. We can see that $d(x, \overline{X}) = 4$ and $d(y, \overline{X}) = 5$. The projection of x (resp. y) on \overline{X} is $\{b\}$ (resp. $\{d\}$). We have $\Pi_{\overline{X}}^e(x) = \{a, b, c\}$, and $\Pi_{\overline{X}}^e(y) = \{d, b, c\}$.

We are now ready to introduce the definition of our discrete λ -medial axis.

Definition 2. *Let X be a finite subset of \mathbb{Z}^n , and let $\lambda \in \mathbb{R}^+$. We define the function \mathcal{F}_X which associates, to each point x of X , the value $\mathcal{F}_X(x) = R(\Pi_{\overline{X}}^e(x))$. The discrete λ -medial axis (or DLMA) of X is the set of points x in X such that $\mathcal{F}_X(x) \geq \lambda$.*

Back to our illustration in Fig. 2, the function \mathcal{F}_X is displayed in Fig. 2c and the 4-medial axis of X (threshold of \mathcal{F}_X for $\lambda = 4$) is displayed in Fig. 2d. Notice that the function \mathcal{F}_X can be computed once and stored in an n -dimensional array (a grayscale image), and that any DLMA of X is obtained by thresholding the values of this array at a particular value λ .

In Fig. 2e, we show how the function \mathcal{F}_X would look like if we omit condition 1 from the definition of the extended projection. In this case, the extended projection of x would be $\{a, b, c, d\}$ instead of $\{a, b, c\}$, this explains the higher value of $\mathcal{F}_X(x)$. Fig. 2f shows how the 4-medial axis of X would be.

In Fig. 3(b,c), we show two examples of DLMA of a bigger shape. More examples are given in Fig. 4 and Fig. 5. Notice that DLMA in general does not exhibit the same topological characteristics as the original shape.

3. The integer medial axis

The integer medial axis was introduced by [22]. Conceptually, this notion is close to the DLMA, and also defined in the framework of discrete grid \mathbb{Z}^n .

Let $X \subset \mathbb{Z}^n$, and $x \in X$. We define $\Pi_{\overline{X}}^l(x)$ as the element of $\Pi_{\overline{X}}^e(x)$ that is the smallest with respect to the lexicographic ordering of its coordinates.

Definition 3 (Hesselink and Roerdink [22]). *Let X be a finite subset of \mathbb{Z}^n , and let $\gamma \in \mathbb{R}^+$. The γ -integer medial axis (or GIMA) of X is the set of points x in X such that at least one $y \in N(x)$ verifies :*

i) $d(\Pi_{\overline{X}}^l(x), \Pi_{\overline{X}}^l(y)) > \gamma$, and

ii) $d(m - \Pi_{\overline{X}}^l(y)) \leq d(m - \Pi_{\overline{X}}^l(x))$, where $m = \frac{x+y}{2}$.

The integer medial axis of X is the γ -integer medial axis of X for $\gamma = 1$.




Figure 2: (a): An object X (white pixels). We have $\Pi_X^c(x) = \{a, b, c\}$, and $\Pi_X^c(y) = \{d, b, c\}$. (b): The squared Euclidean distance map of X . (c): The function \mathcal{F}_X (truncated to the first decimal). (d): The λ -medial axis of X , for $\lambda = 4$ (black disks). (e, f): Same as (c, d) assuming that condition 1 is omitted from the definition of the extended projection. (g): The map, for each $x \in X$, of the minimal value λ such that $x \in \lambda'$ -medial axis of X . (h): The λ' -medial axis of X , for $\lambda' = 4$ (black disks). (i, j): Same as (g, h) for the integer medial axis.




Figure 3: (a): The function \mathcal{F}_X superimposed to the shape X . Darkest colors represent highest values of $\mathcal{F}_X(x)$. Any DLMA of X is a threshold of this function at a particular value λ . (b): Discrete 5-medial axis. (c): discrete 10-medial axis of X .

Notice that the GIMA is not invariant under coordinate permutation.

There are indeed some links between the GIMA and the DLMA. In 2D, in the case of a point x that is, conceptually, a “regular medial axis point” (neither a branch extremity nor a branch junction), condition i) is similar to the condition $R(\Pi_{\overline{X}}(x)) \geq \lambda$ in Def. 1.

Condition ii) plays a role analogous to condition $d(y, \overline{X}) \leq d(x, \overline{X})$ in the definition of the extended projection, that is to get a thinner axis.

However, these definitions differ, leading to sensible differences in the results of these two transformations (see Fig. 5). For example, in definition 3, $\Pi_{\overline{X}}^l(x)$ is the only element of the projection $\Pi_{\overline{X}}(x)$ that is taken into account into the computation of the integer medial axis of X , while all the elements of $\Pi_{\overline{X}}^e(x)$ are taken into account into the computation of the lambda medial axis. Moreover, in definition 3, the result depends on the longest distance between two elements of the (reduced) projection, while in definition 2, the result depends on the size of the smallest ball including the whole projection.

Generally speaking, one can say that the lambda medial axis “takes into account” more elements than the integer medial axis. We analyse quantitatively the consequences of these differences in Sec. 7 and Sec. 8. These differences lead to different axes, as illustrated on Fig. 2.

4. Topology preservation issues

The illustration on Fig. 5 (right) is sufficient to demonstrate that a DLMA or a GIMA of a given shape X may have a homotopy type different from the one of X .

Nearly identical methods for filtering a medial axis while preserving homotopy were proposed by several authors [18, 40, 39, 31]. They consist of iteratively removing simple points [25, 6, 16] from the shape X , that do not belong to the set M of its medial axis points (called constraint set). A priority function is needed in order to specify which points must be considered at each step of the




Figure 4: Top-right: The function \mathcal{F}_X superimposed to the top-left shape X . Red color represents highest values of $\mathcal{F}_X(x)$, and blue color represents lowest ones. Bottom left: discrete 11-medial axis. Bottom right: discrete 15-medial axis of X .

thinning. In the general case, the choice of this priority function is not obvious [39, 17].

An example is shown on Fig. 6, where the filtered DLMA of an object is used as a constraint set during skeletonization. Choosing the exact Euclidean distance map as a priority function for removing simple points from the object may lead to geometric distortions [39]. In some cases, "extra branches" may even appear (see Fig. 6b). Choosing the map \mathcal{F}_X as priority function yields more satisfying results (see Fig. 6c), as it guides the thinning process towards elements that belong to the different nested discrete λ -medial axes.

5. Implementation issues

In this section, we discuss the algorithms that should be used in order to efficiently implement DLMA, and we indicate their complexity. We also introduce




Figure 5: Results of DLMA (first row) and GIMA (second row) for parameter values yielding similar reconstruction ratios (see Sec. 6).




Figure 6: (a): The original object X (in white). (b): Result of the skeletonization using the distance map as a priority function, and a filtered DLMA as constraint set. (c): Result of the skeletonization using \mathcal{F}_X as a priority function, and a filtered DLMA as constraint set (same set than in b).

a variant of the DLMA that can be computed in linear time.

The set of all projections $\Pi_{\overline{X}}(x)$, for all $x \in X$, can be computed in optimal time and space (that is, in $O(n)$ where $n = \sum_{x \in X} |\Pi_{\overline{X}}(x)|$), thanks to an algorithm introduced by Coeurjolly in [17]. After this, to compute the extended projection of a particular point x , we just have to perform the union of the projections of all points y in $N(x)$ that satisfy condition 1.

In order to avoid computing the map $\Pi_{\overline{X}}$, which requires a data structure like an array of lists to store all the sets $\Pi_{\overline{X}}(x)$, we propose a variant of the DLMA that consists of replacing $\Pi_{\overline{X}}$ by $\Pi'_{\overline{X}}$ in the definition of the extended projection. We refer to this variant as the discrete linear λ -medial axis (DLLMA): it is illustrated on Fig. 2g and h. Notice that the DLLMA, as the GIMA, is not invariant under a permutation of the coordinates (on the contrary of the DLMA). One advantage of this variant is that the map $\Pi'_{\overline{X}}$ can be computed in linear time and stored in an array of integers [13, 29, 22]. We included DLLMA in our experimental comparative evaluation (see Sec. 7 and Sec. 8).




Figure 7: Results of homotopic thinning of X constrained by the DLMA for parameter values 5, 10, 20 and 30, and with \mathcal{F}_X as priority function.

The smallest d -dimensional ball enclosing a given set of n points in \mathbb{R}^d can be computed in expected $O(n)$ time if d is considered as a constant, thanks to an algorithm proposed by Welzl [41]. This algorithm is simple and easy to implement.

For each object point x , the computational cost is in $O(n)$ where $n = |\Pi_X^\varepsilon(x)|$ (see Def. 2). In the variant DLLMA, we have $n \leq 5$, since only one projection point is considered for x and each of its four direct neighbors, at worse. Hence, computing the radius of the smallest enclosing d -dimensional ball can be performed in constant time.

Overall, the DLLMA procedure runs in linear time with regard to the number of object points.

Using local characterisations, in 2D and 3D, checking whether a point is simple or not can be done in constant time. The homotopic thinning described in section 4 may be implemented to run in $O(n \log n)$ complexity, where n is the number of object points, using *e.g.* a balanced binary tree to store and retrieve the candidate points and their priority.

6. Methodology for the comparison of medial axes

In order to compare the robustness of DLMA, DLLMA and GIMA, we measure the dissimilarity between the medial axis of a shape X and the one of a shape X' derived from X by a small perturbation. In Sec. 7 this perturbation is obtained through a random process altering the contour of X , and in Sec. 8 it is obtained via a discrete rotation. We also include in this comparison the filtered *Euclidean Medial Axis (EMA)* defined as the set of the centers of maximal balls that have a radius greater than a parameter ρ .

We first introduce some notions that will help us to ensure a fair comparison between methods, for which the filtering parameter does not have the same meaning.

Let X be a finite subset of \mathbb{Z}^n . Let Y be a subset of X , we set $REDT_X(Y) = \bigcup_{y \in Y} \mathring{B}_{d(y, \bar{X})}(y)$. The transformation $REDT_X$ is sometimes called *reverse Euclidean distance transform* [14, 15].

It is well known that any object can be fully reconstructed from its medial axis, more precisely, we have $X = REDT_X(M)$ whenever M is the (exact and

non-filtered) medial axis of X . This property holds if M is the set of the centers of all maximal balls of X (EMA of X with $\rho = 0$), the DLMA of X with $\lambda = 1$, or the integer medial axis of X with $\gamma = 1$. However, it is no longer true if we consider filtered medial axes, *e.g.* EMA, DLMA or GIMA with arbitrary ρ , λ or γ .

Then, it is interesting to measure how much information about the original object is lost when we raise the filtering parameter. We set

$$\mathcal{L}_X(\lambda) = \frac{|X \setminus REDT_X(LM_X(\lambda))|}{|X|}$$

where $LM_X(\lambda)$ is the DLMA of X . In other words, $\mathcal{L}_X(\lambda)$ is the area of the difference between X and the set reconstructed from its DLMA, divided by the area of X . We call $\mathcal{L}_X(\lambda)$ the (*normalised*) *residuals* of the DLMA of X filtered at value λ . The same way, we define $\mathcal{L}'_X(\lambda')$ (resp. $\mathcal{I}_X(\gamma)$, $\mathcal{M}_X(\rho)$) as the (normalised) residuals of the DLLMA (resp. GIMA, EMA) of X filtered at value λ' (resp. γ , ρ).

Fig. 8 shows the evolution of the residuals for the EMA, GIMA, DLMA and DLLMA for a given shape. Since differences are not negligible, to ensure a fair evaluation we compare the results of methods for approximately equal values of their residuals, rather than for equal values of their parameters.




Figure 8: Residuals $\mathcal{L}_X(\lambda)$ (2 variants) and $\mathcal{I}_X(\gamma)$, for the set X depicted in Fig. 3. Horizontal axis: the value of the parameter (λ or γ). Vertical axis: the value of the residual. Notice that curves corresponding to DLMA and DLLMA are superimposed.

For comparing shapes or medial axes, we use the Hausdorff distance (see below), and also a dissimilarity measure proposed by Dubuisson and Jain [19]. The drawback of Hausdorff distance for measuring shape dissimilarity is its extreme sensibility to outliers, the latter measure performing better in this respect.

Let X, Y be two subsets of \mathbb{R}^n . We set

$$H(X|Y) = \max_{x \in X} \{ \min_{y \in Y} \{ d(x, y) \} \},$$

and $d_H(X, Y) = \max\{H(X|Y), H(Y|X)\}$ is the *Hausdorff distance between X and Y*. We set

$$D(X|Y) = \frac{1}{|X|} \sum_{x \in X} \min_{y \in Y} \{d(x, y)\},$$

and $d_D(X, Y) = \max\{D(X|Y), D(Y|X)\}$ is the *Dubuisson and Jain's dissimilarity measure between X and Y* (called *dissimilarity* in the sequel for the sake of brevity).

7. Stability

In this section, we compare the stability of DLMA, DLLMA, GIMA and EMA with respect to contour noise.

We based our 2D experiments on Kimia's database of 216 shapes [35]. For the 3D case, we used 20 three-dimensional objects gathered on the Internet or that we created ourselves. Fig. 9 shows a sample of these shapes .




Figure 9: A sample of the 216 shapes of Kimia's database, and of the 3d shapes from our database.

Medial axes are notoriously sensitive to border noise. Since the λ -medial axis is supposed to cope reasonably well with shape deformation, it is useful to test how it fares in practice.

To introduce noise to the boundary of an object we propose deforming it using a process derived from the Eden's (accretion) process [20]. The Eden's process is an iterative random cellular automaton which, in its simplest form, attributes an equal probability to all the outer border points to be set to 1 at each step. That is, at each step, a neighbour of the object is chosen randomly and added to the shape. In spite of its simplicity, the Eden's process exhibits good asymptotic isotropy [26]. In our process, we required accretion steps to deal only with simple points. This way at each step, the object's homotopy type remains unchanged.

We denote by $E(X, n)$ the result of applying n steps of this process to the shape X . In this experiment, we compare the (filtered) medial axis of an original shape with the one of a deformed shape.




Figure 10: Illustration of the border deformation process, after 0 (original image), 50, 100, 200 and 400 steps. The function \mathcal{F}_X is superimposed to each shape X .

	2D				3D			
Noise	3%	6%	9%	all	3%	6%	9%	all
EMA	6.97	8.78	9.88	7.58	7.07	7.42	7.72	6.92
GIMA	6.30	8.05	8.97	6.77	7.75	8.46	8.77	7.56
DLMA	5.61	7.57	8.32	6.23	5.17	6.01	6.78	5.36
DLLMA	5.89	8.05	9.00	6.66	4.92	5.97	7.72	5.61

Table 1: Average Hausdorff distance between $M(X)$ and $M(E(X, n))$. Noise level (parameter n) is expressed as a percentage of object area, results shown are averages of all tested residuals. Lowest values are highlighted.

	2D				3D			
Noise	3%	6%	9%	all	3%	6%	9%	all
EMA	1.01	1.42	1.75	1.21	0.39	0.52	0.60	0.44
GIMA	1.32	1.83	2.20	1.53	0.80	0.99	1.11	0.85
DLMA	1.02	1.51	1.75	1.23	0.28	0.37	0.43	0.31
DLLMA	1.12	1.66	1.94	1.36	0.28	0.36	0.44	0.31

Table 2: Average dissimilarity between $M(X)$ and $M(E(X, n))$. Noise level (parameter n) is expressed as a percentage of object area, results shown are averages of all tested residuals. Lowest values are highlighted.

In Table 1 (resp. Table 2) we give the average Hausdorff distance (resp. dissimilarity) between $M(X)$ and $M(E(X, n))$ on the 216 shapes of Kimia’s database and the 20 shapes of our 3d database for different definitions of M (filtered EMA, GIMA, DLMA or DLLMA of X). Filtering parameters ρ, γ, λ and λ' were chosen in order to give a residual varying from 5% to 30% in 2d, and a residual varying from 3% to 10% in 3d. Results are given for values of n equal to 3%, 6% and 9% of object’s surface/volume (values shown are averages of results obtained for all tested residuals), and the last column indicates all the results averaged together.

Note that the tested residual values are not the same in the 2d and 3d case. Indeed, in 3d, choosing a residual higher than 10% gives a medial axis which contains very little information about the original object (surfacic parts are mostly deleted). We constrained our study to medial axes holding enough information on the original objects, and ignored residual superior to 10% in the 3d case.

On Fig. 11, we show the evolution of the Hausdorff distance and the dissimilarity between $M(X)$ and $M(E(X, n))$ against the quantity of noise added to the shape, for DLMA, DLLMA, GIMA and EMA algorithms. The results shown are average of the results obtained on the 216 shapes of Kimia's 2D image database. Fig. 12 shows the same experimental results on the 20 shapes of our 3d database.

As we expected, DLMA fares better on the average than other axes, with regard to Hausdorff distance criterion. Its faster variant DLLMA ranks second. According to the dissimilarity criterion, DLMA and DLLMA yield the best results in 3d, and in 2d, DLMA performs nearly as well as EMA.

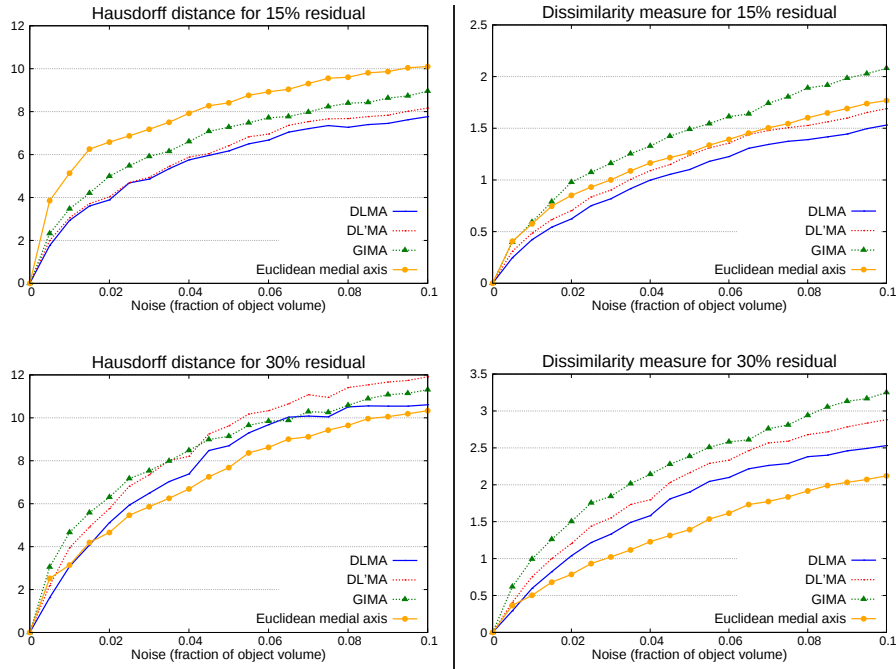


Figure 11: Hausdorff distance (**on the left**) and dissimilarity measure (**on the right**) between $M(X)$ and $M(E(X, n))$ on the 216 shapes of Kimia's 2D image database for residual values 15% and 30%. Noise level (parameter n) is expressed as a percentage of object area.

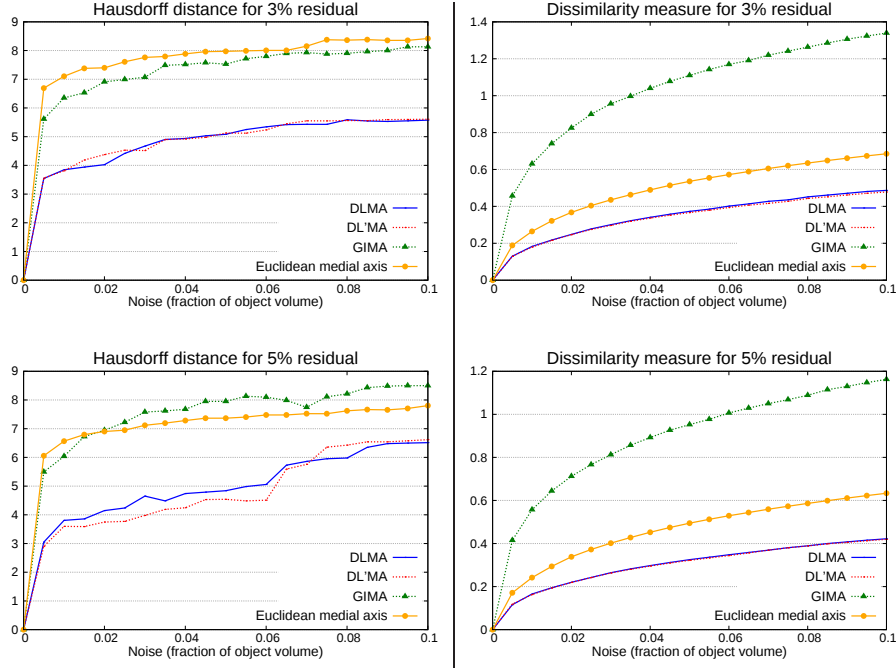


Figure 12: Hausdorff distance (**on the left**) and dissimilarity measure (**on the right**) between $M(X)$ and $M(E(X, n))$ on the 20 shapes of our 3D image database for residual values 3%, 5%. Noise level (parameter n) is expressed as a percentage of object area. When the DLMA curve is not visible, it is superimposed with DLLMA curve.

8. Sensitivity to rotations

In the continuous framework, it is well known that the medial axis is a rotation-invariant shape descriptor. More precisely, if we denote by R_θ the rotation of angle θ about the origin, and by M the medial axis transform, the rotation invariance property states that $M(R_\theta(X)) = R_\theta(M(X))$, whatever the shape X and the angle θ .

In a discrete framework, this property holds only in particular cases (*e.g.*, when θ is a multiple of 90 degrees). Nevertheless, we can experimentally measure the dissimilarity between $M(R_\theta(X))$ and $R_\theta(M(X))$ for different instances, and different definitions of the medial axis. The lower this dissimilarity, the more stable under rotation the method is.

For each shape of Kimia's database and of our database, we computed the Hausdorff distance and the dissimilarity between $M(R_\theta(X))$ and $R_\theta(M(X))$ for values of the parameters yielding 5%, 10%, 15%, 20%, 25% and 30% residuals, and for rotation angles varying from 0 to 90 degrees by 3 degrees steps. In 3d, rotations were successively performed around the Y and Z axis. Moreover, in 3d, only parameters yielding 3%, 5% and 10% residuals were tested: higher

residual values yielded skeletons that were not keeping enough information from the original object (see discussion in Sec. 7).

Fig. 13 shows detailed results of such experiment in 2d, while Fig. 14 and Fig. 15 show results in 3d. For both cases, an inverse truncated real rotation algorithm was used to perform images rotation.

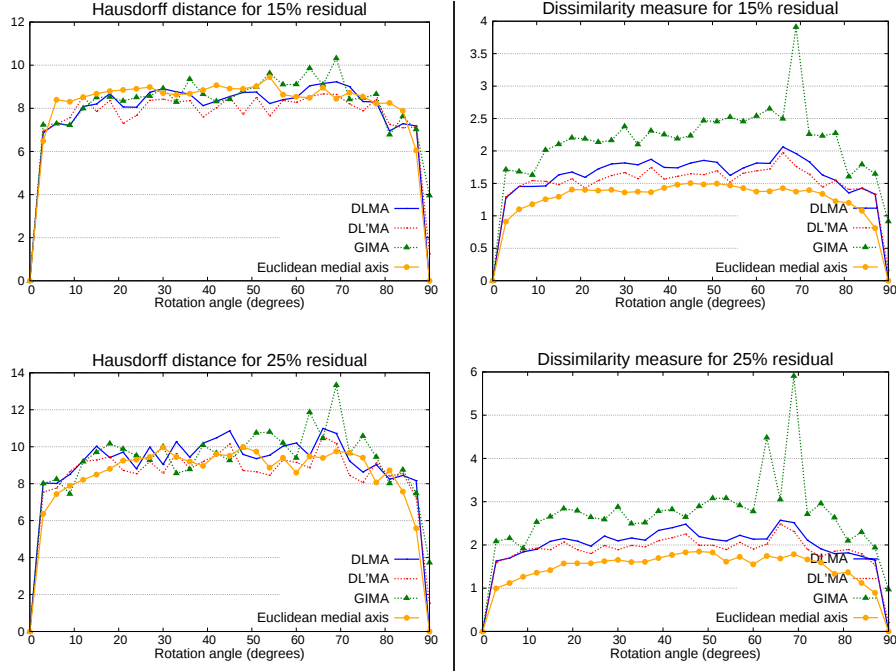


Figure 13: Hausdorff distance (**on the left**) and dissimilarity measure (**on the right**) between $M(R_\theta(X))$ and $R_\theta(M(X))$ on the 216 shapes of Kimia’s 2D image database for residual values 15% and 25%. Noise level (parameter n) is expressed as a percentage of object area.

The results are summarised in Tables 3 and 4. In Table 3 (resp. Table 4) we give the average Hausdorff distance (resp. dissimilarity) between $M(R_\theta(X))$ and $R_\theta(M(X))$ on the 216 shapes of Kimia’s database and on the 20 3d shapes of our database, for residual values 5%, 10%, 15%, 20%, 25% and 30% (for 3d shapes, only 3%, 5% and 10% were tested) and for θ varying from 0 to 89 degrees (in 3d shapes, rotation was performed independently around Y and Z axes). We show, in the first three columns, the average value obtained for all tested rotation angles for a given residual, while the fourth column indicates the average value obtained for all residuals and all rotation angles.

In 3d, DLMA and DLLMA yield, in average, the best results for both measures. In 2d, they are better than other axes with regards to Hausdorff distance, and are only outperformed by EMA with regards to dissimilarity.

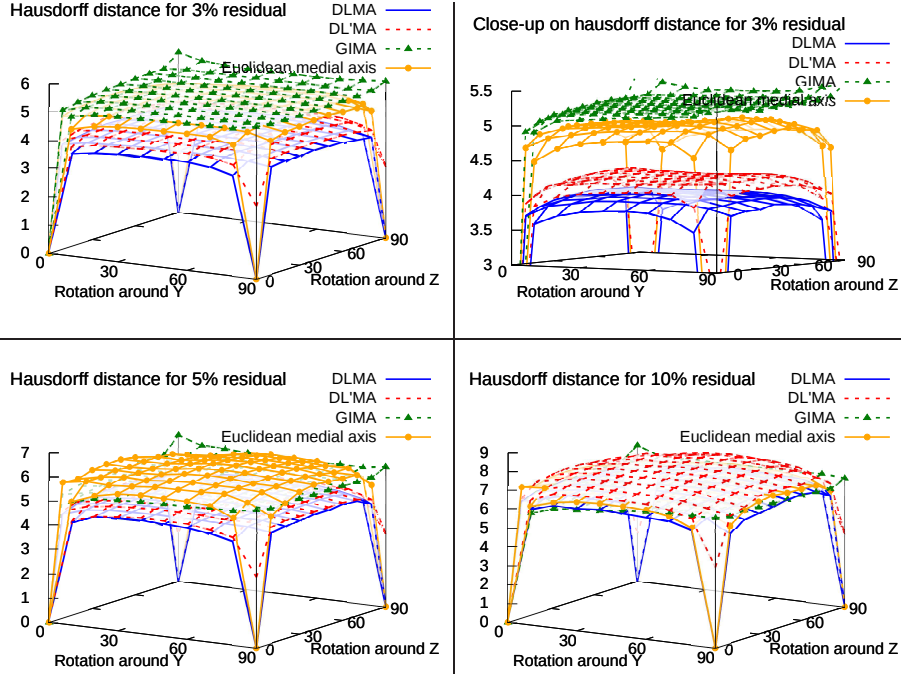


Figure 14: Hausdorff distance (**on the left**) and dissimilarity measure (**on the right**) between $M(R_\theta(X))$ and $R_\theta(M(X))$ on the 20 shapes of our 3D image database for residual values 3%, 5% and 10% (the horizontal axis is proportional to the parameter n).

9. Computational cost

In addition to complexity analysis reported in Sec. 5, we made some experiments to compare actual computation costs of the presented methods.

In Fig. 16, we show the results of computing time measurements that we performed on an Intel Core 2 Duo processor at 1.83 GHz. Computing times for the GIMA are the lowest, but DLLMA is only slightly slower (and also linear).

10. Conclusion

We introduced the discrete λ -medial axis (DLMA), and its variant DLLMA, a transposition in discrete grids of the continuous λ -medial axis recently proposed and known for its stability property. The experimental study reported in this paper shows that both DLMA and DLLMA are robust w.r.t. boundary noise and rotations in practice. For our experimental study of rotational invariance and border noise robustness, both in 2D and 3D, we introduced an original methodology that ensures a fair comparison between different methods, under the mere assumption that their result decreases in size under the control of a

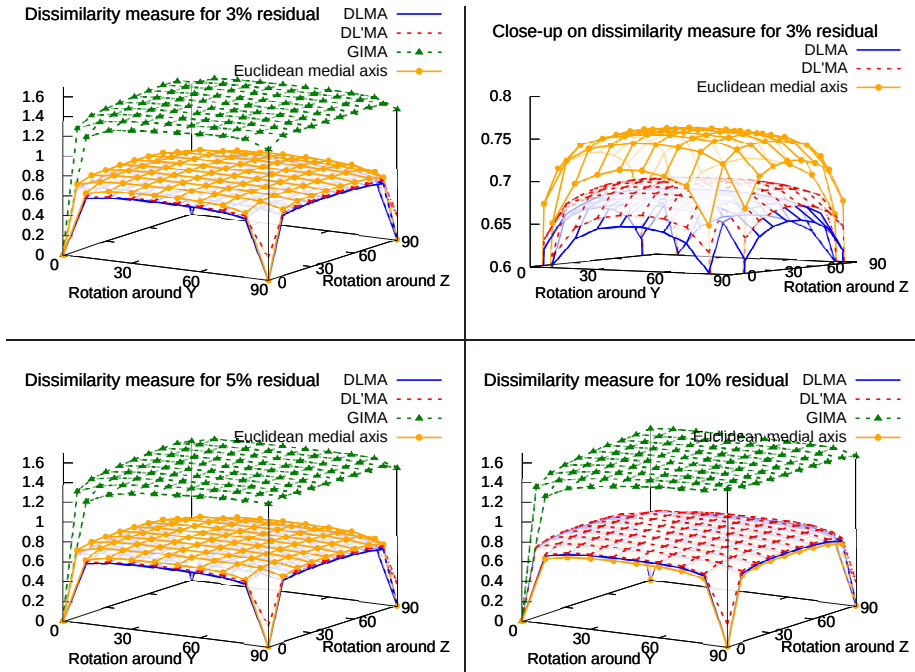


Figure 15: Hausdorff distance (**on the left**) and dissimilarity measure (**on the right**) between $M(R_\theta(X))$ and $R_\theta(M(X))$ on the 20 shapes of our 3D image database for residual values 3%, 5% and 10% (the horizontal axis is proportional to the parameter n).




Figure 16: Computing times (in seconds, vert. axis) versus image sizes (horiz. axis).

single parameter. This methodology could be of interest for comparing other medial axis filtering approaches.

Residual	2D				3D			
	10%	20%	30%	all	3%	5%	10%	all
EMA	8,23	8,13	10,71	8.55	4,86	5,71	6,89	5.82
GIMA	6,94	9,00	13,17	8.65	5,24	5,48	6,64	5.79
DLMA	7,60	9,07	10,97	8.28	3,86	4,65	6,59	5.03
DLLMA	7,41	8,67	10,62	8.02	4,14	4,80	7,26	5.40

Table 3: Average Hausdorff distance between $M(R_\theta(X))$ and $R_\theta(M(X))$ (in 3d, rotations shown were performed around the Y and around the Z axis). Results shown are averages of all tested angles. Lowest values are highlighted.

Residual	2D				3D			
	10%	20%	30%	all	3%	5%	10%	all
EMA	1,21	1,34	3,88	1,71	0,72	0,72	0,73	0,72
GIMA	1,50	2,41	6,09	2,27	1,40	1,43	1,52	1,45
DLMA	1,46	1,89	3,72	1,94	0,66	0,66	0,76	0,69
DLLMA	1,37	1,78	3,60	1,85	0,67	0,68	0,78	0,71

Table 4: Average dissimilarity between $M(R_\theta(X))$ and $R_\theta(M(X))$ (in 3d, rotations shown were performed around the Y and around the Z axis). Results shown are averages of all tested angles. Lowest values are highlighted.

References

- [1] Attali, D., Boissonnat, J.-D., Edelsbrunner, H., 2009. Stability and computation of the medial axis — a state-of-the-art report. In: Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration. Springer-Verlag, pp. 109–125.
- [2] Attali, D., Lachaud, J., 2001. Delaunay conforming iso-surface, skeleton extraction and noise removal. Computational Geometry: Theory and Applications 19, 175–189.
- [3] Attali, D., Montanvert, A., 1996. Modelling noise for a better simplification of skeletons. In: ICIP. Vol. 3. pp. 13–16.
- [4] Attali, D., Sanniti di Baja, G., Thiel, E., 1995. Pruning discrete and semi-continuous skeletons. In: Conference on Image Analysis and Processing. Vol. 974 of LNCS. Springer, pp. 488–493.
- [5] Bai, X., Latecki, L., Liu, W., 2007. Skeleton pruning by contour partitioning with discrete curve evolution. IEEE Trans. on PAMI 29 (3), 449–462.
- [6] Bertrand, G., 1994. Simple points, topological numbers and geodesic neighborhoods in cubic grids. Pattern Recognition Letters 15, 1003–1011.

- [7] Blum, H., 1961. An associative machine for dealing with the visual field and some of its biological implications. *Biol. prototypes and synthetic systems* 1, 244–260.
- [8] Blum, H., 1967. A transformation for extracting new descriptors of shape. In: *Models for the Perception of Speech and Visual Form*. MIT Press, pp. 362–380.
- [9] Borgefors, G., Ragnemalm, I., di Baja, G. S., 1991. The euclidean distance transform: finding the local maxima and reconstructing the shape. In: *7th Scandinavian Conference on Image Analysis*. Vol. 2. pp. 974–981.
- [10] Chaussard, J., Couprie, M., Talbot, H., 2009. A discrete lambda-medial axis. In: *15th Discrete Geometry for Computer Imagery (DGCI'09)*. Vol. 5810 of *Lecture Notes in Computer Science*. Springer-Verlag, pp. 421–433.
- [11] Chazal, F., Cohen-Steiner, D., Lieutier, A., 2009. Normal cone approximation and offset shape isotopy. *Computational Geometry: Theory and Applications* 42, 566–581.
- [12] Chazal, F., Lieutier, A., 2005. The λ -medial axis. *Graphical Models* 67 (4), 304–331.
- [13] Coeurjolly, D., 2002. Algorithmique et géométrie discrète pour la caractérisation des courbes et des surfaces. Ph.D. thesis, Université Lyon II, France.
- [14] Coeurjolly, D., 2003. d-dimensional reverse Euclidean distance transformation and Euclidean medial axis extraction in optimal time. In: *procs. DGCI, LNCS, Springer Verlag*. Vol. 2886. pp. 327–337.
- [15] Coeurjolly, D., Montanvert, A., 2007. Optimal separable algorithms to compute the reverse euclidean distance transformation and discrete medial axis in arbitrary dimension. *IEEE Trans. Pattern Anal. Mach. Intell.* 29 (3), 437–448.
- [16] Couprie, M., Bertrand, G., Apr. 2009. New characterizations of simple points in 2D, 3D and 4D discrete spaces. *IEEE Trans. on PAMI* 31 (4), 637–648.
- [17] Couprie, M., Coeurjolly, D., Zrour, R., 2007. Discrete bisector function and Euclidean skeleton in 2D and 3D. *Image and Vision Computing* 25 (10), 1543–1556.
- [18] Davies, E., Plummer, A., 1981. Thinning algorithms: a critique and a new methodology. *Pattern Recognition* 14, 53–63.
- [19] Dubuisson, M.-P., Jain, A., 1994. A modified hausdorff distance for object matching. In: *12th ICPR*. Vol. 1. pp. 566–568.

- [20] Eden, M., 1961. A two-dimensional growth process. In: Fourth Berkeley Symposium on Mathematical Statistics and Probabilities. Vol. 4. Univ. of Calif. Press, pp. 223–239.
- [21] Ge, Y., Fitzpatrick, J., 1996. On the generation of skeletons from discrete euclidean distance maps. *IEEE Trans. on PAMI* 18 (11), 1055–1066.
- [22] Hesselink, W., Roerdink, J., 2008. Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform. *IEEE Trans. on PAMI* 30 (12), 2204–2217.
- [23] Hesselink, W., Visser, M., Roerdink, J., 2005. Euclidean skeletons of 3d data sets in linear time by the integer medial axis transform. In: 7th ISMM. Vol. 30 of *Computational Imaging and Vision*. Springer, pp. 259–268.
- [24] Kimmel, R., Shaked, D., Kiryati, N., Bruckstein, A. M., 1995. Skeletonization via distance maps and level sets. *Computer Vision and Image Understanding* 62, 382–391.
- [25] Kong, T. Y., Rosenfeld, A., 1989. Digital topology: introduction and survey. *Computer Vision, Graphics and Image Processing* 48, 357–393.
- [26] Lee, T., Cowan, R., 1994. A stochastic tessellation of digital space. In: *Proceedings of 2nd ISMM*. Kluwer, Fontainebleau, pp. 218–224.
- [27] Malandain, G., Fernández-Vidal, S., 1998. Euclidean skeletons. *Image and Vision Computing* 16, 317–327.
- [28] Matheron, G., 1988. Examples of topological properties of skeletons. Vol. 2. Academic Press, pp. 217–238.
- [29] Maurer, C., Qi, R., Raghavan, V., 2003. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans. on PAMI* 25 (2), 265–270.
- [30] Ogniewicz, R., Kübler, O., 1995. Hierarchic voronoi skeletons. *Pattern Recognition* 28 (33), 343–359.
- [31] Pudney, C., 1998. Distance-ordered homotopic thinning: a skeletonization algorithm for 3D digital images. *Computer Vision and Image Understanding* 72 (3), 404–413.
- [32] Rémy, E., Thiel, E., 2005. Exact medial axis with euclidean distance. *Image and Vision Computing* 23 (2), 167–175.
- [33] Samozino, M., Alexa, M., Alliez, P., Yvinec, M., 2006. Reconstruction with voronoi centered radial basis functions. In: *Symposium on Geometry Processing (SGP'06)*. pp. 51–60.
- [34] Serra, J., 1982. *Image analysis and mathematical morphology*. Academic Press.

- [35] Sharvit, D., Chan, J., Tek, H., Kimia, B., 1998. Symmetry-based indexing of image databases. *Journal of Visual Communication and Image Representation* 9 (4), 366–380.
- [36] Siddiqi, K., Bouix, S., Tannenbaum, A., Zucker, S., 1999. The hamilton-jacobi skeleton. In: *International Conference on Computer Vision (ICCV)*. pp. 828–834.
- [37] Soille, P., 1999. *Morphological image analysis*. Springer-Verlag.
- [38] Svensson, S., Sanniti di Baja, G., 2003. Simplifying curve skeletons in volume images. *Computer Vision and Image Understanding* 90 (3), 242–257.
- [39] Talbot, H., Vincent, L., 1992. Euclidean skeletons and conditional bisectors. In: *VCIP'92, SPIE*. Vol. 1818. pp. 862–876.
- [40] Vincent, L., 1991. Efficient computation of various types of skeletons. In: *Medical Imaging V, SPIE*. Vol. 1445. pp. 297–311.
- [41] Welzl, E., 1991. Smallest enclosing disks (balls and ellipsoids). In: *Results and New Trends in Computer Science*. Vol. 555 of LNCS. Springer-Verlag, pp. 359–370.