Fast and Accurate Approximation of Digital Shape Thickness Distribution in Arbitrary Dimension

David Coeurjolly

▶ To cite this version:

Fast and Accurate Approximation of Shape Thickness Distribution in Arbitrary Dimension

David Coeurjolly

Université de Lyon, CNRS, LIRIS, UMR5205, F-69622, France
david.coeurjolly@liris.cnrs.fr

Abstract

We present a fast and accurate approximation of the Euclidean thickness distribution computation of a binary shape in arbitrary dimension. Thickness functions associate a value representing the local thickness for each point of a binary shape. When considering with the Euclidean metric, a simple definition is to associate with each point x, the radius of the largest ball inscribed in the shape containing x. Such thickness distributions are widely used in many applications such as medical imaging or material sciences and direct implementations could be time consuming. In this paper, we focus on fast algorithms to extract such distribution on shapes in arbitrary dimension.

1. Introduction

In many shape analysis applications, the evaluation of the thickness of an object is crucial. For example, thickness based morphometrical analyses of bones on 3D medical imaging systems are relevant for many disease investigations such as osteoporosis [11, 10, 3]. Similar applications can be found in Material Sciences [14] or even in shape description in morphological image processing [26, 19]. As discussed in [11], several options exist to define a thickness measurement on shapes. A first candidate is based on the distance transform of the shape which consists in labelling at each point of the shape with the minimum distance to its boundary [21, 15]. In this approach, the thickness at a point x can be defined by its distance transformation value that can be seen as the radii of the largest ball centered at the point and contained in the shape (see Fig. 1). Such distance information are relevant for local shape analyses but do not match with the common sense dealing with mean thickness or thickness distribution for each point of the shape. For instance, one can see that using this definition the mean thickness of an Euclidean ball in the plane of radius r is exactly $\frac{r}{3}$. The counter-intuitive bias can also be observed in higher order statistics of the thickness distribution. Another option is to define the thickness at a boundary point y in dimension 2 as the length of the Euclidean segment perpendicular to the tangent at y lying inside the shape (see Fig. 1 and [13]). If this definition makes sense on some pure analytical objects (balls, cylinders, plates, . . .), drawbacks are two-fold: First it requires an estimation of the tangent at each point of the shape boundary, which can be sensitive to noise. Second, in some situations,

it can lead to overestimation of the mean thickness. For instance, if one considers a cylinder with length l and radius r, points on the cylinder surface have a thickness of r whereas points on the cylinder caps would have a thickness of l which can be arbitrarily large. Furthermore, such thickness measurement is only defined on boundary points which can be problematic in some applications such as percolation path evaluation in Material Sciences [28, 20]. An interesting definition has been proposed in [24] which consists in computing a topological skeleton of the shape and define the mean thickness as the average of distance transformation values on skeletal points. To more precise, the author considers distance transform based a fuzzy metric [23] and includes in the summation a constant quantity depending on the image resolution. Although the use of fuzzy distance could be an interesting approach in some practical applications, such approach suffers from the high instability of the topological skeleton in presence of noise and this may have deep impact on the estimated quantity (see Sect. 4 for details). Last but not least, such an approach does not provide thickness information at every point contained in the shape which is important for shape filtering by thickness or the construction of a thickness distribution.

Beside these definitions from geometry processing, many related notions exist in the field of granulometric analysis of binary images in mathematical morphology [16, 26, 19]. In this context, the idea is to iterate on opening filters of the input binary shape. At each step, the structuring element size is increased by a dilation process and we count the number of remaining points. In some sense, this process can be viewed as a sieving process with increasing hole size. Using such an analysis, we can construct a function, called *opening function*, which assigns to each point of the input shape, a kind of thickness information parametrized by the considered structuring element (see details in Sect. 2.2). Existing algorithms in this field approximate the Euclidean metric with kinds of fixed neighborhood metrics in which unit balls are squares or diamonds for instance [29].

When dealing with the Euclidean metric, the Euclidean opening function or thickness distribution [11] can be simply defined as follows: we assign to each grid point x of a shape X the radius of the largest ball B contained in X with $x \in X$. We are thus facing a geometrical problem in which the discrete medial axis of a binary shape plays a crucial role (see [7] for an overview on Digital Euclidean Medial Axis). When we consider large binary images in 3D the straightforward implementation of the Euclidean thickness computation leads to expensive computational costs. In this paper, we investigate fast approximation algorithms defined in arbitrary dimensions based on a discrete version of the well-known Power Diagram (or Laguerre Diagram) from Computational Geometry [1]. Figure 3 sketches the important steps of the proposed technique. Since such a discrete power diagram mapping can be obtained in optimal time in arbitrary dimension [7], the proposed thickness distribution computation algorithm can also be generalized to handle binary objects as subset of $\mathbb{Z}^d$. Furthermore, we demonstrate that the proposed approach is still consistent with the mathematical morphology framework.

The paper is organised as follows: First, we introduce the fundamental notions that will be used to construct the algorithm (Sect. 2.2). In Sect. 3 we define the fast approximation algorithms which are evaluated and discussed in Sect. 4. A preliminary version of this paper has been presented in [6].
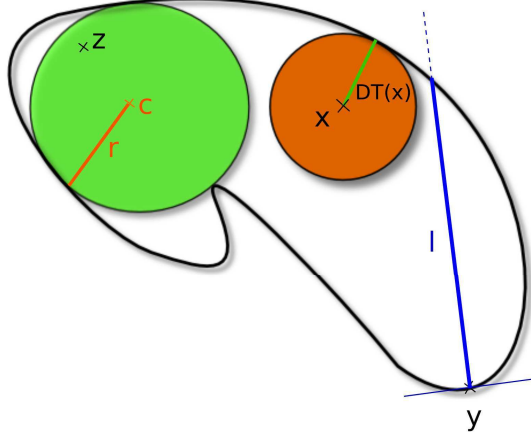
Figure 1: Illustration of different thickness definitions in dimension 2: distance transformation based thickness ($\tau_{DT}(x) = DT(x)$), thickness as the length of inner segment (see text for details, $\tau_{normal}(y) = \|l\|$) and Euclidean thickness as the radius of the maximal ball containing the point ($\tau(z) = r$).

2. Preliminaries

2.1. Notations

In the following and without loss of generality, we consider a binary object X in dimension d defined by a mapping from $[0, \ldots, n-1]^d \subset \mathbb{Z}^d$ to values $\{0, 1\}$. Grid points with value 0 are considered as background points whereas grid points with value 1 define the object points (set X).

As discussed in the introduction, we are interested in a definition of thickness as a function $\tau : X \to \mathbb{R}^+$ which maps each object point $x \in X$ to a thickness value. Whatever the thickness definition we consider, relevant measurements and statistics can be extracted. Indeed, if we consider x as uniformly distributed random variable with realization in X and its associated thickness measurement, we can define a discrete *thickness distribution function F*:

$$F(t) = \frac{|X_t|}{|X|}, \quad X_t = \{x \in X | \tau(x) \leq t\} \tag{1}$$

from which we can derive several quantities such as the mean thickness of a shape $\bar{\tau}$, the standard deviation or even high order statistics using central moments defined on the density function (see Fig. 6 for an illustration of $F(t)$ on a binary shape). In [11], the authors discuss about the use of such high order information for shape classification. Beside these high order parameters, [10] and [3] discussed about the interest of the mean thickness value to extract morphological parameters on bone structures.

2.2. Granulometries

The notion of granulometry and granulometric functions was first introduced by Matheron in 1967 [16] in order to study images of porous materials. First of all, a

3

generic definition of granulometry can be formalized as follows [26].

Definition 1. *Let $\Phi = (\phi_\lambda)_{\lambda \geq 0}$ be a family of image transformations. Φ is a granulometry if and only if for all $\lambda, \mu \geq 0$:*

 (i) *ϕ_λ is increasing;*

 (ii) *ϕ_λ is anti-extensive ($\phi_\lambda(X) \subseteq X$);*

(iii) *$\phi_\lambda \phi_\mu = \phi_{\max(\lambda,\mu)}$.*

From the granulometry, the opening function can be defined as follows: The *granulometry function or opening function G_Φ^X* of a binary image X for granulometry Φ maps each pixel $x \in X$ to the size of the smallest λ such that $x \notin \phi_\lambda(X)$. In the context of the paper, such opening function G_Φ^X can be interpreted as a thickness function. Beside *granulometries generated by a structuring element* widely used in Mathematical Morphology, [17, 29, 26, 19] we consider in the following a thickness function based on the Euclidean metric using perfect increasing disks as structuring elements. We will demonstrate that this function still have the properties of granulometric functions. Designing a thickness measurement which is consistent with these properties is fundamental in order to be able to use this operator in mathematical morphology pipelines and to obtain properties on the result [19].

2.3. Euclidean Medial Axis and Euclidean Thickness Function

In Digital geometry, distance transformation and discrete medial axis extraction are classical problems from decades [21, 4, 27, 23, 15]. Using an exact Euclidean metric, optimal in time algorithms exist to compute without error such a distance transformation whatever the dimension is [22, 5, 12, 18]. The interest for considering the Euclidean metric has been widely discussed in the literature. For short, exact Euclidean metric allows us to reduce the anisotropy of the digital grid and to fill the gap between digital tools and structures from computation geometry such as Voronoi Diagram and Power Diagram (see below). As demonstrated in Sect. 4, the isotropy of the Euclidean metric has a deep impact of the quality if the thickness measurement (see for example Fig. 11 experiment). Among the wide set of measurements that can be computed from the distance transformation, the extraction of the discrete medial axis of a shape X is probably the most important one. Indeed, the Discrete Medial Axis can be defined as the set of maximal balls of X: A ball $B \subset X$ is maximal in X if there is no ball $B' \subset X$ such that $B \subset B'$. In [7], we have presented an algorithm to extract the set of Euclidean maximal balls in linear time in arbitrary fixed dimension. More precisely, if $X \subset [0..n]^d$, the discrete medial axis $MA(X)$ of X is obtained in $O(n^d)$ (see Fig. 2-a in 2-D). We focus here on a thickness function based on the Euclidean metric defined as:

$$\tau(x) = \max(\{r \mid \forall B(c,r) \subset X, x \in B(c,r) \cap X\}), \qquad (2)$$

where $B(c,r)$ denotes an Euclidean open ball with center $c \in X$ and radius r. In other words, we want to assign to each point x in X, the maximal radius among the set of balls inscribed in X that contains x. From this function, we can define a set of transformations $\Phi = \{\phi_\lambda\}$ such that $\phi_\lambda(X) = X_{r_{max}-\lambda}$ which is a granulometry

(with r_{max} being the maximal ball radius). As discussed in [11], Eq. (2) can be optimized since it is sufficient to only consider MA balls $\{B_i(x_i, r_i)\}$. Indeed, Eq. (2) is equivalent to

$$\tau(x) = \max(\{r_i \mid \forall i, x \in B_i(c_i, r_i) \cap X, \ B_i \in MA(X)\}). \tag{3}$$

If we define the height function $h_i(x)$ as follows

$$h_i(x) = \begin{cases} r_i^2 & \text{if } r_i^2 - \|c_i - x\|^2 > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{4}$$

Eq. (3) is thus equivalent to

$$\tau(x) = \sqrt{\max_{B_i(c_i, r_i) \in MA(x)} (h_i(x))}. \tag{5}$$

Figures 2-c and 2-e illustrate these objects as circles and surfaces $\{(x, h_i(x))\}$ with $x \in \mathbb{R}^2$. On a computational point of view, the overall computational cost of the algorithm derived from Eq. (3) or Eq. (5) is $O(n^d + \sum_{B_i \in MA(X)} |B_i \cap X|)$. Indeed, $O(n^d)$ is required to obtain the MA, then we have to scan all MA ball grid points and store the maximal radius at each point (see Fig. 3-b). As illustrated in Sect. 4, such a brute-force algorithm is very expensive since many balls of the MA may overlap and thus their associated grid points would be scanned several times. In the following section, we investigate fast approximation algorithms to compute such distribution function with the help of discrete power diagram mapping.

3. Thickness Distribution Computation Algorithm

3.1. Power Diagram as a Thickness Function

In [7], we have also demonstrated links between the MA extraction and the computation of a *Power diagram* [1]. Note that using other medial axis definitions such as Chamfer based ones [4], or local distance transform value maxima in a fixed size neighborhood [23], the power diagram construction would have been much more complex to extract. First, let us consider a set of N sites $\mathcal{S} = \{c_i\}_{1...N}$ such that each point c_i is associated with a radius r_i. The power $\sigma_i(x)$ of a point x in $\mathbb{R}^d$ according to the site c_i is given by $\sigma_i(x) = \|x - c_i\|^2 - r_i^2$. If $\sigma_i(x) < 0$, x belongs to the disk of center c_i and radius r_i. If $\sigma_i(x) > 0$, x is outside the ball. The *power diagram* is a kind of Voronoi diagram based on the metric σ. Hence, the power diagram $\mathcal{V}_\mathcal{S}$ is a decomposition of the space into open cells $F = \{f_i\}_{1...N}$ associated with each site c_i such that $f_i = \{x \in \mathbb{R}^d \mid \forall j \in \{1...N\}, j \neq i, \sigma_i(x) < \sigma_j(x)\}$. Note that cell f_i associated with site c_i may be empty, otherwise, f_i is a convex polytope (see Figure 3-c in dimension 2) [1].

As discussed above, a Power Diagram Mapping $\Pi_X(x)$ of MA balls can be obtained during the MA extraction without changing the overall computational cost. The mapping assign to each discrete point $x \in X$ the ball $B_i(c_i, r_i)$ such that x belongs to cell f_i (see Fig. 2-b). If x belongs to a cell boundary, we arbitrarily choose to associate
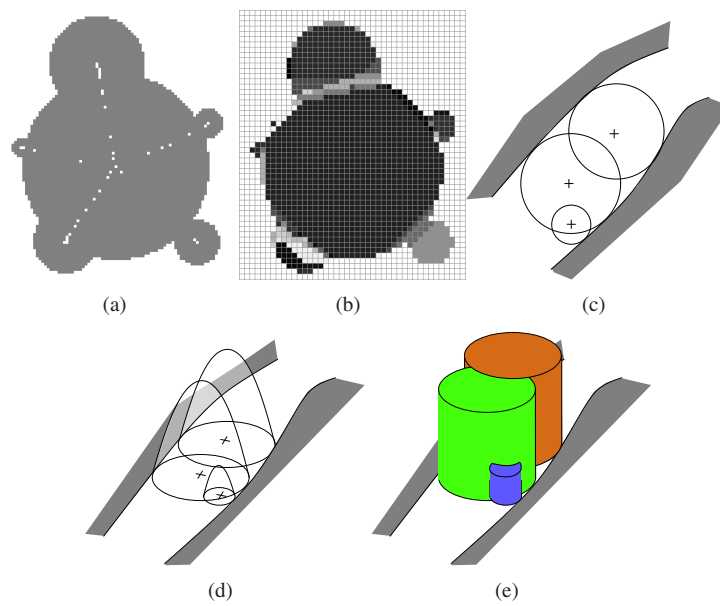
Figure 2: Discrete medial axis (a) and discrete power diagram of a 2D shape $(b$, one gray value color per power cell). Given a subset of three maximal balls $i \in \{1, 2, 3\}$ (c), illustration of $\{x, \sigma_i'(x)\}$ (d) and $\{x, h_i(x)\}$ (e) surfaces.

6

x with one of the balls of its adjacent cells. In the following, the index i is used to identify a ball $B_i(c_i, r_i)$ or a power diagram face f_i.

If we consider now a power function σ' such that $\sigma'_i(p) = -\sigma_i(p)$, we do not change the power diagram cells geometry and we can interpret the diagram construction as the upper envelope of a set of elliptic paraboloids in 2D (see Fig. 2-d).

Let us consider the power diagram of balls in $MA(X)$ (*i.e.* in the following, $N = |MA(X)|$). For $x \in X$, we define the function τ_Π such that $\tau_\Pi(x) = r_i$ if $\Pi_X(x) = i$. More formally,

$$\tau_\Pi(x) = \operatorname*{argmax}_{r_i}\{\sigma'_i(x)\} \text{ with } B_i(c_i, r_i) \in MA(X). \tag{6}$$

In other words, we associate with each point x of the object, the radius of the ball associated with the discrete power diagram cell i containing x. From this definition, we have

Lemma 1. *Given an object X in dimension d, τ_Π induces a granulometry function and $\tau_\Pi(x) \leq \tau(x)$ for each $x \in X$.*

Proof. From $\tau_\Pi(x)$, we can construct a set of transformations $\Phi = \{\phi_\lambda\}$ such that $\phi_\lambda(X) = X_{r_{max}-\lambda} = \{x \in X \,|\, \tau_\Pi(x) \leq r_{max} - \lambda\}$ with $\lambda \in [0, r_{max}]$ (r_{max} is the maximal ball radius in $MA(X)$). Hence it is clear that functions ϕ_λ satisfy the statements (i), (ii), (iii) of granulometry functions in Definition 1. To prove that τ_Π underestimates the thickness, let us consider a ball $B_i(c_i, r_i)$. From Eq.(4) and σ' definition, we have $\sigma'_i(x) \leq h_i(x)$ for all $x \in X$. Since $\tau_\Pi(x)$ and $\tau(x)$ are defined at a point x as the ball radius r_i maximizing respectively the quantities $\{\sigma'_j(x)\}_{1...N}$ and $\{h_j(x)\}_{1...N}$, we have $\tau_\Pi(x) \leq \tau(x)$. $\qquad\square$

Beside the properties presented in Lemma 1, we can see that $\tau_\Pi(x)$ is an *intuitive local thickness* measurement on classical shapes as defined in [11]. Indeed, if we consider regular objects (spheres, cylinders or plates), the estimated mean thickness using $\tau_\Pi(x)$ matches with the intuition we have of the thickness of such objects. Furthermore, $\tau_\Pi(x) = \tau(x)$ if x is in general position, *i.e.* far from plates and cylinders edges or extremities. As detailed above, an implementation of the τ_Π map computation can be designed in $O(n^d)$ for objects in dimension d. As illustrated in Fig. 3-c, the differences between $\tau_\Pi(x)$ and $\tau(x)$ are located on balls intersections (see further error analysis in Sect. 4). Even if τ_Π underestimates the thickness, the geometrical structure provided by the power diagram mapping will allow us to design a quasi-linear and fast thickness approximation algorithm. In the following section, we present an algorithm which starts from an initial τ_Π mapping and which optimize the cell boundaries to have a better approximation of the thickness values.

3.2. Fast Quasi-linear Approximation

The main idea of this algorithm is to start from the τ_Π mapping and then to deform the power diagram boundaries between two balls taking into account the relative order between their radii (cf Fig. 3-d). Let us define the *active border* of a discrete power diagram cell.
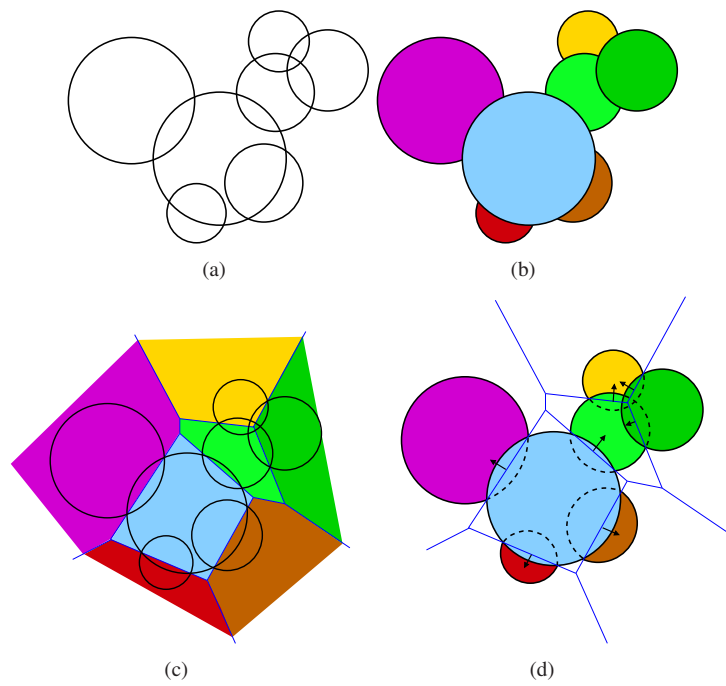
Figure 3: Overview of the proposed method: (a) an input set of maximal balls, (b) its Euclidean thickness function, (c) a first approximation based on the power diagram of balls, and (d) the approximation based on active border propagation.

Definition 2. *Given a Power diagram mapping* Π_X, *the* active border $\mathcal{N}(i)$ *of a cell* f_i *of* Π_X *is such that*

$$\mathcal{N}(i) = \{x \in X \mid \Pi_X(x) = i \text{ and } \exists y \in N(x), \text{ such that } \Pi_X(y) = j \ (j \neq i) \quad (7)$$
$$\text{and } r_j < r_i\} \quad (8)$$

where $N(x)$ *denotes the set of direct neighbors*[1] *of* x *in* X.

The first statement of this definition (Eq. (7)) characterizes points belonging to the discrete boundary of the set $\{f_i \cap \mathbb{Z}^d\}_{1...N}$. The second statement (Eq. (8)) filters the boundary points to only consider those for which the thickness propagation detailed below makes sense.

In the following, such sets will be considered as queues with *Pop_element* and *Push_element* methods.

The fast approximation algorithm can be sketched as follows (see Alg. 1 for details): we first scan the grid points to create active border queues. At this point, we also create a mapping $\tau_{\Pi'}(x) = \tau_{\Pi}(x)$ that will contain the new thickness values. Then, we process active borders by decreasing associated radii and use a *breadth-first* approach to *inflate* active border of a cell i. More precisely, when a grid point x belonging to the list $\mathcal{N}(i)$ is processed, for each neighbor $y \in N(x)$, we push y to $\mathcal{N}(i)$ if the following conditions are fulfilled: y must belong to the ball B_i, y must belong to a cell $j \neq i$ and $r_j < r_i$. If y is added to $\mathcal{N}(i)$, we mark this point as processed and we update its associated thickness $\tau_{\Pi'}(y) = \tau_{\Pi'}(x)$.

Lemma 2. *Given an object* X *in dimension* d, *Alg. 1 computes the thickness function* $\tau_{\Pi'}$ *in* $O(n^d + N \cdot \log N)$ *with* N *the number of balls in* $MA(X)$. *Furthermore,* $\tau_{\Pi'}$ *induces a granulometry function and* $\tau_{\Pi}(x) \leq \tau_{\Pi'}(x) \leq \tau(x)$.

Proof. Whatever the dimension of the shape is, the computational cost of Alg. 1 is quasi-linear since it runs in $O(n^d + N \cdot \log N)$ with N the number of balls in the $MA(X)$ ($N \ll n^d$). Indeed, as discussed above, both $MA(X)$ and τ_{Π} are obtained in $O(n^d)$. Then we have to sort the balls which can be done in $O(N \cdot \log N)$. It is clear that the overall size of the active border queues ($\{\mathcal{N}(i)\}_{i=1...N}$) is bounded by the size of X. In fact, since $\{\mathcal{N}(i)\}_{i=1...N}$ contains only boundary points of Π_X cells, its size is generally of an order on magnitude lower than $|X|$. To conclude the complexity analysis, we can see that if a point x in $\mathcal{N}(i)$ is popped from $\mathcal{N}(i)$, and since we process the queues by decreasing radii, the point will never be considered anymore (thanks to the constraint $r_j < r_i$ in the inner **if** test).

Using similar arguments as in the proof of Lemma 1, $\tau_{\Pi'}$ induces a granulometry function. Furthermore, since we initialize the values of $\tau_{\Pi'}$ by τ_{Π}, and since Alg. 1 only changes the value $\tau_{\Pi'}(x)$ at x if it increases (inner **if** test), it is clear that $\tau_{\Pi}(x) \leq \tau_{\Pi'}(x)$. Furthermore, since the set of balls remains the same, we still have $\tau_{\Pi'}(x) \leq \tau(x)$. $\square$

---

[1] In arbitrary dimension d, $x \in X$ is a direct neighbor of $y \in X$ iff $\|x - y\|_1 = 1$ for the l_1 norm.

Algorithm 1: Euclidean Thickness Function Approximation.

Input: a digital object X in dimension d
Output: the map $\tau_{\Pi'}$
Compute $MA(X)$ and Π_X;
$\tau_{\Pi'} \leftarrow \tau_{\Pi}$;
Initialize the $\mathcal{N}(i)$ empty queues;
foreach *grid point* $x \in X$ **do**
    if x *is an active border* **then**
        Push_element($x, \mathcal{N}(\Pi_X(x))$)

Sort the balls in $MA(X)$ by decreasing radii;
foreach *ball with index i in the sorted list* **do**
    while $\mathcal{N}(i)$ *is not empty* **do**
        $x \leftarrow$ Pop_element($\mathcal{N}(i)$);
        foreach $y \in N(x)$ **do**
            if $\Pi_X(x) = i \neq \Pi_X(y) = j$ **and** $y \in B_i(x_i, r_i)$ **and** $r_j < r_i$ **then**
                $\tau_{\Pi'}(y) = r_i$;
                Remove the element y from $\mathcal{N}(j)$;
                Push_element($y, \mathcal{N}(i)$);

The thickness $\tau_{\Pi'}$ is still an approximation of τ (and thus $\tau_{\Pi'}(x) \leq \tau(x)$) as illustrated in the example depicted in Fig. 4: Let us consider three balls in $\mathbb{Z}^2$ and focus on the index of balls which generates the thickness value τ, τ_{Π} and $\tau_{\Pi'}$ on the row of grid points contained in the round rectangle. Power diagram boundaries are depicted with dashed line in Fig. 4-left. On the right, we have depicted the 1D representation of the $\sigma'_i(x)$ functions (parabolas in this case). If we apply Alg. 1, we start from ball a and its active border set $\mathcal{N}(a)$ contains the point with abscissa 5. During the 1D propagation, the ball a will rewrite the values assigned to pixels with abscissa 6 and 7. Hence, when processing ball b, no active border points in $\mathcal{N}(b)$ must be processed on this row and b cannot propagate to c its radii leading to an incorrect result at $x = 8$ since $\tau_{\Pi'}(x) < \tau(x)$.

As illustrated in Fig. 8, experiments confirm that errors between $\tau_{\Pi'}$ and τ are located on thin *lenses* and *lunes* of the ball arrangement.

4. Experiments and Discussion

In this section we provide a complete experimental evaluation of the proposed algorithms in dimension 3. Table 1 details the overall results on a set of digital objects presented in Fig. 5. We have computed the elapsed time for each algorithm and the speed-up of the proposed approximation algorithm. To focus on the core of the thickness computation algorithms, we have decomposed the computational cost into two parts: the initialization step which is roughly similar for each method and the main loop. During the initialisation step, we basically compute the medial axis and discrete
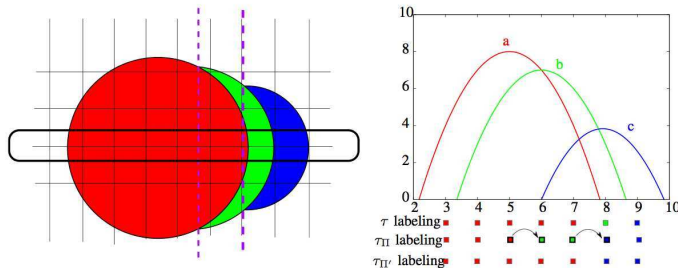
Figure 4: Illustration of a 2D configuration in which Alg. 1 introduces thickness underestimation.

power diagram. For the τ function, the main loop consists in processing all balls one by one and labeling the ball points with the maximum radius (Eq. (3)). For the τ_Π mapping, the main loop consists in scanning the discrete power diagram mapping Π_X and associating the corresponding radii to the grid points. Finally, the main loop for $\tau_{\Pi'}$ computation corresponds to Alg. 1 without computations of $MA(X)$ and Π_X.

Table 2 indicates the Mean Square Error between the τ thickness map and both τ_Π and $\tau_{\Pi'}$ (see Fig. 8 for an illustration of the error location). In our experiment we have considered a comparison with a skeleton based approach similar to Saha's one [24]: we first construct a one-dimensional topological skeleton with extremities (see for example [2] or [25]) and compute the mean thickness as the average of Euclidean distance transform values on the skeleton grid points. Compared to Saha's original approach, we have considered here an exact Euclidean DT instead of a fuzzy one which enhances isotropy and may improves the results (in the following, we denotexs `TopoThickness` such algorithm). In fact, the mean thickness computation is more affected by the high instability of the skeleton than the DT information itself. In the original approach, the authors also consider a correction of the thickness measure by adding a quantity ($0.38 \cdot \rho$ where ρ is the grid resolution) proportional to voxel resolution if such information is available from the acquisition device depending the application. Such correction could also be used in our framework to "correct" τ, τ_Π and $\tau_{\Pi'}$ in a similar way (simply adding the same quantity). In our experiment and without loss of generality, we have chosen to consider un-corrected versions of all thickness estimators and mean thickness values.

Examples of the $\tau_{\Pi'}$ thickness mapping can be found in Fig. 7. Differences between τ, τ_Π and $\tau_{\Pi'}$ can be observed in the thickness distribution $F(t)$ as defined in Eq. (1) (Fig. 6).

To summarize the experimental evaluation, we can first observe that τ_Π is the fastest thickness approximation algorithm. Indeed, once the map Π_X is computed, the computational cost of the main loop is low. Then, we can see that the additional cost we have in $\tau_{\Pi'}$ algorithm with the management of queues and border propagation is low. Indeed, with the τ_Π approximation, we are between 1.2 and 12.23 times faster than the naive algorithm τ. If we add the border propagation step, we are still between 1.06 and 9.88 times faster than τ. It means that both the size of the input object and the ball overlapping discussed at the end of Sect. 2.3 are real bottlenecks for τ computation
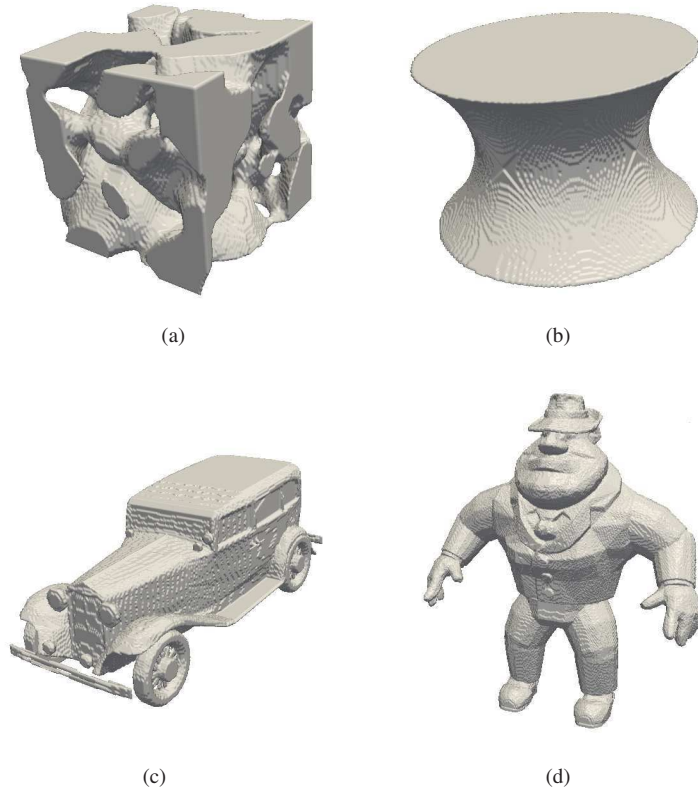
11

(a)

(b)

(c)

(d)

Figure 5: Objects used in the experiments: (a) Snow (Courtesy of GAME-CNRM/CEN and ESRF) (b) catenoid, (c) Dodge and (d) Al.

Table 1: Experimental results in dimension 3.

Object	$	X	$	$\|MA(X)\|$	Time (in sec.)									Speed-up (w.r.t. τ)	
			τ			τ_Π			$\tau_{\Pi'}$				
			Init.	Loop	Total	Init.	Loop	Total	Init.	Loop	Total	τ_Π	$\tau_{\Pi'}$
Snow	111^3	32165	0.4	2.35	2.76	0.22	2.08	2.3	0.22	2.4	2.61	**1.2**	**1.06**
cat.	200^3	22994	1.56	9.86	11.42	1.5	1.91	3.41	1.6	2.46	4.07	**3.34**	**2.81**
Dodge	350^3	84856	7.57	40.9	48.48	6.83	1.44	8.27	7	3.2	10.21	**5.86**	**4.75**
Al	450^3	96947	18.63	196.5	215.1	16.55	1.03	17.58	17.06	4.72	21.78	**12.23**	**9.88**

12

Table 2: Mean Square Error of τ_Π and $\tau_{\Pi'}$ thickness function with respect to the τ values.

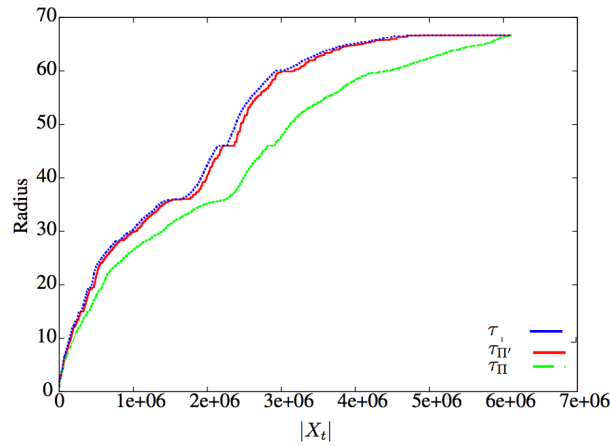| Object	Mean Square Error (w.r.t. τ)		Mean Thickness Values			
	τ_Π	$\tau_{\Pi'}$	$\bar{\tau}$	$\bar{\tau}_\Pi$	$\bar{\tau}_{\Pi'}$	TopoThickness ($|Skel|$)
Snow	5.25	**0.34**	12.192	10.501	11.955	6.78 (3800)
cat.	75.62	**1.44**	36.0931	29.8142	35.4918	12.77 (1966)
Dodge	58.74	**1.46**	39.8798	34.8921	39.4218	10.26 (13247)
Al	65.1	**3.1**	51.094	45.005	50.4822	20.44 (6029)

Figure 6: Thickness distribution $F(t)$ of the object Al: in ordinate, we have considered the radii values (t) and in abscissa we have the cumulative distribution which corresponds to $|X_t|$.

which can be avoided in the $\tau_{\Pi'}$ optimisation. For example, if we consider the Snow object, the overall geometry is closed to a network of tubular structures which limits the number of overlapping balls and thus reduces the differences between τ and $\tau_{\Pi'}$.

From Table 2 and with the help of a mean square error analysis, we can see that the active border propagation step in $\tau_{\Pi'}$ considerably increases the accuracy of the approximation (cf Fig. 8). Furthermore, the accuracy can also be observed when computing statistics (mean thickness in Table 2 and distribution function $F(t)$ in Fig. 6) on the thickness distribution.

In Fig. 9, we have considered the object Al with increasing resolutions ($n \in \{50, 100, 150, \ldots, 450\}$) in order to see how the performance behave while scaling the size of the input data. If we perform a least-square fitting analysis in log-scale space, we can observe that the behavior of the τ computation time curve is in $O(n^4)$ whereas the behavior the $\tau_{\Pi'}^X$ curve stays in $O(n^3)$. Again, the ball overlapping bottleneck can be avoided by $\tau_{\Pi'}$ while preserving good thickness measurements.

We have also designed an experiment to evaluate the stability of the mean thickness value with respect to noise. In Table 3, we have considered an Euclidean digital ball of radius 20 and several copies of this ball with different noise parameter. Hence,
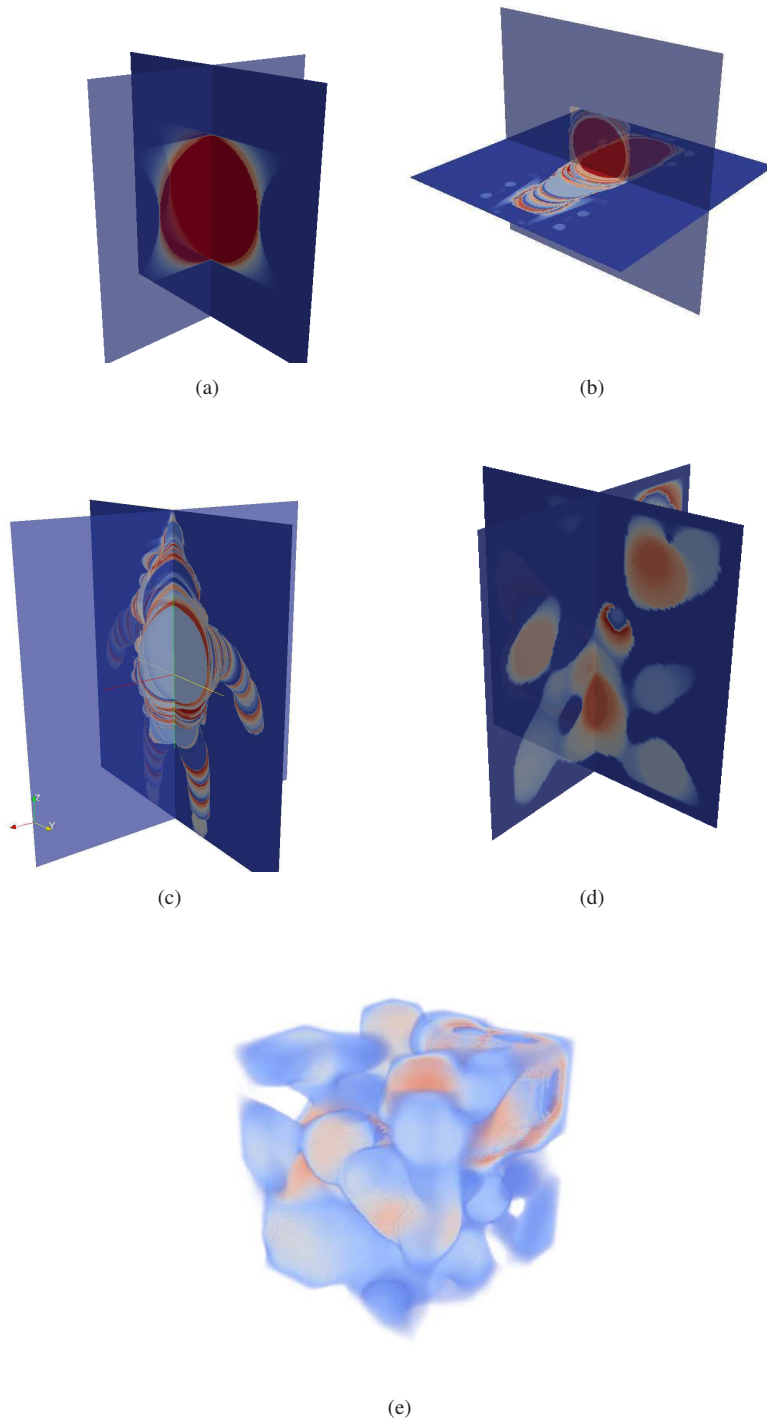
13

Figure 7: Results of the $\tau_{\Pi'}$ thickness mappings with the help of volume slices: (a) result for the `Cat` object, (b) result for the `Dodge` object, (c) result for the `Al` object and $(d - e)$ for object `Snow` (with slices and volumetric rendering).
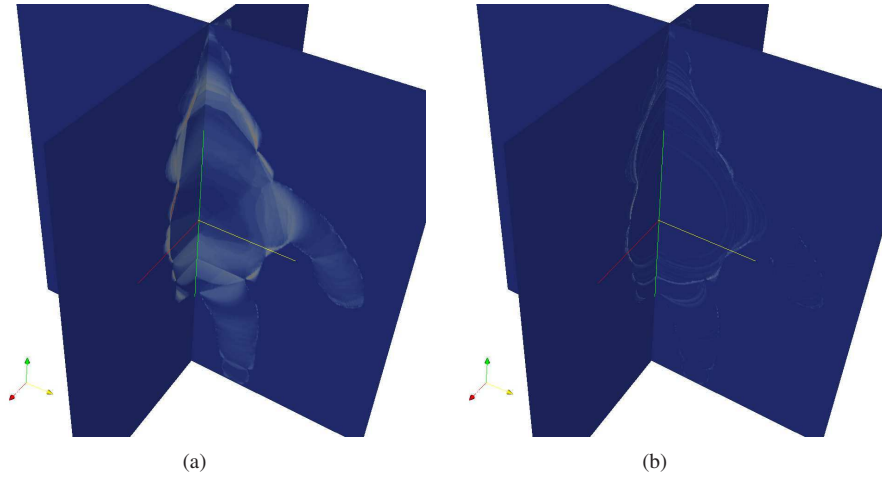
Figure 8: Thickness error locations on object `Al`: thickness differences between τ_Π and τ (a) and $\tau_{\Pi'}$ and τ (b).

expected mean thickness value is 20. In this experiment, we have considered a topological equivalent noise model to be consistent with comparison with [24]: given a noise parameter p, a voxel of the object is removed with probability p only if the voxel removal does not change the object topology (*i.e.* the voxel is simple [25, 2], see Fig. 10). In Table 3, we can observe that all mean thickness values based on Eq. (2) ($\bar{\tau}$, $\bar{\tau}_\Pi$ and $\bar{\tau}_{\Pi'}$) are stable: Even if the Euclidean medial axis is highly distorted by the noise on the boundary, (Fig. 10-(i)), both the power cell boundary propagation and the fact that the mean thickness is computed from all point thickness values in the shape, reduce the impact of the noise on $\bar{\tau}$, $\bar{\tau}_\Pi$ and $\bar{\tau}_{\Pi'}$ measurments. Furthermore and as discussed earlier, both $\bar{\tau}_\Pi$ and $\bar{\tau}_{\Pi'}$ underestimate $\bar{\tau}$. Since noise on the object boundary induces many branches from grid points near the surface toward the object interior in the topological skeleton considered in `TopoThickness` (Fig. 10-(i)) and since the mean value is given by averaging skeletal point values, the error on the mean thickness computation is important (DT values decrease when sampled close to the boundary).

Finally, Fig. 11 presents stability evaluation with respect to rotations around the z-axis of a binary object. Since we use an exact Euclidean metric, distance transformation information are stable according to such rotations. Hence, $\bar{\tau}$, $\bar{\tau}_\Pi$ and $\bar{\tau}_{\Pi'}$ measurements have small variations (*e.g.* only 2.5% of maximal variation for $\bar{\tau}_{\Pi'}$). However, due to instability of the topological skeleton, `TopoThickness` thickness measurements underestimate the quantities and have larger variations (31%).

5. Conclusion

In this paper, we have investigated algorithms for fast and accurate approximation of the Euclidean thickness distribution of binary shapes in $\mathbb{Z}^d$. To summarize the contributions, we have first fill the gap between mathematical morphology and medial
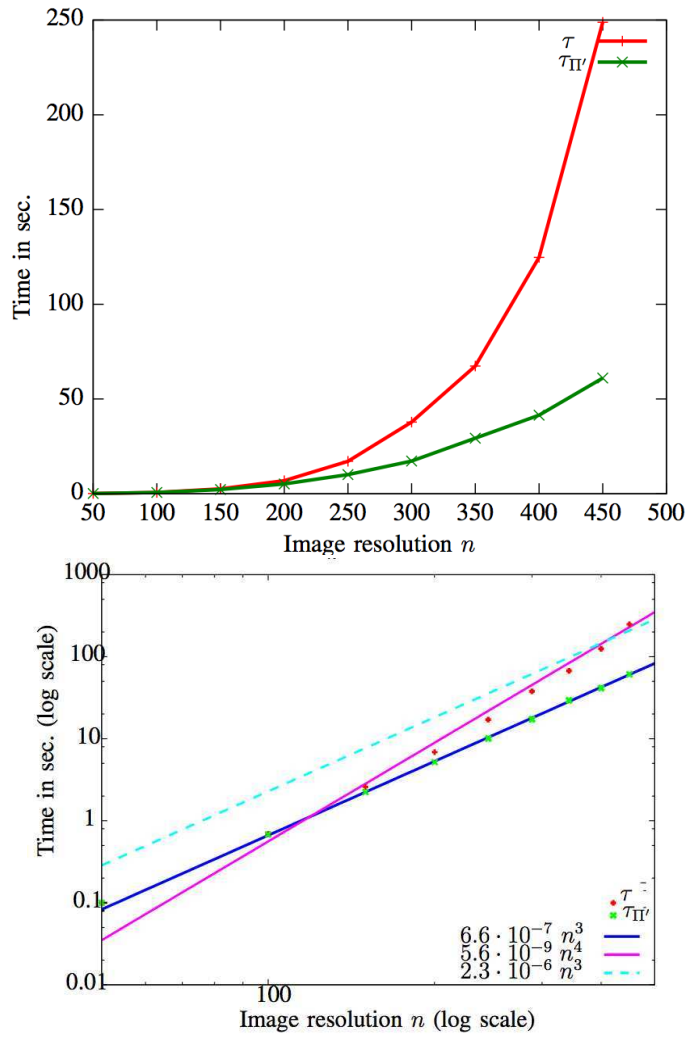
Figure 9: Thickness function computation behavior on the object Al both with normal in log-scale axis. In the log-scale graph, we have depicted two linear least-square fittings in log-scale with a model of order 4 for τ and of order 3 for $\tau_{\Pi'}$. The dashed-line illustrates the high error in the linear approximation when we try to fit a order 3 model to τ computation time data.
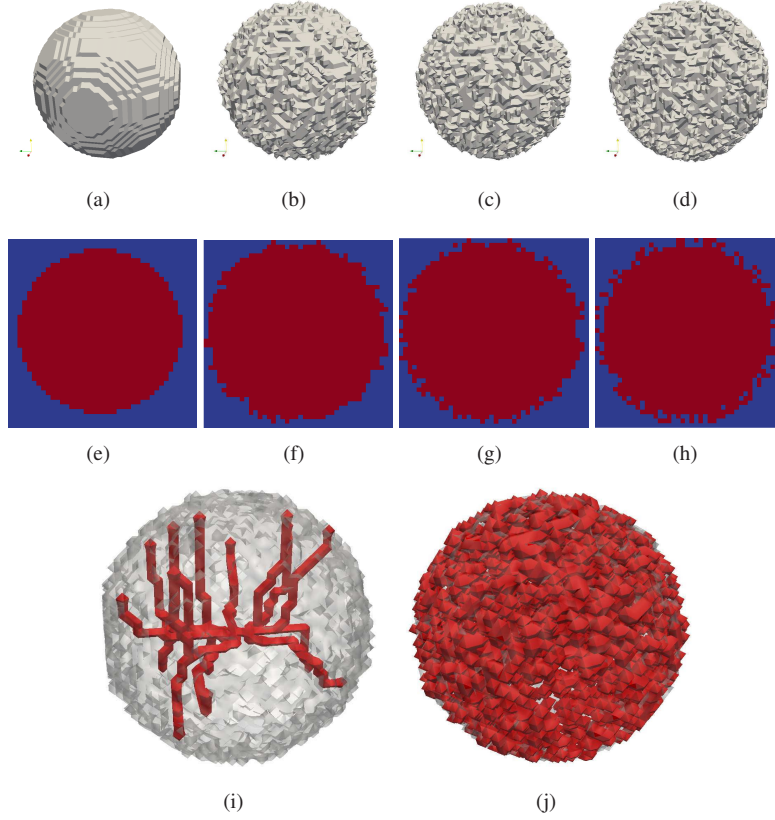
16

Figure 10: Objects used for the noise stability analysis: the original shape (a) and noisy versions from (b) to (d) with parameters 0.2, 0.3 and 0.6. Figures (e) to (h) display a slice containing the ball center for each object. Based on object (b), figure (i) (resp. (j)) illustrates the effect of boundary noise on the topological skeleton (resp. the medial axis extracction).

Table 3: Stability with respect to noise on a ball of radius 20.

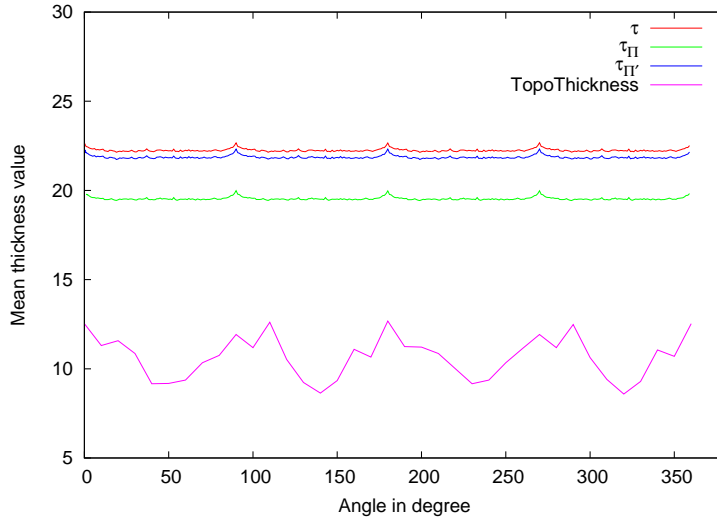|	Orig.	Noise parameter					
		0.1	0.2	0.3	0.4	0.5	0.6
$\bar{\tau}$	20	15.86	15.53	15.38	15.3	15.28	15.28
$\bar{\tau}_{\Pi}$	20	14.75	14.7	14.77	14.82	14.92	14.97
$\bar{\tau}_{\Pi'}$	20	15.52	15.27	15.2	15.16	15.17	15.2
TopoThickness	20	8.11	7.85	6.73	6.01	5.67	5.28
$(|Skel|)$	(1)	(174)	(374)	(698)	(988)	(1511)	(1932)

Figure 11: Stability of thickness measurements on object Al (200^3) according to rotations in degree around the z axis.

ball based thickness computation algorithms. Then, we have demonstrated that using a digital version of a well-known structure from Computational Geometry (the Power diagram), we can obtain a first approximation of the thickness function in linear time with respect to input size. Finally, we have proposed an algorithm which improves the first approximation with a active border propagation mechanism maintaining the computational efficiency. Hence, the proposed algorithm has a computational cost in $O(n^d + |MA(X)| \log |MA(X)|)$ (with $|MA(X)| \ll |X|$ in the general case), for digital shapes in dimension d embedded in a $[0, \ldots, n-1]^d$ domains. This should be compared to the straightforward algorithm which is in $O(n^d + \sum_{B_i \in MA(X)} |B_i \cap X|)$ and we have demonstrated in the experiments that its efficiency could be one order of magnitude higher that the proposed algorithm in our data set. Furthermore, we have demonstrated that this thickness function is still consistent with the granulometry function theory from Mathematical Morphology.

As illustrated in the experiments, our main application of this work is to perform shape measurements on microtomographic snow images (object $Snow$ and [9, 8]). More precisely, we plan to include this thickness measurement in a global percolation path evaluation framework. A very specific aspect of this application is that we have to handle high resolution images (up to 2048^3). In this context, our fast algorithm is very relevant but an important future work would be to have parallel implementation. A good point is that very efficient parallel implementation of both the medial axis extraction and the power diagram mapping computation exist. The remaining bottleneck would be to implement the boundary propagation in an efficient way too. Beside this project on high resolution images, an interesting future work would be to consider fuzzy quantities in the process such as fuzzy distance transform [23] but also fuzzy

maximal balls which are crucial for our thickness definition. We thus expect to design fuzzy thickness measurements in images where segmenting the object of interest is problematic.

References

References

[1] F. Aurenhammer. Power Diagrams: Properties, Algorithms, and Applications. *SIAM Journal on Computing*, 16:78–96, 1987.

[2] G. Bertrand, G. Malandain A new characterization of three-dimensional simple points. *Pattern Recognition Letters*, 15(2):169–175, 1994.

[3] A. Bonnassie, F. Peyrin, and D. Attali. A new method for analyzing local shape in three-dimensional images based on medial axis transformation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 33(4):700–705, 2003.

[4] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344–371, 1986.

[5] H. Breu, J. Gil, D. Kirkpatrick, and M. Werman. Linear time Euclidean distance transform algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):529–533, 1995.

[6] D. Coeurjolly. Fast and Accurate Approximation of the Euclidean Opening Function in Arbitrary Dimension, International Conference on Pattern Recognition, IEEE Computer Society, 2010.

[7] D. Coeurjolly and A. Montanvert. Optimal separable algorithms to compute the reverse euclidean distance transformation and discrete medial axis in arbitrary dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(3):437–448, 2007.

[8] F. Flin, J.-B. Brzoska, B. Lesaffre, C. Coléou, P. Lamboley, D Coeurjolly, O Teytaud, G Vignoles, and J F Delesse. An adaptive filtering method to evaluate normal vectors and surface areas of 3D objects. Application to snow images from X-ray tomography. *IEEE Transactions on Image Processing*, 14:585–596, 2005.

[9] F. Flin, J.-B. Brzoska, B. Lesaffre, C. Coléou, and R. A. Pieritz. Full three-dimensional modelling of curvature-dependent snow metamorphism: first results and comparison with experimental tomographic data. *Journal of Physics D: Applied Physics*, 36(10A):A49–A54, May 2003.

[10] T. Hildebrand, A. Laib, R. Müller, J. Dequeker, and P. Rüegsegger. Direct three-dimensional morphometric analysis of human cancellous bone: microstructural data from spine, femur, iliac crest, and calcaneus. *Journal of bone and mineral research : the official journal of the American Society for Bone and Mineral Research*, 14(7):1167–74, July 1999.

[11] T. Hildebrand and P. Ruegsegger. A new method for the model-independent assessment of thickness in three-dimensional images. *Journal of Microscopy*, 185(1):67–75, January 1997.

[12] T. Hirata. A unified linear-time algorithm for computing distance maps. *Information Processing Letters*, 58:129, 1996.

[13] E. B. Jensen, H. J. Gundersen, and R. Osterby. Determination of membrane thickness distribution from orthogonal intercepts. *Journal of Microscopy*, 115(1):19–33, January 1979.

[14] T. U. Kaempfer and M. Schneebeli. Observation of isothermal metamorphism of new snow and interpretation as a sintering process. *Journal of Geophysical Research*, 112(D24):1–10, December 2007.

[15] R. Klette and A. Rosenfeld. *Digital Geometry: Geometric Methods for Digital Picture Analysis*. Series in Computer Graphics and Geometric Modelin. Morgan Kaufmann, 2004.

[16] G. Matheron. *Eléments pour une Théorie des Milieux Poreux*. Masson, Paris, 1967.

[17] G. Matheron. *Random sets and integral geometry*. John Wiley & Sons, New York, 1975.

[18] A. Meijster, J. Roerdink, and W.H. Hesselink. A general algorithm for computing distance transforms in linear time. *Computational Imaging and Vision*, 18:331–340, 2000.

[19] L. Najman and H. Talbot. *Mathematical Morphology: From Theory to Applications*. ISTE Ltd and John Wiley & Sons Inc, 2010.

[20] D. J. Pringle, J. E. Miner, and H. Eicken. Pore space percolation in sea ice single crystals. *Journal of Geophysical Research*, pages 1–51, 2009.

[21] A Rosenfeld and J Pfaltz. Distance functions on digital pictures. *Pattern Recognition*, 1:33–61, 1968.

[22] T. Saito and J.I. Toriwaki. New algorithms for Euclidean distance transformations of an n-dimensional digitized picture with applications. *Pattern Recognition*, 27, 1551-1565, 1994.

[23] P.K. Saha, F.W. Wehrli and B.R. Gomberg. Fuzzy Distance Transform: Theory, Algorithms, and Applications. *Computer Vision and Image Understanding*, 86:171–190, 2002.

[24] P.K. Saha. Measurement of Trabecular Bone Thickness in the Limited Resolution Regime of In Vivo MRI by Fuzzy Distance Transform. *IEEE Transactions on Medical Imaging*, 23(1):53–62, 2004.

[25] P.K. Saha and B.B. Chaudhuri. Detection of 3-D simple points for topology preserving transformations with application to thinning *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 1028-1032, 1994.

[26] P. Soille. *Morphological Image Analysis*. Springer-Verlag, Berlin, Heidelberg, New York, Etats-unis, 1999.

[27] E. Remy and E. Thiel Exact Medial Axis with Euclidean Distance. *Image and Vision Computing*, 23, 167–175, 2005.

[28] S. Torquato. *Random Heterogeneous Materials: : Microstructure and Macroscopic Properties*. Springer Verlag, 2002.

[29] L. Vincent. *Granulometries and Opening Trees*, pages 57—-90. IOS Press, 40 edition, 2000.