Efficient geodesic attribute thinnings based on the barycentric diameter

Vincent Morard, Etienne Decencière, Petr Dokládál

► To cite this version:

Vincent Morard, Etienne Decencière, Petr Dokládál. Efficient geodesic attribute thinnings based on the barycentric diameter. *Journal of Mathematical Imaging and Vision*, 2013, 46 (1), pp.128-142. 10.1007/s10851-012-0374-7 . hal-00834415

HAL Id: hal-00834415

<https://hal-mines-paristech.archives-ouvertes.fr/hal-00834415>

Submitted on 14 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient geodesic attribute thinnings based on the barycentric diameter

Vincent Morard · Etienne Decencière · Petr Dokládál

Received: date / Accepted: date

Abstract An *attribute opening* is an idempotent, anti-extensive and increasing operator, which removes from an image connected components which do not fulfil a given criterion. When the increasingness property is dropped, we obtain a – more general – *attribute thinning*. In this paper, we propose efficient grey scale thinnings based on geodesic attributes.

Given that the geodesic diameter is time consuming, we propose a new geodesic attribute, the *barycentric diameter* to speed up the computation time. Then, we give the theoretical error bound between these two attributes, and we note that in practice, the barycentric diameter gives very similar results in comparison with the geodesic diameter. Finally, we present the algorithm with further optimisations, to obtain a $60\times$ speed up.

We illustrate the use of these thinnings in automated non-destructive material inspection: the detection of cracks. We discuss the advantages of these operators over other methods such as path openings or the supremum of openings with segments.

Keywords Geodesic attributes · geodesic diameter · barycentric diameter · elongation · tortuosity · thinnings · thickenings · mathematical morphology

1 Introduction

Automated optical inspection for non-destructive testing is an economically important, fast developing domain. One of its main tasks, the detection of cracks,

requires efficient image processing methods. The detection of cracks (usually long, thin and randomly tortuous structures) in the presence of noise is a challenging task.

Generally speaking, filtering of unwanted objects (e.g. noise) while preserving the desired information is a frequent pre-processing step. Mathematical morphology [12, 25, 26] is based on a set approach and classically uses structuring elements to obtain information on the morphology of the objects. In their overview of morphological filtering, Serra and Vincent [27] noticed that simple openings and closings with square, disk or hexagonal structuring elements, are often good enough for the filtering task. However, noise reduction and feature enhancement properties can be improved by adapting the shape and size of the structuring elements to the image content (see e.g. [2], [3] and [14]). Openings and closings by reconstruction are also considered as a part of shape-adaptive morphology. This led Vincent [33] to propose area openings and later Breen and Jones more general attribute openings and thinnings [4]. Later, Urbach and Wilkinson [32, 34] proposed thinnings based on the inertia to detect elongated objects in the image.

In order to enhance thin structures of random orientation, one can use the supremum of linear openings over all orientations. Although this suffices when the structures have a bounded local curvature, it fails for randomly tortuous thin cracks. In order to relax the curvature limit, Heijmans et al. [5] introduced *path openings*. Instead of rigid linear structuring elements, they use flexible paths inferred from an underlying connectivity graph. The paths are kept if, and only if, their length is longer than a given constant. Nonetheless, very tortuous structures still cannot be entirely detected.

The geodesic diameter, initially proposed by Lantuéjoul and Beucher [8], is particularly useful to measure the length of thin structures. We associated, in Morard

Centre de Morphologie Mathématique
Mathématiques et Systèmes, MINES ParisTech;
35, rue Saint-Honoré, 77305 Fontainebleau CEDEX - France
E-mail: Vincent.Morard@mines-paristech.fr
E-mail: Etienne.Decenciere@mines-paristech.fr
E-mail: Petr.Dokladal@mines-paristech.fr

et al. [13], attribute thinnings and geodesic attributes to build a new family of operators. These geodesic attribute thinnings increase the shape flexibility by removing any constraint on the tortuosity. Albeit developed to detect cracks, they can be used to detect other kinds of similar fibrous structures. Nevertheless, the geodesic diameter being computationally expensive, it is unsuitable for time-critical industrial applications.

This paper extends [13], by introducing and deeply analysing a new, fast and accurate geodesic attribute: *the barycentric diameter*. We also propose further experiments, results and an efficient algorithm, which is faster than path openings and readily usable for time-critical industrial applications.

We start by reviewing attribute thinnings and geodesic binary attributes in sections 2 and 3. Section 4 introduces the barycentric diameter. Section 5 explains the method to construct geodesic attribute thinnings. Section 6 illustrates their interests through two applications and provides a detailed comparison with path openings, initially developed to solve similar problems. Lastly, section 7 focuses on an efficient implementation, the complexity and the timings.

2 Background: attribute thinnings

Even if most definitions used in this paper can be given in a continuous domain, for practical reasons we consider a discrete domain.

Let $I : D \rightarrow \{0, 1\}$ be a binary image, with D a finite (typically rectangular) subset of $\mathbb{Z}^2$. A *binary operator* (or operator, when there is no ambiguity) is a function that transforms a binary image into another binary image. The set X contained in I is $X = \{x \in D \mid I(x) = 1\}$ and we denote X^c its complement in D . We associate to D a local neighbourhood describing the connection between adjacent pixels. In this study, each pixel is connected to its eight nearest neighbours. With this 8-connectivity, we introduce the collection of the connected components (*CC*) of X as $\{X_i\}$. From now on, by *object* we understand one connected component from the collection $\{X_i\}$.

2.1 Connected components and attributes

We wish to keep or delete the objects in an image according to intuitive *attributes* like length, tortuosity, elongation or circularity. Given some connected component X_i , these attributes will be respectively denoted $L(X_i)$, $T(X_i)$, $E(X_i)$ and $C(X_i)$. Their definitions will be given later. These attributes allow to define *criteria*

like “longer than or equal to” ($L(\cdot) \geq \lambda$), or “less tortuous than” ($T(\cdot) < \lambda$), with some $\lambda \in \mathbb{R}^+$. Formally, a criterion χ is a function mapping the set of connected components of D into $\{0, 1\}$, where 0 can be interpreted as *false* and 1 as *true*.

The binary operator based on criterion χ is then defined as:

$$\psi_\chi(X_i) = \begin{cases} X_i & \text{if } \chi(X_i) \text{ is true} \\ \emptyset & \text{otherwise,} \end{cases} \quad (1)$$

for all *CCs* X_i included in D .

2.2 Binary attributes thinnings

Based on the binary operator ψ_χ , the corresponding attribute thinning of X is

$$\rho_\chi(X) = \bigcup_i \psi_\chi(X_i). \quad (2)$$

This is equivalent to scanning, one by one, the different *CCs* of X , and either preserving them intact or removing them, depending on the criterion χ .

Attribute thinnings are anti-extensive and idempotent (see [4] for the proof). Moreover, if the criterion is increasing, then the corresponding attribute thinning is also increasing and the operator becomes an attribute opening.

The dual transform of a thinning is called a thickening, and is obtained using the complementation:

$$\delta_\chi(X) = [\rho_\chi(X^c)]^c. \quad (3)$$

In what follows, we restrict our study to thinnings for simplicity, since the computation of thickenings is straightforward with equation 3.

In the next section, we will discuss the extension of thinnings to grey level images.

2.3 Grey level attribute thinnings

A grey level image f is a mapping from D into $V = \{0, \dots, M-1\}$. Image f can be decomposed into a collection of sets obtained by thresholding. Hereafter, $X^h(f)$ denotes the threshold of f above $h \in V$:

$$X^h(f) = \{x \mid x \in D, f(x) \geq h\}, \quad (4)$$

and $X_i^h(f)$ the i -th connected component of $X^h(f)$.

Any increasing binary operator, such as an opening, can be generalised to grey level images by applying it to all the threshold sets $X^h(f)$, and stacking the results to recompose the grey-level image again [11]:

$$(\rho_\chi^{\text{direct}}(f))(x) = \vee \{h \in V \mid x \in \rho_\chi(X^h(f))\}. \quad (5)$$

Fig. 1 Results of an attribute thinning with different filtering rules. From left to right: input signal with two CCs that fulfil the criterion χ (in grey), min, max, direct and subtractive rules.

However, binary attributes are not necessarily increasing, therefore their extension to grey level images is not straightforward. Some filtering rules are reported in the literature to construct this extension [4, 23, 32], see Fig. 1:

1. **Min rule:** A connected component X_i is removed if $\chi(X_i) = 0$ or if it exists a connected component X_j such that $X_i \subset X_j$, which is removed.
2. **Max rule:** A connected component X_i is removed if $\chi(X_i) = 0$ and if all the CCs X_j such that $X_j \subset X_i$ are also removed.
3. **Direct rule:** This is the simplest rule available to extend a binary operator to grey scale images. It uses the same equation as for the opening, eq. 5.
4. **Subtractive rule:** A CC X_i is removed if $\chi(X_i) = 0$. All the other connected components such that $X_j \subset X_i$ are lowered by the value of the contrast of X_i .

Other rules have been introduced for more specific treatments: the median rule [22] is a pre-filter to improve the robustness of the decision, the Viterbi rule [23] solves an optimisation problem for the same purpose, finally, the k -subtractive and the k -absorption rules [19] work with k -flat zones. All these extensions lead to the same result when applied to increasing operators, such as openings.

The choice of the rule depends on the application. However, it is shown in [32, 31], that the subtractive rule is preferable. This is the only rule that fulfils two intuitive requirements: (i) after filtering, all the structures that do not meet the criterion are removed; and (ii) the difference image $f - \rho_\chi(f)$ contains only the structures that do not meet the criterion.

Furthermore, our final application requires the extraction of the cracks from the background and this rule is perfectly suitable for this task. Therefore, we are using the *subtractive rule* in this paper, to extend thinnings to grey scale images.

Geodesic attributes and the barycentric diameter are introduced in the two following sections. They will be used afterwards to build new attribute thinnings.

Fig. 2 Illustration of possible definitions of the length of an object. (a) and (b): the length of the segment between its ends points. (c) and (d), the measurement of its skeleton. These definitions are not always suitable.

3 Geodesic attributes

Geodesic attributes can be defined in both continuous and discrete domains. In this paper, the computations are done in $\mathbb{Z}^2$, with the 8-connectivity.

3.1 Geodesic diameter

Lantuéjoul and Maisonneuve, in [9] asked the question: “What is the length of an object?” The first idea is to measure the length of the segment connecting its end points, Fig. 2(a). However, this definition is not satisfactory, since this segment is not always included in the object, Fig. 2(b). Moreover, the definition of the end points is not trivial. Another possible measurement is the cardinal of the set of points of a homothetic skeleton of the object, Fig. 2(c). However, even though the skeleton is included in the object, it is not continuous in the sense that a slight modification of the shape can introduce important variations to the skeleton (and to its length), Fig. 2(d). Moreover, the skeleton does not necessarily span over the entire object (the skeleton of a disk is its centre).

We sum up below the exposition by Lantuéjoul and Maisonneuve and we present the generalisation of the geodesic diameter, proposed by Soille in [28].

Let X be an object (i.e. a connected component) and let $x, y \in X$. Path connectivity implies that there exists some paths between these two points, Fig. 3 (a). Among them, those with the minimal length are the *geodesic arcs* from x to y within X , Fig. 3 (b). By definition, they have all the same *geodesic length* written $d_X(x, y)$. For a given x belonging to X , we can compute the length of the longest geodesic arcs starting at x :

$$l_x(X) = \sup_{y \in X} d_X(x, y), \quad (6)$$

which corresponds to the distance to the farthest points from x , within X (for the geodesic distance).

We can compute $l_x(X)$ for all $x \in X$. Any point x of X where $l_x(X)$ reaches its global minimum is a *geodesic centre* of X . On the contrary, any couple of points where $d_X(x, y)$ reaches its global maximum are the *geodesic extremities* of X .

Fig. 3 (a) Two paths between x and y ; (b) geodesic arc between these two points; (c) longest geodesic arc of X , whose length is the geodesic diameter of X ; (d) generalised geodesic distance, longest geodesic arc.

The *geodesic diameter* of X is defined as the length of its longest geodesic arc (Fig. 3 (c)):

$$L(X) = \sup_{x,y \in X} l_x(X), \quad (7)$$

which corresponds to the largest geodesic distance between any two points of X :

$$L(X) = \sup_{x,y \in X} d_X(x, y). \quad (8)$$

The geodesic diameter gives, in our opinion, a nice estimation of the intuitive length of an object. It has, in addition, three advantages [9]: it is a general definition valid for every object; it is continuous, since a small change in the shape of the object (excluding any change of topology) will cause, at most, a small change of the measure of the geodesic diameter; finally, the computation of $L(X)$ leads to other attributes, such as the geodesic elongation and the geodesic tortuosity.

Soille in [28], has extended these notions with the generalised geodesic distance. First, he uses the Euclidean distance inside the object, inverts this distance map and finally, computes propagations with the geodesic time [29] from all the boundary points of the object. Therefore, the geodesic diameter tends to follow the medial axis of the object, like in Fig. 3 (d). This increases the accuracy of the measurement, since it measures the same length before and after bending a fibre. In practice though, very similar results are obtained for the detection of thin objects like cracks (see Fig. 8 for a comparison of both methods on a fingerprint image), but the computation times are higher ($\times 2$ at least, table 2). Therefore, we will not use a generalised geodesic distance to compute the length of the objects, unless high precision is required.

It should also be noted that the geodesic distance, and therefore the geodesic diameter, are defined in the same way in a n -dimensional space; their generalisation to higher dimensions is straightforward.

3.2 Geodesic elongation

The geodesic diameter gives a satisfactory definition of the length of an object but it does not tell much about

its shape. By combining it with the area $S(X)$ of X , we obtain some information on its elongation. The geodesic elongation, introduced by Lantujoul and Maisonneuve [9], is computed as follows:

$$E(X) = \frac{\pi L^2(X)}{4S(X)} \quad (9)$$

The longer and narrower the object is, the higher is the elongation. The lower bound is reached with the disk¹. Note that this definition can naturally be generalised to higher dimensions [20].

3.3 A new geodesic attribute: the geodesic tortuosity

We propose a new feature, derived from the geodesic diameter: the *geodesic tortuosity*. The pair of points $\{x, y\}$ is a pair of geodesic extremities of X if, and only if $d_X(x, y) = L(X)$. Note that some objects may have more than one pair of geodesic extremities (e.g. a disk). Let $\Xi(X) = \{\{x_0, y_0\}, \{x_1, y_1\}, \dots\}$ be the set of geodesic extremities of X . Let $L_{Eucl}(X)$ denote the minimal Euclidean distance between geodesic extremities:

$$L_{Eucl}(X) = \min_{\{x,y\} \in \Xi(X)} \|x, y\|. \quad (10)$$

We define the geodesic tortuosity as the ratio between the geodesic diameter and $L_{Eucl}(X)$:

$$T(X) = \frac{L(X)}{L_{Eucl}(X)}. \quad (11)$$

In the case of the two-dimensional Euclidian space, the geodesic tortuosity of convex objects is equal to one.

3.4 Geodesic attributes properties and comments

Provided a rotation-invariant metric is used for the distances, e.g. the $\mathcal{L}_2$, all these attributes are rotation invariant in a Euclidean space; the geodesic elongation and tortuosity attributes are moreover scale invariant.

One can derive other attributes from the geodesic diameter, such as the geodesic circularity, which is the inverse of the geodesic elongation: $C(X) = \frac{1}{E(X)}$. It is also scale and rotation invariant. It reaches its upper bound 1 with the disk. Moreover, we can combine two attributes together to emphasise different structures. For instance, the product between the geodesic diameter and the geodesic tortuosity emphasises long and tortuous structures: $LT(X) = L(X)T(X)$.

¹ The normalisation factors are such that in a two-dimensional Euclidian space, the geodesic elongation of a disk is equal to 1

However, the computation of the geodesic diameter is computationally intensive and we are interested in efficient implementation of the thinnings based on geodesic attributes. Therefore, we introduce in the next section the *barycentric diameter* – a new attribute, approximating the geodesic diameter, but much faster to compute.

4 Barycentric diameter

The aim of the barycentric diameter is to replace the geodesic diameter by a similar measure with a lower complexity. We first describe the available algorithms for computing the geodesic diameter. Then, we introduce the new attribute.

4.1 Review on the computation of the geodesic diameter

The direct approach to compute the geodesic diameter of an object X consists of computing, for each point of the border, the geodesic distance within X to all other points of X (starting from one pixel, this operation is called a *propagation*). The highest geodesic distance computed, is the geodesic diameter. The timings of this algorithm are high and depend on both the area of the object and the number of boundary pixels. We notice that there are many redundant propagations, leading to an inefficient implementation. Schmitt [24] showed that it is enough to consider a subset of the border points as starting points for the propagations. However, despite the important speed-up thus achieved, still too many propagations remain to compute.

Maisonneuve and Lantuéjoul designed an efficient parallel algorithm to compute the geodesic diameter in a hexagonal grid [9]. Let X be a non-porous connected component (without “holes”), and Y the set of border points of X . Using a particular propagation in the hexagonal grid, starting from Y , the algorithm gives the geodesic diameter in a single propagation. However, this algorithm requires that X is not porous. Otherwise, the propagation never stops, turning infinitely around the holes in X . This characteristic makes the algorithm inadequate for our application. Indeed, a group of cracks can represent a porous connected object, e.g. Fig. 10(a).

Given that we have not found a method to efficiently compute the geodesic diameter, we propose what we consider a convenient - an efficient - approximation of the geodesic diameter: the barycentric diameter.

4.2 Definition of the barycentric diameter

For any point x of X , we can compute the length $l_x(X)$ of the longest geodesic arc starting at x in a single propagation. Could this value replace $L(X)$? Is it a good approximation? How to choose a convenient x ?

Let us consider a maximal geodesic arc of X , and let y_1 and y_2 be its extremities. This means that $d_X(y_1, y_2)$ is equal to $L(X)$. Given that the geodesic distance is a distance [9], thanks to the triangular inequality we have, for any x belonging to X :

$$d_X(y_1, y_2) \leq d_X(y_1, x) + d_X(x, y_2). \quad (12)$$

Moreover, by definition $l_x(X)$ is larger than both $d_X(y_1, x)$ and $d_X(x, y_2)$, and smaller than $L(X)$, therefore we finally obtain:

$$L(X) \geq l_x(X) \geq \frac{L(X)}{2}. \quad (13)$$

Thus, the relative error $(L(X) - l_x(X))/L(X)$ obtained when approximating $L(X)$ with $l_x(X)$ is smaller than 50%.

In order to improve this approximation, it is tempting to iterate the propagation, starting this time from the farthest points from x . If Y is the subset of X containing the farthest points from x :

$$Y = \{y \mid y \in X, d_X(x, y) = l_x(X)\}, \quad (14)$$

we can introduce the iterated maximal geodesic distance starting from x , defined as:

$$l_x^2(X) = \sup_{y \in Y} l_y(X). \quad (15)$$

One can easily show that:

$$L(X) \geq l_x^2(X) \geq l_x(X) \geq \frac{L(X)}{2}, \quad (16)$$

and more generally that $l_x^n(X)$ converges, $\forall n \in \mathbb{Z}^+$:

$$L(X) \geq l_x^n(X) \geq l_x^{n-1}(X) \geq \frac{L(X)}{2}. \quad (17)$$

Nonetheless, as it will be seen below, the limit is not necessarily equal to $L(X)$. In practice though, we observe that the convergence is fast, and that going beyond $l_x^2(X)$ is not interesting. We will come back to this point at the end of the section. Fig. 4 provides an example where $l_x^2(X)$ gives a much better approximation than $l_x(X)$, and where further iterations do not improve the result.

In the following, the approximations of $L(X)$ will be based on $l_x^2(X)$. But we still have to choose x .

Does the choice of x within X have an influence on the quality of the approximation? Experiments have shown that indeed it is the case. Several strategies have been tested:

Table 1 Summary of the relative error between the geodesic diameter and 5 approximations of this geodesic diameter. These statistics are computed on 51400 binary shapes on total. The automatic generation of these shapes is presented in the Appendix.

Methods		L_{Bar}	$L_{BarNearest}$	$L_{GeoCentre}$	$L_{GeoCentreNearest}$	L_{Random}
Convex objects	Mean (%)	0.24	0.24	0.47	0.56	1.92
	Std (%)	0.85	0.84	1.41	1.53	3.36
	Max (%)	10.04	11.20	15.9	17.24	22.59
Pixel aggregation	Mean (%)	0.43	0.40	0.74	0.85	2.68
	Std (%)	1.09	1.03	1.58	1.71	3.59
	Max (%)	10.76	9.27	12.87	12.87	24.59
Ball aggregation	Mean (%)	0.22	0.22	0.23	0.59	1.21
	Std (%)	0.72	1.26	1.52	2.31	4.53
	Max (%)	7.17	12.75	13.84	16.85	28.83
Random walk	Mean (%)	0.13	0.28	0.24	0.53	1.31
	Std (%)	0.66	1.03	1.08	1.61	3.27
	Max (%)	9.97	21.7	21.48	21.48	29.96
Perlin noise	Mean (%)	0.12	0.16	0.21	0.40	0.76
	Std (%)	0.74	0.80	1.21	1.50	2.61
	Max (%)	10.87	10.14	16.55	13.61	31.90
Database MPEG7	Mean (%)	0.25	0.43	0.34	0.58	0.66
	Std (%)	1.12	1.93	1.44	1.91	2.38
	Max (%)	20.36	22.31	19.16	19.21	30.11
All Objects	Mean (%)	0.23	0.29	0.37	0.58	1.42
	Std (%)	0.86	1.49	1.37	1.76	3.29
	Max (%)	20.36	22.31	21.48	21.48	31.90

Fig. 4 A black and white cherry: illustration of the barycentric diameter. Point B is barycentre of the CC and the starting point of the first propagation (x) is the farthest point from B. Then, the first propagation gives l_x and the second propagation, starting from y , leads to a much better approximation of the geodesic diameter: l_x^2 .

- When x is one of the farthest points from the barycentre of X , we obtain what we call the barycentric diameter $L_{Bar}(X)$. Note that the barycentre of X does not necessarily belong to X , so we use the Euclidean distance when looking for these farthest points. By taking one of the farthest points from the barycentre, we suppose that it will be close to an extremity of X (even though it is not always the case for non-convex shapes). Alternatively, instead of taking one of the farthest points from the barycentre, we can take one of the closest, thus obtaining another approximation: $L_{BarNearest}(X)$.
- A geodesic centre of X can also be used as reference point, instead of the barycentre. We obtain then two other approximations, $L_{GeoCentre}(X)$ and

$L_{GeoCentreNearest}(X)$. Note that this strategy is only proposed for comparison; in practice, computing the geodesic centre is computationally intensive.

- Finally, we have also considered as starting point the first point of the object found with a raster scan, called $L_{Random}(X)$.

Therefore, five methods are available to approximate the geodesic diameter with only two propagation steps. To compare them, we apply them to six different sets of binary objects and we compute the geodesic diameter as well as the approximations defined above. For each object, we store $L(X)$ and $l_x^2(X)$ and we finally compute some statistics for each database : the mean relative error $(L(X) - l_x^2(X))/L(X)$, its standard deviation, and its maximum.

Five sets correspond to realisations from a random connected component model. Each one contains 10000 realisations. We give in the Appendix further explanations on the automatic generation of these objects. The sixth set is a standard database of 1400 binary objects (MPEG7 CE Shape-1 Part B database²).

Table 1 summarises the results. A first general observation concerns the fact that relative errors are, in practice, much lower than the 50% theoretical limit. All mean errors are in fact smaller than 1%. Even the maximal error is far from reaching the theoretical limit. Note also that the barycentric diameter achieves the smallest

² Database available in www.imageprocessingplace.com/root_files_V3/image_databases.htm

error for most sets of objects: in average, the error is at least reduced by 20% compared to the other methods.

The maximal relative error on all 51400 objects obtained with the barycentric diameter is only 20%. Can the 50% relative error theoretical limit be reached? In Fig. 5 we exhibit an example, where $L_{Bar}(X)$ can be arbitrarily close to $\frac{L(X)}{2}$. Note also that this counter-example proves that $l_x^n(X)$ does not necessarily converge towards $L(X)$. Indeed, in this case the limit is $L_{Bar}(X)$, reached after only one iteration.

Fig. 5 Example of an object X , where the barycentric diameter does not converge towards the geodesic diameter. The geodesic diameter is equal to the distance $d_X(E, B) + d_X(B, C) + d_X(C, D) + d_X(D, F) = 4r + 2\epsilon$ whereas the barycentric diameter is equal to $d_X(A, B) + d_X(B, C) = 2r$. The relative error is : $\lim_{\epsilon \rightarrow 0} \frac{L - L_{Bar}}{L} = \lim_{\epsilon \rightarrow 0} \frac{r - \epsilon}{2r - \epsilon} = 0.5$

Concerning the convergence speed of the different approximations, we note that in practice, the barycentric diameter shows also the best performance, as illustrated in Fig. 6. This figure presents the convergence rate of the different approximations on the objects of the MPEG7 database. Note also that these curves show that, as stated previously, the second iteration brings an important improvement, but that further steps do not seem necessary.

We observe that a high relative error is reached when the propagation starts from a geodesic centre and converges towards another geodesic centre. Now, using the farthest point from the barycentre usually avoids a geodesic centre as starting point of the propagation. This observation provides us with some clues to understand why the barycentric diameter is, in practice, close to the geodesic diameter.

To conclude this section, the barycentric diameter appears to be, in practice, a very good and efficient replacement of the geodesic diameter. Hereafter, geodesic elongations and geodesic tortuosities are computed using L_{Bar} instead of L . In the next section, new thinnings, based on geodesic attributes, are introduced.

Fig. 6 Convergence of the geodesic diameter approximations. The relative error (in % and in log scale) is considerably reduced with 2 propagations (from l_x to l_x^2). However increasing the number of iterations beyond 2 does not bring a huge improvement on the accuracy of the measurement. Therefore, we choose l_x^2 with x being the farthest point from the barycentre. These curves are computed using all the objects of the MPEG7 database.

5 Geodesic attribute thinnings

To the extent of the authors knowledge, this is the first time that the geodesic diameter, its approximation L_{Bar} , as well as the derived attributes (geodesic elongation and tortuosity) are used to build attribute thinnings.

We recall that ρ_χ stands for the attribute thinning using criterion χ . We will consider criteria such as $L(X) \geq \lambda$, where λ is some real value, chosen according to the application.

5.1 Results on binary images

Fig. 7 is a toy example containing thin, tortuous objects, and other non-thin objects. We want to separate these structures with the following criteria:

- Suppress objects whose barycentric diameter value is smaller than 80 pixels and smaller than 120 pixels (second column);
- Suppress objects which are not elongated, e.g. whose geodesic elongation is smaller than 4 and smaller than 7 (third column);
- Suppress all non tortuous objects, i.e. whose geodesic tortuosity is smaller than 1.3 and smaller than 2 (fourth column);
- Suppress all objects having a circularity measurement smaller than 0.4 and smaller than 0.6 (fifth column).

We can observe through this simple experiment that geodesic attribute thinnings are intuitively parametrised, and provide an interesting tool to classify objects.

Fig. 7 Filtering with geodesic attributes criteria: (first column) initial image, (second column) barycentric diameter thinnings, (third column) geodesic elongation thinnings, (fourth column) geodesic tortuosity thinnings and (fifth column) geodesic circularity thinnings.

Fig. 8 Illustration of the thinnings based on geodesic attributes on a grey scale image. The first line uses classical geodesic distance and the second line presents the results of thinnings based on the generalised geodesic distance [28], written with a * (e.g. L^* , E^* , T^* and C^*)

5.2 Grey level images

We present some results on the application of geodesic attribute thinnings to a grey scale image. Fig. 8 shows a fingerprint with long and tortuous structures, and the resulting images after application of thinnings with length, elongation, tortuosity and circularity criteria, using the subtractive rule. Different structures are enhanced and we can make the distinction between long, elongated or tortuous structures with these filters. Finally, the circularity attribute naturally erases most structures from the image.

In section 3, we have presented both the classical geodesic distance [9] and the generalised geodesic distance [28]. We compare, in Fig. 8, thinnings with both methods and we note the similarity of the results; the main visible difference corresponds to a fingerprint ridge

around the center of the tortuosity image. The generalised geodesic distance is more accurate, but the computation times are longer (see Tab. 2). According to whether one prefers speed or accuracy, one may choose the classical or the generalised geodesic distance.

6 Results

First, we propose an application³ where geodesic attribute thinnings are particularly adapted. Fig. 9(a) shows the image of a DNA molecule acquired with an electron microscope⁴. These structures are long, thin,

³ A demonstration version is available <http://cmm.enscm.fr/~morard/DemoGeoThinnings.html>

⁴ Image from the Institute for Molecular Virology. University of Wisconsin - Madison <http://www.biochem.wisc.edu/faculty/inman/empics/dna-prot.htm>

and extremely tortuous. In order to evaluate the results we have visually compared them with two other methods, classically used for this sort of application:

1. The supremum of openings (Fig. 9(b), 60-pixel long segments oriented every 2 degrees) only preserves straight parts, and fails with tortuous structures. Note that the length λ is very low compared to the other methods. Longer segments (100-pixel long) would completely remove the molecule. A smaller length would preserve the background noise.
2. The path opening, see Heijmans et al. [5], (Fig. 9(c), length 160 pixels), tolerates some tortuosity. It yields better result than the previous method. The noise is considerably reduced but some parts of the molecule are also discarded. This operator is not able to follow the entire structure, and therefore underestimates its length.
3. The barycentric diameter thinning (Fig. 9(d), size 600 pixels) yields a better result, since the molecule is correctly extracted. It relaxes any constraints on the tortuosity, and offers a better flexibility than path openings.
4. The geodesic elongation thinning (Fig. 9(e), elongation superior to 50), filters out all the noise and offers a very efficient detection. Any other non elongated structures are similarly deleted.
5. The geodesic tortuosity thinning (Fig. 9(f), tortuosity superior to 2.5), filters out every structure that is not tortuous. Hence, this molecule is easily extracted with this tool.

These attribute thinnings have been used in an industrial non-destructive material-inspection application to detect long, narrow and randomly tortuous cracks. Fig. 10 offers a crack enhancement example with the same operators. Similar conclusions on the results can be pointed out: geodesic thinnings yield the best detection. The tortuosity thinning (Fig. 10(f)) enhances the crack as well as some noise in the background image. However, the noise has a lower grey level value than the crack and a simple threshold suffices to extract the crack.

7 Algorithm, practical considerations and optimisation

We describe the algorithm used to build geodesic attribute thinnings. We compare it with the algorithm previously used by the authors [13], where the geodesic diameter is computed exhaustively from the contour points. Furthermore, we propose a comparison with path openings [5].

7.1 Geodesic attribute thinnings algorithm

Attribute filters are often implemented using a tree representation of the image called *max-tree* (see e.g. Salembier et al. [23]). The max-tree creation relies on a recursive flooding procedure starting from the lowest pixel in the image. Its worst-case time complexity is quadratic $\mathcal{O}(N \times M)$, with M the cardinal of the set of values and N the number of pixels of the image, see [15]. However, the worst case is rare and in practice, Salembier *et al.* algorithm is faster than the algorithm of Najman and Couprie [15] for 8 bits images, even if the construction of max-tree is done in quasi linear time with the size of the image, using an union find approach. Other efficient implementations exist, see e.g. Ngan et al. [16].

A max-tree structure is adapted for attributes that can be *updated* each time a new pixel is aggregated to a CC. This is not the case of the barycentric diameter which needs to be *recomputed* and needs to access all the pixels of a connected component. This requires some modifications of the max-tree algorithm and prohibits the recursion. The algorithm Najman and Couprie [15] maintains several tree-like structures at a time and accessing the pixels is uneasy either.

Instead, we start with the relief completely submerged by water, and let the water progressively sink. As soon as appears the first (global) maximum, its connected component is progressively reconstructed and tested on χ . When other local maxima appear, they are not process yet and we reconstruct all connected component at lower threshold sets that are supersets of the global maxima (see Fig. 11 component 1). Finally, other local maxima are processed in the same way but the aggregation stops when a CC is already processed (Fig. 11 components 2 and 3).

Fig. 11 Illustration of the algorithm principle for a one-dimensional signal with three local maxima. The components in light grey (marked by one) are analysed first, followed by the components marked by 2 and 3.

The Alg. 1 simulates this principle. It uses a priority queue *HQueue* that supports operations (modifying the content), and queries (not modifying the content): the operation *Push*(x, p) inserts an element x with a given priority p , and $(x, p) = \text{Pop}()$ retrieves the currently highest-priority element x and its priority p . The

Fig. 9 DNA molecule extraction: we compare five different methods to detect this thin structure. We note that path openings and segment opening completely underestimate the length of this tortuous structure. By opposition, geodesic attribute thinnings easily enhance this molecule and correctly evaluate its length.

Fig. 10 Crack detection: to detect these thin structures, we use five different methods. Geodesic attribute thinnings yield the best enhancement.

queries: $p = \text{Prio}()$ returns the currently highest priority, and $\text{Empty}()$ tests the emptiness of $HQueue$.

In the beginning, the output image g is initialised by the input f , and the local maxima of f are extracted by a function FindLocalMax (code lines 4 and 5). The list $LocalMax$ contains one pixel per every maximum, and is sorted by decreasing order of the level.

The *for* cycle (lines 8 to 28) runs over all the local maxima to process. Line 9 reads from $LocalMax$ the position p of next maximum to process. The *repeat-until* cycle (lines 14 to 28) simulates the decreasing water level to progressively reconstruct the associated component associated, and tests it against χ (line 26). The extraction of the component uses two auxiliary flag arrays $State : D \rightarrow \mathcal{Z}$ and $Level : D \rightarrow V$ (lines 2 to 3), initialised to 0 for all image pixels (lines 6 to 7). $State$ receives -1 whenever a pixel is pushed in the queue (line 11), and later the index of the local maximum to which component it belongs to (line 18), cf. Fig. 11.

We initialise the region growing algorithm by pushing p in the hierarchical queue with the priority $f(p)$ (line 10), by marking this pixel (lines 11 to 12) and by initialising CC to the empty set. Then, a region-growing algorithm (code lines 15 to 25) extracts a connected component CC to which the pixel p belongs to. Criteria χ based on the barycentric diameter L_{Bar} need to access the connected component as a set. When the region growing at this level is finished (line 25), CC contains all pixels of this connected component such that $f(x) \geq h, \forall x \in CC$. Then, if the test $\chi(CC)$ is false, we apply the subtractive rule by removing the contrast of this connected component for all pixels belonging to CC . The function $\text{contrast}(CC)$ returns the difference between the level h and the highest grey level of the neighbours of CC .

The same process is repeated until all regional maxima are processed.

The Fnct. L_{Bar} illustrates how to compute the barycentric diameter when χ is based on this attribute. First, it computes the barycentre $\bar{p}$ of CC , i.e. the mean of the coordinates (code line 1). Then, it searches one of the farthest point from the barycenter with the Euclidean distance (line 2). Finally, it computes twice the geodesic distance inside CC (see Eq. 6, and lines 3 and 4), and returns the result l_x^2 of the second iteration.

Note : Algorithm 1 is set up to compute the subtractive rule. However, if one wants to apply the direct rule, only two modifications of the algorithm are required: the output image is initialised with 0 and if $\chi(CC)$ is true, we write the CC in g by taking the maximum value between h and g .

Algorithm 1: $g \leftarrow \text{GEOATTRTHINNING}(f, \chi)$

```

Input:  $f : D \rightarrow V$  - input image
           $\chi$  - criterion
Result:  $g : D \rightarrow V$  - output image

/* Declare auxiliary variables */
1  $HQueue = \{\}$  /* Init  $HQueue$  to empty */
2  $State : D \rightarrow \mathcal{Z}$ 
3  $Level : D \rightarrow V$ 

/* Initialize */
4  $g = f$ 
5  $LocalMax = \text{FindLocalMax}(f)$  /* get sorted list of local maxima */
6  $State(:) = 0$ 
7  $Level(:) = 0$ 

/* Run over all local maxima */
8 for  $i = 1 \dots \text{Card}(LocalMax)$  do
9    $p = LocalMax(i)$ 
10   $HQueue.Push(p, f(p))$ 
11   $State(p) = -1$  /* pixel enqueued */
12   $Level(p) = f(p)$ 
13   $CC = \{\}$ 

/* Simulate decreasing water level */
14 repeat
15   /* Extract a CC to test on  $\chi$  */
16   repeat
17      $(q, h) = HQueue.Pop()$ 
18      $CC = CC \cup \{q\}$ 
19      $State(q) = i$ 
20     foreach  $n \in \text{Neighbour}(q)$  do
21       if  $f(n) > Level(n)$  and  $State(n) \neq i$  then
22          $hmin = \min(h, f(n))$ 
23          $HQueue.Push(n, hmin)$ 
24          $State(n) = -1$  /* pixel enqueued */
25          $Level(n) = hmin$ 
26   until  $h < HQueue.Prio()$ 
27   if not  $\chi(CC)$  then
28      $g(CC) = g(CC) - \text{contrast}(CC)$ 
29 until  $HQueue.Empty()$ 

```

Function $l_x^2 \leftarrow L_{Bar}(CC)$

```

Input:  $CC$  - A connected component
Result:  $l_x^2$ : Barycentric diameter of  $CC$ 
1  $l_x^2 \leftarrow L_{Bar}(1)l_x^2 \leftarrow L_{Bar}$ 
2  $b = \bar{p}, p \in CC$ 
3  $x = \arg \max_{y \in CC} ||y, b||$ 
4  $x^1 = \arg \sup_{y \in CC} d_{CC}(x, y)$ 
5  $x^2 = \arg \sup_{y \in CC} d_{CC}(x^1, y)$ 
6  $l_x^2 = d_{CC}(x^1, x^2)$ 
7 return  $l_x^2$ 

```

7.2 Complexity

The complexity analysis is split into two parts; the complexity of the CC's extraction (Alg 1) and the complexity of the computation of the attribute.

Table 2 Running times for different images of size 256x256 for a geodesic diameter thinning or thickening for $\lambda = 20$. See section 7.4 for details. Timings are in seconds. Laptop computer: Intel Core2 Duo T7700 @ 2.40GHz

Images	From [13]		From this paper	
	$L_{no\ opt}$	L_{opt}	L_{Bar}	L_{Bar}^*
Coffee	36	0.079	0.0028	0.0062
Eutectic	37	0.015	0.0029	0.0069
Grains	3650	5.4	0.041	0.328
Macula	3850	2.7	0.031	0.220
Relief	1024	0.99	0.022	0.139
Retina	1820	1.56	0.024	0.136

Fig. 12 Images used to build table 2. They have all the same size: 256×256 pixels.

First, we detail the complexity of Alg. 1. The function *FindLocalMax* fills the array *LocalMax* using a region growing process with a queue, and sorts the array *LocalMax* with a linear complexity radix sort [17], which gives a complexity of $\mathcal{O}(N)$ for this function ($N = \text{Card}(D)$). Finally, from the local maxima, we use a relief emerging approach using a hierarchical queue, which adds $\log(M)$ in the complexity. Every local maximum is processed sequentially and the worst case is a highly nested signal. Then, for the worst signal, a pixel is analysed $N/4$ times in average, which leads to a complexity of $\mathcal{O}(N^2 \log(M))$, just like the algorithm presented by Breen and Jones [4]. We note however, that the computation times are very far from this upper bound⁵ and in practice, we observe a linear evolution of the computation time with the number of pixels.

All geodesic attributes are based on the length measurement. Therefore, all attributes are computed with the complexity of $\text{Fcnt } l_x^2 \leftarrow L_{Bar}$.

For a connected component having n pixels, the computation of the attribute value is done in 4 steps. The barycenter and the farthest point both require one scan over all pixels, giving $\mathcal{O}(n)$. Then, two propagations are performed to get the barycentric diameter. For the propagation step, according whether $\mathcal{L}_1$ or $\mathcal{L}_2$ is used for the distance, one needs a queue (or a priority queue, resp.) giving a complexity of $\mathcal{O}(1)$ (or $\mathcal{O}(\log(m))$ resp.) per pixel, with m the mean number of elements in the priority queue.

This is an improvement in comparison with the complexity of the geodesic diameter. With a propagation from all boundary points, the complexity is $\mathcal{O}(n^2)$ (the worst case is a thin line, which needs n propagations). With the Schmitt's method, the number of boundary points is reduced however, the worst case still requires $\sqrt{n}$ propagations (for a disc).

⁵ We made an experiment on more than 100 natural images showing that a pixel is processed less than two times in average.

7.3 Optimisations

The first and probably, the most important optimisations is that we can stop the propagation step whenever the criterion value is reached. This acceleration is only available for the barycentric diameter, or the geodesic elongation and it does not change the complexity of the algorithm.

The acceleration also comes from other optimisations, which are computer's tricks:

- We compute simple attributes such as the barycentre or the area during the region growing process, in order to avoid useless scan of all the pixels of CC in $\text{Fcnt } l_x^2 \leftarrow L_{Bar}$.
- We also compute the contrast of the CC during the region growing process.
- In $\text{Fcnt } l_x^2 \leftarrow L_{Bar}$, we remove the distance map during the propagation step. Instead, we enqueue the distance value with its position whenever a pixel is stored in a queue or in the hierarchical queue.
- We manage our own queue, which is specialised for this algorithm. Therefore, we can remove all the elements from a queue in only one assignment.
- For 8 bits images, we can merge the buffers *State* and *Level*. Then, we use the first byte to store the buffer *Level* and the other bytes for the variable *State*. Then, we get the value of two variables in only one access of the memory.

These optimisations does not change the complexity of this algorithm, however, the timings have been reduced by a very large factor, as shown in the next section.

7.4 Timings

We compare in this section all the benefits brought by the optimisations, accelerations presented in section 7

and the barycentric diameter (section 4). Table 2 collects the timings for the images of Fig. 12, where we compare four methods: a thinning based on the geodesic diameter with no optimisation ($L_{no\ opt}$), the geodesic diameter with the stop of the propagation when the front wave is larger than λ (L_{opt}), the barycentric diameter with all the discussed optimisations (L_{Bar}) and the barycentric diameter using the generalised geodesic distance (L_{Bar}^*).

Timings have been reduced by a large factor between thinnings using L_{opt} and L_{Bar} (around 60 in average for these 6 images). Moreover, using a generalised geodesic distance increases the computation time by at least a factor of 2, in comparison with L_{Bar} . The overhead is introduced by the computation of the distance map for every connected component. This distance map also cuts down the benefits of the optimisation, where the propagation step is stopped if the criterion is fulfilled during the propagation.

7.5 Efficiency comparison with path openings

Path openings [5] have been developed to solve applications similar to those treated in this paper. In section 6, it has been shown that they provide indeed interesting results, even if they cannot detect objects as tortuous as those detected by geodesic attributes thinnings. For the comparison to be complete, we now compare their computational performance.

In Fig. 13, we plot a comparison between the thinnings based on geodesic attributes (barycentric diameter) and an efficient implementation of path openings [30, 6]. The timings were computed on the image presented in Fig. 10(a). On this crack image (768x576 pixels), thinnings are faster than path openings and this tendency is the same for other images.

Finally, we can state that the resulting implementation of geodesic attribute thinnings is at least as efficient as other operators that have been designed with similar objectives, which makes them interesting for a large number of industrial applications.

8 Conclusion and future work

A new geodesic attribute is introduced, the barycentric diameter to approximate the costly geodesic diameter. We give the theoretical upper error bound of this approximation, and show that the typical error is negligible.

This, as well as other shape attributes, such as the geodesic length, elongation and circularity, is combined

Fig. 13 Benchmark on a crack image (768x576 pixels, 8 bits) with regard to the size of the attribute, for path openings and thinnings based on the barycentric diameter with the subtractive rule.

with attribute thinnings to obtain a new family of operators. They are ideal for characterising long, elongated or tortuous objects. We show that their shape flexibility is superior to that of classical linear openings or path openings.

We use the subtractive filtering rule, proposed by Urbach and Wilkinson [32], and we provide an efficient algorithm for computing the resulting operators. We report competitive timings, which allow using these operators in time-critical, industrial applications. The geodesic attribute thinnings are not only more flexible than path openings, but also faster. Furthermore, geodesic attribute thinnings readily generalise to 3-dimensional images, by only changing the connectivity.

At present, geodesic thinnings fail to extract thin, curvilinear structures disconnected by noise. Indeed, their length will be underestimated. A possible solution shall consist in using a generalised connectivity, like the second generation connectivity introduced by Ouzounis and Wilkinson [18].

Future work will concern the generalisation of the proposed algorithm to efficiently compute other operators based on geodesic attribute thinnings. On the one hand, morphological pattern spectra [10] estimate the size and shape distribution of the searched structures. On the other hand, ultimate openings [1, 7] extract structures with the highest contrast. Typically, these operators require the computation of a family of thinnings of increasing size. Using this algorithm, we can compute pattern spectra and ultimate thinnings within only one “relief emerging” process.

Acknowledgements This work was made possible thanks to the support of the “Pôle ASTech” and the “Pôle Nucléaire

de Bourgogne”, and has been financed by the French “Département de Seine et Marne”

The authors are grateful to Ms Raviart, working in the “Centre des matériaux, MINES ParisTech”, for help with the optical microscope.

The authors also wish to thank the anonymous reviewers for their comments.

References

1. Beucher, S.: Numerical residues. *Image and Vision Computing* **25**(4), 405–415 (2007)
2. Beucher, S., Bosseville, J.M., Lenoir, F.: Traffic spatial measurements using video image processing. *Intelligent Robots and Computer Vision, Proc. SPIE* **848**, 648–655 (1987)
3. Braga-Neto, U.M.: Alternating sequential filters by adaptive-neighborhood structuring functions. In: *Mathematical Morphology and Its Applications to Image and Signal Processing*, vol. 5, pp. 139–146. Maragos, Schafer and Butt, Atlanta (1996)
4. Breen, E.J., Jones, R.: Attribute openings, thinnings, and granulometries. *Computer Vision and Image Understanding* **64**(3), 377–389 (1996)
5. Heijmans, H., Buckley, M., Talbot, H.: Path openings and closings. *Journal of Mathematical Imaging and Vision* **22**(2), 107–119 (2005)
6. Hendriks, C.L.L.: Constrained and dimensionality-independent path openings. *Image Processing, IEEE Transactions on* **19**(6), 1587–1595 (2010)
7. Hernández, J., Marcotegui, B.: Ultimate attribute opening segmentation with shape information. *Mathematical Morphology and Its Application to Signal and Image Processing* pp. 205–214 (2009)
8. Lantuejoul, C., Beucher, S.: On the use of the geodesic metric in image analysis. *Journal of Microscopy* **121**(1), 39–49 (1981)
9. Lantuejoul, C., Maisonneuve, F.: Geodesic methods in quantitative image analysis. *Pattern Recognition* **17**(2), 177–187 (1984)
10. Maragos, P.: Pattern spectrum and multiscale shape representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **11**(7), 701–716 (1989)
11. Maragos, P., Ziff, R.D.: Threshold superposition in morphological image analysis systems. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **12**(5), 498–504 (1990)
12. Matheron, G.: *Random sets and integral geometry*, vol. 1. Wiley New York (1975)
13. Morard, V., Decencire, E., Dokldal, P.: Geodesic attributes thinnings and thickenings. In: *Mathematical Morphology and Its Applications to Image and Signal Processing*, vol. 6671, pp. 200–211. Springer (2011)
14. Morard, V., Decencire, E., Dokldal, P.: Region growing structuring elements and new operators based on their shape. In: *Signal and Image Processing*, pp. 78–85. ACTA Press (2011)
15. Najman, L., Couprie, M.: Building the component tree in quasi-linear time. *Image Processing, IEEE Transactions on* **15**(11), 3531–3539 (2006)
16. Ngan, N., Dokldalov, E., Akil, M., Contou-Carrre, F.: Fast and efficient FPGA implementation of connected operators. *Journal of Systems Architecture* **57**(8), 778 – 789 (2011)
17. Nilsson, S.: Radix sorting and searching. Ph.D. thesis, Department of Computer Science, Lund University, Lund, Sweden (1996)
18. Ouzounis, G.K., Wilkinson, M.H.F.: Mask-based second-generation connectivity and attribute filters. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **29**(6), 990–1004 (2007)
19. Ouzounis, G.K., Wilkinson, M.H.F.: Hyperconnected attribute filters based on k-flat zones. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **33**(2), 224–239 (2011)
20. Parra-Denis, E.: *Analyse morphologique 3D de particules de forme complexes: application aux intermtalliques dans les alliages d’aluminium*. Ph.D. thesis, Universit Jean Monnet de Saint-Etienne (2007)
21. Perlin, K.: Course in advanced image synthesis. In: *ACM SIGGRAPH Conference*, vol. 18 (1984)
22. Salembier, P., Oliveras, A., Garrido, L.: Motion connected operators for image sequences. In: *European Signal Processing Conference*, vol. 96, pp. 1083–1086 (1996)
23. Salembier, P., Oliveras, A., Garrido, L.: Antiextensive connected operators for image and sequence processing. *Image Processing, IEEE Transactions on* **7**, 555–570 (1998)
24. Schmitt, M.: *Des algorithmes morphologiques l’intelligence artificielle*. Ph.D. thesis, Ecole Nationale Suprieure des Mines de Paris (1989)
25. Serra, J.: *Image analysis and mathematical morphology*, vol. 1. Academic Press, London (1982)
26. Serra, J.: *Image analysis and mathematical morphology*, vol. 2 Theoretical Advances. Academic Press, London (1988)
27. Serra, J., Vincent, L.: An overview of morphological filtering. *Circuits Syst. Signal Process.* **11**(1), 47–108 (1992)
28. Soille, P.: Generalized geodesic distances applied to interpolation and shape description. *Mathematical morphology and its applications to image processing* pp. 193–200 (1994)
29. Soille, P.: Generalized geodesy via geodesic time. *Pattern Recognition Letters* **15**(12), 1235–1240 (1994)
30. Talbot, H., Appleton, B.: Efficient complete and incomplete path openings and closings. *Image and Vision Computing* **25**(4), 416–425 (2007)
31. Urbach, E.R., Roerdink, J.B.T.M., Wilkinson, M.H.F.: Connected shape-size pattern spectra for rotation and scale-invariant classification of gray-scale images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **29**(2), 272–285 (2007)
32. Urbach, E.R., Wilkinson, M.H.F.: Shape-only granulometries and grey-scale shape filters. In: *Proc. Int. Symp. Math. Morphology (ISMM)*, vol. 2002, pp. 305–314 (2002)
33. Vincent, L.: Grayscale area openings and closings, their efficient implementation and applications. In: *First Workshop on Mathematical Morphology and its Applications to Signal Processing*, pp. 22–27 (1993)
34. Wilkinson, M.H.F., Westenberg, M.A.: Shape preserving filament enhancement filtering. *Medical Image Computing and Computer-assisted Intervention, lecture notes in computer science* **2208**, 770–777 (2001)

Appendix: random connected component models

Many strategies exist to generate random binary shapes. The results of the comparison between the geodesic diameter and its approximations are linked to the method used to generate these shapes. Therefore, we use five different methods to have a high variety of objects. For each method, 5 realisations are presented in Fig. 15. The size of the support of these random shapes is a 500 by 500 pixels square.

Convex shape

A random number of points (between 10 and 100) are randomly and uniformly picked on D . The final connected component is the convex hull of these points.

Pixel aggregation

This method is used to generate relatively dense objects, which are almost convex. The set is initialised with a single point. At each iteration, a randomly chosen neighbour point is added to the set. The procedure is iterated a random number of times.

Ball aggregation

This method uses the same process as the pixel aggregation method, except that we iteratively aggregate a ball instead of a point to the set. The ball radius is chosen randomly between 5 and 40 pixels, for each ball. The generated shapes are much more complex than the shapes generated using the pixel aggregation method.

Random walk

We start from a ball in the centre of the domain. Then, we use a Brownian motion to choose the next location of the ball. At each iteration, the radius of the ball is chosen randomly.

Perlin noise

Perlin noise [21] is a procedural texture primitive. It has a pseudo random appearance that is highly controllable and multiscale. Fig. 14 provides an illustration of a realisation of this noise. By thresholding this image, we get a set of objects and we select the biggest CC of this set (Fig. 14). Some resulting objects can be very smooth, whereas others can have a high tortuosity.

Fig. 14 Object generation with Perlin noise. Left to right: Perlin noise, thresholding, selection of the biggest CC .

Fig. 15 Examples of random connected components, obtained with five different models.