nD skeletonization: a unified mathematical framework

Antoine Manzanera †,* Thierry M.Bernard †,
Françoise Prêteux ‡, and Bernard Longuet ⋄.

† Ecole Nat. Sup. de Techniques Avancées, Lab. d'Electronique et d'Informatique, Paris - France
‡ Institut National des Télécommunications, Unité de Projet ARTEMIS, Evry - France
⋄ EADS / Aérospatiale-Matra-Missiles, Chatillon - France

Abstract

We present a skeletonization algorithm defined by *explicit Boolean conditions* which are dimension-independent. The proposed procedure leads to new thinning algorithms in 2D and 3D. We establish the mathematical properties of the resulting skeleton refered to as the *MB skeleton*. From a topological point of view, we prove that the algorithm *preserves connectivity* in 2D and 3D. From a metric point of view, we show that the MB skeleton is located on a *median hypersurface* (MHS) that we define. This MHS does not correspond to the standard notion of median axis/surface in 2D/3D, as it combines the various distances associated with the hypercubic grid. The MHS specificities prove to make the skeleton robust with respect to noise and rotation. Then we present the *algorithmic properties* of the MB skeleton: Firstly the algorithm is *fully parallel*, which means that no spatial subiterations are needed. This property, together with the symmetry of the Boolean n-dimensional patterns leads to a perfectly *isotropic* skeleton. Secondly, we emphasize on the *extreme conciseness* of the Boolean expression, and derive the computational efficiency of the procedure.
KEYWORDS: *Skeleton, Fully Parallel Algorithm, nD, Discrete Topology, Boolean Complexity, Median hypersurface*

1 INTRODUCTION

Faithfully representing a shape with a small amount of information is a challenging issue in computer vision. A skeleton-based approach is one way to achieve this goal. A skeleton is faithful if it preserves the *topology* (connected components, holes...), as well as the *geometry* (elongated parts, ramifications...) and the *location* of the original shape. The skeleton-based representation is a common preprocessing step in pattern recognition algorithms. It has motivated intensive research activities [1], [2], mainly for 2D images, and more recently for 3D objects. The skeleton of a binary image is usually obtained through a procedure called *thinning*: border points are iteratively removed if they are not significant for the topology, geometry and location. The remaining shape is then called the skeleton. The available knowledge about digital skeletons suffers from two major limitations. From a theoretical point of view, 2D characterizations are not easily generalized to a 3D framework. This points out the need for a unifying concept. From a practical point of view, such characterizations are generally very difficult to turn into algorithm. Furthermore, as far as nD skeletonization is concerned (with $n > 3$), no satisfying answer currently exists in the literature (yet computing an nD skeleton is useful e.g. to provide the safest trajectories of a robot in a multi-parameter space).

In this paper, we propose a new thinning algorithm, called MB. It is defined with respect to the hypercubic grid, and is dimension-independent. The Boolean characterization of the thinning process is based on one removal and one preserving condition. We explicitely provide the corresponding patterns for dimension 1, 2 and 3. Expressed under this form, our procedure can be easily compared with its competitors, and evaluated with respect to a compactness criterion. In 2D, the overall number of elementary Boolean operations needed to compute the MB skeleton proves to be one of the smallest, as stated by

*This work was mainly done at the *Centre Technique d'Arcueil* (DGA/DCE) and the *ARTEMIS Project Unit* (INT Evry), with the financial support of *EADS Aérospatiale-Matra*

studying the literature. In 3D, the conciseness of the Boolean definition also makes the MB algorithm straightforward to implement. In order to state the topological properties of the MB algorithm, we provide a framework that unifies the theorems of topology preservation in 2D and 3D. In this context, we prove that the MB skeleton preserves the topology in 2D and 3D. From a geometrical point of view, we put in evidence the relationship between the MB-skeleton and the so-called *mixed median hypersurface*.

This paper is organized as follows. Section 2, dealing with preliminaries, introduces the theoretical background useful for understanding the other sections. The proposed thinning algorithm is then presented in Section 3, and illustrated by some results. The MB algorithm properties are presented and discussed in the next sections: topological properties in Section 4, geometrical properties in Section 5, and algorithmic properties in Section 6.

2 PRELIMINARIES

2.1 The hypercubic grid

Let $\mathbb{Z}$ be the set of integers, $\mathbb{N}$ the set of positive integers, $\mathbb{R}$ the set of real numbers. Let $\mathcal{P}(A)$ be the collection of all subsets of A, $\mathcal{P}^*(A) = \mathcal{P}(A) \setminus \emptyset$. Let $n \in \mathbb{N}$, $\mathbb{Z}^n$ is the n-dimensional discrete space. The *cubic grid* is defined by plunging $\mathbb{Z}^n$ into the affine space associated with $\mathbb{R}^n$ as follows:

$$\begin{array}{rccl} \Phi: & \mathbb{Z}^n & \longrightarrow & \mathcal{P}(\mathbb{R}^n) \\ & z & \longmapsto & \Phi(z) = \prod_{i=1}^{n} [z_i - \frac{1}{2}; z_i + \frac{1}{2}]. \end{array}$$

Point z of the discrete space is identified to an *hypercube* of the quantified affine space, i.e. to the Cartesian product of the closed segments centered around z, z_i being the i-th coordinate of z in the canonical basis.

2.2 Discrete hypercubic topologies

Let Ψ be the function defined on $\mathcal{P}^*(\mathbb{R}^n)$ such that $\Psi(P)$ is the linear manifold generated by P (linear manifolds of dimension 0, 1, and 2 are respectively points, lines and planes). Then n kinds of adjacency relations (and then topologies as well) can be defined on the hypercubic mesh, as follows:

Definition 1 *Let z and z' be two points of $\mathbb{Z}^n$ such that $\Phi(z) \cap \Phi(z') \neq \emptyset$.
z and z' are said to be* k-adjacent *(resp.* k-neighbors*), with $0 \leq k \leq n$, if and only if:*
$$dim(\Psi(\Phi(z) \cap \Phi(z')) = k \text{ (resp. } dim(\Psi(\Phi(z) \cap \Phi(z')) \geq k).$$

The adjacency of two points corresponds to the non-empty intersection of two hypercubes, and the level of adjacency to the dimension of the manifold generated by the intersection. Please note the difference with the notion of neighbor, which matches the definitions usually given in the literature. Examples are shown on Figure 1.

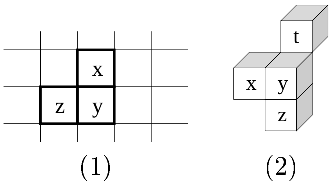

Figure 1: Connectivity relations: (1) in 2D: x and y are 1-adjacent (but not 0-adjacent). x and z are 0-adjacent. x and y are 1-neighbors, therefore 0-neighbors. (2) in 3D: x is 2-adjacent to y, 1-adjacent to z and 0-adjacent to t. x and y are 2-, 1- and 0-neighbors. x and z are 1- and 0-neighbors.

Let a *binary image* I be a subset of $\mathbb{Z}^n$. Let us define the notion of *interior* point of I.

Definition 2 *Let $I \subset \mathbb{Z}^n$, and $z \in I$. Point z is a k-*interior *point of I if and only if: $\forall z' \in \mathbb{Z}^n, z$ and z' are k-neighbors $\Rightarrow z' \in I$.*

Definition 3 *Let A and B be two subsets of $\mathbb{Z}^n$. A and B are* k-connected *if there exists $a \in A$ and $b \in B$ such that a and b are k-neighbors. $X \subset \mathbb{Z}^n$ is a* k-connected component *(denoted by k-cc) if there exists no partition of X in two subsets that are not k-connected.*

Now let us count the number of neighbors for any connectivity level and space dimension. Considering the origin O of $\mathbb{Z}^n$, a point is k-adjacent to O if and only if it has $n-k$ coordinates in the set $\{-1, +1\}$, the others being equal to 0. Let us denote by $A(n,k)$ (resp. $V(n,k)$), the number of k-adjacent points (resp. the number of k-neighbors, the point itself excluded), in dimension n. We get:

$$A(n,k) = 2^{n-k}\binom{n}{k} \text{ and } V(n,k) = \sum_{i=k}^{n-1} A(n,i), \text{ using the usual notation } \binom{n}{k} = \frac{n!}{k!(n-k)!}.$$

Table 1: Numbers of k-neighbors in nD, for $n \leq 4$.

k \ n	1	2	3	4
3	-	-	-	8
2	-	-	6	32
1	-	4	18	64
0	2	8	26	80

Table 1 gives the $V(n,k)$ numbers for n lower than 4. These numbers are used in the literature to indicate the topology level (e.g. 4- or 8-connectivity in 2D). Nevertheless, for the sake of homogeneity and concision, we shall use the single k- prefix (e.g. 1- or 0- connectivity in 2D).

2.3 Discrete distances and median hypersurfaces

For $0 \leq k < n$, let δ_k^n denote the distance induced by the k-topology in dimension n (or δ_k when there is no ambiguity regarding the dimension). Let $X \subset \mathbb{Z}^n$. Let X^c denote $\mathbb{Z}^n \setminus X$, the background of X.

Definition 4 *The* distance map *associated with X and related to δ_k is the function associating $\delta_k(x, X^c)$ with any x in $\mathbb{Z}^n$.*

Definition 5 *Let $r \in \mathbb{N}$. Let $x \in \mathbb{Z}^n$. Let $\mathcal{B}_k(x,r) = \{y \in \mathbb{Z}^n; \delta_k(x,y) \leq r\}$ be the ball with centre x and radius r. Let $X \subset \mathbb{Z}^n$. $\mathcal{B}_k(x,r)$ is a* maximal ball *of X if and only if:*
$$\forall y \in X, \forall q \in \mathbb{N}, \mathcal{B}_k(x,r) \subset \mathcal{B}_k(y,q) \subset X \Rightarrow (y,q) = (x,r).$$

Definition 6 *Let $S_k(X)$ denote the collection of all the centres of maximal balls related to δ_k. $S_k(X)$ is called the* median hypersurface *(MHS) associated with distance δ_k.*

Property 1 *Let $x \in X$. x belongs to $S_k(X)$ if and only if for every y k-neighbor of x, $\delta_k(y, X^c) \leq \delta_k(x, X^c)$.*
In other words, the collection of the centres of maximal balls is equal to the set of local maxima of the corresponding distance map.

These notions are useful to express how to preserve both shape geometry and location. Constraining the skeleton to contain the median hypersurface related to a given distance guarantees that it lies right "at the middle" of the shape for this distance. Let us recall that the first historical definition of the 2D skeleton exactly coincided with that of the median axis [3]. Yet, it does not match the modern notion of skeleton since, in the general case, X and $S_k(X)$ do not have the same topology.

3 THINNING PROCEDURE

Using the mathematical framework previously introduced, we now define a new thinning procedure expressed by explicit Boolean conditions. These conditions are given through two functions of the discrete space in itsef. The first one, called function *Alpha*, expresses a *removing* condition. The second one, called function *Beta*, specifies a *preserving* condition. The MB algorithm consists in removing iteratively, until convergence, the points matching function *Alpha* but not function *Beta*.

Before giving explicitely those two functions, defined from $\mathcal{P}(\mathbb{Z}^n)$ to $\mathcal{P}(\mathbb{Z}^n)$, we shall explain their semantics, that can be roughly understood by looking at Figure 2, displaying the binary patterns corresponding to functions Alpha and Beta, for dimensions from 1 to 3.

Function Alpha computes the set of points having along one direction of their neighborhood, a digital ball belonging to the object on one side, and a set of neighbors belonging to the background on the opposite side. A point matching Alpha is then k-adjacent to an (n-1)-interior point, and it has $n-k$ (n-1)-adjacent neighbors *opposite* to this interior point and belonging to the background. It is somehow, a generalization of the notion of *perfect points* introduced by Eckhardt[4]. Perfect points, defined in 2D, correspond to points matching pattern α_1 in 2D. The authors of [4] have proved that all the perfect points preserving locally the topology could be removed simultaneously from the shape to get a connected skeleton. We show in this paper that this property extends to function Alpha.

Function Beta detects the smallest configuration in which the topology may differ when considering different adjacency levels of the cubic grid. As we will see in section 6, it is related to the notion of *well-shaped* images of Latecki [5].

Definition 7 Function Alpha
Let $I \subset \mathbb{Z}^n$. The set $\alpha(I)$ is defined as the set of all points $z \in I$, such that:
There exists z' (n-1)-interior point and 0-neighbor of z, such that the following condition holds:
Any point t (n-1)-adjacent to z and such that: $\mathcal{S}_z(\Phi(z) \cap \Phi(z')) \subset (\Phi(z) \cap \Phi(t))$ belongs to I^c.
($\mathcal{S}_z$ denotes the symmetry of centre z).

The former definition can be explained as follows: a point z of I belongs to $\alpha(I)$ only if it is a 0-neighbor of a (n-1)-interior point z' of I. If this condition holds, z and z' have a non-empty intersection, generating a linear manifold of dimension k ($k < n$). Let us now consider the image of this manifold by the symmetry of centre z. It is a parallel linear manifold of dimension k, which can be uniquely obtained by intersecting $n-k$ affine hyperplanes, each of which being perpendicular to a vector of the canonical basis of $\mathbb{R}^n$. Then $z \in \alpha(I)$ if all the points (n-1)-adjacent to z, and whose intersection with z generates one of those hyperplanes, belong to the background of I.

Definition 8 Function Beta
Let $I \subset \mathbb{Z}^n$. The set $\beta(I)$ is defined as the set of all points $z \in I$ for which there exists $k, 0 \leq k \leq n-2$, and two couples of k-adjacent points (a,b) and (c,d) such that a,b,c,d are all k-neighbors of z, and such that the following two conditions hold:
(1) $\Phi(a) \cap \Phi(b) = \Phi(c) \cap \Phi(d)$,
(2) $\{a,b\} \subset I$ and $\{c,d\} \subset I^c$.

In other words, a point z of I belongs to $\beta(I)$ iff it contains in its k-neighborhood two couples of k-adjacent points that have the same intersection, and such that one couple belongs to the image, and the other to the background. Note that z itself might play the role of a or b. In fact, it is always so in 2D (but not necessarily in 3D).

Finally the complete algorithm is defined as follows:

Definition 9 The MB algorithm
Let $X \subset \mathbb{Z}^n$. Let $mb(X) = \alpha(X) \setminus \beta(X)$. Let $X^0 = X$, $X^{n+1} = X^n \setminus mb(X^n)$.
The MB skeleton of X is $X^\infty = X^{\min\{n; X^{n+1} = X^n\}}$.

In order to ease the reading of the prooves in the following sections, we shall refer to $\alpha_i(X)$ (resp. $\beta_i(X)$) as the set of the points of X that match pattern α_i (resp. that contain pattern β_i within their i-neighborhood).

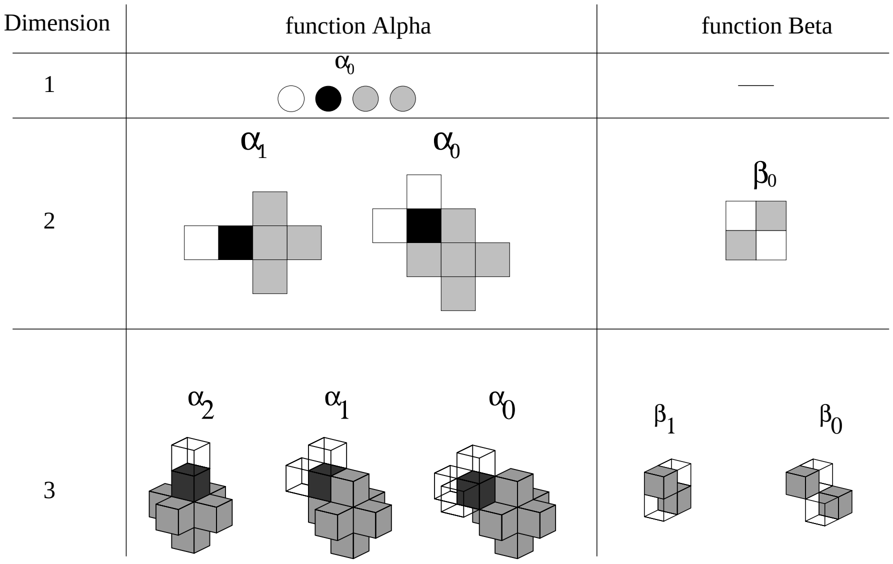

Figure 2: The MB patterns in dimension 1 to 3. Function Alpha is the removing condition, and function Beta the preserving one. The grey hypercubes correspond to points in X, the white ones to points in X^c. The black hypercube, belonging to X, represents the origin of the pattern, that is the matched point. Thus in dimension n, a point is removed iff: (1) it matches one of the patterns α_i (or one of the $\pi/2$ rotated version around any vector of the canonical basis). (2) it does not contain any pattern β_i within its i-neighborhood.

Some results of the MB algorithm are displayed in 2D in Figure 3, and in 3D in Figure 4. They illustrate the properties and behavior of the thinning procedure, that we develop in the following sections.

4 TOPOLOGICAL PROPERTIES

Removing points from an object without changing its topology means preserving the connectivity relations that exist in the object and in its background. To this end, points have to be removed without disconnecting any connected component, nor creating or deleting a "hole". Now in a discrete topology, special care must be taken to deal with these notions. For instance, a connected component of the background may run across a piece of surface of the object only if there is a hole in it ! To keep this "natural" property, however, it is necessary to choose two different connectivity levels, one for the object and another one for the background. In this paper, we always consider the 0-connectivity for the object (two hypercubes of the object are neighbors as soon as they share a point) and the (n-1)-connectivity for the background (two hypercubes of the background are neighbors only if they share a hyperface). We use the (0,n-1)-prefix to recall the choice of the topology.
In this section, we do not intend to prove that the MB-algorithm preserves the topology in n-D. Anyway, to our knowledge, no characterization in n-D has been recognized by the discrete topology community. Nevertheless, the formalism that we use in the following results is completely generic as for the dimension, and the results are true for dimensions 2 and 3.
We must first give some meaning to "preserving the topology". In the 2D and 3D continuous spaces, topology can be represented by means of the fundamental group, which is the group of the equivalent classes of homotopic arcs (two arcs are said to be homotopic if there exists a bicontinuous function mapping one to the other). Then, two subsets of $\mathbb{R}^2$ have the same topology if and only if their fundamental groups are isomorphic. In $\mathbb{R}^3$, this property is no longer sufficient, as fundamental groups of the

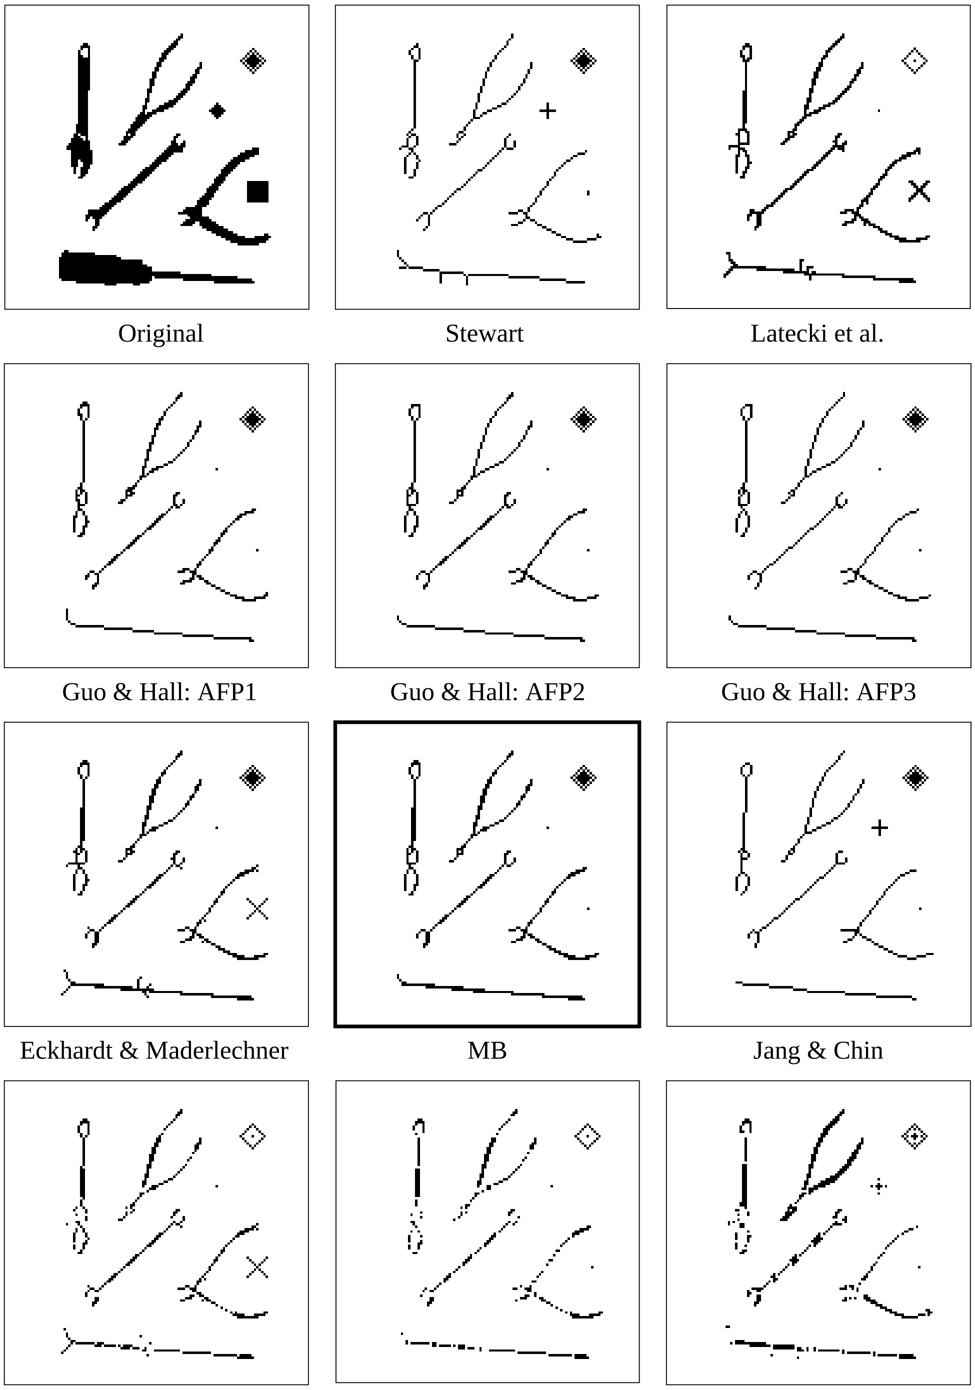

S_1^1 S_1^0 S_0^0

Figure 3: Some results of the MB-2D algorithm. For an objective comparison, we also displayed the results of the algorithms of Stewart[6], Latecki et al.[5], the three algorithms of Guo and Hall[7], Eckhardt and Maderlechner[4], and Jang and Chin[8]. The last line shows the different median axis, as defined in Section 5.

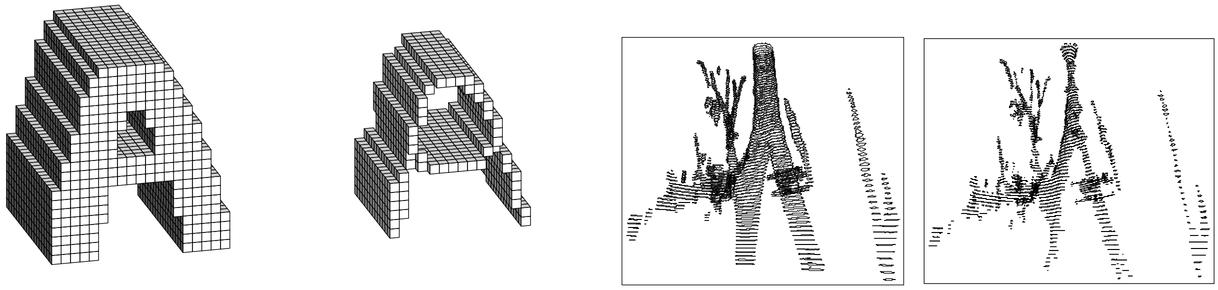

Original MB-3D Original MB-3D

(1) (2)

Figure 4: Results of the MB-3D algorithm: on a simple home-made 3D object (1) and on a three-dimensional segmented image of lung (2).

backgrounds must be isomorphic too. In higher dimensions, homotopy of higher order is necessary to characterize the topology.

In discrete spaces, the central notion regarding topology preservation is *simplicity*. A subset A of a discrete image X is said to be simple if X and $X \setminus A$ have the same topology. For a point $x \in X$, "x is simple" means that $\{x\}$ is simple. In 2D and 3D there are explicit characterizations to decide whether a point is simple or not. A major result is that the decision can be made locally, examining only a finite neighborhood. We now give the characterization, for the (0,n-1)-connectivity model:

Theorem 1 *Let $X \subset \mathbb{Z}^n$ be a binary image. Let $x \in X$. Let X_0^x denote the set of all 0-neighbors of x in X, except x itself, and $\overline{X_{n-2}^x}$ the set of all (n-2)-neighbors of x in X^c. x is (0,n-1)-*simple *in X if and only if the following two conditions hold:*

- *$\{x\}$ is 0-connected to a single 0-connected component of X_0^x, and*
- *$\{x\}$ is (n-1)-connected to a single (n-1)-connected component of $\overline{X_{n-2}^x}$.*

In 2D, this result corresponds to the *0-connectivity number* defined by Yokoi[9], and used earlier by Hilditch[10]. In 3D, this result is due to Bertrand and Malandain[11]. Thus, Theorem 1 proposes a unified expression of those two fundamental results. The use of the (n-2)-connectivity level in the set $\overline{X_{n-2}^x}$ can be explained by the fact that the (n-2)-neighbors of x is the smallest set in which a subset of X^c may be (n-1)-connected.

Unfortunately, a union of simple points is not a simple set, in general, and this is the major difficulty for designing parallel thinning algorithms. The first problem is to characterize simple sets. Ronse[12] did it first for 2D images. This work was then generalised by Kong[13] for higher-dimensional images. In these papers, it is shown that a set is simple for the image X if and only if it can be ordered in a sequence of points $\{x_1, \dots, x_n\}$ such that for every i in $\{1, \dots, n\}$, x_i is individually simple with respect to $X \setminus \{x_1, \dots, x_{i-1}\}$. From this property, Ronse proposed[14] very efficient sufficient conditions to prove the soundness of parallel thinning algorithms in 2D. This result has been extended to the 3D case by Ma[15]. In this case we may also give a unified expression of these results as follows: Let a *unit lattice element of dimension k*, $0 \leq k \leq n$, be a set of 2^k points of $\mathbb{Z}^n$ such that every pair of points is a pair of (n-k)-neighbors (see Figure 5).

Theorem 2 *(Ronse 88, Ma 94)*
Let $X \subset \mathbb{Z}^n$ be a binary image. An algorithm that removes points in parallel from a binary nD shape X preserves (0,n-1)-connectivity if the two following conditions are satisfied:

- *Every subset of X that is contained in a unit lattice element of dimension $(n-1)$ and that is removed by the algorithm is simple, and*

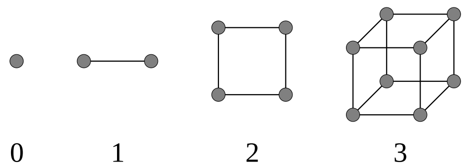

Figure 5: The unit lattice elements of dimension 0 to 3.

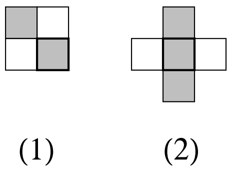

Figure 6: The (0,1)-non-simplicity patterns in 2D.

- *No connected component of X contained in a unit lattice element of dimension n can be completely removed.*

This theorem allows to prove the soundness of a parallel thinning algorithm by checking only a limited number of configurations. We may now formulate a topological property about the MB algorithm:

Proposition 1 *(Topological property)*

MB-nD preserves the (0,n-1)-topology.

We prove this proposition for n=2 and n=3.

For n=2, we have to show that

(1) MB-2D only removes simple points,

(2) Every pair of 1-adjacent points that is removed is simple, and

(3) No connected component contained in a 2×2 square can be completely removed.

To prove assertion (1), we deduce from theorem 1 in 2D that a point 1-adjacent to a point of the background and not (0,1)-simple, corresponds to a point matching one of the two patterns represented in Figure 6. A point x such that $x \in \alpha(X)$ cannot match pattern (2), and pattern (1) is identical to pattern β_0. Then if $x \in mb(X)$, x is (0,1)-simple.

To prove assertion (2), let us notice that two 1-adjacent points, $\{x, y\}$, simultaneously removed by the MB-2D algorithm correspond to one of the three configurations shown on Figure 7. If x and y both belong to $\alpha_1(X)$ (resp. $\alpha_2(X)$), then the only possible configuration is that of Figure 7(1) (resp. 7(2)). In both cases, none of the two points belongs to the pattern of the other, the two points would have been removed even if treated sequentially, hence $\{x, y\}$ is (0,1)-simple. If $x \in \alpha_2(X)$ and $y \in \alpha_1(X)$, then the only possibility is given by Figure 7(3). In that case, x does not belong to pattern α_1 of y. Thus y is simple in $X \setminus \{x\}$, and the pair $\{x, y\}$ is (0,1)-simple.

Assertion (3) is obviously met, since no pattern α_i fits in a 2×2 square. This proof establishes that the MB-2D algorithm preserves the (0,1)-topology. □

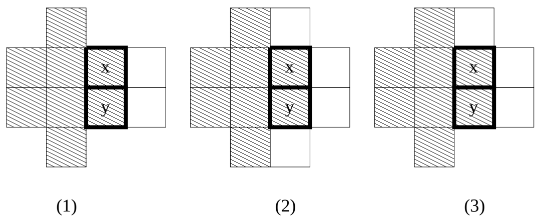

Figure 7: The tree possible cases of 1-adjacent pair removed by MB-2D.

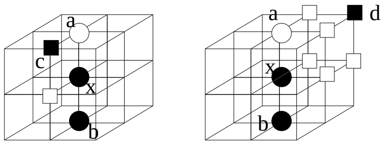

(1) (2)

Figure 9: Lemma 1: geometric interpretation.

For n=3, we need to prove that

(1) Every subset of X contained in a 2×2 square is simple, and

(2) No connected component contained in a $2 \times 2 \times 2$ cube can be completely removed.

If $x \in X$, we use the two sets X_0^x and $\overline{X_1^x}$ defined in Theorem 1. The proof is based on five lemmae. Lemmae 1 to 3 deal with the 0-topology preservation of objects, whereas lemmae 4 and 5 deal with the 2-topology preservation of the background. Lemmae 1 and 4 prove that one iteration of the MB-3D algorithm removes only simple points. Lemmae 1 and 2 are used to prove lemma 3. Lemma 4 is used to prove lemma 5. Lemmae 3 and 5 prove that any pair of 2-adjacent points removed by the MB-3D algorithm is a simple set. Finally, the proof is completed in proposition 2.

Lemma 1 *Let $x \in X$, between two 2-neighbors a and b, with $a \notin X$ and $b \in X$ (cf Figure 9). If $\{x\}$ is 0-connected to more than one 0-cc of X_0^x, then either x is contained in pattern β_1, or x is contained in the pattern λ represented in Figure 8.*

PROOF

If $\{x\}$ is 0-connected to more than one 0-cc of X_0^x, then there exists a point y in X_0^x which is not 0-neighbor of b. Here, y cannot be a 2-neighbor of x, but it may be a 1-neighbor, as illustrated by c in Figure 9(1). In that case, since c and b are not in the same 0-cc, x matches β_1. If there is no such c, then y is only a 0-neighbor of x, as illustrated by d in Figure 9(2). In that case, x matches λ □

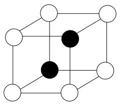

Figure 8: Pattern λ.

Corollary 1 *Any point removed by one iteration of the algorithm fulfils condition 1 of theorem 1.*

Any point matching pattern α_2 or α_1 is necessarily between two 2-neighbors: one in X, the other in the background. The same holds for a point that matches α_0, and not β_1. Then Lemma 1 applies and, since pattern λ is a particular case of pattern β_0, the point is 0-neighbor of only one 0-cc of X_0^x.

Lemma 2 *Let $x \in X$. Let Y be a subset of X such that $Y \subset mb(X)$ and $Y \cup \{x\}$ is contained in a 2×2 square. Then $x \in \lambda(X \setminus Y)$ implies $x \in \beta_0(X)$.*

PROOF

Let us consider $x \in \lambda(X \setminus Y)$. Remember that $\lambda(X \setminus Y)$ stands for the set of points in $X \setminus Y$ that match pattern λ. If $x \in \lambda(X)$, then $x \in \beta_0(X)$. If not, the situation is that of Figure 10(1), where $Y \subset \{y_1, y_2, y_3\}$. Note that the three points represented by squares belong either to Y or to X^c. If $y_1 \in X^c$ or $y_3 \in X^c$, then obviously $x \in \beta_0(X)$. If not, $\{y_1, y_3, z\} \subset X$. It follows that y_2 may match an α_i only with an interior point within the cube drawn in Figure 10(1). But for each of the seven possibilities, one can easily check that this is not possible. Then $y_2 \notin Y$, so $y_2 \in X^c$, and the four points $\{x, t, y_2, z\}$ make up a β_0 pattern □

Lemma 3 *Let x and y be two 2-neighbors such that $\{x, y\} \subset mb(X)$. Then $\{x\}$ is 0-connected to only one 0-cc of $(X \setminus \{y\})_0^x$.*

PROOF

Under the premises of lemma 3, it can easily be checked that whatever the α_i it matches, x is always between two 2-neighbors such that one belongs to $X \setminus \{y\}$ and the other to X^c. Now suppose that $\{x\}$ is 0-connected to more than one 0-cc of $(X \setminus \{y\})_0^x$. From lemma 1, x must match one of the two patterns β_1 or λ within $(X \setminus \{y\})$. But lemma 2 shows it cannot be λ since x would have matched β_0 before the removal of y, in contradiction with x being removed by MB-3D. So x matches β_1 within $(X \setminus \{y\})$;

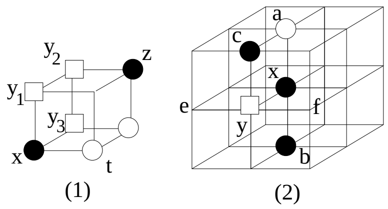

Figure 10: Lemmae 2 and 3: geometric interpretation.

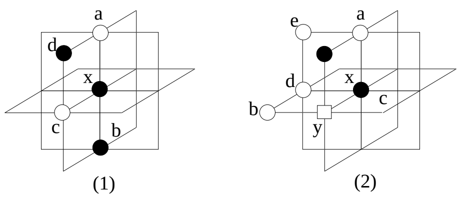

Figure 11: Lemmae 4 and 5: geometric interpretation.

more precisely, the situation of x is that of Figure 9(1), with c and b in distinct 0-ccs. Since x does not match β_1 within X, y as a removed point, is part of β_1, as shown on Figure 10(2). Besides, e and f must both belong to X^c. But then, y could not have matched an α_i pattern, which is in contradiction with its removal by the MB-3D algorithm $\square$

Lemma 4 *Let $x \in X$, between two 2-neighbors a and b, with $a \notin X$ and $b \in X$. If $\{x\}$ is 2-connected to more than one 2-cc of $\overline{X_1^x}$, then x is contained in pattern β_1.*

PROOF

See Figure 11(1). If there exists $c \notin X$ such that a and c belong to two distinct 2-ccs of $\overline{X_1^x}$, then point d such that $d \neq x$, d 2-adjacent to both a and c must belong to X. So x matches pattern β_1 $\square$

Corollary 2 *Any point removed by one iteration of the algorithm fulfils condition 2 of Theorem 1.*

Lemma 5 *Let x and y be two 2-neighbors such that $\{x, y\} \subset mb(X)$. Then $\{x\}$ is 2-connected to only one 2-cc of $\overline{(X \setminus \{y\})_1^x}$.*

PROOF

The premises of lemma 5 (identical to those of lemma 3), imply that x is between two 2-neighbors such that one belongs to $X \setminus \{y\}$ and the other to X^c. Now suppose that $\{x\}$ is 2-connected to more than one 2-cc of $\overline{(X \setminus \{y\})_1^x}$. From lemma 4, x must match β_1 within $(X \setminus \{y\})$. See Figure 11(2), where a and y belong to distinct 2-ccs of $\overline{(X \setminus \{y\})_1^x}$. If b and c both belong to X, then y could not have matched an α_i pattern, so b or c belong to X^c. Let us suppose it is b. Since x is removed, it does not match pattern β_0, and so $d \in X^c$. Since x does not match pattern β_1, $e \in X^c$ also, and finally a and y belong to the same 2-cc. That leads to a contradiction $\square$

We may now give the main proposition.

Proposition 2 *The MB-3D algorithm preserves the (0,2) topology.*

PROOF

As mentioned earlier, lemmae 1 and 4 prove that one iteration of the MB-3D algorithm removes only simple points. Now let $\{x_1, x_2\}$ be a pair of 2-adjacent points, simultaneously removed by MB-3D. Lemmae 3 and 5 prove that $\{x_1, x_2\}$ is a simple set. More generally, let Y be a set of points such that $Y \subset mb(X)$ and Y is contained in a 2×2 square. Let $x \in Y$ such that x is not simple in $(X \setminus (Y \setminus \{x\}))$. Then lemmae 1 and 4 show that x matches pattern β_1 or λ, but the latter is forbiden by lemma 2. Then x matches β_1 within $(X \setminus (Y \setminus \{x\}))$. Now let us consider $\{x_1, x_2\} \subset mb(X)$ a pair of 1-adjacent points. It is easy to see that if $x_1 \notin \beta_1(X)$, then $x_1 \notin \beta_1(X \setminus \{x_2\})$. So x_1 is simple in $(X \setminus \{x_2\})$, and then

$\{x_1, x_2\}$ is a simple set. Let $\{x_1, x_2, x_3\} \subset mb(X)$ be a triplet of points contained in 2×2 square such that x_1 and x_2 are 2-adjacent. Then $\{x_1, x_2\}$ is simple, and it is easy to see that if $x_3 \notin \beta_1(X)$, then $x_3 \notin \beta_1(X \setminus \{x_1, x_2\})$, so $\{x_1, x_2, x_3\}$ is a simple set. Let $\{x_1, x_2, x_3, x_4\} \subset mb(X)$ be the four corners of a 2×2 square. $\{x_1, x_2, x_3\}$ is a simple set, and if $x_4 \notin \beta_1(X)$, then $x_4 \notin \beta_1(X \setminus \{x_1, x_2, x_3\})$, so $\{x_1, x_2, x_3, x_4\}$ is a simple set. Thus we have proved that any set contained within a 2×2 square is a simple set. At last, it is obvious that an iteration of the MB-3D algorithm cannot entirely remove a connected component contained in a $2 \times 2 \times 2$ cube, since no α_i fits into this elementary cube. So we have proved that MB-3D is a parallel reduction operator that fulfils conditions (1) and (2) of theorem 2. Therefore the MB-3D algorithm preserves (0,2)-topology □

5 GEOMETRICAL PROPERTIES

In this section, we are concerned with the non-topological side of skeletonization, that is, the geometry preservation issue. Ideally, the skeleton lies at the middle of the shape if the local maxima of the Euclidean distance belong to the skeleton. As our prime purpose was conciseness, we have only allowed ourselves using the canonical distances of the hypercubic grid. Still, we show in this section that the MB algorithm is based on a special metrics that gives to the MB skeleton good properties with repect to noise immunity and rotation invariance. Furthermore, if a predefined metrics is needed, such as Euclidean or pseudo-Euclidean distance, we show that the MB algorithm lends itself well to *constrained thinning*, in order to base the skeleton on another kind of median surface.
Let $p \leq m \leq n$. Let us define the (m,p)-*median hypersurface* ((m,p)-MHS).

Definition 10 *Let $X \subset \mathbb{Z}^n$. The (m,p)-*median hypersurface *of X is the following set:*

$$S_m^p(X) = \{x \in X; \forall y \; p-neighbor \; of \; x, \delta_m(y, X^c) \leq \delta_m(x, X^c)\}$$

Note that the (m,m)-MHS of X corresponds to the standard $S_m(X)$ defined in Section 2.3. The behavior of the MB algorithm with respect to MHS preservation is characterized by the following proposition:

Proposition 3 *(Non-topological property)*
Let $n \in \mathbb{N}$. Let X be a subset of $\mathbb{Z}^n$ such that the points of X are deleted in the order induced by the (n-1)-distance. Then the following property holds:

$$\boxed{\text{The MB-nD skeleton of } X \text{ contains the set } S_{n-1}^0(X).}$$

Let us consider patterns $\alpha_i, 0 \leq i \leq n-1$ of Figure 2. If a point x is removed, it has within its 0-neighborhood a (n-1)-interior point y. If the points are examined in the order of the (n-1)-distance function, $\delta_{(n-1)}(y, X^c) > \delta_{(n-1)}(x, X^c)$ and thus the points of $S_{n-1}^0(X)$ necessarily belong to the skeleton. □

It is worth observing that, for usual images, points are examined according to the order induced by the distance function. Nevertheless, there are exceptions, like those corresponding to pathological images (Figure 14). Such images correspond to a configuration that would "protect" a piece of surface, inhibiting the thinning of a thick volume, as in the shape present in the upper right corner of the image in Figure 3.

The property of MHS preservation is illustrated in 2D in Figure 12 and in 3D in Figure 13. In Figure 12, the upper pictures show the distance functions for the two different connectivity models, where the value is represented by the grey level. On the centre, the left and the right images show the corresponding local maxima set, and at the middle, the peculiar median axis on which the MB skeleton (bottom, centre) is built. For an objective comparison, other skeletons are also shown on the left and on the right, based on the local maxima for the 1- and the 0- distance respectively. In Figure 13, it can be seen as well that different 3D skeletons can be obtained, depending on the median surface. In Figure 13(b), the restriction of the Alpha function of the MB thinning algorithm to the α_2 pattern leads to a skeleton based on the local maxima of the 2-distance. In Figure 13(c), restricting to patterns α_2 and α_1 leads to a skeleton based on the $(2,1)$-median surface. Finally, the whole function Alpha is applied in Figure 13(d), and the resulting skeleton is based on the $(2,0)$-median surface.

The possibility of changing the shape of the MHS by restricting function Alpha turns the MB skeleton into a versatile tool, useful for modelizing complex objects.

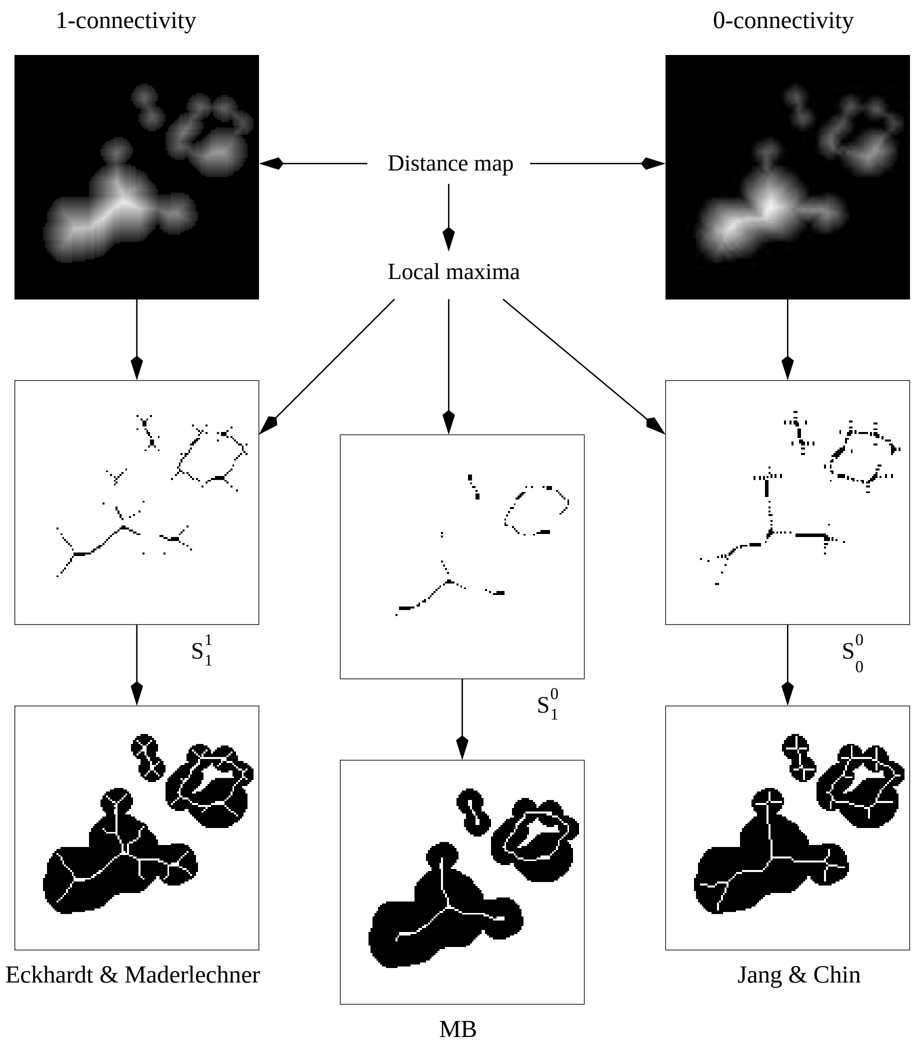

Figure 12: Different median axis leading to different skeletons in 2D.

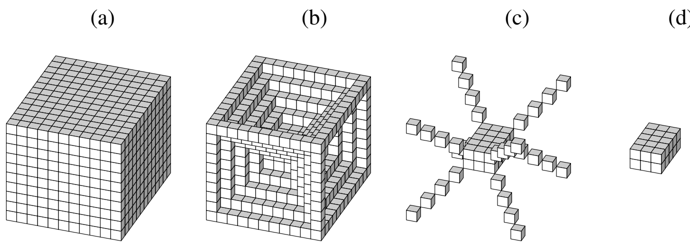

Figure 13: Different median surfaces in 3D: (a) Original (b) S_2^2 (c) S_2^1 (d) S_2^0.

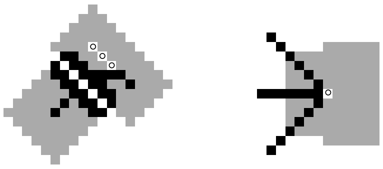

Figure 14: Pathological images in 2D: the skeleton (in black) does not contain the whole median axis (the points marked by a circle form the difference)

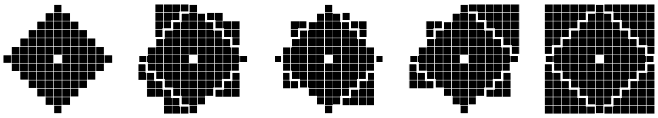

Figure 15: Some fuzzy balls of radius 7 in 2D, and their skeleton reduced to their centre (superimposed in white).

The MB skeleton is based on a special kind of MHS, able to homogeneously process the different kinds of distances in the cubic grid. In particular, the thinning operator does not distinguish the balls of the different distances, producing always one single point for any kind of ball. It even produces a single point for a family of sets that are bounded (for the inclusion) by two balls of different distances, and with the same radius. Let us call these sets *Fuzzy balls*. We show in Figure 15 some fuzzy balls of radius 7 in 2D, and their skeleton reduced to their centre (superimposed as a white dot). In 2D, it can be proved that a fuzzy ball B (formally defined as a set whose (1,0)-median axis is reduced to a point, called its centre c) is a set that is included between a 1-ball and a 0-ball, such that for all x in B, B contains the smallest rectangle containing x and c.

The behavior of the MB skeleton with respect to fuzzy balls is the cause of its good properties regarding noise immunity and rotation invariance. This is shown in Figure 16, where the results are compared to other algorithms, previously referenced in Figure 3. In return, MB only allows partial reconstruction. This fact is illustrated in Figure 18: the shape is reconstructed from the weighted skeleton (here, the value of the distance function on the skeleton is represented by the grey level). Unlike the left skeleton (the Eckhardt and Maderlechner one[4]), based on the (1,1)-median axis, and achieving an *exact* reconstructibility, the MB-algorithm, represented here on the right, does not. We show here an *approximate* reconstruction, based on octagons of the corresponding radius.

In spite of the properties of the (n-1,0)-MHS, the MB skeleton remains characterized by the (n-1)-distance, because it removes points in the order induced by the (n-1)-distance function. This implies that the skeleton cannot be completely rotation invariant (i.e. up to the discretization process). This is particularly apparent in the thickest parts of the objects, as it can be seen on Figure 19(1), which displays the result of MB-2D on the same object with different rotations. It can be seen that the eye of the lizard generates a skeletal branch in one case only, and that the hierarchy of the branches may vary sensitively (see for example the shoulders of the lizard).

This typical problem of thinning algorithms has motivated recent studies on alternate skeletons in order to be as close as possible to the Euclidean median surface. The first type of approach is based on a continuous framework. For example in [16], the authors compute the median surface through a differential geometry model, while ensuring the homotopy at a discrete level. The second type of approach is fully discrete. For example, the skeleton of Costa [17] can be computed directly from local configurations of the Euclidean distance transform, as long as the distance function is computed in its full dynamics, which

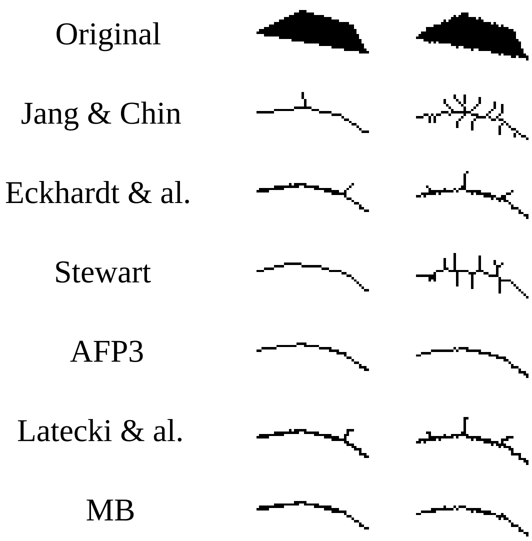

Figure 16: Behavior of the MB algorithm in 2D, with respect to noise.

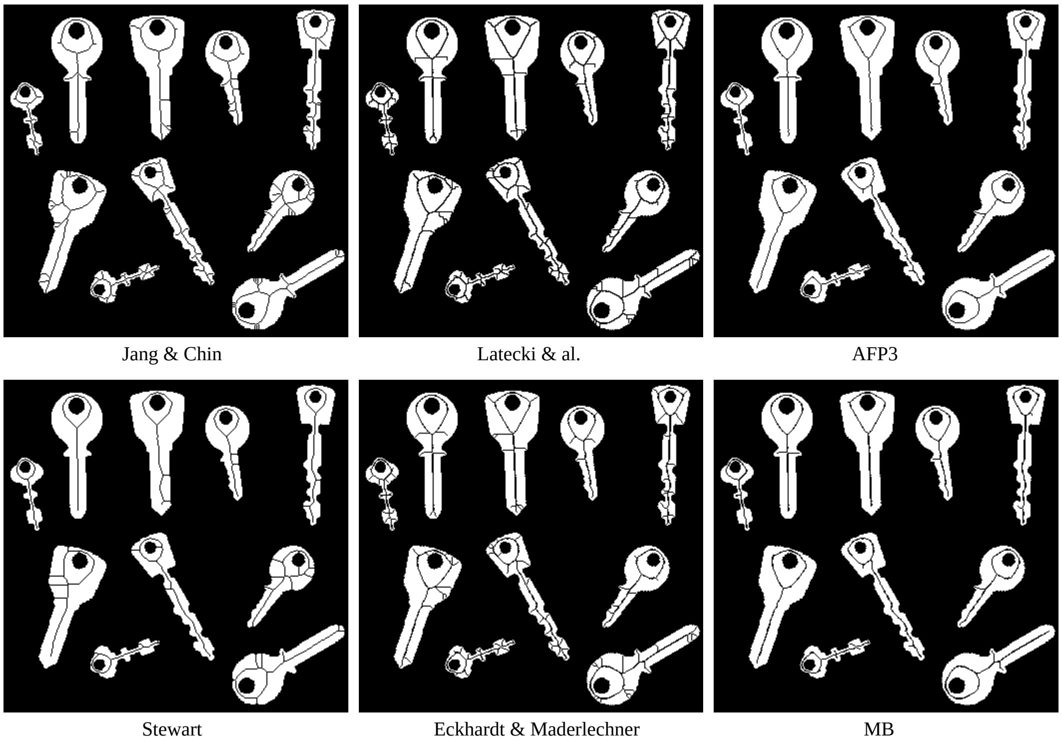

Figure 17: Behavior of the MB algorithm in 2D, with respect to rotation.

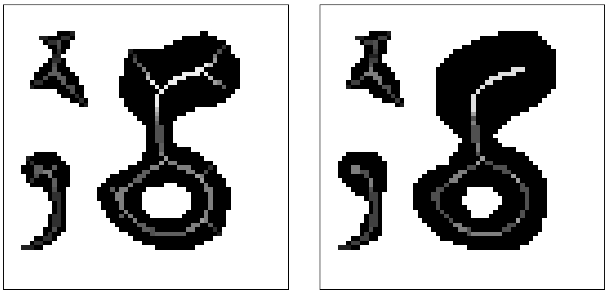

Figure 18: Exact reconstruction for the skeleton based on the 1-distance balls (left), approximate reconstruction for the MB-skeleton based on the fuzzy balls (right).

is done thanks to the concept of *exact dilation* [18] involving a special data structure.

One simple method for a thinning algorithm to cope with this problem is to constrain the thinning process to respect the order induced by a pre-processed distance function. Nevertheless, if a thinning algorithm is too strongly characterized by its underlying metrics, it remains rotation sensitive. On the contrary, the MB algorithm is neutral enough to tolerate the conditioning by a given distance function. This is illustrated on Figure 19, where the MB-2D thinning algorithm is applied by simply examining the points in the order of some distance functions computed before. The chamfer distances [19] are used as simple approximation on the Euclidean distance. So, the MB-skeleton is by no means a "Euclidean" skeleton like the Costa skeleton [17] (which appears on Figure 19(4)) because it does not distinguish the fuzzy balls, but it can achieve, thanks to the conditioning, similar rotation invariance.

6 ALGORITHMIC PROPERTIES

6.1 Isotropy and full parallelism

In the hypercubic grid, *isotropy* means that all the directions represented by the n axis must be treated the same way, as well as the two ways on every axis. Isotropy is one of the fundamental principles on which our algorithm is built. A natural outcome is the presence of two-pixel thick surfaces in skeletons (Figures 3 and 4). In the definitions of the two Boolean functions, no direction gets a special treatment. This property clearly appears on the pattern matching form of the procedure, as all the patterns are completely symmetrical. According to this property, isotropy is guaranteed, as long as all the iterations have exactly the same action. This is a major property of the MB algorithm: the procedure is *fully parallel*, i.e. all the iterations act exactly in the same manner, removing points in the $2n$ cardinal directions at the same time. As a consequence, the number of iterations before convergence equals the radius of the biggest ball corresponding to the strongest adjacency relation.

6.2 Computational speed

Does conceptual conciseness imply computational speed ? We have argued in an earlier work[20] that the most fundamental way to compare computational efficiency of data parallel algorithm is based on the so-called *Shannon measure*, i.e. the overall number of elementary operations per pixel needed to perform the algorithm. The nature of the elementary operations depends on the computer. However, we get a sound general measure by choosing the most elementary operation to be computed in a digital algorithm: the Boolean function of two variables. For any digital algorithm, the cost of every elementary Boolean operation is found either at a software level, where it has an effect on the computation time, or at a hardware level, under the form of an equivalent two-entries logic gate, where it has an effect on the expense in hardware resources. Table 2 displays the Shannon measure of the MB algorithm in 2D and in 3D. In 2D, the measure is compared to some of the most recent thinning algorithms found in the literature.

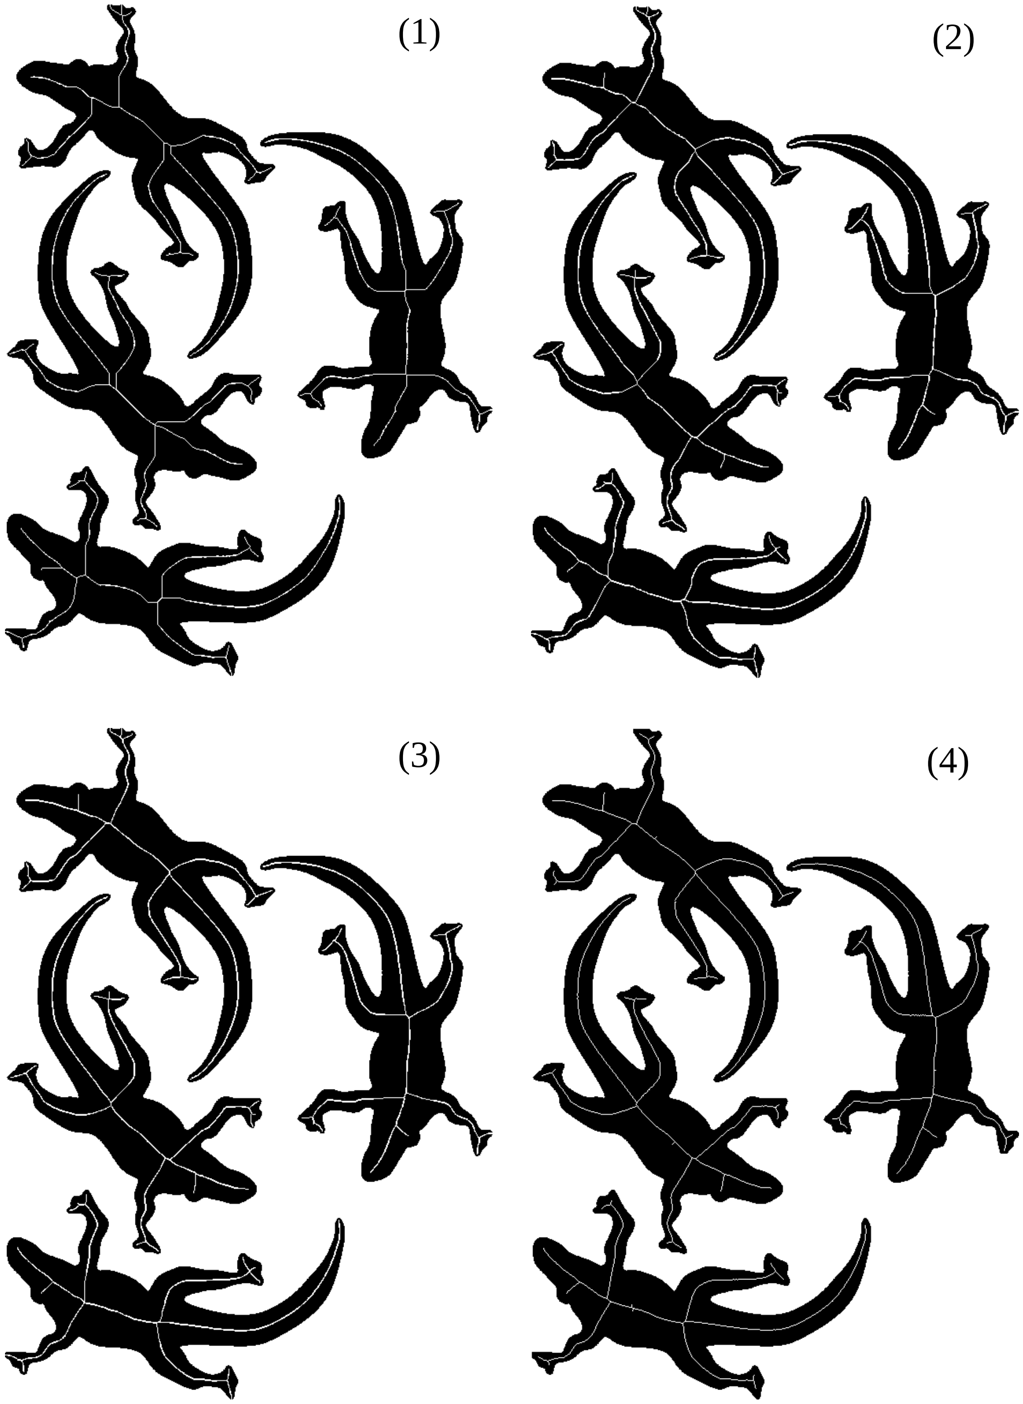

Figure 19: Constraining the MB thinning with a given distance function. (1) The 1-distance function (in this case the result is strictly identical to the unconstrained algorithm) (2) The 3-4 chamfer distance function (3) The 5-7-11 chamfer distance function (4) For comparison purposes, the Euclidean skeleton computed using the exact dilation framework [18], (reproduced with kind permission of Pr. da Fontoura Costa).

Table 2: Comparing the cost of parallel thinning algorithms. r_1 (resp. r_0) stands for the radius of the biggest 1-ball (resp. 0-ball) contained in the image. The number of elementary operations given for the other algorithms is sometimes only an estimation, as the quoted papers did not always provide a logic minimization of their algorithm.

Algorithm	Size of neighborhood examined	Number of elementary operations required
Jang & Chin[8]	7	$32 \times r_0$
Cardoner & Thomas[21]	7	$40 \times r_0$
Latecki & al.[5]	11	$16 \times r_1$
Stewart[6]	19	$64 \times r_1$
Wu & Tsai[22]	11	$60 \times r_1$
Guo & Hall (AFP1)[7]	11	$73 \times r_1$
Guo & Hall (AFP2)[7]	11	$81 \times r_1$
Guo & Hall (AFP3)[7]	11	$91 \times r_1$
Eckhardt & Maderlechner[4]	13	$18 \times r_1$
MB-2D	21	$28 \times r_1$
MB-3D	81	$148 \times r_1$

These comparative results show that the computational cost of MB-2D is among the lowest, although its behavior is very close to algorithms that are among the best. The work of logic minimization that we have done in 2D has shown that the most compact algorithms are those of Eckhardt and Latecki. The algorithm of Eckhardt[4] correspond to the MB-2D algorithm by restricting function Alpha to pattern α_1. Thus it is entirely computable from the two patterns α_1 and β_0. The algorithm of Latecki[5] is defined from the same patterns exactly, but its principle is different. For Latecki, an image is *well-shaped* if it does not contain pattern β_0. This corresponds to images whose (0,1)-topology and (1,0)-topology are the same. In this restricted framework, Latecki computes a (1-connected) skeleton, by removing points matching pattern α_1, as long as the removal does not create a β_0 pattern. Because of an anisotropy that is necessary to obtain a thin skeleton, it only computes one half of the rotated versions of pattern β_0, and so it is slightly more concise than Eckhardt's procedure. Thus it is the parallel thinning algorithm with the lowest Shannon measure we know of.

7 CONCLUSION

The MB thinning algorithm is, to our knowledge, the first thinning algorithm proposed within nD framework. Further researches in discrete topology for higher dimension meshes will allow to check its overall validity. More generally, we may hope that this work will suggest some research tracks. In particular, we have proposed a unified formalism to express the theorems of homotopy in n-D propositions: the Yokoi/Bertrand and Malandain theorem for the simple point characterization, the Ronse/Ma theorem for the topology preservation sufficient conditions. Do these propositions make sense for dimensions higher than 3 ? May our work motivate some investigations in this direction... Nevertheless, we think one of the best achievements of the MB algorithm is its great simplicity, as the n-D thinning procedure is defined through two binary Boolean functions, and only $(2n-1)$ Boolean patterns. We have shown in 2D that this simplicity had led to computational efficiency, and we can justifiably expect the same advantage in 3D. Finally we have characterized the median hypersurface on which the skeleton is based, and shown that it provides some noise immunity to the skeleton. This median hypersurface implies a new kind of shape descriptors: fuzzy balls, that we have characterized in 2D. Further work should provide a more general view of the latter issue.

ACKNOWLEDGMENTS

This work has taken advantage of many helpful comments from Pr. *Luciano da Fontoura Costa*, *Leandro Estrozi*, Pr. *Longin J. Latecki* and Pr. *Michel Schmitt*. The authors are also grateful to *Bertrand Collin* and *Damien Mercier*, from CTA/GIP, for software support. The volumetric medical image was acquired in the Service de Radiologie Central de l'Hôpital de la Salpétrière (Professor Ph. Grenier). The lung was segmented by *Catalin Fetita* from the ARTEMIS Project Unit (INT Evry).

References

[1] T. Kong and A. Rosenfeld, "Digital Topology : Introduction and Survey," *Computer Vision, Graphics and Image Processing* **48**, pp. 357–393, 1989.

[2] L. Lam, S.-W. Lee, and C. Y. Suen, "Thinning methodologies : A Comprehensive Survey," *IEEE Transactions on pattern analysis and machine intelligence* **14**, pp. 869–885, 1992.

[3] H. Blum, "A tranformation for extracting new descriptors of shape," in *Proc. Symposium Models for the perception of speech and visual form*, pp. 62–380, W.Wathen-Dunn ed. M.I.T. Press Cambridge MA, 1967.

[4] U. Eckhardt and G. Maderlechner, "Invariant thinning," *International Journal of Pattern Recognition and Artificial Intelligence* **7-5**, pp. 1115–1144, 1993.

[5] L. J. Latecki, U. Eckhardt, and A. Rosenfeld, "Well-composed sets," *Computer Vision and Image Understanding* **61-1**, pp. 70–83, 1995.

[6] A. Stewart, "A one-pass thinning algorithm with interference guards," *Pattern Recognition Letters* **15**, pp. 825–832, 1994.

[7] Z. Guo and R. W. Hall, "Fast fully parallel thinning algorithms," *Computer Vision, Graphics and Image Processing* **55-3**, pp. 317–328, 1992.

[8] B.-K. Jang and R. T. Chin, "Analysis of thinning algorithms using mathematical morphology," *IEEE Transactions on pattern analysis and machine intelligence* **12-6**, pp. 514–551, 1990.

[9] S. Yokoi, J. Toriwaki, and T. Fukumura, "Topological properties in digitized binary pictures," *Systems, Computers, Controls* **4-6**, pp. 32–39, 1973.

[10] C. Hilditch, "Linear skeletons from square cupboards," *Machine Intelligence* **4**, pp. 403–420, 1969.

[11] G. Bertrand and G. Malandain, "A new characterization of three-dimensional simple points," *Pattern Recognition Letters* **15**, pp. 169–175, 1994.

[12] C. Ronse, "A topological characterization of thinning," *Theoretical Computer Science* **43**, pp. 31–41, 1986.

[13] T. Kong, "On the problem of determining whether a parallel reduction operator for n-dimensional binary images always preserve topology.," in *SPIE Conference on Vision Geometry*, (Boston - MA), 1993.

[14] C. Ronse, "Minimal test patterns for connectivity preservation in parallel thinning algorithms for binary digital images," *Discrete Applied Mathematics* **21**, pp. 67–79, 1988.

[15] C. M. Ma, "On topology preservation in 3d thinning," *CVGIP: Image Understanding* **59-3**, pp. 328–339, 1994.

[16] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker, "The Hamilton-Jacobi skeleton," in *International Conference on Computer Vision*, pp. 828–834, (Kerkyra - Greece), Sept. 1999.

[17] L. da Fontoura Costa and L. F. Estrozi, “Multiresolution shape representation without border shifting,” *Electronics Letters* **35(21)**, pp. 1829–1830, 1999.

[18] L. da Fontoura Costa, “Robust skeletonization through exact Euclidean distance transform and its application to neuromorphometry,” *Real-Time Imaging* **6**, pp. 415–431, 2000.

[19] G. Borgefors, “Distance transformations in digital images,” *Computer Vision, Graphics and Image Processing* **34**, pp. 344–371, 1986.

[20] A. Manzanera, T. M. Bernard, F. Prêteux, and B. Longuet, “Ultra fast skeleton based on an isotropic fully parallel algorithm,” in *8th Discrete Geometry for Computer Imagery*, vol. 1568, pp. 313–324, Lecture Notes in Computer Science - Springer Verlag, (Marne-La-Vallée - France), Mar. 1999.

[21] R. Cardoner and F. Thomas, “Residuals + directional gaps = skeletons,” *Pattern Recognition Letters* **18**, pp. 343–353, 1997.

[22] R.-Y. Wu and W.-H. Tsai, “A new one-pass parallel thinning algorithm for binary images,” *Pattern Recognition Letters* **13**, pp. 715–723, 1992.