Skeletal representations of orthogonal shapes

Jonàs Martínez, Núria Pla, Marc Vigo

Dept. de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya
Edifici ETSEIB, Diagonal 647, 8a planta
E - 08028 Barcelona, SPAIN

Abstract

In this paper we present two skeletal representations applied to orthogonal shapes of $\mathbb{R}^n$: the cube axis and a family of skeletal representations provided by the scale cube axis. Orthogonal shapes are a subset of polytopes, where the hyperplanes of the bounding facets are restricted to be axis aligned. Both skeletal representations rely on the L_∞ metric and are proven to be homotopically equivalent to its shape. The resulting skeleton is composed of $n-1$ dimensional facets. We also provide an efficient and robust algorithm to compute the scale cube axis in $\mathbb{R}^2$ and compare the resulting skeleton with other skeletal representations.

1. Introduction

Skeletal representations attempt to capture geometric and topological properties of shapes and attempt to understand how these structures encode local and global features, while reducing the dimensionality of the shape. Intuitively, the skeleton of a shape is a thin version of that shape equidistant to its boundary. However, there is not a unique way to define a skeleton. The skeleton usually emphasizes geometrical and topological properties of the shape, such as its connectivity, topology, length and direction. A main challenge is the simplification of the skeleton while retaining the required information. Furthermore, several applications require a skeleton stable to boundary perturbations and shape noise.

A polytope is a bounded region of the n-dimensional space enclosed by a finite number of facets. Orthogonal shapes are a subset of polytopes, where the facets hyperplanes are restricted to be axis aligned. The constrained structure of orthogonal shapes has enabled advances on complex or unsolved problems for arbitrary shapes [44]. In the main, computation of orthogonal shape properties is robust and is less complex compared with more general shape representations.

We present the cube axis and the scale cube axis, two novel skeletal shape representations for the restricted case of orthogonal shapes. The scale cube axis is based upon the scale axis transform [27]. The scale axis transform defines a simplification scheme where shape features are ignored first if they are small relative to their neighbourhood. We demonstrate that the cube axis preserves the homotopy of the original shape and that

the scale cube axis is homotopy equivalent to the shape, for some constant value of the scale factor. We also provide an efficient and robust algorithm for computing the scale cube axis of two dimensional orthogonal polygons in $O(n \log n)$ time and $O(n)$ space, where n is the number of vertices of the input polygon or a scaled polygon from the input. In addition we show the practical application of the scale cube axis with a robust and fast implementation in 2D and compare the scale cube axis of several shapes with the medial axis and the straight skeleton.

2. Previous work

The most well-known skeletal structure is the medial axis [12]. The medial axis of a shape can alternatively be defined as the locus of centres of maximally inscribed balls or as the set of points which have at least two closest points on the shape boundary. It is known that any bounded open subset of arbitrary dimension is homotopy equivalent to its medial axis [29] and finite for subanalytic open sets [18]. Unfortunately, small modifications in a shape can induce large modifications in its medial axis and the exact computation of the medial axis is difficult except for certain classes of shapes [5]. Furthermore, the medial axis of polyhedra has a high combinatorial complexity [38], its components have a high algebraic degree and its exact computation is difficult. We refer the reader to the book of Pizer et al. [39] for a detailed review of medial representations and surveys in [5, 11, 35]

The high combinatorial complexity and instability of the medial axis lead to the development of approximate skeletons. The problem of estimating the medial axis from a set of boundary points was first pointed by Brandt and Algazi [14] who showed that the Delaunay triangulation of a sufficiently dense set of samples contains a reconstruction of the boundary as a subset of its edges. Given a shape in $\mathbb{R}^2$, the Voronoi diagram defined by a sample of points of the boundary approximates its medial axis [15]. Attali and Montanvert [7] showed that a three dimensional shape can be approximated by a finite union of balls. Amenta and Kolluri [4] extended this result and pointed that given a sample of points on the boundary shape, the union of a subset of the Voronoi balls approximates the original shape.

There have been also attempts to linearize and reduce the complexity of the medial axis by simplifying the underlying Voronoi diagram. These strategies try to obtain an approximated skeletal representation only composed by flat surfaces and straight segments. Canny and Donald [17] define a Voronoi diagram based on a measure of distance which is not a true metric. Dey and Zhao [23] approximate the medial axis by computing a subcomplex of the Voronoi diagram that converges to the medial axis as the sampling density increases.

Pruning methods shorten peripheral branches of the medial axis trying to capture its stable part [6, 37]. Pruning can either be performed implicitly as a post processing step or implicitly integrated in the skeleton computation. Sud et al. [41] propose an homotopy preserving algorithm to compute a simplified medial axis. They first compute the medial axis and prune it using the separation angle formed by connecting a point on the medial axis to the closest points on the boundary as a measure of the stability of the medial axis

at that point. Foskey et al. [26] compute a simplified medial axis that rely on the angle formed by the medial axis and its closest neighbours on the surface.

Sud et al. [41] compute a simplified medial axis that removes unstable features of the original medial axis and retain its topological structure. Chazal and Lieutier [19] proposed the λ-medial axis, that contains the set of medial axis points whose closest neighbors on the boundary cannot be enclosed in a ball smaller than a global threshold parameter λ, and proved its geometric stability with respect to small perturbations in terms of Hausdorff distance. Recently, Giesen et al. [27] introduced the scale axis transform that is based on the medial axis transform and the simplification of the shape under multiplicative scaling in order to capture the relevant shape features. Miklos et al. [33] introduced a discrete scale axis transform for 3D shapes.

Chuang et al. [20] apply the generalized potential field model [2] in order to identify the skeleton as potential valleys using a Newtonian potential model in place of the distance function. A one dimensional skeletal representation that often relies on the computation of the medial axis is the curve skeleton. Curve-skeleton can be formulated as a subset of the medial axis with the help of a medial geodesic function [24] or may involve several skeletal representations [21]. Stolpner and Whitesides [40] compute a sample of points, with a guarantee that each of them is within a specified tolerance from the medial axis. Ward and Hamarneh [45] develop a statistical skeletonization framework to identify stable medial points of a group of shapes. Au et al. [8] compute a curve-skeleton by performing geometry contraction using constrained Laplacian smoothing. Sudhalkar et al. [42] introduced the box skeleton of discrete 2D and 3D shapes that identifies the set of centers of all maximal boxes included in the shape and relies on the L_∞ distance instead of the Euclidean one.

The straight skeleton was first proposed as an skeletal representation for simple polygons [3]. While the medial axis is defined using a distance function, the straight skeleton relies on a shrinking process in which the edges of the polygon are moved inwards parallel to themselves at a constant speed. A major drawback of straight skeletons arises from the non-locality effects of reflex vertices: a reflex vertex may collide with a far element. The linear axis [43] is a straight skeleton that approximates the medial axis of a polygon trying to reduce the effects of non-locality of reflex vertices. Eftekharian and Ilieş[25] presented an approach that can generate the medial axis and the straight skeleton within the same formulation using constructive representations of shapes. The straight skeleton can be extended to polyhedra [22]. In three dimensions, the straight skeleton is only composed by points, edges and planar faces. The straight skeleton of polyhedra has a lower combinatorial complexity than the medial axis.

3. Previous definitions

For a point $x \in \mathbb{R}^n$, we denote by x_i the coordinate $i \in \{1 \ldots n\}$. The Euclidean distance between two points $x, y \in \mathbb{R}^n$ is denoted by $d_E(x, y)$ while the L_∞ distance is:

$$d(x, y) = \max_i \left(|x_i - y_i| \right) \tag{1}$$

Notation	Description
$\|A\|$	Cardinality of the set A
$\mathring{A}$	Interior of the set A
$\overline{A}$	Closure of the set A
∂A	Boundary of the set A
A^c	Complement of the set A
$\mathrm{sgn}\,(x)$	Signum function of $x \in \mathbb{R}$ with domain $\{0, 1, -1\}$
$[x, y]$	Closed segment between points x and y
$\|\|x\|\|$	Euclidean norm of vector of $x \in \mathbb{R}^n$

Table 1: Basic notation

The open n-dimensional cube $\mathbb{C}_{c,r}$ with center c and radius r is the set:

$$\mathbb{C}_{c,r} = \{x \in \mathbb{R}^n : d\,(c, x) < r\} \tag{2}$$

Definition 1. *A polytope is an open bounded region of n-dimensional space enclosed by a finite number of hyperplanes. The boundary of an n-dimensional polytope is composed of polytope elements of dimension $\{0 \ldots n - 1\}$. A polytope is bounded by a number of facets, that are polytope elements of dimension $(n - 1)$. These facets are in turn bounded by $(n - 2)$-dimensional ridges of the original polytope and so on.*

Note that our definition of polytope is not restricted to the convex case. Along this paper f denote a facet and e denotes a polytope element of arbitrary dimension. We denote by $\dim\,(e)$ the dimension of the polytope element e. The set $\tau\,(e)$ defines the set of coordinates aligned with e:

$$\tau\,(e) = \{i \in \{1 \ldots n\} : \forall x, y \in e, x_i = y_i\} \tag{3}$$

If $i \in \tau\,(e)$, we denote by e_i the constant value of the coordinate i for all the points in e. The predicate $\mathrm{hyp}\,(e_1, e_2)$ is true if the two polytope elements e_1 and e_2 lie on a common orthogonal hyperplane and false otherwise. The L_∞ distance between a point x and a polytope element e is:

$$d\,(x, e) = \min\,\{d\,(x, y) : \forall y \in e\} \tag{4}$$

The L_∞ distance between two polytope elements is defined by the smallest n-dimensional cube in contact with both elements. The L_∞ bisector between two elements is the locus of points equidistant to the two elements. Table 1 shows other basic notation employed along this paper.

We denote as x_f the orthogonal projection of $x \in \mathcal{S}$ onto the hyperplane containing $f \in \partial\mathcal{S}$. We denote as $\mathcal{S}_f$ the intersection of $\mathcal{S}$ and the hyperplane cointaining f.

4. The cube axis

It is known that the medial axis can be computed exactly for some classes of shape. For example, the medial axis of a union of balls in $\mathbb{R}^n$ can be derived from the Apollonius diagram of the balls or from convex hulls of finitely many balls in $\mathbb{R}^n$ [9, 13]. Under the L_∞ metric the natural counterpart of the union of balls is the union of n dimensional cubes. The cube axis is defined over a restricted class of shapes in $\mathbb{R}^n$ called orthogonal shapes that are well suited for its computation.

Definition 2. *The set $\mathcal{S}$, which we will refer to as orthogonal shape, is an open bounded subset of $\mathbb{R}^n$ limited by axis aligned hyperplanes.*

In fact $\mathcal{S}$ is an open n-dimensional polytope bounded by axis aligned facets. The formulation of a skeletal representation under the L_∞ metric becomes more simple if we consider the shape $\mathcal{S}$ than if we consider an arbitrary open bounded subset of $\mathbb{R}^n$. In addition, the L_∞ bisector between two or more polytope elements of $\partial \mathcal{S}$ can be robustly computed because the polytope elements are restricted to be axis aligned. Our definition of facet is slightly different compared with the traditional definition:

Definition 3. *A facet of $\partial \mathcal{S}$ corresponds to the intersection between $\partial \mathcal{S}$ and a bounding axis aligned hyperplane of $\mathcal{S}$.*

Note that in our setting a facet of $\partial \mathcal{S}$ may be a disconnected set. Moreover if $e \in \partial \mathcal{S}$ we have that $\dim(e) = n - |\tau(e)|$. For $n \geq 2$, the dihedral angle between two neighbouring facets of $\partial \mathcal{S}$ along a ridge is either $\frac{\pi}{2}$ or $\frac{3\pi}{2}$. A reflex ridge of $\partial \mathcal{S}$ is bounded by two facets with dihedral angle of $\frac{3\pi}{2}$. In general, a reflex polytope element e of dimension $0 \geq \dim(e) \geq n - 2$ is bounded by two polytope elements of dimension $\dim(e) + 1$ with dihedral angle $\frac{3\pi}{2}$. The L_∞ bisector of a polytope element is induced by its two bounding polytope elements of higher dimension. The strictly positive L_∞ distance between a point $x \in \mathcal{S}$ and the boundary $\partial \mathcal{S}$ is:

$$d(x, \partial \mathcal{S}) = \min \{d(x, f) : \forall f \in \partial \mathcal{S}\} \tag{5}$$

The cube axis has a tight connection with the L_∞ Voronoi diagram induced by the facets of $\partial \mathcal{S}$.

Definition 4. *For each facet f of $\partial \mathcal{S}$, the Voronoi region of f is the set of points in $\mathcal{S}$ that are closer to f than to other facets of $\partial \mathcal{S}$. The Voronoi diagram is the space partition induced by Voronoi regions. We denote as $E_\infty(\mathcal{S})$ the set of points in $\mathcal{S}$ that are L_∞ equidistant to two or more facets that separate two or more Voronoi regions.*

The set $E_\infty(\mathcal{S})$ may be regarded as an skeletal representation of $\mathcal{S}$. However, $E_\infty(\mathcal{S})$ may not be homotopically equivalent to $\mathcal{S}$ or may not reduce the dimension of $\mathcal{S}$. Consider the left orthogonal shape in $\mathbb{R}^2$ of Figure 1. There exists a set of points in $E_\infty(\mathcal{S})$ that are equidistant to the two collinear segments of the polygon and that are contained in the triangular area drawn in black. In this case $E_\infty(\mathcal{S})$ does not effectively reduce the

5

dimension of $\mathcal{S}$. Papadopoulou et al. [34] consider in $\mathbb{R}^2$ only the outermost boundary of the region of $E_\infty(\mathcal{S})$ induced by two or more collinear elements according to a lexicographic order as proposed in [10]. However, this simple criterion defines a subset of $E_\infty(\mathcal{S})$ that may not be unique as it depends on the lexicographic order.

In the case of the right orthogonal shape in $\mathbb{R}^3$ of the Figure 1, the homotopy equivalence is not preserved because the Voronoi region associated to the face with a hole disconnects the upper connected set and the lower connected set of $E_\infty(\mathcal{S})$. This problem arises when $\mathcal{S}$ have a dimension up to two. Sudhalkar et al. [42] observed the same problem associated with the L_∞ infinity metric when defining the skeleton of a 3D discrete voxelization and proposed a thinning method based on simplicial complexes in order to preserve the homotopy equivalence in these regions.

Note that if we consider that both collinear segments of the left orthogonal polygon of Figure 1 lie on the same segment, by following our definition of facet of $\mathcal{S}$, the set $E_\infty(\mathcal{S})$ will be also disconnected in the same way as the three dimensional example of the shape in the right of Figure 1. Thus, it is necessary to find a skeletal representation that does not only rely on the set of closest facets from a point of $\mathcal{S}$.
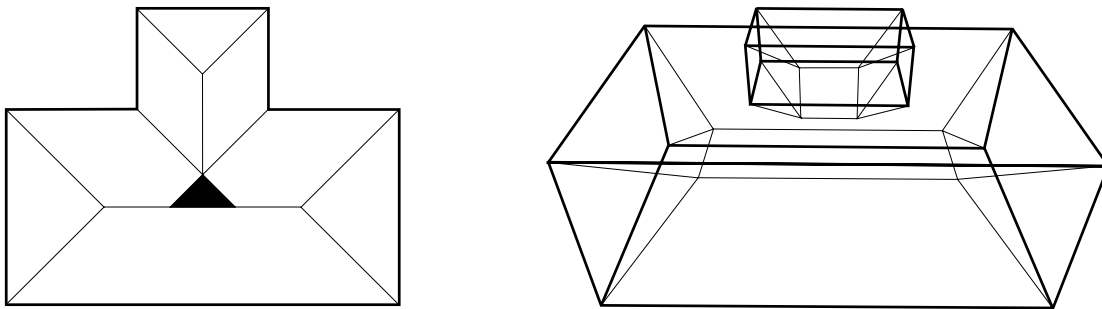

Figure 1: Set of points $E_\infty(\mathcal{S})$ of an orthogonal polygon on the left and of an orthogonal polyhedron on the right.

Although some authors observed that the set $E_\infty(\mathcal{S})$ of orthogonal polygons coincide with the straight skeleton, this is only true if the orthogonal polygon does not have collinear segments. Furthermore, in $\mathbb{R}^3$ the straight skeleton of orthogonal polyhedra may not coincide with $E_\infty(\mathcal{S})$ even if the polyhedron does not have coplanar faces.

If we consider the Euclidean distance, given a point of $\mathcal{S}$ there exists a set of closest points in $\partial\mathcal{S}$ and the number of closest points is used to define the medial axis. However, if we consider the L_∞ distance it does not make sense to define a skeletal representation by considering the number of closest points in $\partial\mathcal{S}$ under the L_∞. As any polytope element of $\partial\mathcal{S}$ is axis aligned, the set of nearest points in $\partial\mathcal{S}$ from a point $x \in \mathcal{S}$ is in general infinite. For this reason, we define the set of closest polytope elements instead of the set of closest boundary points in order to define the cube axis.

Definition 5. *For* $x \in \mathcal{S}$ *the set* $\epsilon(x)$ *is:*

$$\epsilon(x) = \{e \in \partial\mathcal{S} : d(x,e) = d(x,\partial\mathcal{S})\} \tag{6}$$

Because $\partial \mathcal{S}$ is compact, $\epsilon(x)$ is a non empty compact set. Note that the polytope elements are open. As a consequence, we define an appropiate set of polytope elements by discarding the polytope elements that exclusively contain in its boundary any nearest point to x. That is, we only select the polytope elements in which we take the L_∞ distance to a point of a element, and not as a limit.

Property 1. *For $n \geq 2$, if x lies on the L_∞ bisector induced by a reflex polytope element e then the set $\epsilon(x)$ can not contain any of the two polytope elements e_1, e_2 of dimension $\dim(e) + 1$ bounding e. By definition it is not possible to select any point y in the set e_1 or e_2 such that $d(x, y) = d(x, \partial \mathcal{S})$.*

In addition, we need to *filter* the set $\epsilon(x)$ in order to obtain a subset of representative polytope elements. Given a set $\epsilon(x)$ we may have several polytope elements belonging to a common orthogonal hyperplane. The polytope elements of lower dimension are contained in the clausure of elements of higher dimension. In this case, we want to consider only the polytope elements with maximal dimension:

Definition 6. *For $x \in \mathcal{S}$ the set $\sigma_p(x)$ defines a subset of $\epsilon(x)$ such that if several polytope elements of $\epsilon(x)$ lie on a common orthogonal hyperplane, only the polytope elements with the highest dimension are considered to belong to $\sigma_p(x)$, namely:*

$$\sigma_p(x) = \{e_1 \in \epsilon(x) : \forall e_2 \in \epsilon(x), \text{hyp}(e_1, e_2) \Rightarrow \dim(e_1) \geq \dim(e_2)\} \tag{7}$$

Consider that we define a skeletal representation simply as the set of points in $\mathcal{S}$ that define two or more polytope elements of $\sigma_p(x)$. In this case, the corresponding skeletal representation could not be homotopically equivalent to $\mathcal{S}$. Consider the Figure 2, that shows with a continuous line the points such that $|\sigma_p(x)| \geq 2$. There is a set of points such that the set $\sigma_p(x)$ defines only a single segment and that are necessary to maintain the homotopy equivalence.
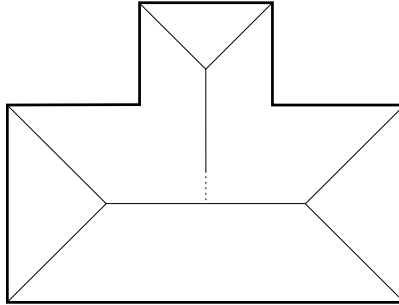

Figure 2: Simple orthogonal polygon. In continuous line the set of points such that $|\sigma_p(x)| \geq 2$. In dotted line, set of points that in addition meet the condition $|\sigma(x)| \geq 2$.

Property 2. *For $x \in \mathcal{S}$, $\sigma_p(x) = \{f\}$ and $x_f \in \mathcal{S}_f$, if $\exists e \in \partial \mathcal{S}_f$ such that $d(x_f, e) = d(x_f, \partial \mathcal{S}_f)$ then we also must have that $d(x, e) = d(x, \partial \mathcal{S})$.*

If we recursively find the closest polytope element in $\mathcal{S}_f$ we can define an extension of the set $\sigma_p(x)$ that identifies the closest polytope elements in any lower dimensional subset of $\mathcal{S}$ (see Property 2). We can refine appropiately the search of the closest polytope elements and erase the ambiguity induced by the L_∞ distance by selecting which polytope elements contained in f are closer to x, when x is projected onto $\mathcal{S}_f$.

Definition 7. *For $x \in \mathcal{S}$ the set $\sigma(x, \mathcal{S})$ defines the set of polytope elements with dimension $\{0 \ldots n-1\}$ such that:*

$$\sigma(x, \mathcal{S}) = \begin{cases} \sigma(x_f, \mathcal{S}_f) & \text{if } \sigma_p(x) = \{f\} \text{ and } x_f \in \mathcal{S}_f \\ \sigma_p(x) & \text{otherwise} \end{cases} \tag{8}$$

We simply denote as $\sigma(x)$ the set $\sigma(x, \mathcal{S})$. Note that at each recursion level, the set $\sigma_p(x)$ applies to the current set $\mathcal{S}$ that may be a lower dimensional subset of the original orthogonal shape. The recursive definition of $\sigma(x, \mathcal{S})$ distinguishes two cases. If $\sigma_p(x, \mathcal{S})$ defines a single facet f of $\partial \mathcal{S}$ and $x_f \in \mathcal{S}_f$, the closest polytope elements are found in $\partial \mathcal{S}_f$. In this case, the closest polytope element may be a ridge or a lower dimensional boundary element of $\partial \mathcal{S}$. The orthogonal projection is done in order to distinguish which boundary polytope elements of f may be closer to x when projected onto $\mathcal{S}_f$. Note that this recursive procedure can only be done if $\dim(\mathcal{S}) \geq 2$. Otherwise, the set $\sigma(x, \mathcal{S})$ is just equal to the set of closest polytope elements $\sigma_p(x)$.

Property 3. *For $x \in \mathcal{S}$ the set $\sigma(x)$ is non empty and defines polytope elements of $\partial \mathcal{S}$.*

Definition 8. *(**Cube axis**) The cube axis CA of $\mathcal{S}$ is the set of points that have at least two closest polytope elements of $\partial \mathcal{S}$:*

$$CA = \{x \in \mathcal{S} : |\sigma(x)| \geq 2\} \tag{9}$$

The recursive definition of $\sigma(x)$ extends naturally to any dimension and guarantees that the cube axis is *centered* in the regions where the implicit ambiguity of L_∞ distance may arise. For example, in the left orthogonal polygon of the Figure 1 instead of arbitrarily selecting a segment of the triangular region drawn in black, the points of the cube axis correspond to the central bisector contained in the triangular region. Thus, the cube axis is unique for each different input orthogonal shape. Note that the set of points that only lie on the L_∞ bisector induced by one reflex polytope element do not belong to the cube axis (see Property 1). For example, although in the left orthogonal polygon of the Figure 1 the bisector induced by the two reflex vertices are included in $E_\infty(\mathcal{S})$ this does not happen in the case of Figure 2. The set $\sigma(x)$ extends the characterization given by the set of polytope elements of $\sigma_p(x)$ in order to obtain and homotopical equivalent skeleton with respect $\mathcal{S}$ (see Theorem 1) that reduces the dimension of $\mathcal{S}$ (see Lemma 1).

While the medial axis is defined by taking the set of closest boundary points from a point, the cube axis is defined by taking the set $\sigma(x)$ of closest polytope elements from a point x. The cube axis together with the distance to the boundary can also be used as a complete shape descriptor of $\mathcal{S}$. The cube transform is the set of n-dimensional cubes whose center lie on the CA and with radius equal to $d(x, \partial \mathcal{S})$.
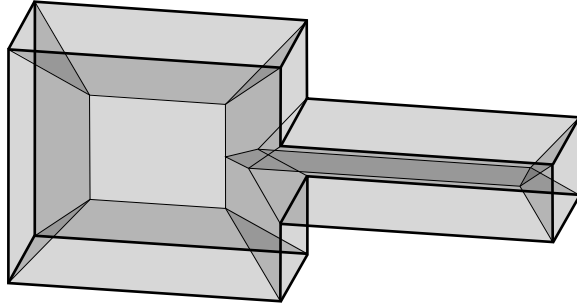
Figure 3: Cube axis of a simple orthogonal polyhedron.

Definition 9. *(**Cube transform**) The cube transform CT of $\mathcal{S}$ is the set of n-dimensional cubes associated to each point of the cube axis:*

$$CT = \left\{ \mathbb{C}_{x,d(x,\partial \mathcal{S})} : x \in CA \right\} \tag{10}$$

In general, the proof of the important properties of the cube axis takes advantage of a piecewise partition that may be induced by the polytope elements of $\partial \mathcal{S}$ or the cube axis. Thus, we are able to separately analyze each part, which is more simple than trying to derive global properties, and then proceed to analyze the boundary between the different parts. By following this simple strategy we are able to derive global results.

Lemma 1. *The cube axis is composed of unoriented $n-1$ dimensional facets that can be axis aligned or rotated $\frac{\pi}{4}$ with respect an orthogonal axis.*

Proof. As the polytope elements of $\partial \mathcal{S}$ are restricted to be axis aligned its induced L_∞ bisector must be constrained and as a consequence the structure of the cube axis. We analyze first the case of the L_∞ bisector induced by two facets of $\partial \mathcal{S}$. Consider that $\sigma(x) = \{f_1, f_2\}$. We have two different cases depending on the relative position between f_1 and f_2. If f_1 and f_2 are parallel then x must belong to a facet of CA also parallel to f_1 and f_2. For $n \geq 2$, if f_1 and f_2 are not parallel then x must belong to a facet of CA contained in the L_∞ bisector of f_1 and f_2 that is rotated $\frac{\pi}{4}$ with respect an orthogonal axis.

For $n \geq 2$, consider secondly that $\sigma(x) = \{e_1, e_2\}$ such that e_1 and e_2 have the same dimension $\{0 \ldots n-2\}$ and lie on a common intersection $\mathcal{S}_f$ during the recursion process of $\sigma(x, \mathcal{S})$. The L_∞ bisector induced by those polytope elements of a dimensional space lower than $n-1$ is defined in the same manner previously explained for the base case of two facets of dimension $n-1$. Thus, the dimension reduction property of CA remains for any lower dimension.

Finally, if a point of CA does not meet any of the other two cases it must belong to the boundary of one or more facets of CA. If $|\sigma(x)| > 3$, then it is clear that x belongs to the boundary. If $\sigma(x)$ defines two polytope elements of different dimension, the polytope element with the lowest dimension is contained in two or more polytope elements with dimension equal to the polytope element of $\sigma(x)$ with the highest dimension. This happens when the element of lowest dimension is reflex. $\square$

Definition 10. *For $x \in \mathcal{S} \setminus CA$ and $e \in \partial\mathcal{S}$ the function $v(x,e)$ and $v(x)$ define an n-dimensional vector such that:*

$$v(x,e) = \{v \in \mathbb{R}^n : i \in \tau(e), j \in \{1 \ldots n\} \setminus \tau(e), v_i = x_i - e_i, v_j = 0\} \tag{11}$$

$$v(x) = \{v(x,e) : \sigma(x) = \{e\}\} \tag{12}$$

The function $v(x)$ is constructed in order to properly define the deformation retract. The function $v(x,e)$ is employed to demonstrate the continuity of $v(x)$. Note that the set $\sigma(x)$ always defines a single polytope element if $x \in \mathcal{S} \setminus CA$. As the set $\sigma(x)$ is well defined in $\mathcal{S} \setminus CA$ the function $v(x)$ is also well defined and $||v(x)|| > 0$. We define as $\hat{v}(x) = \frac{v(x)}{||v(x)||}$ the unit vector of $v(x)$.

Definition 11. *For $e \in \partial\mathcal{S}$ we define the set $\lambda(e)$ as:*

$$\lambda(e) = \{x \in \mathcal{S} \setminus CA : \sigma(x) = \{e\}\} \tag{13}$$

The set $\lambda(e)$ contains all the points of $\mathcal{S} \setminus CA$ whose associated vector $v(x)$ is defined by the same polytope element e. Note that the set $\lambda(e)$ can be empty.

Lemma 2. *For $x \in \mathcal{S} \setminus CA$ the function $v(x)$ is continuous.*

Proof. If $x \in \lambda(\overset{\circ}{e})$ the function $v(x)$ is continuous because the term e_i of $v(x,e)$ (see Equation 12) remains constant. Therefore, we need to demonstrate that continuity is also maintained on the boundary of $\lambda(e)$, that is in $\lambda(e) \cap \partial\lambda(e)$. Consider that e is a facet f at some moment of the recursion of $\sigma(x)$ (see Equation 8) with dimension $\{1 \ldots n-1\}$. If we have that $x_f \notin \mathcal{S}_f$ and at the same time $xf \in \partial\mathcal{S}_f$, then x must belong to $\lambda(f) \cap \partial\lambda(f)$ (see Figure 4). That is, if $x \in \lambda(f) \cap \partial\lambda(f)$, the neighbouring region of $\lambda(e)$ is defined by a poltyope element of the boundary of e. Let e_1 be the polytope element of $\partial\mathcal{S}_f$ containing x_f, with dimension $\{0 \ldots n-2\}$. The point x is contained in all the orthogonal hyperplanes passing through e_1 that are aligned with a coordinate of $\tau(e_1) \setminus \tau(f)$. Once we have identified the points in $\mathcal{S} \setminus CA$ where the continuity must be preserved, we evaluate $v(x)$ in the two sides. If the equality $v(x,f) = v(x,e_1)$ is always true, then we can conclude that $v(x)$ is continuous in $\mathcal{S} \setminus CA$. This is clearly seen, as $\forall i \in \tau(e_1) \setminus \tau(f)$ we always have that the term x_i of the equations $v(x,f)$ and $v(x,e_1)$ are equal because x lies in all the hyperplane that pass through e_1 and that are aligned with $\tau(e_1) \setminus \tau(f)$. $\square$

Property 4. *Let $e \in \partial\mathcal{S}$ and $l \in \mathbb{R}^*$. If $x_1, x_2 \in \lambda(e)$ such that $x_2 = x_1 + lv(x_1)$ then we have that $\hat{v}(x_2) = \hat{v}(x_1)$. In other words, $v(x)$ has the same direction in a line contained in $\lambda(e)$.*

Due to the Property 4, we can associate to each point $x \in \mathcal{S} \setminus CA$ a single point of CA by extending the ray with direction $v(x)$ and origin x to the closest point of CA.
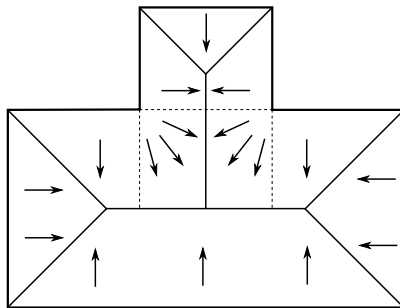
Figure 4: Simple example in $\mathbb{R}^2$. Cube axis shown in continuous line. Direction of $v(x)$ shown in with arrows. The points of $\mathcal{S} \setminus CA$ such that $\lambda(f) \cap \partial\lambda(f)$ (see Lemma 2) are shown in dotted lines.

Definition 12. *For $x \in \mathcal{S} \setminus CA$ and $l \in \mathbb{R}^+$ being the minimal Euclidean distance from x to CA through the ray following the direction of $v(x)$ (see Property 4), the function $\psi(x)$ maps x to a point of the cube axis such that:*

$$\psi(x) = \{y \in CA : y = x + l\hat{v}(x)\} \tag{14}$$

For $x \in CA$ the function $\psi(x)$ is just:

$$\psi(x) = x \tag{15}$$

Lemma 3. *For $x \in \mathcal{S} \setminus CA$ the function $\psi(x)$ is continuous.*

Proof. We first show that the function $\psi(x)$ is well defined, that is for every point $x \in \mathcal{S} \setminus CA$ the function $\psi(x)$ describes a unique point. For $e \in \partial\mathcal{S}$ and $x \in \lambda(e)$, namely assume that there exists a point $p \in CA$ with minimal distance from x along the direction $v(x)$, such that $\psi(x) = p$. The maximal n-dimensional cube $\mathbb{C}_{p,d(p,\partial\mathcal{S})}$ is empty, and contains every n-dimensional cube centered at a point on the segment $[p, x]$, hence all of those are empty as well, and must belong to $\lambda(e)$. So there can not be other point of CA in $[p, x]$.

Let analyze the continuity of $\psi(x)$. We know that $v(x)$ is continuous, so it remains to see that l is bounded. Let consider the set of points in $\mathcal{S} \setminus CA$ such that $\psi(x)$ are contained in a common facet f of CA. It is clear that distance l is bounded for the set of points that are contained in $\mathring{f}$. It remains to show that l is still bounded when $\psi(x)$ belongs to a boundary element ∂f. As we have seen, the points of the boundary must be defined with the two closest polytope elements defining f (see Lemma 1), and therefore there can not be a discontinuity in the boundary of f. Thus, we proof the continuity of $\psi(x)$ in $x \in \mathcal{S} \setminus CA$.

$\square$

Theorem 1. *The cube axis is a deformation retract of S.*

Proof. We follow a similar retraction as the one defined by Wolter [46] to demonstrate the homotopy equivalence of the medial axis with respect to C^2-smooth shapes. The continuous map:

$$F : \mathcal{S} \times [0, 1] \to \mathcal{S} \tag{16}$$

is a deformation retraction of a space $\mathcal{S}$ onto a subspace CA if for every $x \in \mathcal{S}$, $t \in [0, 1]$ and $a \in CA$ we have:

$$F(x, 0) = x, \; F(x, 1) \in CA \text{ and } F(a, t) = a \tag{17}$$

In other words, a deformation retraction is a homotopy between a retraction and the identity map on $\mathcal{S}$ [32]. The subspace CA is called a deformation retract of $\mathcal{S}$. A deformation retract is a special case of homotopy equivalence. We define the homotopy $F(x, t)$ by:

$$F(x, t) = (1 - t) x + t \psi(x) \tag{18}$$

Note that $F(x, t)$ is a linear interpolation between x and $\psi(x)$. We know that $v(x)$ and $\psi(x)$ are continuous functions when $x \in \mathcal{S} \setminus CA$ (see Lemma 2 and 3) so we have that $F(x, t)$ is continuous in $x \in \mathcal{S} \setminus CA$. It remains to show that $F(x, t)$ is continuous when $x \in CA$. By definition of $\psi(x)$ it is clear that when $x \in CA$, $\psi(x) = x$. For $x_0 \in CA$, $t_0 \in [0, 1]$, let (x_n, t_n) be any sequence in $\mathcal{S} \times [0, 1]$ converging to (t_0, x_0). Then, we have that $F(x, t)$ is continuous in CA if $F(x_n, t_n)$ converges to $F(x_0, t_0)$. As $\psi(x) = x$ for $x \in CA$ and $\psi(x)$ is continuous in $\mathcal{S} \setminus CA$ the sequence $t_n \psi(x)$ converges to 0. Therefore the sequence $F(x_n, t_n)$ must converge to x_0.

By definition of $F(x, t)$ it is clear that $F(x, 0) = 0$. We also have by definition that for $a \in CA$, $\psi(a) = a$. Thus the term $d_e(a, \psi(x))$ is equal to zero and then we have that for every $a \in CA$ and $t \in [0, 1]$, $F(a, 1) = a$. Finally, for $F(x, 1)$ we directly have that $F(x, t) = \psi(x) \in CA$. Therefore the proof of Theorem 1 is now complete.

$\square$

Theorem 2. *The orthogonal shape $\mathcal{S}$ can be reconstructed from its cube transform CT.*

Proof. It is clear that:

$$\mathcal{S} \supset \bigcup_{x \in CA} \mathbb{C}_{x, d(x, \partial \mathcal{S})} \tag{19}$$

It remains to show that:

$$\mathcal{S} \subset \bigcup_{x \in CA} \mathbb{C}_{x, d(x, \partial \mathcal{S})} \tag{20}$$

That is, for every point $y \in \mathcal{S}$ there exists a point x such that $y \in \mathbb{C}_{x, d(x, \partial \mathcal{S})}$. If $y \in CA$ it is obvious that we can select $x = y$. If $y \notin CA$, we show that if we select $x = \psi(y)$ the next condition met:

$$\mathbb{C}_{y,d(y,\partial\mathcal{S})} \subset \mathbb{C}_{\psi(y),d(\psi(y),\partial\mathcal{S})} \tag{21}$$

As we have seen in Lemma 3 the maximal cube associated to $\psi(y)$ is empty and contains any maximal cube associated to a point of the segment $[y, \psi(y)]$. Therefore we prove Equation 21 and the proof of Theorem 2 is now complete.

$\square$

5. The scale cube transform

We give analogous definitions of the scale cube transform based upon the scale axis transform [27]. The multiplicative distance of a point $x \in \mathbb{R}^n$ to the n-dimensional cube $\mathbb{C}_{c,r}$ is:

$$d_m(x, \mathbb{C}_{c,r}) = \frac{d(x, c)}{r} \tag{22}$$

The multiplicative distance of a point $x \in \mathbb{R}^n$ to the set $\mathcal{S}$ is the minimum of all multiplicative distances to the set of n-dimensional cubes CT:

$$d_m(x, CT) = \min\{d_m(x, \mathbb{C}_{c,r}) : \forall \mathbb{C}_{c,r} \in CT\} \tag{23}$$

Definition 13. *For any point $x \in \mathcal{S}$, we denote by $\sigma_m(x)$, the set of closest n-dimensional cubes under the multiplicative distance:*

$$\sigma_m(x) = \{\mathbb{C}_{c,r} \in CT : d_m(x, \mathbb{C}_{c,r}) = d_m(x, CT)\} \tag{24}$$

Definition 14. *For $s \geq 0$, the multiplicatively s-scaled shape is:*

$$\mathcal{S}_s = \bigcup_{\mathbb{C}_{c,r} \in CT} \mathbb{C}_{c,sr} \tag{25}$$

Definition 15. *(Scale cube axis) The s-scale cube axis CA_s is the cube axis (see Definition 8) of the shape $\mathcal{S}_s$.*

Definition 16. *(Scale cube transform) For $s \geq 1$ the scale cube transform of $\mathcal{S}$ is:*

$$SCT_s = \{\mathbb{C}_{c,r/s} : \mathbb{C}_{c,r} \in CT_s\} \tag{26}$$

Giesen et al. [27] apply the property that the multiplicative euclidean distance is semi-concave [16] in order to demonstrate the properties of the scale axis transform. However, the multiplicative distance under the L_∞ distance is not semiconcave and we require another technique to demonstrate the properties of the SCT. As in the proof of the properties of the cube axis, we take advantage of the simpler structure of orthogonal shapes in order to prove the properties of the scale cube axis.

For $x \in \mathbb{R}^n$ the open orthogonal halfspace $\mathcal{H}(x, i, r)$ partitioning the coordinate i in the orientation $r \in \{+1, -1\}$ is defined as:

$$\mathcal{H}(x, i, r) = \{y \in \mathbb{R}^n : ry_i > rx_i\} \tag{27}$$

Definition 17. *For $\mathbb{C}_{c,r}$ and $x \in \mathbb{R}^n \setminus \{c\}$, we define the set of halfspaces $H(x, \mathbb{C}_{c,r})$ as:*

$$H(x, \mathbb{C}_{c,r}) = \{h \subset \mathbb{R}^n : i \in \{1 \ldots n\}, |x_i - c_i| = d(x,c), h = \mathcal{H}(x, i, \operatorname{sgn}(x_i - c_i))\} \quad (28)$$

Definition 18. *For $\mathbb{C}_{c,r}$ and $x \in \mathbb{R}^n \setminus \{c\}$ we define the set $\mathcal{E}(x, \mathbb{C}_{c,r})$ as:*

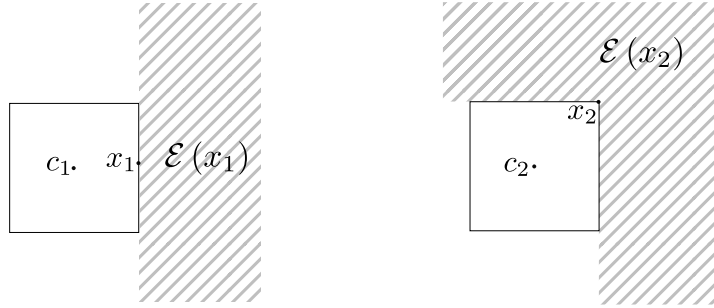$$\mathcal{E}(x, \mathbb{C}_{c,r}) = \bigcup_{h \in H(x, \mathbb{C}_{c,r})} h \quad (29)$$

Figure 5: Two squares and its associated region $\mathcal{E}(x)$ shown in diagonal dotted lines.

Property 5. *For $\mathbb{C}_{c,r}$ and $x \in \mathbb{R}^n \setminus \{c\}$, $\forall y \in \mathcal{E}(x, \mathbb{C}_{c,r})$ we have that $d_m(y, \mathbb{C}_{c,r}) > d_m(x, \mathbb{C}_{c,r})$ and $[xy] \cap \mathbb{C}_{c,d(x,c)} = \emptyset$.*

As $\mathcal{E}(x, \mathbb{C}_{c,r})$ is defined as the union of orthogonal halfspaces passing through the point x, the segment defined by x and any point of $\mathcal{E}(x, \mathbb{C}_{c,r})$ never intersects $\mathbb{C}_{c,d(x,c)}$.

Definition 19. *For $x \in \mathbb{R}^n \setminus CA$ we define the set $\mathcal{E}(x)$ as:*

$$\mathcal{E}(x) = \bigcap_{\mathbb{C}_{c,r} \in \sigma_m(x)} \mathcal{E}(x, \mathbb{C}_{c,r}) \quad (30)$$

Definition 20. *A point $x \in \mathbb{R}^n$ is a critical point of $d_m(x, CT)$ if $d_m(x, CT)$ is a maximum in a neighbourhood of x.*

Lemma 4. *A point $x \in \mathbb{R}^n \setminus CA$ is a critical point of $d_m(x, CT)$ if and only if $\mathcal{E}(x) = \emptyset$.*

Proof. We know from Property 5 that if $y \in \mathcal{E}(x, \mathbb{C}_{c,r})$, then the multiplicative distance from y increases. The set $\mathcal{E}(x)$ is defined as the intersection of several $\mathcal{E}(x, \mathbb{C}_{c,r})$ associated to the nearest n-dimensional cubes to $x \in \mathbb{R}^n \setminus CA$. Thus, $\mathcal{E}(x)$ defines the space where the multiplicative distance can increase. It is always possible to select a point in the neihbourhood of x in $\mathcal{E}(x)$ such that the multiplicative distance increases in all the points. Thus, if $\mathcal{E}(x) = \emptyset$, it is not possible to move from x to another point such that the multiplicative distance increases and this means that x is a critical point. Conversely, if $\mathcal{E}(x) = \emptyset$ means that there not exist a neighbourhood on x where the multiplicativ distance can increase so x must be a critical point. □
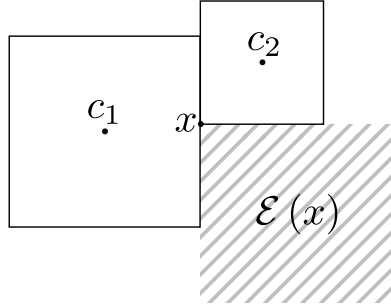
Figure 6: Set $\mathcal{E}(x)$ induced by two squares shown in diagonal dotted lines.

Lemma 5. *Let e_1 and e_2 be two polytope elements of $\partial \mathcal{S}$ and $x \in \lambda(e_2)$. If $e_1 \cap e_2 \neq \emptyset$ we always have that $d(x, e_2) \leq d(x, e_1)$.*

Proof.

- If $e_1 \subset e_2$ we have that $d(x, e_2) \leq d(x, e_1)$.

- If $e_2 \subset e_1$ we have that $d(x, e_2) = d(x, e_1)$.

- As a consequence we always have that $d(x, e_2) \leq d(x, e_1)$.

$\square$

Definition 21. *For $x \in \mathcal{S} \setminus CA$ we define the set $\alpha(x)$ of polytope elements of $\partial \mathcal{S}$ as:*

$$\alpha(x) = \{e \in \partial \mathcal{S} : \mathbb{C}_{c,r} \in \sigma_m(x), e \in \sigma(c), d(x, e) < d(c, \partial \mathcal{S})\} \tag{31}$$

The set $\alpha(x)$ defines the polytope element of $\partial \mathcal{S}$ associated to the set of closest n-dimensional cubes of x under the multiplicative distance.

Lemma 6. *For $x \in \mathcal{S} \setminus CA$ the set $\alpha(x)$ is non empty.*

Proof.

- We know that $\sigma_m(x)$ is never empty.

- Suppose that $\mathbb{C}_{c,r} \in \sigma_m(x)$ and that $e \in \sigma(c)$. Suppose in order to create a contradiction that $\alpha(x) = \emptyset$. Then, we always have that $d(x, e) \geq d(c, \partial \mathcal{S})$. As $d(x, e) = d(x, \partial \mathcal{S})$ we arrive in a contradiction as we have the inequality $d(x, \partial \mathcal{S}) \geq d(c, \partial \mathcal{S})$, that is impossible if $x \in \mathcal{S} \setminus CA$.

$\square$

Lemma 7. *For $x \in \mathcal{S} \setminus CA$ we have that $\forall e_1, e_2 \in \alpha(x), e_1 \cap e_2 \neq \emptyset$.*

Proof.

15

- Suppose that $x \in \lambda(e_2)$. Then we have $d(x, e_2) = d(x, \partial S)$.

- Suppose that $e_1 \in \alpha(x)$ and that $\exists \mathbb{C}_{c_1, r} \in \sigma_m(x)$ such that $e_1 \in \sigma(c_1)$. As $e_1 \in \alpha(x)$ we have that $d(x, e_1) < d(c_1, \partial S)$.

- Suppose, in order to create a contradiction, that $e_1 \cap e_2 = \emptyset$. Then, according to Lemma 5 we can have that $d(x, e_2) > d(x, e_1)$ and equivalenty $d(x, \partial S) > d(x, e_1)$.

- So we have the next inequations:

$$d(x, e_1) < d(x, \partial S)$$

$$d(x, e_1) < d(c_1, \partial S)$$

As a consequence we can have that $d(c_1, \partial S) < d(x, \partial S)$ and we arrive in a contradiction as it is not possible if $x \in S \setminus CA$.

$\square$

Definition 22. *For $x \in S \setminus CA$, $\alpha'(x)$ defines the intersection of all the polytope elements of $\alpha(x)$:*

$$\alpha'(x) = \bigcap_{e \in \alpha(x)} e \tag{32}$$

Property 6. *For $x \in S \setminus CA$, $\alpha'(x)$ define a single polytope element of ∂S.*

Lemma 8. *The distance function $d_m(x, CT)$ has no critical points in $S \setminus CA$.*

Proof. $\square$

Definition 23. *We define the weak feature size wfs analogously to [19] in our setting as:*

$$wfs(S) = \min\{d_m(x, CT) : x \in \mathbb{R}^n \setminus CA, x \text{ is a critical point}\} \tag{33}$$

Theorem 3. *For $s < wfs(S)$, the s-scale cube axis CA_s is homotopy equivalent to S.*

Proof. $\square$

6. Computing the SCT of orthogonal shapes in $\mathbb{R}^2$

We present an efficient and robust algorithm to compute the SCT in $\mathbb{R}^2$. In the subsequent of this section, S denotes an orthogonal shape of $\mathbb{R}^2$ and ∂S denotes its associated orthogonal polygon. Our algorithm requires as input the set of contours representing ∂S and the SCT scale value $s > 1$. The general steps are as follows:

1. Compute the cube axis CA of S.
2. Retrieve a finite set of rectangles necessary to compute the s-scaled shape S_s.
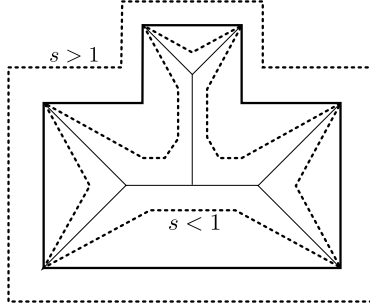3. Compute the s-scale cube axis CA_s of S_s.

Figure 7: Orthogonal polygon and its cube axis shown in continuous line. The associated s-scaled shape is shown in dotted lines for $s > 1$ and $s < 1$.

For $\mathbb{R}^2$, a facet of $\partial\mathcal{S}$ corresponds to an edge and a ridge corresponds to a vertex bounding an edge. The cube axis is composed of line segments. There are two kinds of segments of CA depending on its L_∞ bisector (see Lemma 1): segments whose bisector is defined by two segments with the same orthogonal direction, and those whose bisector is defined by two segments with different orthogonal direction.

Lemma 9. *(MARTINEZ et al. [31]) The L_∞ Voronoi diagram of an orthogonal polygon has $O(n)$ vertices and can be robustly computed in $O(n \log n)$ time and $O(n)$ space, where n is the number of vertices of the orthogonal polygon.*

Property 7. *The cube axis of an orthogonal polygon can be directly derived from its the L_∞ Voronoi diagram.*

The approach presented in [31] is able to compute the L_∞ Voronoi diagram of orthogonal polygons that contain geometrical degeneracies or coplanar edges. The cube axis can be derived from the L_∞ Voronoi diagram as the vertices and segments of both representations have a direct correspondence. Note that this property does not hold for higher dimensions. The bisector emanating from reflex vertices is just erased. As the colinear segments are considered to belong to a different segment, the L_∞ ambiguity (see left shape of Figure 1 is resolved by selecting the central bisector of the area, that exactly corresponds to the cube axis.

Once we computed the cube axis we extract a subset of rectangles that cover the cube transform CT and that when are grown represent the shape $\mathcal{S}_s$, for $s > 1$.

Lemma 10. *(GÜTING [28]) The contour of the union of n iso-oriented rectangles can be computed in $O(n \log n + p)$ time and $O(n)$ space, where p is the number of edges of the orthogonal polygon.*

Lemma 11. *For $\mathbb{R}^2$ and $s > 1$, the s-scaled shape $\mathcal{S}_s$ can be computed in $O(n \log n + p)$ time and $O(n)$ space, where n is the number of vertices of $\mathcal{S}$ and p is the number of edges of $\mathcal{S}_s$.*

Proof. We first show that it is possible to represent the set of cubes CT by a finite set of rectangles that are grown. Consider first that the bisector of a segment of CA is defined by two segments with the same orthogonal direction. The set of squares of CT associated to the segment can be bounded by a rectangle defined by the two squares associated to the endpoints of the segment. Growing this rectangle is equivalent to grow all the associated squares. Let consider the second case, where the segment of CA is defined by two segments with different orthogonal direction. In this case, it suffices to consider the biggest square between the two squares associated to the enpoinds of the segment in order to bound all the squares associated to the segment.

There are $O(n)$ segments in CA (see Lemma 9), each of them defining a single bounding rectangle so the number of rectangles necessary to cover the s-scaled shape is $O(n)$. Thus, we can compute the contour of $\mathcal{S}_s$ by using the algorithm of Güting [28], that computes the contour of the union of rectangles in $O(n \log n + p)$ time and $O(n)$ space. Note that $\mathcal{S}_s$ is also an orthogonal polygon when $s > 1$. $\qquad\square$
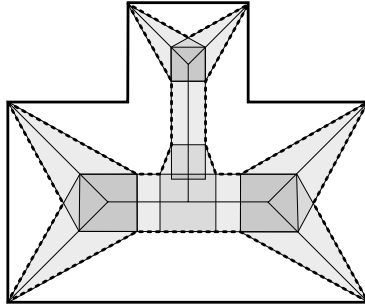

Figure 8: Orthogonal polygon and its cube axis shown in continuous line. The associated s-scaled shape is shown in dotted lines for $s < 1$ and is found by obtaining the contour of the union of rectangles and convex polygons.

Property 8. *For $\mathbb{R}^2$, the maximum number of segments of CA converging in a common point of CA is four.*

Lemma 12. *For $\mathbb{R}^2$ and $s < 1$, the s-scaled shape $\mathcal{S}_s$ has $O(n)$ vertices and can be computed in $O(n)$ time and space, where n is the number of vertices of $\mathcal{S}$.*

Proof. We first show that is possible to shrink all the cubes associated to a segment of CA without requiring any discretization step. Then we show that the union of every shrunken element has $O(n)$ boundary vertices. Consider first that the bisector of a segment of CA is defined by two segments with the same orthogonal direction. In this case, it suffices to shrink the rectangle defined by the two squares associated to its endpoints (as when $s > 1$) in order to cover the result of shrinking all the squares associated to the segment. If the bisector of the segment of CA is defined by two segments with different orthogonal

18

direction, the convex polygon with at most six edges containing the two squares of the endpoints covers all the shrunken squares induced by the segment.

Given an segment of CA, its associated shrink convex polygon (see Figure 8) can only cross the convex polygons associated to its neighbouring segments in CA. As the maximum number of neighbouring segments is constant (see Property 8), thus the number of crossings between the convex polygons remains linear. By traversing every segment associated to a different edge of the orthogonal polygon, we are able to straightforwardly construct the contour of $\mathcal{S}_s$ in $O(n)$ time and space.

$\square$

Lemma 13. *For $\mathbb{R}^2$ and $s > 1$, the SCT of $\mathcal{S}$ can be computed in $O(n \log n)$ time and $O(n)$ space, where n is the number of vertices of $\mathcal{S}$ or $\mathcal{S}_s$.*

Proof. The *SCT* requires two different computations: computation of the L_∞ Voronoi diagram of $\mathcal{S}$ and computation of the s-scaled shape $\mathcal{S}_s$. We have seen that the L_∞ Voronoi diagram can be computed in $O(n \log n)$ time (see Lemma 9). For $s > 1$ the s-scaled shape can be computed in $O(n \log n + p)$ time and $O(n)$ (see Lemma 11). As we may have that $p > n$, the final complexity is dominated by the number of vertices of the scaled shape $\mathcal{S}_s$ or of the original shape $\mathcal{S}$ (the one with more vertices). $\square$

7. Results

Although the introduced skeletal representations are applied to orthogonal shapes, it is possible to extend our approach to any kind of shape by considering the approximation framework described in [5]. We first find an orthogonal shape approximating an input shape. Secondly, we compute is cube axis and simplify it with the multiplicative scaling, providing an stable skeletal approximation of the input shape. Our approach does not require any filtering of the skeleton and its complexity is bounded by the number of vertices of the input orthogonal polygon or the grown orthogonal polygon.

Approximating a shape by an orthogonal shape is straightforward and efficient. Vigo et al. [44] proposed efficient algorithms to extract the boundary of 2D and 3D orthogonal shapes from several input shape representations. A fast way to extract an orthogonal shape is to first discretize the input shape and then extract the boundary of the discretized data. The discretization can be greatly accelerated with the help of modern GPUs [36] when dealing with 3D shapes and is straightforward in the 2D case.

We have computed the medial axis, straight skeleton and scale cube axis of several polygons in order to compare their skeletons. The medial axis and the straight skeleton of polygons have been implemented using the C++ CGAL library [1]. The algorithm to extract the scale cube axis of orthogonal polygons has been implemented in C++ and the source code consists of about 2000 lines of code. Table 2 shows the results on a PC Intel E6600 2.40 Ghz with 3.2 GB RAM for four sample objects that are shown in Figure 9. In the case of the straight skeleton and the cube axis, the segments of the skeleton emanating from the vertices of the input polygon are erased. Note that the *SCT* of polygons requires to approximate the input polygon by an orthogonal polygon and we have to select an

orientation of the Cartesian coordinates for the approximation. In order to present the results, we define a metric to quantify the similarity between the original shape and the approximated orthogonal shape. We define a normalized hausdorff distance between two sets (d_{HN}) as the hausdorff distance between the two sets divided by the diameter of the bounding box of both sets. For all the models in the Table 2, the approximation of the input polygon with an orthogonal polygon is accurate $(d_{HN} < 0.001)$. Figure 9 demonstrates that the scale cube axis is able to simplify the skeleton while retaining its important topological features. In addition, the time to compute the scale cube axis is comparable the performance of CGAL library and is also slightly faster than the straight skeleton computation, although the computation of the union of rectangles is done with a greedy algorithm. Thus, the computation time of the SCT can be improved more.

Figure 10 illustrates the fact that a less accuracy of the shape approximation produces skeletons composed of less vertices and less accurate. We also see that as the scale factor s increases both the scale cube axis and the s-scaled shape are simplified. In fact, by taking the the s-scaled shape the SCT may be also employed for shape simplification.

Shape	Vertices	Medial axis		Straight skeleton		Scale cube axis		
		Elements	Time	Vertices	Time	Vertices	Time	s-scale
(a)	533	530	0.07	528	0.75	674	0.19	1.3
(b)	573	581	0.05	1112	0.7	1169	0.23	1.1
(c)	521	537	0.06	517	0.68	1390	0.22	1.2
(d)	1414	1418	0.19	1412	1.45	1006	0.26	1.2

Table 2: Running time of each skeletal representation computation in seconds.

8. Conclusions

We have presented two new skeletal representations of shapes defined according the L_∞ metric and applied to the orthogonal shapes. We also introduced an algorithm to compute the scale cube axis of orthogonal polygons and shown its ability to simplify the skeletal structure. The scale cube axis may serve as an alternative to the scale axis transform in applications where the skeleton is no restricted to be *centered* in the sense of euclidean metric. In addition, geometric computations applied to orthogonal shapes are robust thanks to the restricted alignment of n dimensional faces and approximating an orthogonal shape from other classes of shape representations is efficient. The scale cube axis allows a tradeoff between the skeletal representation resolution and the computation time.

As a future work, we plan to develop and algorithm to compute the scale cube axis in $\mathbb{R}^3$ by using the algorithm presented in [30]. As with the scale axis transform, for large values of s the cube axis may lie outside $\mathcal{S}$ as the scale cube axis is not a subset of the cube axis. That property is not desirable and we will investigate how to avoid it.
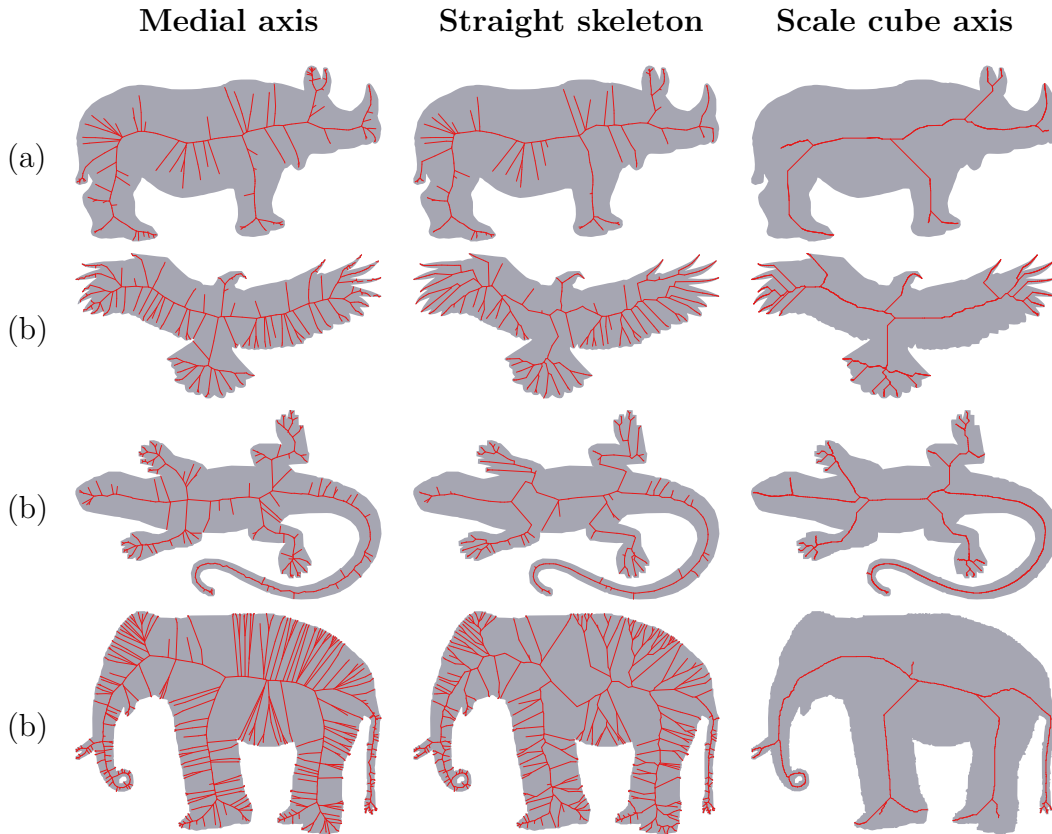
| Medial axis | Straight skeleton | Scale cube axis |

Figure 9: Sample objects. Input polygon shown in gray. Skeleton shown in red.

9. Acknowledgements

[1] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org.

[2] N. Ahuja and Jen-Hui Chuang. Shape representation using a generalized potential field model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:169–176, 1997.

[3] O. Aichholzer, F. Aurenhammer, D. Alberts, and B. Gärtner. A novel type of skeleton for polygons. *Journal of Universal Computer Science*, 1:752–761, 1995.

[4] Nina Amenta and Ravi Krishna Kolluri. The medial axis of a union of balls. *Computational Geometry Theory and Applications*, 20:25–37, 2001.

[5] Dominique Attali, Jean Daniel Boissonnat, and Herbert Edelsbrunner. Stability and computation of medial axes - a state-of-the-art report. In *Mathematical Foundations*

of Scientific Visualization, Computer Graphics, and Massive Data Exploration, Mathematics and Visualization, pages 109–125. Springer Berlin / Heidelberg, 2009.

[6] Dominique Attali, Gabriella di Baja, and Edouard Thiel. Pruning discrete and semi-continuous skeletons. In *Image Analysis and Processing*, volume 974 of *Lecture Notes in Computer Science*, pages 488–493. Springer Berlin / Heidelberg, 1995.

[7] Dominique Attali and Annick Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding*, 67:261–273, 1997.

[8] Oscar Kin-Chung Au, Chiew-Lan Tai, Hung-Kuo Chu, Daniel Cohen-Or, and Tong-Yee Lee. Skeleton extraction by mesh contraction. *ACM Transactions on Graphics*, 27:44:1–44:10, 2008.

[9] F. Aurenhammer and H. Imai. Geometric relations among Voronoi diagrams. *Geometriae Dedicata*, 27(1):65–75, 1988.

[10] Franz Aurenhammer. Voronoi diagrams. a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23:345–405, 1991.

[11] Silvia Biasotti, Dominique Attali, Jean-Daniel Boissonnat, Herbert Edelsbrunner, Gershon Elber, Michela Mortara, Gabriella Sanniti Baja, Michela Spagnuolo, Mirela Tanase, and Remco Veltkamp. Skeletal structures. In *Shape Analysis and Structuring*, pages 145–183. Springer Berlin Heidelberg, 2008.

[12] Harry Blum. A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, pages 362–380. MIT Press, 1967.

[13] J.D. Boissonnat and M.I. Karavelas. On the combinatorial complexity of Euclidean Voronoi cells and convex hulls of d-dimensional spheres. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 305–312, 2003.

[14] Jonathan W. Brandt and V. Ralph Algazi. Continuous skeleton computation by Voronoi diagram. *CVGIP: Image Understanding*, 55:329–338, 1991.

[15] J.W. Brandt. Convergence and continuity criteria for discrete approximations of the continuous planar skeleton. *CVGIP: Image Understanding*, 59:116–124, 1994.

[16] Piermarco Cannarsa and Carlo Sinestrari. *Semiconcave functions, Hamilton-Jacobi equations, and optimal control.* Birkhäuser, 2004.

[17] J. Canny and B. Donald. Simplified Voronoi diagrams. *Discrete and Computational Geometry*, 3(1):219–236, 1988.

[18] F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems*, 10:149–170, 2004.

[19] Frédéric Chazal and André Lieutier. The λ medial axis. *Graphical Models*, 67:304–331, 2005.

[20] Jen-Hi Chuang, Chi-Hao Tsai, and Min-Chi Ko. Skeletonization of three-dimensional object using generalized potential field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:1241–1251, 2000.

[21] Nicu D. Cornea and Deborah Silver. Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 13:530–548, 2007.

[22] Erik Demaine, Martin Demaine, Jeffrey Lindy, and Diane Souvaine. Hinged dissection of polypolyhedra. In *Algorithms and Data Structures*, volume 3608 of *Lecture Notes in Computer Science*, pages 205–217. Springer Berlin / Heidelberg, 2005.

[23] Tamal K. Dey and Wulue Zhao. Approximate medial axis as a Voronoi subcomplex. In *Proceedings of the seventh ACM symposium on Solid modeling and applications*, pages 356–366, 2002.

[24] T.K. Dey and J. Sun. Defining and computing curve-skeletons with medial geodesic function. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, pages 143–152, 2006.

[25] Ata A. Eftekharian and Horea T. Ilieş. Distance functions and skeletal representations of rigid and non-rigid planar shapes. *Computer-Aided Design*, 41:865–876, 2009.

[26] Mark Foskey, Ming C. Lin, and Dinesh Manocha. Efficient computation of a simplified medial axis. In *Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 96–107, 2003.

[27] Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. The scale axis transform. In *Proceedings of the 25th annual symposium on Computational geometry*, pages 106–115, 2009.

[28] R. Hartmut Güting. An optimal contour algorithm for iso-oriented rectangles. *Journal of Algorithms*, 5(3):303–326, 1984.

[29] André Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. In *Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 65–75, 2003.

[30] J. Martinez, M. Vigo, and N. Pla-Garcia. Skeleton computation of orthogonal polyhedra. *Computer Graphics Forum*, 30(5):1573–1582, 2011.

[31] Jonas Martinez, Marc Vigo, Nuria Pla-Garcia, and Dolors Ayala. Skeleton computation of an image using a geometric approach. In *Proceedings of Eurographics*, pages 13–16, 2010.

[32] William S. Massey. *Algebraic topology, an introduction.* Springer-Verlag, 1967.

[33] Balint Miklos, Joachim Giesen, and Mark Pauly. Discrete scale axis representations for 3D geometry. *ACM Transactions on Graphics*, 29:101:1–101:10, 2010.

[34] E. Papadopoulou and D. T. Lee. The L_∞ Voronoi diagram of segments and VLSI applications. *International Journal of Computational Geometry and Applications*, 11:503 – 508, 2001.

[35] Stephen M. Pizer, Kaleem Siddiqi, Gabor Székely, James N. Damon, and Steven W. Zucker. Multiscale medial loci and their properties. *International Journal of Computer Vision*, 55:155–179, 2003.

[36] Michael Schwarz and Hans-Peter Seidel. Fast parallel surface and solid voxelization on GPUs. *ACM Transactions on Graphics*, 29(6):179:1–179:10, 2010.

[37] D. Shaked and A.M. Bruckstein. Pruning medial axes. *Computer Vision and Image Understanding*, 69:156–169, 1998.

[38] M. Sharir and P.K. Agarwal. *Davenport-Schinzel sequences and their geometric applications.* Cambridge University Press, 1995.

[39] Kaleem Siddiqi and Stephen Pizer. *Medial Representations: Mathematics, Algorithms and Applications.* Springer Publishing Company, Incorporated, 2008.

[40] Svetlana Stolpner and Sue Whitesides. Medial axis approximation with bounded error. In *Proceedings of the 2009 Sixth International Symposium on Voronoi Diagrams*, pages 171–180, 2009.

[41] Avneesh Sud, Mark Foskey, and Dinesh Manocha. Homotopy-preserving medial axis simplification. In *Proceedings of the 2005 ACM symposium on Solid and physical modeling*, pages 39–50, 2005.

[42] Atul Sudhalkar, Levent Gürsöz, and Friedrich Prinz. Box-skeletons of discrete solids. *Computer-Aided Design*, 28(6–7):507 – 517, 1996.

[43] Mirela Tanase and Remco C. Veltkamp. A straight skeleton approximating the medial axis. In *European Symposium on Algorithms*, volume 3221 of *Lecture Notes in Computer Science*, pages 809–821. Springer Berlin / Heidelberg, 2004.

[44] Marc Vigo, Nuria Pla, Dolors Ayala, and Jonas Martinez. Efficient algorithms for boundary extraction of 2D and 3D orthogonal pseudomanifolds. *Graphical Models*, 74(3):61 – 74, 2012.

[45] Aaron D. Ward and Ghassan Hamarneh. The groupwise medial axis transform for fuzzy skeletonization and pruning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32:1084–1096, 2010.

[46] Franz-erich Wolter. Cut locus and medial axis in global shape interrogation and representation. In *MIT Design Laboratory Memorandum 92-2 and MIT Sea Grant Report*, 1992.
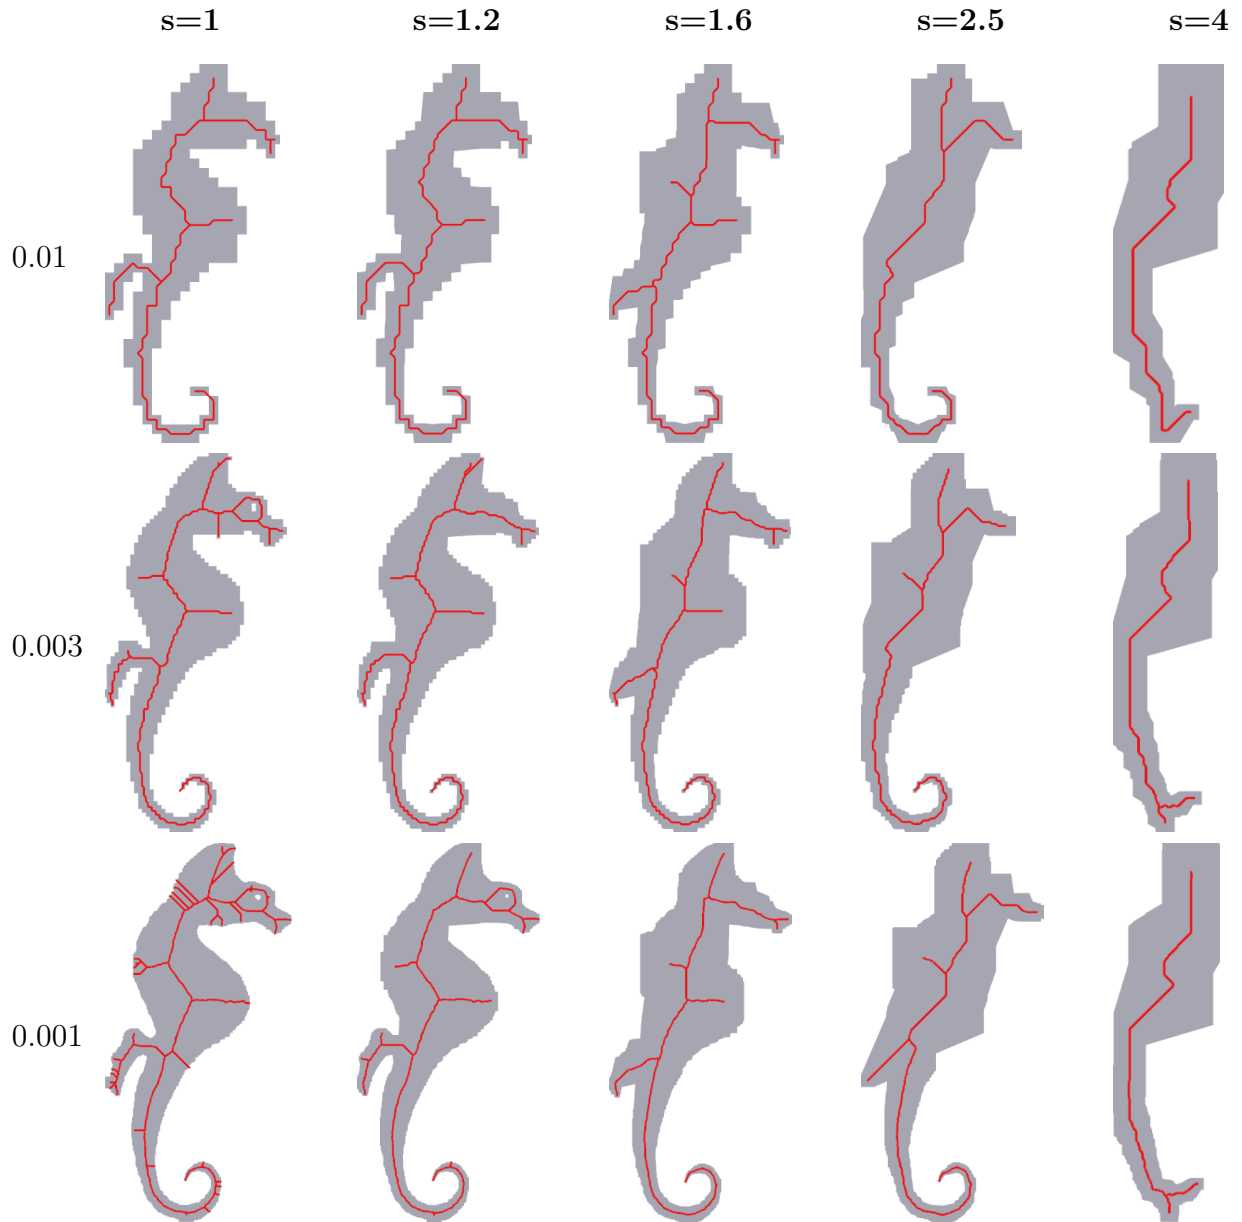
Figure 10: First column indicates the d_{HN} of the orthogonal polygons with respect the original shape associated to each row. First row shows the scale factor values associated to each column. The s-scaled shape $\mathcal{S}_s$ drawn in gray and the scale cube axis drawn in red are shown together.