The Scale Axis Transform

Joachim Giesen
Friedrich Schiller University
Jena, Germany
giesen@informatik.uni-jena.de

Balint Miklos*
Applied Geometry Group
ETH Zürich
balint@inf.ethz.ch

Mark Pauly
Applied Geometry Group
ETH Zürich
pauly@inf.ethz.ch

Camille Wormser
Applied Geometry Group
ETH Zürich
cwormser@inf.ethz.ch

ABSTRACT

We introduce the *scale axis transform*, a new skeletal shape representation for bounded open sets $\mathcal{O} \subset \mathbb{R}^d$. The scale axis transform induces a family of skeletons that captures the important features of a shape in a scale-adaptive way and yields a hierarchy of successively simplified skeletons. Its definition is based on the medial axis transform and the simplification of the shape under multiplicative scaling: the s-scaled shape $\mathcal{O}_s$ is the union of the medial balls of $\mathcal{O}$ with radii scaled by a factor of s. The s-scale axis transform of $\mathcal{O}$ is the medial axis transform of $\mathcal{O}_s$, with radii scaled back by a factor of $1/s$. We prove topological properties of the scale axis transform and we describe the evolution $s \mapsto \mathcal{O}_s$ by defining the multiplicative distance function to the shape and studying properties of the corresponding steepest ascent flow. All our theoretical results hold for any dimension. In addition, using a discrete approximation, we present several examples of two-dimensional scale axis transforms that illustrate the practical relevance of our new framework.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms and Problems**]: Geometrical problems and computations; I.3.5 [**Computational Geometry and Object Modeling**]: Curve, surface, solid, and object representations

General Terms

Theory, Algorithms

Keywords

medial axis, skeleton, topology

*Corresponding author.

SCG'09, June 8–10, 2009, Aarhus, Denmark.

1. INTRODUCTION

Skeletal representations of shapes are an important tool in digital geometry analysis and processing. A central goal is to define skeletal structures that capture geometric and topological properties of shapes and to understand how these structures encode local and global features. Common applications include shape matching, segmentation, simplification, and meshing. One of the most prominent examples of a skeletal structure is the medial axis transform that describes a shape as the union of all maximal open balls contained in the shape [3]. A well-known disadvantage of the medial axis transform is its inherent instability under small perturbations: two similar shapes can have very different medial axis transforms. A common approach to deal with this issue is to extract the stable parts of the medial axis by filtering irrelevant branches using a suitable stability criterion. Typically, these criteria are based on angle [9, 2, 16], distance [7], or area [17] measures computed at a medial axis point and its closest neighbors on the surface. Such methods rely on a threshold computed at one single medial ball to prune unwanted parts of the medial axis and do not naturally adapt to the local scale of the shape. Hence they are best suited for removing branches corresponding to surface noise, but are not designed to simplify the medial axis based on a comparison of features in a scale-adaptive way.

Perception studies have shown that the human visual system relies on skeletal structures to understand and identify shapes. More specifically, empirical evidence in [18] confirms that *"[...] the visual system represents simple spatial regions by their medial axes [...]"* and that *"[...] the medial representation arises in a scale-specific way [...]"*. We formalize this concept of scale for skeletal computations and introduce the *scale axis transform*, a new mathematical structure that provides a systematic treatment of spatial adaptivity based on a relative measure of feature importance.

This measure relies on a size comparison between a feature and the surrounding ones, effectively defining a simplification scheme where features are ignored first if they appear small relative to their neighborhood (see Figures 8 and 1). To construct a skeleton representing the significant features of a shape, we compute the medial axis of the shape after an evolution that simplifies less significant features. Using this construction, we overcome the limitations of approaches that filter the medial axis based on a measure computed at one single medial ball. Our evolution, called multiplicative scal-

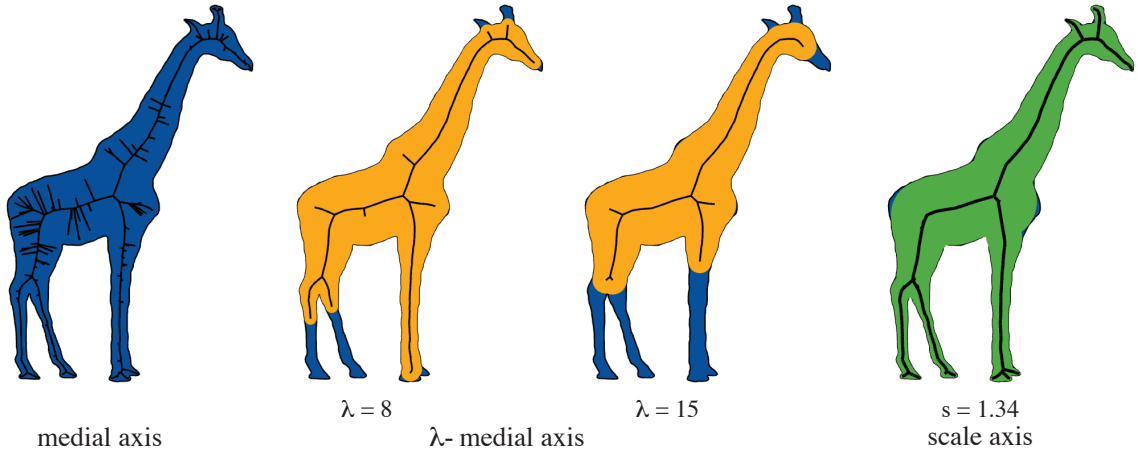

Figure 1: Comparison of the medial axis, the λ-medial axis filtration, and our new proposed scale axis. Blue represents the shape, orange the shape reconstructed from the λ-medial axis approximation, green is the reconstruction from the scale axis approximation.

ing, formalizes the notion of *local contrast of size* by growing or shrinking the shape to detect relative feature size. The scaling leads us to study the properties of a function that we call *multiplicative distance*. From its steepest ascent flow we deduce topological properties of both the scaling process and the scale axis transform. More specifically, we use non-smooth critical point theory to show that homotopy is preserved under multiplicative shrinking. Consequently, we obtain a lower bound on when topological changes occur for the scale axis transform. These results, and the various proofs and methods we introduce, lay the ground for further study of this new framework. In addition, we demonstrate the practical relevance of our approach with our prototype implementation and show different examples that illustrate potential applications of the scale axis transform.

1.1 Related work

The medial axis transform is a fundamental construction in computational geometry and computer vision and has been studied extensively, see [3, 14] for recent surveys. A multitude of medial axis filtration algorithms rely on the angle formed by the medial axis and its closest neighbors on the surface. Methods have been presented that explicitly enforce topology preservation during simplification [16]. Many angle based approaches are successful in removing noisy branches in practical datasets [9, 2], but no geometric stability guarantee has been presented for angle based filtration methods. The λ-medial axis [7] contains the set of medial axis points whose closest neighbors on the boundary cannot be enclosed in a ball smaller than the global threshold parameter λ. Chazal and Lieutier [7] showed that the λ-medial axis preserves topology for a restricted range of values of λ and proved geometric stability with respect to small perturbations in terms of Hausdorff distance. Such an absolute global λ threshold is well suited for describing and removing surface noise, but it is not designed to identify the important skeletal branches for shapes with features on different scales, the problem we are addressing in this paper (see Figure 1). For a comparison of the skeleton families of our approach and related medial axis filtrations, see more videos and images at `http://people.agg.ethz.ch/balintmiklos/scale-axis/`

1.2 Basic definitions

Let $d(x, y)$ be the Euclidean distance between $x, y \in \mathbb{R}^d$, $[xy]$ the closed segment between x and y, and (xy) the line connecting them. We use the notation $\langle u, v\rangle$ for the dot product of two vectors and $\|u\|$ for the norm of u. A function f is Lipschitz if $\exists K \geq 0$ such that $\|f(x)-f(y)\| \leq K\|x-y\|$ for all x, y. The notation $|A|$ is used for the cardinality of a set A. We denote with $\oplus$ the Minkowski sum, i.e. for sets A and B we have $A \oplus B = \{a + b \mid a \in A,\ b \in B\}$. The open ball $B(c, r)$ is the set $B(c, r) = \{x \in \mathbb{R}^d \mid d(c, x) < r\}$. The set $\mathcal{O}$, which we will often refer to as *shape*, is an open bounded subset of $\mathbb{R}^d$. We denote by $\overline{\mathcal{O}}$ the closure of $\mathcal{O}$, while $\partial\mathcal{O} = \overline{\mathcal{O}} \setminus \mathcal{O}$ is the boundary of $\mathcal{O}$. The boundary of a ball $\partial B(c, r)$ is a $(d-1)$-sphere, and its intersection with some halfspace forms a subset of $\partial B(c, r)$ which we call a spherical cap. The medial axis $\mathrm{M}(\mathcal{O})$ of $\mathcal{O}$ is the set of points in $\mathcal{O}$ with at least two closest points in $\partial\mathcal{O}$. The medial axis transform is the set of maximal balls inside $\mathcal{O}$ centered at the medial axis: $\mathrm{MAT}(\mathcal{O}) = \{B(x, d(x, \partial\mathcal{O})) \mid x \in \mathrm{M}(\mathcal{O})\}$. Balls in $\mathrm{MAT}(\mathcal{O})$ are called medial balls. Medial balls of $\mathcal{O}$ cannot cover each other and their union is exactly $\mathcal{O}$.

2. THE SCALE AXIS TRANSFORM

The definition of the scale axis transform is based on the multiplicative scaling operation that is designed to eliminate locally small features. Using the medial axis transform, a shape can be considered simply as a union of balls, where every ball contributes to the description of the shape. Therefore, the task of finding locally small features can be posed as the problem of finding locally small balls, i.e. balls that have a significantly larger ball close to them. We detect such configurations using a very simple construction: scale the radius of every ball by a factor $s > 1$ (see Figure 2). As the example of the red ball illustrates, small balls will be covered by nearby large balls for a small scaling factor, hence they can be considered irrelevant for the union of all balls. We define the multiplicative scaling as this construction for any $s > 0$ and open set $\mathcal{O} \subset \mathbb{R}^d$:

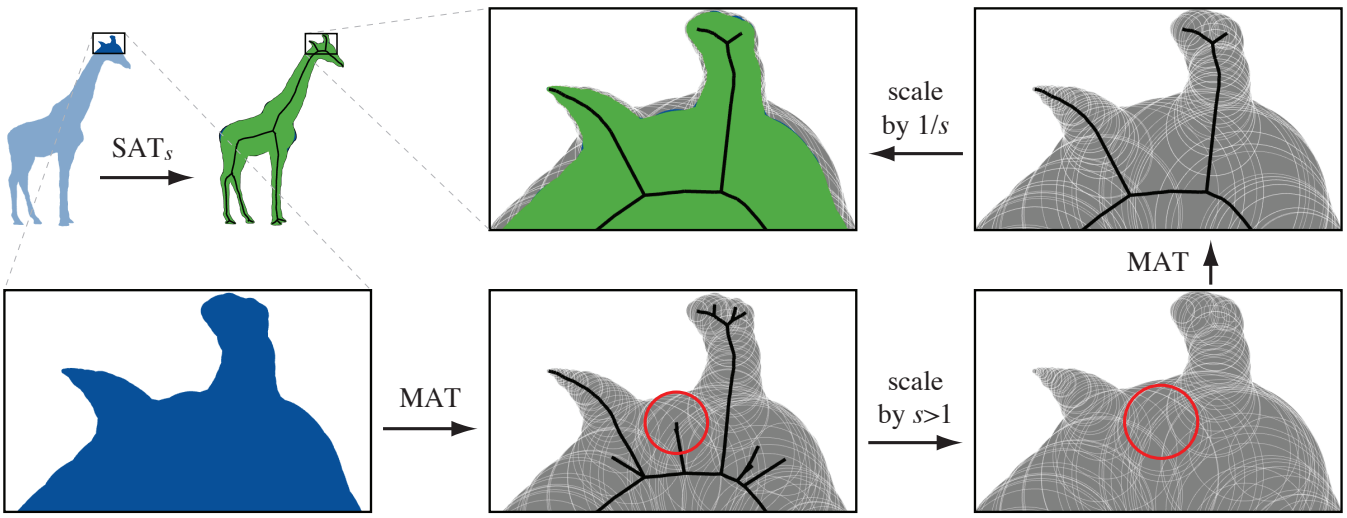

Figure 2: The construction of the scale axis transform. For illustration purposes we show only a representative subset of the medial balls.

DEFINITION 1. *For an open set $\mathcal{O}$ and $s > 0$, the multiplicatively s-scaled shape is* $\mathcal{O}_s = \bigcup_{B(c,r)\in \mathrm{MAT}(\mathcal{O})} B(c, sr)$.

We have $\mathcal{O}_1 = \mathcal{O}$, and $s \mapsto \mathcal{O}_s$ is an increasing function for the inclusion. For $s < 1$ we say we shrink the shape, for $s > 1$ we grow it. Because the growing eliminates locally small balls, we can use the medial axis of the grown shape as the significant skeleton of the original shape. Since growing is only used for detecting this contrast of sizes, we compensate for the overall growth of the shape and define the scale axis transform as the set of scaled-back medial balls of the grown shape.

DEFINITION 2. *For $s \geq 1$, the* s-scale axis transform *of an open set $\mathcal{O} \subset \mathbb{R}^d$ is* $\mathrm{SAT}_s(\mathcal{O}) = \{B(c, r/s) \mid B(c, r) \in \mathrm{MAT}(\mathcal{O}_s)\}$. *We call the set of centers of the balls in* $\mathrm{SAT}_s(\mathcal{O})$ *the* s-scale axis.

For $s = 1$, the scale axis is identical to the medial axis. With increasing s, the scale axis gradually ignores less important features of $\mathcal{O}$, leading to a successive simplification of the skeletal structure (see Figure 8). We consider the union $(\mathcal{O}_s)_{1/s} = \bigcup_{B\in \mathrm{SAT}_s} B$ to be the a simplified version of $\mathcal{O}$ at scale s. Note that similar shrinking constructions were used by Edelsbrunner for the definition of skin surfaces, where the "convex hull" of a finite set of balls is shrunken to describe a surface [10]. In our case, the set of medial balls is shrunken to bring back the shape to the original scale.

We study the behavior of the scale axis transform and the evolution $s \mapsto \mathcal{O}_s$ by defining a scalar function whose sublevelsets are the scaled shapes. Let $\mu_{B(c,r)}$ denote the Minkowski functional of $B(c, r)$, i.e., $\mu_{B(c,r)}(x)$ is the factor by which B has to be scaled so that x lies on its boundary: $\mu_{B(c,r)}(x) = d(c, x)/r$. This is a special case of the multiplicative distance, which we define for any open set:

DEFINITION 3. *The* multiplicative distance *of a point $x \in \mathbb{R}^d$ to a an open set $\mathcal{O} \subset \mathbb{R}^d$ is the infimum of all multiplicative distances to the medial balls of $\mathcal{O}$:* $\mu_{\mathcal{O}}(x) = \inf\{\mu_B(x) \mid B \in \mathrm{MAT}(\mathcal{O})\} = \inf\{d(c, x)/r \mid B(c, r) \in \mathrm{MAT}(\mathcal{O})\}$.

This definition directly yields that each sublevelset of $\mu_{\mathcal{O}}$ is a specific scaled shape: $\mathcal{O}_s = \mu_{\mathcal{O}}^{-1}([0, s)) = \{x \mid \mu_{\mathcal{O}}(x) \in [0, s)\}$. We illustrate the behavior of multiplicative scaling with a simple cone-like planar shape $\mathcal{C}$: the union of a ball $B(c, r)$ and the region bounded by the two tangent segments to B drawn from an outside point t (see Figure 3.a). The medial axis of $\mathcal{C}$ is the segment open at one endpoint $[ct]\backslash\{t\}$. As we move along the medial axis towards t, the radii of the medial balls decrease linearly and converge to 0. If we multiplicatively shrink $\mathcal{C}$, its medial axis stays unchanged. The shape simply becomes a "thinner" cone as illustrated by the isolines. As we grow $\mathcal{C}$, the medial axis changes only in the instant when all medial axis balls simultaneously become tangent to the tip of the cone. At that moment, the medial axis collapses into a point and the grown shape becomes a disk. The evolution is described by the multiplicative distance $\mu_{\mathcal{C}}$ shown in Figure 3.a-b. The point t is a discontinuity point of $\mu_{\mathcal{C}}$. Along the medial axis the limit of $\mu_{\mathcal{C}}$ at t is 0, but the function limit at t along the boundary curve is 1. Moreover, all the isolines in the range of $[0, \mu_{\mathcal{C}}(t)]$ meet in the limit at the point t, and the function value is the largest value corresponding to all such isolines. Thus in a steepest ascent flow induced by $\mu_{\mathcal{C}}$, the point t would not move for a period of time until the largest ball reaches it and starts governing its movement. This example illustrates an interesting property of the scale axis transform: features are not necessarily smoothed out gradually. Locally small and sharp features can be preserved until the supporting neighborhood is considered unimportant as a whole.

Before studying the details of the multiplicative distance and its induced steepest ascent flow, we illustrate the nature of this flow with another 2D example of two intersecting balls, a smaller ball B_1 and a larger ball B_2 (see Figure 3.c). Let us consider the minimum of the two multiplicative distances induced by each ball. The set of points equidistant to B_1 and B_2, $\{x \mid \mu_{B_1}(x) = \mu_{B_2}(x)\}$, is a circle $M_{1,2}$ called the Möbius bisector, that passes through the intersection points of the boundaries of the two balls. Every point in the interior of the ball bounded by $M_{1,2}$ is closer to B_1, hence the steepest ascent vector points away radially from the center

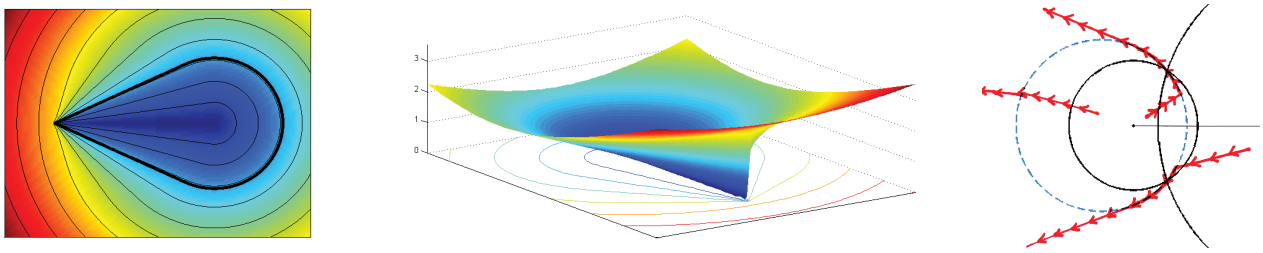

a) b) c)

Figure 3: The multiplicative distance of a cone-like shape in a 2D and a surface plot. On the right the bisector and multiplicative flow induced by the two black balls.

of B_1. Similarly, the steepest ascent vector of a point outside of the ball bounded by $M_{1,2}$ is the vector pointing away from the center of B_2. Interestingly, the flow of a point on the bisector follows the bisector as long as the steepest ascent of B_2 points into the interior of the ball bounded by $M_{1,2}$, leading to the flow lines shown in Figure 3.c. This flow is closely related to the multiplicatively weighted Voronoi diagram which is a special case of the Möbius diagram (see [4] for a description of such diagrams). In contrast to this, the flow of the Euclidean distance to a set of points is simpler as it has piecewise linear flow curves, and is closely related to the Voronoi diagram as discussed in [5]. Note that, different to this discrete setting, the multiplicative distance $\mu_{\mathcal{O}}$ has an infinite number of balls defining the distance to $\mathcal{O}$. Figure 4 shows a comparison between the Euclidean and the multiplicative distance.

3. OVERVIEW OF PROOF TECHNIQUES

To show topological properties of the scale axis transform, we use the multiplicative distance $\mu_{\mathcal{O}}$ to study the evolution of a shape for a varying scale parameter s. While $\mu_{\mathcal{O}}$ is not smooth, we show that it is locally semiconcave (see [6, 13] for a presentation of semiconcave functions). As a consequence, we can define a gradient vector field as the steepest ascent vector field, and this vector field is known to be integrable into a flow. We use these tools to prove homotopy equivalence between the shape $\mathcal{O}$ and its scaled version $\mathcal{O}_s$, for $s < 1$. Note that topology may change for $s > 1$, and we determine when these changes happen.

Our strategy is similar to the one Lieutier [12] used to prove homotopy equivalence of a shape and its medial axis, which can in fact be reformulated in the semiconcave setting. In [12], the gradient flow of the Euclidean distance to the boundary is used to construct a deformation retract. In contrast, our shape scaling is described by a weighted distance to the medial axis whose gradient flow expands the shape (see Figure 5). To construct a deformation retraction, we have to somehow reverse this gradient flow of a non-smooth function. For this purpose, we use standard regularization techniques in variational analysis (see the chapter on mollifiers in [15], and a typical use of such techniques by Grove [11]). Central to the proof of homotopy equivalence is our theorem that insures that no critical points of the multiplicative distance are encountered during the flow. We show that critical points can be located only on the medial axis and outside the shape. In Lieutier's case, the equivalent result followed directly (and was implicitly used) from the definition of the distance function. These properties allow us to conclude the proof that homotopy is preserved under shrinking. Furthermore, similar to the λ-medial axis [7], we obtain a topology preservation guarantee during the simplification step (i.e. the growing step), assuming that the scaling factor stays below a certain value dependent on the shape. Nevertheless, the scale axis transform is designed such that it induces meaningful topology changes even after this point.

4. PROPERTIES OF THE MULTIPLICATIVE DISTANCE

We now present important properties of the multiplicative distance that lay the ground for the topological results we prove in section 5. We start with some definitions: let us consider the set $\overline{\mathrm{MAT}(\mathcal{O})}$, the closure of the set of the medial balls, which is a set of open balls defined as follows: every open ball $B(c,r) \subset \mathbb{R}^d$ can be mapped to a point in $\mathbb{R}^d \times \mathbb{R}_+^*$ with the radius as its last coordinate. The closure of the set of points in $\overline{\mathbb{R}^d \times \mathbb{R}_+}$ representing the balls in $\mathrm{MAT}(\mathcal{O})$ defines the set $\overline{\mathrm{MAT}(\mathcal{O})}$, a set composed of open balls and points. Now, for all $x \in \mathbb{R}^d$, we can define the set of *closest balls* as $\Gamma_{\mathcal{O}}(x) = \{B \in \overline{\mathrm{MAT}(\mathcal{O})} \mid \mu_B(x) = \mu_{\mathcal{O}}(x)\}$. Defining $\Gamma_{\mathcal{O}}(x)$ as a subset of $\overline{\mathrm{MAT}(\mathcal{O})}$ rather than one of $\mathrm{MAT}(\mathcal{O})$ implies that $\Gamma_{\mathcal{O}}(x)$ is never empty, as a consequence of Lemma 8. But the main reason why we define it that way is that it is needed for Lemma 11 to be true. Let the *radius function* be $r_{\mathcal{O}}(x) = \sup\{r \mid B(c,r) \in \Gamma_{\mathcal{O}}(x)\}$. One can see that for any $x, y \in \mathbb{R}^d$ we have $\mu_{\mathcal{O}}(y) \leq \mu_{\mathcal{O}}(x) + \frac{d(x,y)}{r_{\mathcal{O}}(x)}$. More specifically, this implies the following lemma.

LEMMA 4. *The multiplicative distance is Lipschitz function on any set bounded away from* $\mathrm{M}(\mathcal{O})$. *It is continuous on* $\mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$ *and upper semicontinuous on* $\mathbb{R}^d$.

The closest ball function $\Gamma_{\mathcal{O}}$ has a similar semicontinuity property as the "unweighted" closest point function for the Euclidean distance (see Lemma 4.6 in [12]):

LEMMA 5 (SEMICONTINUITY OF CLOSEST BALLS). *For any* $x \in \mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$ *and any* $\varepsilon > 0$, *exists* $\alpha > 0$, *such that* $\forall y \in B(x,\alpha)$ *and* $\forall B(c_y, r_y) \in \Gamma_{\mathcal{O}}(y)$, $\exists B(c_x, r_x) \in \Gamma_{\mathcal{O}}(x)$ *such that we have* $c_y \in c_x \oplus B(0,\varepsilon)$ *and* $|r_x - r_y| < \varepsilon$.

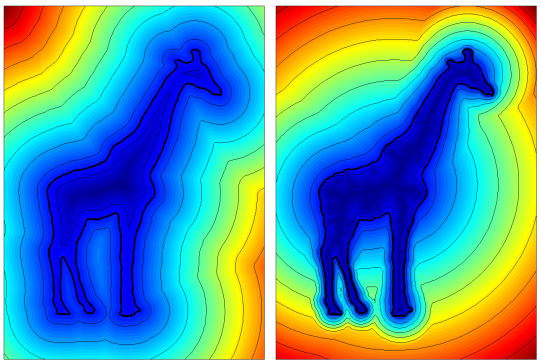

Figure 4: The signed Euclidean distance to the boundary and the multiplicative distance to the shape.

4.1 Semiconcavity

In the following, A always represents an open subset of $\mathbb{R}^d$. Recall that a function $g : A \to \mathbb{R}$ is concave if and only if we have $\alpha g(x) + (1-\alpha)g(y) \leq g(\alpha x + (1-\alpha)y)$ for any $\alpha \in [0,1]$ and x, y such that $[x,y] \subset A$. We now present the definition of locally semiconcave functions. See [6] for a general presentation, and [13] for more properties. Note that the class of functions which are called locally semiconcave with linear modulus by Cannarsa and Sinestrari [6] are referred to as semi-concave functions by Petrunin [13].

Definition 6. *A function $f : A \to \mathbb{R}$ is* semiconcave *with semiconcavity constant λ, if $\lambda > 0$ and $x \mapsto f(x) - \lambda|x|^2$ is concave. Similarly, a function $f : A \to \mathbb{R}$ is* locally semiconcave*, if for each $x \in A$, there exists a neighborhood N_x of x in A such that the restriction of f to N_x is semiconcave.*

The following lemma is a direct consequence of the more general Proposition 2.2.2 from [6]. It follows from bounding the second derivative of the distance function:

Lemma 7. (Semiconcavity of the distance to a point). *The Euclidean distance to a point $p \in \mathbb{R}^d$ is semiconcave on $\mathbb{R}^d \setminus B(p,r)$ with semiconcavity constant $1/r$, and locally semiconcave in $\mathbb{R}^d \setminus \{p\}$.*

The multiplicative distance function to $\mathcal{O}$ is defined by the set of all balls in $\mathrm{MAT}(\mathcal{O})$. Let us now show that in the neighborhood of each point that does not belong to $\overline{\mathrm{M}(\mathcal{O})}$, medial balls with radius smaller than a certain threshold can be ignored from the definition.

Lemma 8 (Radius filtering). *Let $x \in \mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$ and N_x be a bounded neighborhood of x such that $d(N_x, \overline{\mathrm{M}(\mathcal{O})}) = \delta > 0$. Then there exists $r_0 > 0$ such that for any $y \in N_x$ we have:*

$$\mu_{\mathcal{O}}(y) = \inf\left\{\mu_{B(c,r)}(y) \mid B(c,r) \in \mathrm{MAT}(\mathcal{O}) \ , \quad r > r_0\right\}$$

Proof. As the infimum of positive continuous functions on $\mathbb{R}^d$, $\mu_{\mathcal{O}}$ is bounded on the set N_x. Let t be such that $t > \sup \mu_{\mathcal{O}}(N_x)$. If $y \in N_x$, only those balls B may affect the computation of $\mu_{\mathcal{O}}(y)$ for which we have $\mu_B(y) < t$. Such a ball $B(c,r)$ satisfies $t > \frac{d(y,c)}{r} > \frac{\delta}{r}$, and in particular $r > r_0 = \frac{\delta}{t}$. □

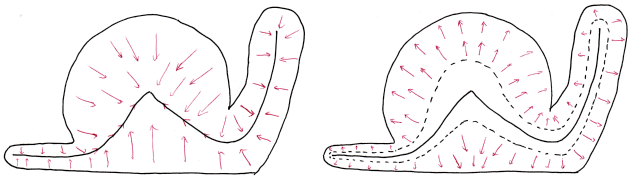

Figure 5: On the left, the flow defined in [7] induced by the distance to the boundary, on the right, our flow which describes the shape scaling.

Lemma 9. (Semiconcavity of the multiplicative distance). *The multiplicative distance $\mu_{\mathcal{O}}$ is semiconcave on a set bounded away from $\mathrm{M}(\mathcal{O})$ and it is locally semiconcave in $\mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$.*

Proof. Let $x \in \mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$. Denote by N_x a bounded neighborhood of x such that $d(N_x, \overline{\mathrm{M}(\mathcal{O})}) = \delta > 0$. Lemma 7 shows that for each ball $B(c,r) \in \mathrm{MAT}(\mathcal{O})$, $x \mapsto d(x,c) - \frac{1}{\delta}x^2$ is concave inside N_x. It follows that $x \mapsto d(x,c)/r - \frac{1}{r\delta}x^2$ is concave too. In other words, $\mu_{B(c,r)}$ is semiconcave with constant $\frac{1}{r\delta}$ in N_x. Lemma 8 shows that there exists $r_0 > 0$ such that for all $y \in N_x$, $\mu_{\mathcal{O}}(y) = \inf\{\mu_{B(c,r)}(y) | \ B(c,r) \in \mathrm{MAT}(\mathcal{O}) \ , \ r > r_0\}$. Since the infimum of concave functions is concave too, it follows that $\mu_{\mathcal{O}}(y)$ is semiconcave with constant $\frac{1}{r_0\delta}$ in N_x. This proves that $\mu_{\mathcal{O}}$ is locally semiconcave in $\mathbb{R}^d \setminus \overline{\mathrm{M}(\mathcal{O})}$. The same arguments yield that $\mu_{\mathcal{O}}$ is semiconcave on a set bounded away from $\mathrm{M}(\mathcal{O})$, since the filtering lemma can provide a global bound in this case. □

4.2 Differential properties

Let us recall some differential properties of a locally semiconcave function $f : A \to \mathbb{R}$. For every point $x \in A$, the directional derivative $\partial f(x,v) = \lim_{\varepsilon \to 0^+} \frac{f(x+\varepsilon v) - f(x)}{\varepsilon}$ is well defined (see Theorem 3.2.1 [6]). Additionally, the set of superdifferentials $D^+f(x) = \{g \in \mathbb{R}^d \mid \forall v \in \mathbb{R}^d \ , \ \partial f(x,v) \leq \langle g, v\rangle\}$ is a nonempty bounded convex set (see Proposition 3.1.5 [6]). One can then define the steepest ascent field, or gradient field, of a locally semiconcave function (see [13]): the gradient vector $\nabla f(x)$ is the unique vector $g \in D^+f(x)$ such that $\partial f(x,g) = \langle g, g\rangle$. One can show that $\nabla f(x)$ is the projection of the origin on $D^+f(x)$ and when non-zero, its direction v is the one such that $\partial f(x,v)$ is maximum among all $\partial f(x,u)$, with $\|u\| = 1$. This gradient vector field has useful properties, in particular, it is possible to construct a continuous flow whose right derivative is equal to ∇f (see [13]).

Definition 10. *For a locally semiconcave function $f : A \to \mathbb{R}$, we call $x \in A$ a* critical point *of f if $\nabla f(x) = 0$.*

In order to understand the relation between the topology of various sublevelsets, we need to determine the location of critical points of $\mu_{\mathcal{O}}$. For a locally semiconcave function f, it follows from $D^+f(x)$ being convex that the critical points of f are the points for which $\forall p \in \mathbb{R}^d, \exists v \in D^+f(x)$ such that $\angle(p,v) \geq \pi/2$. As a consequence, we have the following lemma:

Lemma 11 (Characterization of critical points). *A point $x \in \mathbb{R}^d$ is a critical point of $\mu_{\mathcal{O}}$ if and only if $x \in \mathrm{conv}\{c \mid B(c,r) \in \Gamma_{\mathcal{O}}(x)\}$.*

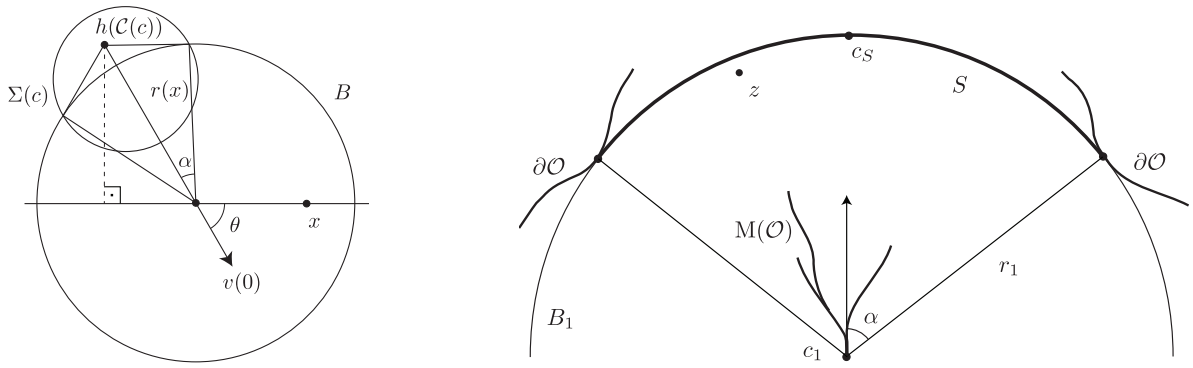

Figure 6: Figure for Lemma 12 on the left, and for Lemma 16 on the right.

5. HOMOTOPY EQUIVALENCE UNDER SHRINKING

We show that $\mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$ does not contain any critical point of $\mu_\mathcal{O}$. Then we can construct a smooth approximation of $\mu_\mathcal{O}$, without any critical point in $\mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$. The gradient flow of this smooth function can be reversed and this reversed flow defines a retraction between sublevelsets, effectively showing the homotopy equivalence of the sublevelsets $\mu_\mathcal{O}^{-1}([0,s))$ for $0 < s \leq 1$.

5.1 Critical points

We first prove that $\mu_\mathcal{O}$ has no critical point in $\mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$. We have seen in Lemma 11 that $z \in \mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$ is a critical point if it belongs to the convex hull of the centers of its closest medial balls with respect to the multiplicative distance. Therefore, a point $z \in \mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$ can only be a critical point if $|\Gamma_\mathcal{O}(z)| \geq 2$. Throughout this section we will use the notation $\Gamma_\mathcal{O}(z) = \{B_i(c_i, r_i)\}_{1\leq i\leq n}$, with $n \geq 2$, and we may assume w.l.o.g. that $\forall i, r_1 \leq r_i$. In order to expose a contradiction to the assumption that z is a critical point, we construct another ball, $B_\star = B(c_\star, r_\star) \in \mathrm{MAT}(\mathcal{O})$ such that $\mu_{B_\star}(z) < \mu_\mathcal{O}(z) = \mu_{B_1}(z)$.

We show that one can find $B_\star = B(c_\star, r_\star)$, with the center $c_\star$ close to c_1. For this construction, we use the flow defined by Lieutier [12], that we denote by ϕ in the following. We describe how $\mu_\mathcal{O}$ evolves along ϕ, and we show that either $\mu_\mathcal{O}$ decreases, or we can construct a path on $\mathrm{M}(\mathcal{O})$ in another direction, so that $\mu_\mathcal{O}$ decreases along this reversed path.

Let us now recall the definition and a few properties of ϕ, which are detailed and proved in [12]. The flow ϕ is defined as the steepest ascent flow of the Euclidean distance to $\partial\,\mathcal{O}$. This distance function is a locally semiconcave function, as well (see Proposition 2.2.2 in [6]). Let $\Gamma(x)$ denote the set of contact points, i.e. the closest points to x on $\partial\,\mathcal{O}$. As shown in Lemma 4.6 in [12], $x \mapsto \Gamma(x)$ has a semicontinuity property: $\forall x \in \mathcal{O}, \forall\varepsilon > 0, \exists\alpha > 0$ such that $y \in B(x,\alpha) \Rightarrow \Gamma(y) \subset \Gamma(x) \oplus B(0,\varepsilon)$. The gradient at point x is defined as the vector $\nabla_x = \frac{x-\Theta(x)}{d(x,\partial\,\mathcal{O})}$, where $\Theta(x)$ is the center of the smallest enclosing ball $\Sigma(x)$ of the points in $\Gamma(x)$. There exists a flow $\phi_x(t)$ for $t \geq 0$ such that $\phi_x(0) = x$ and that admits ∇_x as right derivative at $t = 0$. Importantly, if $c \in \mathrm{M}(\mathcal{O})$, then $\forall t \geq 0, \phi_c(t) \in \mathrm{M}(\mathcal{O})$. Therefore, we introduce the notation $\phi_B(t) = B(c(t), r(t))$, where B is the medial ball centered at c, and $B(c(t), r(t))$ is the medial ball centered at $c(t) = \phi_c(t)$. We call c a critical point of ϕ if $\nabla_c = 0$. If c is critical, ϕ_c is constant and the minimal enclosing sphere $\Sigma(c)$ of the contact points of c is identical to the medial ball centered at c. If c is not a critical point of ϕ, then the minimal enclosing sphere $\partial\Sigma(c)$ of the contact points intersects the medial sphere $\partial B(c,r)$ in $\mathcal{C}(c)$ (two points in dimension 2, a circle in dimension 3, a $(d-2)$-sphere in dimension d). Denote by $h(\mathcal{C}(c))$ the apex of the cone tangent to $B(c,r)$ along $\mathcal{C}(c)$. The following lemma describes how $\mu_\mathcal{O}$ evolves along ϕ_c:

Lemma 12 (Derivative along the flow). *Let $B(c,r) \in \mathrm{MAT}(\mathcal{O})$ and $x \in B(c,r)$. If c is not a critical point of ϕ, we have*

$$\frac{d\mu_{\phi_B(t)}(x)}{dt}\Big|_{t=0^+} < 0 \quad \Leftrightarrow \quad \langle h(\mathcal{C}(c)) - x, x - c\rangle < 0$$

In other words, the derivative of the multiplicative distance along ϕ is negative if and only if $h(\mathcal{C}(c))$ projects before x on line (cx), oriented from c to x.

Proof. Let $v(t)$ be the right derivative of $\phi_c(t)$ and α the half-angle of the cone Y with apex c and generated by $\mathcal{C}(c)$ (see Figure 6 left). The axis of Y is aligned with $v(0)$, and $\cos(\alpha) = \|v(0)\|$. Let θ be the angle between $v(0)$ and $[cx]$. Lemma 4.11 in [12] shows that

$$\frac{dr(t)}{dt}\Big|_{t=0^+} = \|v(0)\|\cos(\alpha), \quad \text{and we have}$$

$$\frac{d\|x - c(t)\|}{dt}\Big|_{t=0^+} = -\|v(0)\|\cos(\theta).$$

It follows that

$$\frac{d\mu_{B(t)}(x)}{dt}\Big|_{t=0^+} = \frac{d\frac{\|x-c(t)\|}{r(t)}}{dt}\Big|_{t=0^+} =$$

$$\frac{\|v(0)\|}{r(0)^2}\Big(-r(0)\cos(\theta) - \|x - c(0)\|\cos(\alpha)\Big) \qquad (1)$$

In particular, this derivative is negative if and only if $-r(0)\cos(\theta) < \|x - c(0)\|\cos(\alpha)$. Since $\cos(\alpha)$ is positive and $-r(0)\cos(\theta)/\cos(\alpha)$ is the position of the projection of $h(\mathcal{C}(c))$ on line (cx), the result follows. □

Now let us consider $z \in \mathcal{O}\setminus\overline{\mathrm{M}(\mathcal{O})}$ such that $|\Gamma_\mathcal{O}| \geq 2$, and let $B_1(c_1, r_1)$ be the smallest ball in $\Gamma_\mathcal{O}(z)$. We know that $z \in B_i$ for every $B_i \in \Gamma_\mathcal{O}(z)$. In the following, H_z denotes the hyperplane passing through z that is orthogonal to $(c_1 z)$,

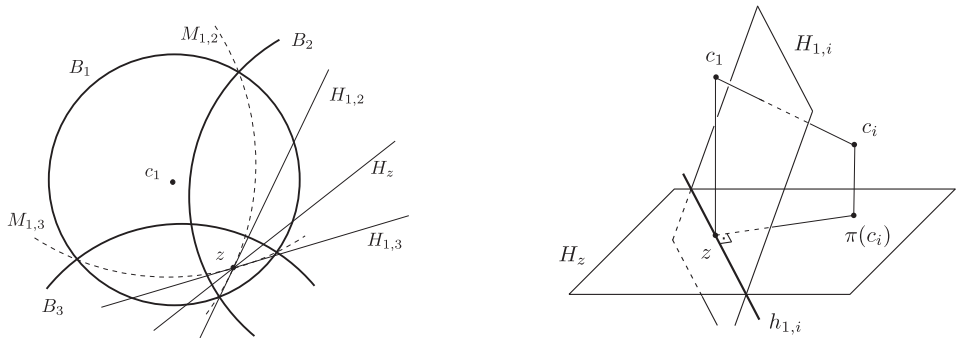

Figure 7: Illustrations for Lemma 14. On the left a planar case showing the H_z, $H_{1,2}$ and $H_{1,3}$. On the right a three dimensional example showing that $h_{1,i}$ is perpendicular to $(z\pi(c_i))$.

and H_z^+ is the halfspace delimited by H_z and not containing c_1. The following lemma gives a sufficient condition for the derivative of $\mu_{\mathcal{O}}$ to be negative along the flow ϕ:

LEMMA 13. *Assume that c_1 is not a critical point of ϕ, and that the spherical cap $\Sigma(c_1) \cap \partial B_1$ does not intersect the spherical cap $H_z^+ \cap \partial B_1$. Then*

$$\frac{d\mu_{\phi_{B_1}(t)}(z)}{dt}\Big|_{t=0^+} < 0.$$

PROOF. $\Sigma(c_1) \cap \partial B_1$ not intersecting $H_z^+ \cap \partial B_1$ implies that $h(\mathcal{C}(c_1))$ projects onto the line $(c_1 z)$ on the same side of H_z as c_1. The result then follows from Lemma 12. □

We have shown that if the conditions of Lemma 13 are fulfilled, by starting at c_1 and moving along the flow ϕ, we find medial balls that are closer to z than B_1 is, in terms of multiplicative distance. Any of these balls is a suitable $B_\star$.

If c_1 is a critical point of ϕ, then ϕ_{c_1} is constant and the derivative is zero. If c_1 is not critical, but $\Sigma(c_1) \cap \partial B_1$ intersects $H_z^+ \cap \partial B_1$, the function $\mu_{\phi_{B_1}(t)}(z)$ may have a non-negative derivative, depending on the position of the projection of $h(\mathcal{C}(c_1))$ along $(c_1 z)$ (note that this cannot happen in dimension 2, but that it is possible in higher dimensions). In these cases, we need another technique to find a suitable $B_\star$.

Let us show that there still exists a direction around c_1, where we can find medial axis points that are suitable centers for $B_\star$. The center of the spherical cap considered in the following lemma will give us a direction around c_1 where $\mu_{\mathcal{O}}$ decreases.

LEMMA 14. (A SPHERICAL CAP WITHOUT CONTACT POINTS). *The spherical cap $H_z^+ \cap \partial B_1$ does not contain any points of $\partial\,\mathcal{O}$.*

PROOF. Note that the case when z belongs to some segment $[c_1 c_i]$ is simple: $H_z^+ \cap \partial B_1$ is then included in $B_1 \cap B_i$, which does not contain any contact point. In the following we assume that z does not belong to such a segment. For each B_i with $i \neq 1$, the Möbius bisector $M_{1,i}$ of B_1 and B_i contains z and it is either a sphere centered on $(c_1 c_i)$ or a hyperplane orthogonal to $(c_1 c_i)$ (see Figure 7 left). Let $H_{1,i}$ be the hyperplane tangent to $M_{1,i}$at z. We know that $M_{1,i}$ delimits a spherical cap of ∂B_1 entirely covered by B_i. Since B_1 is a smallest ball in $\Gamma_{\mathcal{O}}(z)$, this spherical cap contains the one delimited by $H_{1,i}$. The intersection of $H_{1,i}$ and H_z is a line $h_{1,i}$ containing z and orthogonal to $[c_1 c_i]$ (see Figure 7 right). We denote by $H_{1,i}^+$ the halfspace delimited by $H_{1,i}$ and not containing c_1, and we define $h_{1,i}^+ = H_{1,i}^+ \cap H_z$. It is sufficient to show that $H_z^+ \subset \bigcup_{2\leq i\leq n} H_{1,i}^+$, or equivalently, that $H_z = \bigcup_{2\leq i\leq n} h_{1,i}^+$. Let π be the orthogonal projection on H_z. Since $h_{1,i}$ is orthogonal to $[z\pi(c_i)]$ and $\pi(c_i) \subset h_{1,i}^+$ (we excluded at the beginning of the proof the case $z = \pi(c_i)$), this is equivalent to z being inside the convex hull of $\pi(c_2), \ldots, \pi(c_n)$, which follows directly from z being in the convex hull of $c_1, \ldots, c_n$ (recall that $\pi(c_1) = z$ and $c_1 \neq z$). □

Note that the assumption that B_1 is the smallest of the balls is crucial in the previous proof.

Intuitively, in the case where the derivative of $\mu_{\phi_{B_1}(t)}(z)$ is not negative, we want to reverse ϕ in order to find a direction where the derivative of μ is negative. Since ϕ is the gradient flow of $d(\cdot, \partial\,\mathcal{O})$, one cannot simply reverse it. However, we can use of the fact that $d(\cdot, \partial\,\mathcal{O})$ is locally semiconcave on $\mathcal{O}$ and find a suitable direction with a well defined derivative. For this, let us recall Lemma 4.2.5 of [6], page 86. The notation $\partial_d X$ below represents the boundary of X, as defined by the topology of $\mathbb{R}^d$, and not by the induced topology on X (see Remark 3.3.5, page 58 of [6]):

LEMMA 15 (CANNARSA AND SINESTRARI). *Let f be a function that is semiconcave in a neighborhood N_c of c. Fix $p_0 \in \partial_d D^+ f(c)$ and let $q \in \mathbb{R}^d \setminus \{0\}$ be such that $\forall p \in D^+ f(c), \langle q, p - p_0\rangle \geq 0$. Then there exists $\sigma > 0$, and a Lipschitz function $\ell : [0, \sigma] \to N_c$, with $\ell(0) = c$, and such that $\lim \frac{\ell(s)-c}{s} = q$ when s converges to 0^+. Furthermore, $p(s) = p_0 + \frac{\ell(s)-c}{s} - q$ belongs to $D^+ f(\ell(s))$ for all $s \in (0, \sigma]$, and it converges to p_0.*

We now use this lemma for the case where f is the distance function to $\partial\,\mathcal{O}$, and c is c_1, the center of the medial ball B_1, to prove the following result, which concludes our search for $B_\star$:

LEMMA 16. *Assume that c_1 is a critical point of ϕ or the spherical cap $\Sigma(c_1) \cap \partial B_1$ intersects the spherical cap $H_z^+ \cap \partial B_1$. Then one can find $\sigma > 0$ and a Lipschitz function $\ell : [0, \sigma] \to \mathrm{M}(\mathcal{O})$ with $\ell(0) = c_1$, such that $q = \lim \frac{\ell(s)-c}{s}$ exists when s converges to 0^+ and the derivative of the multiplicative distance of z to the medial balls in the direction q is negative.*

PROOF. Let us consider the set $\mathcal{S}$ of maximal spherical caps of B_1 that do not contain any point of $\partial\mathcal{O}$. The set $\mathcal{S}$ is the medial axis transform of $\partial B_1 \setminus \partial\mathcal{O}$ for the intrinsic distance. Let S denote an element of $\mathcal{S}$ containing $H_z^+ \cap \partial B_1$. Such an S exists, because $H_z^+ \cap \partial B_1$ is a spherical cap which does not contain any point of $\partial\mathcal{O}$, as stated in Lemma 14. Let $c_S \in \partial B_1$ be the center of S and Y be the solid cone with apex c_1 and generated by S. See Figure 6. Denote by α the half-angle of Y.

We apply Lemma 15 to the function $f = d(\cdot, \partial\mathcal{O})$: with the notations of the lemma, we choose $c = c_1$. Let us denote the hyperplane $H = \{x \mid \langle x, \frac{c_s - c_1}{\|c_s - c_1\|}\rangle = -\cos(\alpha)\}$. We define $X = D^+f(c_1) \bigcap H$. Since $S \bigcap \partial\mathcal{O} = \emptyset$, the set $D^+f(c_1)$ is contained in the halfspace delimited by H and containing the origin. It follows that $X \subseteq \partial_d D^+f(c_1)$. Furthermore, $|\partial S \bigcap \partial\mathcal{O}| \geq 2$ implies that X contains the convex hull of at least two distinct points with norm 1. As a consequence, X contains points with norm less than 1. Let p_0 be such a point and $q = c_S - c_1$. The hypothesis of the Lemma 15 is now satisfied: we have $\forall p \in D^+f(c), \langle q, p - p_0\rangle \geq 0$, because q is orthogonal to H, which is a supporting hyperplane of the convex set $D^+f(c_1)$.

Hence we obtain a Lipschitz function ℓ, with $\ell(0) = c_1$, and such that $\lim \frac{\ell(s)-c}{s} = c_S - c_1 = q$ when s converges to 0^+. The fact that $p(s) = p_0 + \frac{\ell(s)-c}{s} - q$ converges to p_0 and belongs to $D^+f(\ell(s))$ implies that for s small enough, $D^+f(\ell(s))$ contains a vector with norm less than 1. Note that such a vector is not the gradient of the distance to any of the contact points in $\partial\mathcal{O}$. It follows that for s small enough, $\ell(s)$ has more than one contact point: $\ell(s)$ belongs to the medial axis and $c_S - c_1$ is its tangent vector at c_1.

Let us finally distinguish several cases, based on the value of α, the half-angle of the cone Y. Let us denote with $B(\ell(s))$ the medial ball centered at $\ell(s)$. We show that in all cases the computations in the proof of Lemma 12 yield that the derivative of $\mu_{B(l(s))}(z)$ at $s = 0$ is negative. If $\alpha = \pi/2$, then we can adapt Equation (1) from the proof of Lemma 12 to get the derivative of the multiplicative distance $\mu_{B(l(s))}(z)$ at $s = 0$ in the direction q:

$$\frac{\|c_s - c_1\|}{(c_s - c_1)^2}\left(-\|c_s - c_1\|\cos(c_s, c_1, z) - \|z - c_1\|\cos\left(\frac{\pi}{2}\right)\right)$$

Therefore, the derivative of the multiplicative distance is $-\cos(c_s, c_1, z)$, where the cosine is positive, and the result follows.

If $\alpha \neq \pi/2$, let $h(S)$ be the apex of the cone tangent to B_1 along ∂S. We use the fact that S contains $H_z^+ \cap \partial B_1$ to find the location of the projection of $h(S)$ on the line $(c_1 z)$, oriented from c_1 to z:

- if $\alpha > \pi/2$ then $h(S)$ projects before z
- if $\alpha < \pi/2$ then $h(S)$ projects after z

Notice that q points in such a direction that in both of these cases the derivative of the multiplicative distance to z is negative (see the computations in Lemma 12). This concludes the proof. □

We have shown that if $z \in \mathcal{O} \setminus \overline{\mathrm{M}(\mathcal{O})}$, we can always find a ball $B_\star$ closer to z than the balls $B_i \in \Gamma_{\mathcal{O}}(z)$ in terms of the multiplicative distance. Our main theorem follows:

THEOREM 17. *The function $\mu_{\mathcal{O}}$ has no critical point in $\mathcal{O} \setminus \overline{\mathrm{M}(\mathcal{O})}$.*

5.2 Deformation retraction

Now that we have shown that $\mu_{\mathcal{O}}$ has no critical points in $\mathcal{O} \setminus \overline{\mathrm{M}(\mathcal{O})}$, we would like to use its gradient flow to construct a retraction from sublevelset $\mathcal{O}_s$ to sublevelset $\mathcal{O}_{s'}$, for $s > s'$. However, since $\mu_{\mathcal{O}}$ is not smooth, we cannot directly revert the flow. Instead we first regularize $\mu_{\mathcal{O}}$ into a smooth function, whose gradient is close to the superderivatives of $\mu_{\mathcal{O}}$ using Proposition 4.1 of [8]. To state it, we extend the notation of superderivatives D^+ for sets in the following way: $D^+f(X) = \bigcup_{x \in X} D^+f(x)$.

THEOREM 18 (CZARNECKI AND RIFFORD). *Let A be an open subset of $\mathbb{R}^d$. Let $f : A \to \mathbb{R}$ be a locally Lipschitz function. For every continuous function $\varepsilon : A \to \mathbb{R}_+^*$, there exists a smooth function $g : A \to \mathbb{R}$ such that for every $x \in A$, we have*

- $|f(x) - g(x)| \leq \varepsilon(x)$;
- $\nabla g(x) \subset D^+f(B(x, \varepsilon(x)) \cap A) \oplus B(0, \varepsilon(x))$.

The following theorem ensues from Theorems 17 and 18:

THEOREM 19. *For $0 < s' \leq s \leq 1$ the sets $\mathcal{O}_s$ and $\mathcal{O}_{s'}$ are homotopy equivalent.*

PROOF. We use the characterization of homotopy equivalence recalled in Proposition 3.2 of [12]: if there exists a continuous map $H : [0,1] \times \mathcal{O}_s \to \mathcal{O}_s$ such that

(i) $\forall x \in \mathcal{O}_s, H(0, x) = x$
(ii) $\forall x \in \mathcal{O}_s, H(1, x) \in \mathcal{O}_{s'}$
(iii) $\forall x' \in \mathcal{O}_{s'}, \forall t \in [0,1], H(t, x') \in \mathcal{O}_{s'}$

then $\mathcal{O}_s$ and $\mathcal{O}_{s'}$ are homotopy equivalent.

Let us construct a smooth function ν such that we can use the gradient descent flow of ν as a function H retracting $\mathcal{O}_s$ into $\mathcal{O}_{s'}$, up to some reparametrization of the first parameter. We construct ν by smoothing $\mu_{\mathcal{O}}$. Property (i) follows from the fact that we consider a gradient descent. For Property (iii), we need to show that the gradient descent flow keeps $\mathcal{O}_{s'}$ stable. This stability property is a direct consequence of a transversality property of $\nabla\nu(x)$ that we prove now: $\forall x \in \mathcal{O}_s, \angle(\nabla\mu_{\mathcal{O}}(x), \nabla\nu(x)) < \pi/2$. Similarly, Property (ii) follows from $\|\nabla\nu(x)\|$ being bounded away from zero, a fact that follows from a similar proof.

For any point x in $A = \mathcal{O}_s \setminus \mathcal{O}_{s'}$, we know that $D^+\mu_{\mathcal{O}}(x)$ is a bounded convex set, bounded away from the origin O, and that the projection of O on $D^+\mu_{\mathcal{O}}(x)$ is $\nabla\mu_{\mathcal{O}}(x)$. In particular, there exists $\alpha < \pi/2$ such that $\forall v \in D^+\mu_{\mathcal{O}}(x)$, $\angle(v, \nabla\mu_{\mathcal{O}}(x)) < \alpha$: let L be a finite positive constant such that $D^+\mu_{\mathcal{O}}(x) \subset B(O, L)$ and $\delta = d(O, D^+\mu_{\mathcal{O}}(x)) > 0$. Then, for any $\cos^{-1}(\delta/L) \leq \alpha < \pi/2$, if there exists $v \in D^+\mu_{\mathcal{O}}(x)$ such that $\angle(v, \nabla\mu_{\mathcal{O}}(x)) > \alpha$, one would find an element of $D^+\mu_{\mathcal{O}}(x)$ shorter than $\nabla\mu_{\mathcal{O}}(x)$ on the segment $[\nabla\mu_{\mathcal{O}}(x)\ \ v]$. In the following, we define the angle $\alpha = \frac{1}{2}(\pi/2 + \cos^{-1}(\delta/L)) < \pi/2$.

By using the semicontinuity of the closest balls mapping stated in Lemma 5, we obtain a positive function ε such that $T = \mathrm{conv}(D^+\mu_{\mathcal{O}}(B(x, \varepsilon(x)) \cap A) \oplus B(0, \varepsilon(x)))$ is bounded and bounded away from the origin, too. In fact, one can choose ε to be continuous. Denote by π_T the projection of O to T. Note that for ε small enough, for all $v \in T$, we have $\angle(v, \pi_T) < \frac{1}{2}(\pi/2 + \alpha)$.

Furthermore, we can decrease ε, so that $\angle(\nabla\mu_{\mathcal{O}}(x), \pi_T)$ is smaller than $(\pi/2 - \alpha)/2$. Theorem 18 then implies that we

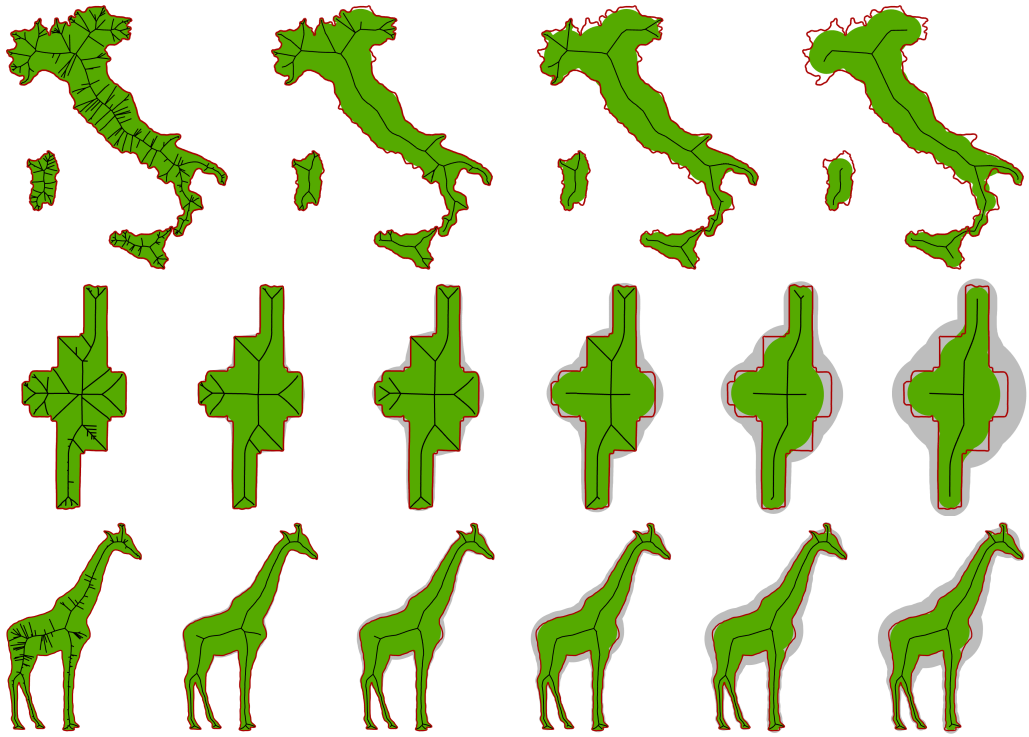

Figure 8: Scale axis transform: a discrete approximation for the map of Italy, a mechanical CAD part, and the giraffe. We color code with gray the grown shape and with green the reconstruction from the scale axis transform.

can find an approximation ν of $\mu_{\mathcal{O}}$ without critical points, such that $\nabla\nu(x) \in D^+\mu_{\mathcal{O}}(B(x,\varepsilon(x)) \cap A) \oplus B(0,\varepsilon(x)) \subset T$. It follows that $\angle(\nabla\mu_{\mathcal{O}}(x), \nabla\nu(x)) < \frac{1}{2}(\pi/2+\alpha) + (\pi/2-\alpha)/2 = \pi/2$. We have thus proved that $\angle(\nabla\mu_{\mathcal{O}}(x), \nabla\nu(x)) < \pi/2$, and the result follows. □

5.3 Homotopy type of the scale axis

The above results yield a topological characterization of the scale axis very similar to the λ-medial axis. To formulate this result we need to define the equivalent of the weak feature size from [7] in our setting.

DEFINITION 20. *The multiplicative weak feature size* mwfs *of an open set $\mathcal{O}$ is defined as:*

$$\text{mwfs}(\mathcal{O}) = \inf\{\mu_{\mathcal{O}}(x) \mid x \in \mathbb{R}^d \setminus \overline{\text{M}(\mathcal{O})},\ \nabla\mu_{\mathcal{O}}(x) = 0\}.$$

Note that Theorem 17 yields that the multiplicative weak feature size cannot be smaller than 1. In order to state the next corollary we assume $\text{mwfs}(\mathcal{O}) > 1$. This condition can be viewed as a minimum regularity condition on the boundary, the same way the condition of positive weak feature size used in [7] for the λ-medial axis.

COROLLARY 21. *For $s < \text{mwfs}(\mathcal{O})$, the s-scale axis of $\mathcal{O}$ is homotopy equivalent to $\mathcal{O}$.*

The proof follows from the same arguments as Theorem 19 and the homotopy equivalence between a shape and its medial axis shown in [12].

6. CONCLUSIONS AND FUTURE WORK

We have introduced the scale axis transform, an extension of the medial axis transform that computes a family of successively simplified skeletons in a scale-adaptive way. The scale axis transform is defined using the medial axis transform of the multiplicatively scaled shape. We have presented a theoretical framework based on non-smooth analysis and the theory of semiconcave functions to prove properties of the multiplicative distance. Most notably, we show that multiplicative shrinking does not change the homotopy type of a shape, and we prove a topological stability guarantee of the scale axis similar to the one existing for the λ-medial axis filtration.

The simplicity of the scale axis construction based on growing and shrinking of medial balls leads to efficient algorithms and thus it has practical relevance in digital shape analysis and processing. We illustrated the potential for applications using a prototype implementation for 2D shapes based on CGAL [1], yet all the theoretical results apply for arbitrary dimensions.

The scale axis is designed such that topological changes provide meaningful simplifications of the shape. Therefore, we envision future theoretical results built on top of our framework that formalize the property that two close shapes have topologically similar simplified shapes under the scale axis transform. Similarly, a geometry stability result of the scale axis should formalize that two close shapes have close scale axis in a *scale-adaptive* way. The appropriate closeness measure for such a stability result is the symmetric Hausdorff version of the multiplicative distance introduced in this paper.

Acknowledgments

We are especially grateful to Frédéric Chazal for pointing us to the theory of semiconcave functions, and to the anonymous reviewer for pointing out an error in the original version of the proof of Lemma 16. We would like to thank to Steve Oudot, Quentin Mérigot, André Lieutier, and Leonidas Guibas for all the insightful discussions and comments. This work has been partially supported by the Swiss National Science Foundation.

7. REFERENCES

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

[2] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *SMA '01: Proceedings of the sixth ACM symposium on Solid modeling and applications*, pages 249–266, New York, NY, USA, 2001. ACM.

[3] Dominique Attali, Jean-Daniel Boissonnat, and Herbert Edelsbrunner. Stability and computation of the medial axis — a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, 2007.

[4] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Curved Voronoi diagrams. In Jean-Daniel Boissonnat and Monique Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, pages 67–116. Springer-Verlag, Mathematics and Visualization, 2006.

[5] Kevin Buchin, Tamal K. Dey, Joachim Giesen, and Matthias John. Recursive geometry of the flow complex and topology of the flow complex filtration. *Comput. Geom. Theory Appl.*, 40(2):115–137, 2008.

[6] Piermarco Cannarsa and Carlo Sinestrari. *Semiconcave Functions, Hamilton-Jacobi Equations, and Optimal Control*. Birkhäuser, Boston, USA, 2004.

[7] Frédéric Chazal and André Lieutier. The λ-medial axis. *Graphical Models*, 67(4):304–331, 2005.

[8] Marc-Olivier Czarnecki and Ludovic Rifford. Approximation and regularization of Lipschitz functions: Convergence of the gradients. *Transactions-American Mathematical Society*, 358(10):4467, 2006.

[9] Tamal Dey and Wulue Zhao. Approximate medial axis as a Voronoi subcomplex. *Computer-Aided Design*, 36(2):195–202, 2004.

[10] Herbert Edelsbrunner. Deformable smooth surface design. *Discrete & Computational Geometry*, 21(1):87–115, 1999.

[11] Karsten Grove and Katsuhiro Shiohama. A generalized sphere theorem. *The Annals of Mathematics*, 106(1):201–211, 1977.

[12] André Lieutier. Any open bounded subset of R n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.

[13] Anton Petrunin. Semiconcave functions in Alexandrov's geometry. In *Surveys in Differential Geometry: Metric and Comparison Geometry*, volume XI. Internationational Press of Boston, 2007.

[14] Stephen M. Pizer, Kaleem Siddiqi, Gabor Székely, James N. Damon, and Steven W. Zucker. Multiscale medial loci and their properties. *Int. J. Comput. Vision*, 55(2-3):155–179, 2003.

[15] R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*. Springer, 1997.

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

[17] Roger C. Tam and Wolfgang Heidrich. Feature-preserving medial axis noise removal. In *ECCV '02: Proceedings of the 7th European Conference on Computer Vision-Part II*, pages 672–686, London, UK, 2002. Springer-Verlag.

[18] Xiaofei Wang and Christina A. Burbeck. Scaled medial axis representation: evidence from position discrimination task. *Vision Research*, 38(13):1947–1958, 1998.