Epsilon-covering: a greedy optimal algorithm for simple shapes

Tuong-Bach Nguyen * Isabelle Sivignon *

Abstract

Unions of balls are widely used shape representations. Given a shape, computing a union of balls that is both accurate in some sense and of small cardinality is thus a challenging problem. In this work, accuracy is ensured by imposing that the union of balls, called *covering*, is included in the shape and covers a parameterized core set (namely the erosion) of the shape. For a family of simple shapes, we propose a polynomial-time greedy algorithm that computes a covering of minimum cardinality for a given shape.

1 Introduction

Unions of balls are common shape representations, useful for instance to describe molecules in biochemistry [4], to quickly detect collisions [3] between shapes or to derive higher-level representations. The ubiquity of unions of balls is largely due to the existence of provably good conversion algorithms that allow us to derive them from various representations such as point clouds and polygonal meshes [6]. However, the union of balls output by the conversion process provides only an approximation of the original shape.

In a previous work [2], we introduced a novel way, called ε-covering, of controlling the geometric error between a given input shape and a selected union of balls. The idea is to impose that the union of balls covers a core set of the shape *and* does not cross over an outer set. This problem falls in the family of geometric set cover problems, where the goal is to minimize the number of balls. We proved that, in the general case, computing an ε-covering of minimum cardinality is an NP-complete problem. Other approaches, related to the maximum k-cover problem, aim at maximising the coverage for a fixed number of balls [4], but the problem is also NP-complete.

In this work, we consider the family of input shapes that are themselves 2D unions of balls with a tree-like structure. An ε-covering of such a shape is a simplified union of balls. We present a polynomial-time algorithm that computes an ε-covering of minimum cardinality for this specific family. To do so, we rely on the medial axis structure of unions of balls, and show how to continuously sweep its pencils in order to get a correct and optimal result.

*Univ. Grenoble Alpes, GIPSA-Lab, F-38000 Grenoble, France CNRS, GIPSA-Lab, F-38000 Grenoble, France
`tuong-bach.nguyen@gipsa-lab.grenoble-inp.fr`
`isabelle.sivignon@gipsa-lab.grenoble-inp.fr`

2 Statement of the result

In this paper, $\mathbb{R}^2$ is endowed with the Euclidean distance. For any point c and real $r > 0$, we denote by $b(c,r)$ the closed ball of center c and radius r. For any subset $S \subseteq \mathbb{R}^2$, we respectively denote its closure, interior, complement, boundary, and medial axis by $\overline{S}$, $\mathring{S}$, S^{c}, ∂S, and $\mathrm{MA}(S)$. Let $\varepsilon > 0$ be a real number. The erosion of S (by ε) is $S^{\ominus\varepsilon} = \{y \mid b(y,\varepsilon) \subseteq S\}$. For any collection of balls $\mathscr{B}$, we write $\bigcup \mathscr{B} = \cup_{b\in\mathscr{B}} b$.

Definition 1 An (inner) **ε-covering** of S is a collection of balls $\mathscr{B}$ such that $S^{\ominus\varepsilon} \subseteq \bigcup \mathscr{B} \subseteq S$.

For given S and ε, there exist many ε-coverings of S, with different cardinalities. We say that an ε-covering is **optimal** if it achieves minimum cardinality. In general, finding such an optimal ε-covering for S is an NP-complete problem [2], but we focus here on simple shapes S and prove the following result:

Theorem 1 There is a polynomial-time algorithm to compute optimal ε-coverings for finite unions of balls whose medial axis is cycle-free.

In Section 3, we present some results on the structure of unions of balls, before describing the principle of our algorithm in Section 4. Section 5 expands on some practical considerations required for the algorithm, and Section 6 is dedicated to proving that it indeed achieves the claimed result.

3 Union of balls

3.1 Medial axis and pencils

In order to specify our algorithm, we must elaborate on the structure of unions of balls, in particular that of their medial axis. Recall that a ball $b \subseteq S$ is a medial ball if its boundary ∂b intersects ∂S at least twice. The medial axis $\mathrm{MA}(S)$ is the collection of the centers of these medial balls. Owing to the structure theorem of the medial axis of a union of balls [1], we know that for a finite union of balls S, $\mathrm{MA}(S)$ is a finite collection of line segments. We also know that each of these line segments coincides with a pencil of balls [4] in the following sense.

Borrowing the terminology used in [7], an elliptic pencil can be characterized by two points $u, v \in \mathbb{R}^2$: it is the family of all balls whose boundary goes through u and v. The collection of their centers forms a line. In this paper, we only manipulate elliptic pencil segments, that are subsets of elliptic pencils whose collection of centers forms a segment instead of a line. From here on, we will not consider any proper line pencil, and thus we refer to these elliptic pencil segments simply as pencils. As such, the pencils we consider always have two endpoint balls b_1 and b_2. We denote the pencil they generate by $[b_1 b_2]$. A basic property of a pencil is that the domain it covers, that is the collection of all points covered by some ball of the pencil, is the union of b_1 and b_2, $\bigcup [b_1 b_2] = b_1 \cup b_2$.

Linking back to the previous remark, $\mathrm{MA}(S)$ is a collection of pencils. Indeed, for each segment of the medial axis, there are two points $u, v \in \partial S$, such that any medial ball centered at a point of that segment contains both u and v in its boundary. Thus, each segment of $\mathrm{MA}(S)$ coincides with a pencil. Hence a union of balls can always be interpreted as a union of pencil domains (see Figure 1 for an illustration).

The medial axis of a closed shape and its erosion always satisfy the below inclusion.

Proposition 2 $\mathrm{MA}(S^{\ominus\varepsilon}) \subseteq \mathrm{MA}(S)$

Proof. Consider $c \in \mathrm{MA}(S^{\ominus\varepsilon})$ and $b = b(c,r) \subseteq S^{\ominus\varepsilon}$ medial in $S^{\ominus\varepsilon}$. Let $b_+ = b(c, r+\varepsilon)$. We prove that b_+ is medial in S which implies that $c \in \mathrm{MA}(S)$.
First, note that $b_+ = \cup_{y \in b} b(y, \varepsilon) \subseteq S$, hence if ∂b_+ intersects ∂S at least twice, it is medial in S. By definition of b, there are at least two points $u \neq v$ in the intersection of ∂b and $\partial(S^{\ominus\varepsilon})$. Since $u, v \in \partial(S^{\ominus\varepsilon})$, there are $u_+, v_+ \in \partial S$ such that $\|u - u_+\| = \varepsilon = \|v - v_+\|$. By triangular inequality we have $\|c - u_+\| \leq \|c - u\| + \|u - u_+\| = r + \varepsilon$, hence $u_+ \in b_+$. Since $u_+ \in \partial S$, necessarily $u_+ \notin \mathring{b}$ and we have $\|c - u_+\| \geq r + \varepsilon$. Therefore $\|c - u_+\| = r + \varepsilon$, thus $u_+ \in \partial b_+ \cap \partial S$. Also, we have equality in a triangular equality, hence c, u and u_+ are aligned. Likewise, $v_+ \in \partial b_+ \cap \partial S$, and c, v and v_+ are also aligned. Thus, if $u_+ \neq v_+$, then b_+ is medial in S. By contradiction, assume that we have $u_+ = v_+$. Then, c, u, v and u_+ must be aligned. Because $u \neq v$, we have the below situation.

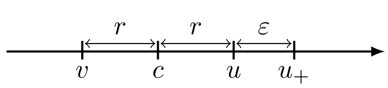

Hence $\|v - v_+\| = \|v - u_+\| = 2r + \varepsilon = \varepsilon$. This implies $r = 0$ and $u = v$, which is impossible. Therefore, $u_+ \neq v_+$ and b_+ is medial in S. $\square$

Thus for unions of balls, $\mathrm{MA}(S^{\ominus\varepsilon})$ splits $\mathrm{MA}(S)$ into two finite collections of line segments: segments that are part of both $\mathrm{MA}(S)$ and $\mathrm{MA}(S^{\ominus\varepsilon})$, and segments that are exclusively part of $\mathrm{MA}(S)$. We respectively refer to them as **eroded** and **non-eroded** pencils.

3.2 Partial ordering on medial balls

$\mathrm{MA}(S)$ being a collection of segments, it can be viewed as the embedding of a graph in $\mathbb{R}^2$. By assumption on the class of shapes considered, $\mathrm{MA}(S)$ is cycle-free, hence it is a forest. Since we can process each tree of the forest independently, we assume without loss of generality that $\mathrm{MA}(S)$ is a tree. By picking any point x of $\mathrm{MA}(S)$ as a root, we obtain an orientation of $\mathrm{MA}(S)$ which induces a partial order on $\mathrm{MA}(S)$. Indeed, we simply have to orient all the edges of $\mathrm{MA}(S)$ from the leaves to the root x. We denote by T the resulting oriented tree. The structure represented by T is at times called anti-arborescence or in-tree, and can also be viewed as a directed acyclic graph with a unique sink. For any $y, z \in \mathrm{MA}(S)$, we say that y is T-smaller than or equal to z, and note $y \leq_T z$, if z belongs to the unique path from y to the root x of T. We also use the usual order symbols and notions such as being T-larger or equal to, $\geq_T$, or also being strictly T-smaller, $<_T$.

Note that this T-order is valid for all points of $\mathrm{MA}(S)$, and not simply vertices of T. Because points of $\mathrm{MA}(S)$ are centers of medial balls of S and $S^{\ominus\varepsilon}$, this T-order extends to medial balls. Specifically, we can T-compare two medial balls of either S or $S^{\ominus\varepsilon}$, but also a medial ball of S with a medial ball of $S^{\ominus\varepsilon}$.

Since T only induces a partial order, we say that two balls that cannot be ordered by T are ***T*-unrelated**. An additional useful notion is that of T-maximal ball for a collection: given a collection of balls $\mathscr{B}$, $b \in \mathscr{B}$ is ***T*-maximal** in $\mathscr{B}$ if for all $b' \in \mathscr{B}$, either $b' \leq_T b$ or b and b' are T-unrelated. Likewise, b is ***T*-minimal** in $\mathscr{B}$ if for all $b' \in \mathscr{B}$, either $b' \geq_T b$ or b and b' are T-unrelated. Finally, we extend the notion of degree for any point $c \in \mathrm{MA}(S)$. By convention, if c is not a vertex of T but an inner edge point, we say that c has degree 2. We denote the degree by $\deg(c)$.

4 Algorithm

4.1 Principle

For our proposed algorithm, the partial ordering we introduced allows the definition of clear start and end points, as well as a measure of progress. Indeed, let $b_0 = b(c_0, r_0)$ be a medial ball of S. Its center c_0 splits $\mathrm{MA}(S)$ into $\deg(c_0)$ connected components. We denote these components by $\mathrm{branch}(c_0, i)$, $1 \leq i \leq \deg(c_0)$. For our purpose, we want to express the domain covered by balls centered at points of these components of $\mathrm{MA}(S)$, but not covered by b_0. First we define the collection of related medial balls, $\mathscr{C}(b_0, i) = \{b(c,r) \text{ medial in } S \mid c \in \mathrm{branch}(c_0, i)\}$. Then the domain of each $\mathscr{C}(b_0, i)$ is $C(b_0, i) = \bigcup \mathscr{C}(b_0, i) \setminus b_0$. With these notations, b_0 also splits S into the different $C(b_0, i)$'s, and $S \setminus b_0 = \cup_{i=1}^{\deg(c_0)} C(b_0, i)$. Unless c_0 is the

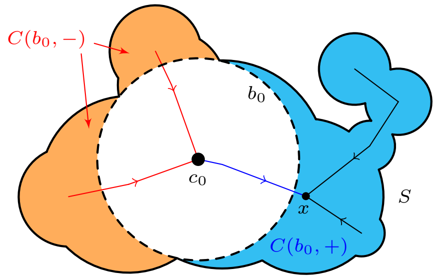

Figure 1: T-large and T-small components of a medial ball. x is the root of T. Red segments are points T-smaller than c_0, blue ones are points T-larger, black ones are T-unrelated.

root of T, one of these domains corresponds to balls T-larger than or T-unrelated to b_0. The other $\deg(c_0) - 1$ domains corresponds to balls T-smaller than b_0. To promote clarity, we refer to the former domain simply as the ***T*-large component**, and denote it by $\boldsymbol{C(b_0,+)}$ (see Figure 1). As for the later domains, we refer to their union as the ***T*-small component** and note $\boldsymbol{C(b_0,-)}$. Hence, we have $S \setminus b_0 = C(b_0,+) \cup C(b_0,-)$. From the definition, we also deduce the following:

Proposition 3 *If $b_1 \leq_T b_2$, then $C(b_1,-) \subseteq C(b_2,-)$ and $C(b_1,+) \supseteq C(b_2,+)$.*

Now assume we want to traverse and sweep S with a medial ball, starting from a leaf of T, toward its root. When we reach a medial ball b_0, then at that moment, $C(b_0,-) \cup b_0$ corresponds to the domain of S that was swept by our medial ball, and $C(b_0,+)$ to the domain of S that was not swept by it. Our approach is based on this particular decomposition of S. We want to use a greedy approach to iteratively compute an ε-covering of the T-small component $C(b_0,-)$. Because T may have several leaves, it is necessary to extend the above definitions of T-small and T-large components to collections $\mathscr{B}$ of medial balls, while preserving the interpretation that $C(\mathscr{B},-) \cup (\bigcup \mathscr{B})$ is the domain of S already processed, and $C(\mathscr{B},+)$ is the domain of S that has not been processed yet. The domain already processed for a collection of balls should thus be the union of the domains already processed by some $b \in \mathscr{B}$. Hence the T-small component of $\mathscr{B}$ is $C(\mathscr{B},-) = \cup_{b\in\mathscr{B}} C(b,-)$. Likewise, the domain that still needs to be processed for $\mathscr{B}$ should be the intersection, over all $b \in \mathscr{B}$, of the domains to be processed for b. Hence the T-large component of $\mathscr{B}$ is $C(\mathscr{B},+) = \cap_{b\in\mathscr{B}} C(b,+)$. Owing to Proposition 3, these definitions emphasize the importance of the T-maximal balls of $\mathscr{B}$. Let $\boldsymbol{T\text{-max}(\mathscr{B})}$ be the collection of these T-maximal balls. Then $C(\mathscr{B},+) = C(T\text{-max}(\mathscr{B}),+)$ and likewise $C(\mathscr{B},-) = C(T\text{-max}(\mathscr{B}),-)$.

To formalize the procedure presented above, we require two more definitions.

Definition 2 Let $\mathscr{B}$ be a collection of medial balls in S. We say that $\mathscr{B}$ is a ***T*-small *ε*-covering** of S if it covers $S^{\ominus\varepsilon}$ in its T-small component $C(\mathscr{B},-)$, that is if $C(\mathscr{B},-) \cap S^{\ominus\varepsilon} \subseteq \bigcup \mathscr{B}$.

Note that every ε-covering is also a T-small ε-covering of S. As such, we employ the term partial T-small ε-covering if we need to distinguish from complete ε-coverings.

Definition 3 Let $\mathscr{B}$ be a partial T-small ε-covering of S, and b_0 be medial in S. We say that b_0 is a **candidate** ball with respect to $\mathscr{B}$, if $\mathscr{B}_0 = \mathscr{B} \cup \{b_0\}$ is also a T-small ε-covering of S, and $S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B}_0 \subsetneq S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B}$.

The strict inclusion $S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B}_0 \subsetneq S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B}$ ensures that $\mathscr{B}_0$ is closer to being a complete ε-covering than $\mathscr{B}$. Hence, for any partial T-small ε-covering, iteratively adding a candidate to the collection ensure that at some point we will obtain a complete ε-covering. Because any partial T-small ε-covering always have an infinity of candidates, we elect to add only T-maximal candidates, that is T-maximal among the collection of candidates.

4.2 Specification

Our algorithm is based on a loop over the collection of all eroded and non-eroded pencils of $\mathrm{MA}(S)$ in a topological order. Since $\mathrm{MA}(S^{\ominus\varepsilon}) \subseteq \mathrm{MA}(S)$, we can simultaneously sweep S and $S^{\ominus\varepsilon}$. With our partial ordering on $\mathrm{MA}(S)$, a topological order of its vertices $(v_1, \ldots, v_{n+1})$ is such that for $i \leq j$, then either $v_i \leq_T v_j$, or v_i and v_j are T-unrelated. Besides the root, each vertex is incident to exactly one pencil composed of T-larger points of $\mathrm{MA}(S)$, hence any topological ordering of vertices induces an ordering of pencils $([v_1\widehat{v}_1], \ldots, [v_n\widehat{v}_n])$, where $\widehat{v}_i \in \{v_1, \ldots, v_n\}$. This ordering of pencils thus satisfies that for $i < j$, then either $v_i <_T \widehat{v}_i \leq_T v_j <_T \widehat{v}_j$, or $\widehat{v}_i$ and v_j are T-unrelated.

As we loop over the pencils, we maintain a collection of medial balls $\mathscr{B}$ which is a T-small ε-covering, while looking for T-maximal candidates to add to said collection. When we process a pencil, we compute a collection of constraints to pass on to the next incident pencil. A constraint is a point or circular arc that any ball T-larger than the pencil (that is T-larger than any ball of that pencil) must contain in order to be a candidate for $\mathscr{B}$. Hence, if no T-maximal candidate is found in the currently processed pencil, the set of constraints it will pass on to its incident T-larger pencil is the collection of all constraints it itself inherited from incident T-smaller pencils, plus new constraints specific to the current pencil. Then, we can compute the T-maximal ball that contains all these constraints. If it is not the T-large endpoint of the pencil, we call it **critical**. We claim that for well chosen constraints, critical balls are T-maximal candidates. The overall approach is summed up in Algorithm 1.

Algorithm 1 Greedy ε-covering

Input: A finite union of balls S
Output: An ε-covering $\mathscr{B}$ of S

```
Compute a topological ordering of MA(S)
𝓑 ← ∅
Loop over all pencils in topological order
    Retrieve incident constraints
    Search for a critical ball in the pencil
    If a critical ball b is found then
        𝓑 ← 𝓑 ∪ {b}
    end If
    Compute the constraints to pass on
end Loop
Return 𝓑
```

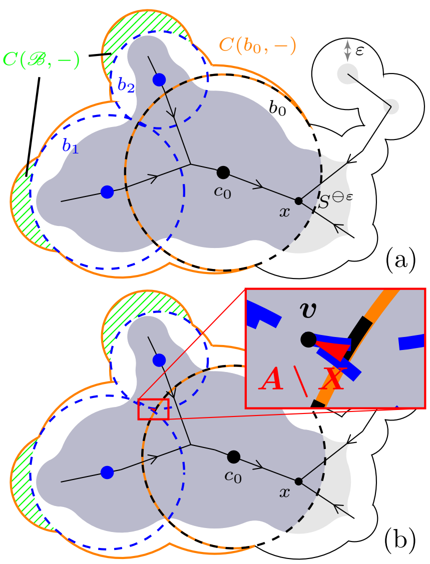

Figure 2: A is the dark shaded area, with $\mathscr{B} = \{b_1, b_2\}$ and $X = b_0 \cup b_1 \cup b_2$. In (a) b_0 is a candidate to $\mathscr{B}$, but in (b) vertex v does not satisfy condition (ii) and $A \setminus X$ (in red) is non empty.

All that remains is to explicit constraints which guarantee that critical balls are indeed T-maximal candidates. Consider a candidate b_0 to the collection $\mathscr{B}$. We can ignore balls which are T-unrelated to b_0, thus let $\mathscr{B}' = \{b \in \mathscr{B},\ b \leq_T b_0\}$. By definition, $\mathscr{B}'_0 = \mathscr{B}' \cup \{b_0\}$ is a T-small ε-covering. The domain of $S^{\ominus\varepsilon}$ covered by $C(\mathscr{B}'_0, -) \cup (\bigcup \mathscr{B}'_0)$ but not by $C(\mathscr{B}', -)$ is contained in $(\bigcup \mathscr{B}'_0) \setminus C(\mathscr{B}', -)$. Let that former domain be

$$A = \Big(\big(C(\mathscr{B}'_0, -) \cup \big(\bigcup \mathscr{B}'_0\big)\big) \setminus C(\mathscr{B}', -)\Big) \cap S^{\ominus\varepsilon}$$
$$= ((C(b_0, -) \cup b_0) \setminus C(\mathscr{B}', -)) \cap S^{\ominus\varepsilon}$$

and the latter be

$$X = \Big(\bigcup \mathscr{B}'_0\Big) \setminus C(\mathscr{B}', -) = b_0 \cup \Big(\bigcup T\text{-max}(\mathscr{B}')\Big).$$

See Figure 2a. A and X both vary depending on $\mathscr{B}$ and b_0, and in light of the previous remark, b_0 is a candidate to $\mathscr{B}$ if and only if $A \subseteq X$. Explicitly ensuring and verifying that we indeed have the inclusion $A \subseteq X$ is non trivial. Instead, we claim that it is sufficient to ensure the two conditions:

(i) $\partial A \subseteq X$,

(ii) $\forall$ vertex $v \in \partial X$, $\exists N_v$ an open neighbourhood of v, such that $N_v \cap A \subseteq X$.

Note that both ∂A and ∂X are finite collections of circular arcs, whose intersections we call vertices of the boundary. Since these boundaries have a finite combinatorial structure, (i) and (ii) are more readily verifiable and can be enforced. Also, as illustrated in Figure 2b, condition (i) by itself is insufficient to ensure that b_0 is a valid candidate. We take as constraints for Algorithm 1 any arc of ∂A not in $\bigcup \mathscr{B}'$, as well as any vertex of $\partial(\bigcup T\text{-max}(\mathscr{B}'))$ whose neighbourhood in A is not fully contained in $\bigcup \mathscr{B}'$. This ensures that critical balls can be computed as per Section 5 and fulfill conditions (i) and (ii). Both conditions are necessary to have the inclusion $A \subseteq X$. We prove in Section 6 that they are also sufficient.

5 Computation of critical balls

As stated previously, each pencil inherits a collection of constraints that are either a singleton point or a circular arc. Because of Proposition 4, which will be stated and explained in detail later, when we sweep a pencil from an endpoint to the other, a point that exited the sweep ball will never re-enter it, and a point that entered the sweep ball will never exit it. Thus, to each constraint corresponds a **single-constraint critical** ball b_{crit}: any ball strictly T-larger than b_{crit} cannot fully contain the constraint, hence cannot be a candidate, while any ball T-smaller than or equal to b_{crit} will always contain the constraint and may be a candidate. Therefore, the critical ball for the overall collection of constraints is the T-maximal ball that is T-smaller than all single-constraint critical balls, that is the unique T-minimal ball amongst all single-constraint critical balls. From here on, we omit the qualifier "single-constraint". Also, because we parameterize balls of a pencil by an interpolation value λ, we call critical λ an interpolation value which corresponds to a critical ball.

The next sections present how to compute a critical ball given a single constraint. Section 5.1 presents a very useful property of pencils and how to handle point constraints, while Section 5.2 deals with arc constraints.

5.1 Point inclusion

Given a ball $b = b(c, r)$ and a point y, there are many ways to test whether y belongs to b or not. By definition, we can compare the distance from y to c, to the radius of b. Here, we rely on the **power** of y with respect to ball b, which by definition is $\mathbf{pow}(\boldsymbol{y},\boldsymbol{b}) = \|y-c\|^2 - r^2$. Hence, y belongs to ball b if and only if $\operatorname{pow}(y,b) \le 0$.

Consider now a pencil of balls $[b_1b_2]$, with $b_i = b(c_i,r_i)$. For $\lambda \in [0,1]$, we denote by $b_\lambda = b(c_\lambda, r_\lambda)$ the ball of pencil $[b_1b_2]$ that is centered at $c_\lambda = \lambda c_1 + (1-\lambda)c_2$. We argue that then, for any point y, the power of y with respect to b_λ is the linear interpolation of its power with respect to b_1 and b_2.

Proposition 4

$$\operatorname{pow}(y,b_\lambda) = \lambda \operatorname{pow}(y,b_1) + (1-\lambda)\operatorname{pow}(y,b_2)$$

The power of a point with respect to balls of the pencil is thus linear in λ and may only change sign once. This justifies our earlier claim that when sweeping a pencil, a point can exit or enter the sweep ball only once.

Proof. By definition, all the balls of a pencil $[b_1b_2]$ share exactly two points $\{u,v\} = \partial b_1 \cap \partial b_2$ on their boundary. We have $r_\lambda^2 = \|u-c_\lambda\|^2$. Hence:

$$\begin{aligned}\operatorname{pow}(y,b_\lambda) &= \|y-c_\lambda\|^2 - \|u-c_\lambda\|^2\\ &= \|y\|^2 - 2\langle c_\lambda,\, y\rangle + \|c_\lambda\|^2\\ &\quad - \left(\|u\|^2 - 2\langle c_\lambda,\, u\rangle + \|c_\lambda\|^2\right)\\ &= 2\langle c_\lambda,\, u-y\rangle + \|y\|^2 - \|u\|^2\end{aligned}$$

Note that this identity also holds for b_1 and b_2 since they coincide with b_λ, for $\lambda \in \{0,1\}$.

$$\begin{aligned}\operatorname{pow}(y,b_\lambda) &= 2\langle \lambda c_1 + (1-\lambda)c_2,\, u-y\rangle + \|y\|^2 - \|u\|^2\\ &= \lambda\left(2\langle c_1,\, u-y\rangle + \|y\|^2 - \|u\|^2\right)\\ &\quad + (1-\lambda)\left(2\langle c_2,\, u-y\rangle + \|y\|^2 - \|u\|^2\right)\\ &= \lambda \operatorname{pow}(y,b_1) + (1-\lambda)\operatorname{pow}(y,b_2)\end{aligned}$$

□

This proof actually extends to all balls of the complete line pencil. From there, given two endpoint balls of a segment pencil and a single constraint point, we can easily compute the value λ_{crit} for which $b_{\lambda_{\text{crit}}}$ will be critical. From the design of Algorithm 1, and assuming that $b_1 \le_T b_2$, we will always have $\lambda_{\text{crit}} \ge 0$. If we happen to have $\lambda_{\text{crit}} > 1$, then the segment pencil does not contain any critical ball for that constraint point, since all balls of the pencil contain that constraint point.

5.2 Arc inclusion

For arc constraints, we use the same approach of computing a critical value for λ. If all of these critical λ's are strictly larger than 1, then the pencil does not contain any T-maximal candidate. If any are less than or equal to 1, then the minimum value yields a T-maximal candidate. All that remains is thus being able to compute such a critical λ for arc constraints. To do so, we first need to compute this critical value for a ball.

5.2.1 Critical λ for balls

Consider a constraint ball $b = b(c,r)$. We want to compute the critical values of λ for which b is fully contained in b_λ. Hence:

$$\begin{aligned} b \subseteq b_\lambda &\iff r_\lambda \ge \|c-c_\lambda\| + r\\ &\iff r_\lambda^2 \ge \|c-c_\lambda\|^2 + r^2 + 2r\|c-c_\lambda\|\\ &\iff -\left(\|c-c_\lambda\|^2 - r_\lambda^2\right) - r^2 \ge 2r\|c-c_\lambda\|\\ &\iff -\operatorname{pow}(c,b_\lambda) - r^2 \ge 2r\|c-c_\lambda\|\\ &\iff \left(-\operatorname{pow}(c,b_\lambda) - r^2\right)^2 \ge 4r^2\|c-c_\lambda\|^2\\ &\qquad \& \quad -\operatorname{pow}(c,b_\lambda) - r^2 \ge 0\end{aligned}$$

Since $\operatorname{pow}(c_0,b_\lambda)$ is linear in λ as per Proposition 4, we introduce two constants A and B such that $A\lambda + B = \operatorname{pow}(c,b_\lambda) + r^2$ to simplify notations. We have

$$\begin{aligned}A &= \operatorname{pow}(c,b_1) - \operatorname{pow}(c,b_2)\\ B &= \operatorname{pow}(c,b_2) + r^2.\end{aligned}$$

We focus on the first inequality, $(A\lambda+B)^2 \ge 4r^2\|c-c_\lambda\|^2$. For the right hand side, the factor $\|c-c_\lambda\|^2$ can be developed as follows

$$\begin{aligned}&\|c_\lambda - c\|^2\\ &= \|\lambda(c_1-c_2) + c_2 - c\|^2\\ &= \lambda^2\|c_1-c_2\|^2 + 2\lambda\langle c_1-c_2,\, c_2-c\rangle + \|c_2-c\|^2\end{aligned}$$

and thus, it is quadratic in λ. Therefore

$$b \subseteq b_\lambda \iff C\lambda^2 + D\lambda + E \ge 0 \quad \& \quad A\lambda + B \le 0$$

where

$$\begin{aligned}C &= A^2 - 4r^2\|c_1-c_2\|^2\\ D &= 2AB - 8r^2\langle c_1-c_2,\, c_2-c\rangle\\ E &= B^2 - 4r^2\|c_2-c\|^2\end{aligned}$$

A, B, C, D and E are constant with respect to λ and thus the inclusion test can be reduced to a sign analysis of polynomials of degree 2 and 1. From the roots of these polynomials, we can thus derive the critical values of λ.

5.2.2 Critical λ for arcs

In order to compute the critical λ value for an arc constraint e, we first require the critical value for the ball b supporting that arc, that is the ball such that $e \subseteq \partial b$. Let λ' be the critical value for that supporting ball.

With λ', we can then compute the tangency point y between b and $b_{\lambda'}$. This particular point, in addition with the endpoints of arc e, are sufficient to compute the critical value of e. Indeed, for $\lambda \leq \lambda'$, the whole ball b is contained in b_λ, hence we also have $e \subseteq b_\lambda$. For $\lambda > \lambda'$, two cases arise. Either $y \in e$ or $y \notin e$. In the former, we have $y \in e \setminus b_\lambda$ hence the critical value for the arc e is In the former, we have $y \in e$ and $y \notin b_\lambda$ hence the critical value for the arc e is λ' itself. In the latter, note that b_λ splits ∂b into two connected components, one which is inside b_λ, and the other outside. By Proposition 4, the outside component will always contain y. For $\lambda > \lambda_{\text{crit}}$, any point of e not in b_λ must be path-connected in ∂b to y. Therefore, some endpoint of e also belongs to the outside connected component. It immediately follows that if both endpoints of e actually belong to b_λ, then the whole arc e is also contained in b_λ. Therefore, the critical value for e is in this case equal to the smallest critical value of its endpoints.

6 Correctness of the algorithm

6.1 Convergence to an ε-covering

Lemma 5 *Let b_0 and $\mathscr{B}$ such that b_0 fulfills both conditions (i) and (ii). Then b_0 is a candidate for $\mathscr{B}$.*

Proof. Let A and X be defined from b_0 and $\mathscr{B}$. Consider $H = A \setminus X = A \cap X^{\mathsf{c}}$. By contradiction assume that $H \neq \varnothing$. First, notice that $\partial H \subseteq \partial X$. Indeed, $\partial H = \partial (A \cap X^{\mathsf{c}}) \subseteq \left(\partial A \cap \overline{X^{\mathsf{c}}}\right) \cup \left(\partial\left(X^{\mathsf{c}}\right) \cap \overline{A}\right)$. By condition (i), $\partial A \subseteq X$ and we have $\partial A \cap \overline{X^{\mathsf{c}}} \subseteq \partial X$. Also, $\partial\left(X^{\mathsf{c}}\right) = \partial X$. Hence $\partial H \subseteq \partial X$. Let H_i be a connected component of H. Since $\partial H_i \subseteq \partial H$, hence $\partial H_i \subseteq \partial X$. We have $H_i \subseteq X^{\mathsf{c}}$ connected, with $\partial H_i \subseteq \partial X$. Thus, H_i is actually a connected component of X^{c}. H_i being bounded, it is commonly called a hole of X. As a hole of X, any vertex of ∂H_i is also a vertex of ∂X. Let v be a vertex of ∂H_i. For all open neighbourhood N_v of v, we have $N_v \cap H \neq \varnothing$. Since $H = A \cap X^{\mathsf{c}}$, we deduce $\varnothing \neq N_v \cap A \nsubseteq X$. This contradicts (ii) and is impossible. Therefore, $H = \varnothing$, $A \subseteq X$, and b_0 is indeed a candidate for $\mathscr{B}$. $\square$

From the above lemma, Algorithm 1 indeed finds candidates and eventually converges to an ε-covering. We now show that it converges in polynomial time.

Lemma 6 *Algorithm 1 converges to an ε-covering in $\mathcal{O}\left(|\mathrm{MA}\,(S)|^2\right)$.*

Proof. Let $n = |\mathrm{MA}\,(S)|$. First, we show that Algorithm 1 outputs a collection $\mathscr{B}_{\text{algo}}$ with size at most $2n$, and then that the complexity is at most quadratic.

For any partial T-small ε-covering $\mathscr{B}$, let $b \in T\text{-max}\,(\mathscr{B})$. There is a unique pencil incident to b that contains balls T-larger than b. Let $b_0 >_T b$ be the T-large endpoint of that pencil. We show that b_0 is always a candidate to $\mathscr{B}$. Though b_0 may not be a T-maximal candidate, this still implies that there can be at most two medial balls from the same pencil in $\mathscr{B}_{\text{algo}}$.
Let $\mathscr{B}_0 = \mathscr{B} \cup \{b_0\}$. Because b and b_0 belong to the same pencil, $C(\mathscr{B}_0, -)$ and $C(\mathscr{B}, -)$ only differ in the domain of S covered by the pencil $[bb_0]$. This implies the equalities $C(\mathscr{B}_0, -) \setminus C(\mathscr{B}, -) = C(b_0, -) \setminus C(b, -) = b \setminus b_0$. Hence, we obtain

$$A = ((b \setminus b_0) \cup b_0) \cap S^{\ominus\varepsilon} \subseteq b \cup b_0,$$

thus b_0 is a candidate for $\mathscr{B}$. Therefore, $|\mathscr{B}_{\text{algo}}| \leq 2n$.

We now analyse the complexity. Computing a topological ordering [5] is linear in n. Enforcing a single constraint takes constant time (see Section 5), hence the time expanded to search for a critical ball depends on the number of constraints. This number cannot exceed the combinatorial complexity of the boundaries of $S^{\ominus\varepsilon}$ and $\bigcup \mathscr{B}_{\text{algo}}$. The former is linear in n. As for the latter, it is linear in the size of $\mathscr{B}_{\text{algo}}$, which is itself linear in n. Thus, Algorithm 1 is quadratic in n. $\square$

6.2 Convergence to an optimal solution

In order to prove that our algorithm reaches an optimal, we rely on several intermediate results. We introduce a sequence of three lemmas, the last of which we reformulate, through two corollaries, in terms of candidate ball to a partial T-small ε-covering. Then, we finally prove Proposition 12.

Lemma 7 *Consider a finite union of balls S, and a ball b such that $b \nsubseteq S$. Let $S' = S \cup b$. Then, $\mathring{S}' = \mathring{S} \cup \mathring{b}$.*

Lemma 8 *Consider a finite union of balls S such that $\mathrm{MA}\,(S)$ is a tree. Then $\overline{S^{\mathsf{c}}}$ is path-connected.*

Proof. Let $\mathscr{S}$ be the collection of medial balls in S which are centered at a vertex of $\mathrm{MA}\,(S)$. $\mathscr{S}$ has finite cardinality and we have $\bigcup \mathscr{S} = S$. We proceed by induction on $n = |\mathscr{S}|$, and henceforth use the notation $S_n = \bigcup \mathscr{S}_n$ for collection of balls with cardinality n. For $n = 1$, S_1 is only a ball, the property is verified. Now consider a collection $\mathscr{S}_{n+1}$, and let $b \in \mathscr{S}_{n+1}$ such that b is centered on a leaf of $\mathrm{MA}\,(S_{n+1})$. Let $\mathscr{S}_n = \mathscr{S}_{n+1} \setminus \{b\}$ and $S_n = \bigcup \mathscr{S}_n$.

Using Lemma 7, we deduce that $\overline{S^{\mathsf{c}}_{n+1}} = \overline{S^{\mathsf{c}}_n} \cap \overline{b^{\mathsf{c}}}$. Consider $y, z \in \overline{S^{\mathsf{c}}_{n+1}}$. By induction assumption, we know that $\overline{S^{\mathsf{c}}_n}$ is path-connected, hence there is a path $\gamma \subseteq \overline{S^{\mathsf{c}}_n}$ connecting y and z. We build from γ another path $\gamma' \subseteq \overline{S^{\mathsf{c}}_{n+1}}$ that connects y and z. If $\gamma \subseteq \overline{b^{\mathsf{c}}}$ we already have $\gamma \subseteq \overline{S^{\mathsf{c}}_{n+1}}$. Otherwise let $\pi_y, \pi_z \in \gamma \cap \partial b$ such that γ does not meet b between y and π_y, and likewise between z and π_z. Because $\pi_y, \pi_z \in \overline{S^{\mathsf{c}}_n}$, they cannot be in $\mathring{S}_n$. Let $e = \partial b \setminus \mathring{S}_n$. e is a path-connected circular arc, and is the contribution of ∂b to ∂S_{n+1}. We have $\pi_y, \pi_z \in e$. Since e is path-connected and contained in $\overline{S^{\mathsf{c}}_{n+1}}$, we can complete γ' by following e to go from π_y to π_z, showing that $\overline{S^{\mathsf{c}}_{n+1}}$ is path-connected. $\square$

Lemma 9 *Consider b_0 a medial ball of S. $C(b_0,+)$ and $C(b_0,-)$ are interior disjoint.*

Proof. By contradiction, assume that $\mathring{C}(b_0,+)$ and $\mathring{C}(b_0,-)$ are not disjoint. From there, we exhibit two paths $\gamma \subseteq \mathring{S}$ and $\delta \subseteq \overline{S^{\mathsf{c}}}$ with non empty intersection, which is impossible.

Let $y \in \mathring{C}(b_0,+) \cap \mathring{C}(b_0,-)$. Without loss of generality, we assume that $y \notin \mathrm{MA}\,(S)$. Necessarily, there are two medial balls of S, b_+ and b_-, such that:

$$\begin{aligned} b_+ &\subseteq C(b_0,+) \cup b_0 \\ b_- &\subseteq C(b_0,-) \cup b_0 \\ y &\in \left(\mathring{b}_+ \cap \mathring{b}_-\right) \setminus b_0. \end{aligned}$$

Let c_+ and c_- be the respective centers of b_+ and b_-. We have segment $[c_+y] \subseteq \mathring{S}$ and likewise segment $[c_-y] \subseteq \mathring{S}$. There is also a path in $\mathrm{MA}\,(S) \subseteq \mathring{S}$ connecting c_+ and c_-. Hence there exists a Jordan curve $\gamma \subseteq \mathrm{MA}\,(S) \cup [c_+y] \cup [c_-y] \subseteq \mathring{S}$. By the Jordan-Brouwer separation theorem, γ has a well defined interior and exterior. Consider now $e = \partial b_0 \cap \overline{C(b_0,+)}$. It is a path-connected circular arc. Its two endpoints are vertices of ∂S, hence we have $\partial e \subseteq \overline{S^{\mathsf{c}}}$. Moreover, one of those endpoints lies in the interior of γ, while the other lies in the exterior of γ. See Figure 3 for a schematic representation. However by Lemma 8, $\overline{S^{\mathsf{c}}}$ is path-connected. Therefore, let $\delta \subseteq \overline{S^{\mathsf{c}}}$ be a path connecting the endpoints of e. Since the interior and exterior of γ are separated, necessarily $\varnothing \neq \gamma \cap \delta \subseteq \mathring{S} \cap \overline{S^{\mathsf{c}}} = \varnothing$, hence the contradiction. $\square$

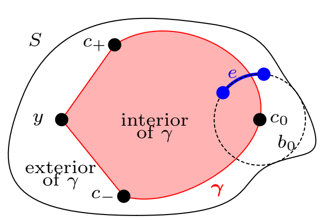

Figure 3: Schematic representation for Lemma 9.

Though Lemmas 7 and 8 are quite remote from the result we want to prove, in essence Lemma 9 implies that whatever medial ball we chose in $C(b_0,+)$, it cannot contribute to cover $S^{\ominus\varepsilon}$ in $C(b_0,-)$. That is why we can process each of these components separately in a greedy way and still achieve a global optimal solution. Formally, we rely on the following corollaries.

Corollary 10 *Let b_0 be a medial ball in S. Then we have the three identities:*

$$\begin{aligned} (C(b_0,-) \cup b_0)^{\mathsf{c}} \cap S^{\ominus\varepsilon} &= C(b_0,+) \cap S^{\ominus\varepsilon} \\ (C(b_0,+) \cup b_0)^{\mathsf{c}} \cap S^{\ominus\varepsilon} &= C(b_0,-) \cap S^{\ominus\varepsilon} \\ (C(b_0,+) \cup C(b_0,-))^{\mathsf{c}} \cap S^{\ominus\varepsilon} &= b_0 \cap S^{\ominus\varepsilon} \end{aligned}$$

Corollary 10 simply states that when restricted to $S^{\ominus\varepsilon}$, the complement of $C(b_0,+)$, $C(b_0,-)$, and b_0, is the union of the other two subsets.

Proof. We only prove the first equality. We have $S^{\ominus\varepsilon} \subseteq \mathring{S}$. Additionally by Lemma 9, $\mathring{S}$ is the disjoint union of $C(b_0,-) \cap \mathring{S}$, $b_0 \cap \mathring{S}$, and $C(b_0,+) \cap \mathring{S}$. Hence,

$$\begin{aligned} S^{\ominus\varepsilon} \cap (C(b_0,-) \cup b_0)^{\mathsf{c}} &= S^{\ominus\varepsilon} \cap \mathring{S} \cap (C(b_0,-) \cup b_0)^{\mathsf{c}} \\ &= S^{\ominus\varepsilon} \cap \mathring{S} \cap C(b_0,+) \\ &= S^{\ominus\varepsilon} \cap C(b_0,+). \end{aligned}$$

$\square$

Corollary 11 *Consider an ε-covering $\mathscr{B}$. Let $\mathscr{B}_- \subsetneq \mathscr{B}$ be a partial T-small ε-covering, and let b_0 be any T-maximal candidate to $\mathscr{B}_-$. Then $\mathscr{B} \setminus \mathscr{B}_-$ contains a candidate to $\mathscr{B}_-$ that is T-smaller than or equal to b_0.*

Proof. Let $\mathscr{B}_+ = \mathscr{B} \setminus \mathscr{B}_-$. First we prove that $\mathscr{B}_+$ always contains candidates to $\mathscr{B}_-$, and then that one of these candidates is T-smaller than or equal to b_0.

By contradiction, assume that $\mathscr{B}_+$ is void of candidate to $\mathscr{B}_-$. Consider $b' \in \mathscr{B}_+$, T-minimal in $\mathscr{B}_+$. By T-minimality of b', for all $b \in \mathscr{B}_+$, $b \subseteq b' \cup C(b',+)$. Thus $\bigcup \mathscr{B}_+ \subseteq b' \cup C(b',+)$. By Corollary 10, we deduce that $S^{\ominus\varepsilon} \cap C(b',-) \subseteq S^{\ominus\varepsilon} \cap (\bigcup \mathscr{B}_+)^{\mathsf{c}} = S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B}_+$. By assumption, b' cannot be a candidate to $\mathscr{B}_-$, hence by Definitions 2 and 3 $(S^{\ominus\varepsilon} \cap C(b',-)) \setminus (b' \cup (\bigcup \mathscr{B}_-))$ is non empty. With the previous inclusion, we get the following development.

$$\begin{aligned} &\left(S^{\ominus\varepsilon} \cap C(b',-)\right) \setminus \left(b' \cup \left(\bigcup \mathscr{B}_-\right)\right) \\ &= \left(S^{\ominus\varepsilon} \cap C(b',-)\right) \setminus \left(b' \cup \left(\bigcup \mathscr{B}_-\right) \cup \left(\bigcup \mathscr{B}_+\right)\right) \\ &= \left(S^{\ominus\varepsilon} \cap C(b',-)\right) \setminus \bigcup \mathscr{B} \\ &\subseteq S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B} \end{aligned}$$

Hence, $S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B} \neq \varnothing$ which is impossible since $\mathscr{B}$ is an ε-covering. Thus, $\mathscr{B}_+$ contains a candidate to $\mathscr{B}_-$.

Because $\mathscr{B}_-$ may have several distinct T-maximal candidates, $\mathscr{B}_+$ may only contain candidates T-unrelated to b_0. By contradiction assume $\mathscr{B}_+$ is void of candidate to $\mathscr{B}_-$ T-smaller than or equal to b_0. Since $\mathscr{B}_+$ contains at least one candidate to $\mathscr{B}_-$, b_0 cannot be centered at the root of T. By Proposition 3, any ball T-smaller than or equal to b_0 is a candidate to $\mathscr{B}_-$. Hence

$\mathscr{B}_+$ only contains balls that are either strictly T-larger than b_0, or T-unrelated to b_0. Therefore, there exists a medial ball $b' >_T b_0$, with b' also strictly T-smaller or T-unrelated to balls in $\mathscr{B}_+$. By T-maximality of b_0, b' cannot be a candidate to $\mathscr{B}_-$. Once again, we have a ball b', T-minimal in $\mathscr{B}'_+ = \mathscr{B}_+ \cup \{b'\}$ which is not a candidate. The same development as above using Corollary 10 yields $S^{\ominus\varepsilon} \setminus \bigcup \mathscr{B} \neq \varnothing$, which is still impossible. Therefore, $\mathscr{B}_+$ must contain a candidate to $\mathscr{B}_-$ that is T-smaller than or equal to b_0. □

Proposition 12 *Algorithm 1 converges to an optimal ε-covering.*

Proof. We denote by $\mathscr{B}_{\text{algo}}$ the ε-covering found by Algorithm 1. We number the balls of $\mathscr{B}_{\text{algo}}$ by $b_1, \ldots, b_k$, such that for $i \leq j$, b_i was found before b_j. Consider any optimal ε-covering $\mathscr{B}_{\text{opt}}$. We assume without loss of generality that $\mathscr{B}_{\text{opt}}$ only contains medial balls. Indeed, S is a finite union of balls, hence any ball $b \in \mathscr{B}_{\text{opt}}$ is wholly contained in a medial ball. Using consecutive substitutions, we want to build a finite sequence of ε-coverings $\mathscr{B}_0, \ldots, \mathscr{B}_k$ that satisfies the properties:

(a) $\mathscr{B}_0 = \mathscr{B}_{\text{opt}}$,

(b) $|\mathscr{B}_{i+1}| = |\mathscr{B}_i|$, $\forall i \in [\![0, k-1]\!]$,

(c) $\{b_1, \ldots, b_i\} \subseteq \mathscr{B}_i$, $\forall i \in [\![1, k]\!]$,

If such a sequence exists, we immediately deduce that $|\mathscr{B}_{\text{algo}}| = |\mathscr{B}_{\text{opt}}|$, and $\mathscr{B}_{\text{algo}}$ is also optimal.

We proceed by induction. Assume that for $0 \leq i < k$, we have built $\mathscr{B}_0, \ldots, \mathscr{B}_i$ with the above properties. Consider b_{i+1}. Let $\mathscr{B}_- = \{b_1, \ldots, b_i\}$. By construction, $\mathscr{B}_- \subsetneq \mathscr{B}_i$. Let $\mathscr{B}_+ = \mathscr{B}_i \setminus \mathscr{B}_-$. Algorithm 1 guarantees that $\mathscr{B}_-$ is a partial T-small ε-covering and that b_{i+1} is a T-maximal candidate for $\mathscr{B}_-$. Hence we can apply Corollary 11 and there is a candidate b to $\mathscr{B}_-$, such that $b \in \mathscr{B}_+$ and $b \leq_T b_{i+1}$. Then, let $\mathscr{B}_{i+1} = (\mathscr{B}_i \cup \{b_{i+1}\}) \setminus \{b\}$. $\mathscr{B}_{i+1}$ satisfies both properties (b) and (c), we must prove that it is also an ε-covering. To do so, it suffices to prove that both $C(b_{i+1}, -) \cap S^{\ominus\varepsilon}$ and $C(b_{i+1}, +) \cap S^{\ominus\varepsilon}$ are contained in $\bigcup \mathscr{B}_{i+1}$. Because b_{i+1} is a candidate to $\mathscr{B}_-$, we have $C(b_{i+1}, -) \cap S^{\ominus\varepsilon} \subseteq b_{i+1} \cup (\bigcup \mathscr{B}_-) \subseteq \bigcup \mathscr{B}_{i+1}$. Also, $b \leq_T b_{i+1}$, hence by Proposition 3 we have $C(b_{i+1}, +) \subseteq C(b, +)$. This implies $C(b_{i+1}, +) \cap S^{\ominus\varepsilon} \subseteq C(b, +) \cap S^{\ominus\varepsilon} \subseteq (\bigcup \mathscr{B}_i) \setminus b \subseteq \bigcup \mathscr{B}_{i+1}$. Thus, $\mathscr{B}_{i+1}$ is an ε-covering, and $\mathscr{B}_{\text{algo}}$ is an optimal ε-covering. □

7 Discussion

Consider the following more general covering definition, where $S^{\oplus\varepsilon'} = \cup_{y \in S} b(y, \varepsilon')$ is the dilation of S (by ε'):

Definition 4 An $\boldsymbol{\varepsilon'\varepsilon}$**-covering** of S is a collection of balls $\mathscr{B}$ such that $S^{\ominus\varepsilon} \subseteq \bigcup \mathscr{B} \subseteq S^{\oplus\varepsilon'}$.

The algorithm presented computes in polynomial time an optimal $\varepsilon'\varepsilon$-covering of any union of balls S such that $\mathrm{MA}(S)$ is a forest and $\mathrm{MA}(S^{\ominus\varepsilon}) \subseteq \mathrm{MA}(S^{\ominus\varepsilon'})$.

An interesting perspective is to build on this work to design a heuristic algorithm where both conditions on the medial axis are relaxed.

References

[1] N. Amenta and R. K. Kolluri. The medial axis of a union of balls. *Computational Geometry*, 20(1):25–37, 2001.

[2] D. Attali, T.-B. Nguyen, and I. Sivignon. Epsilon-covering is NP-complete. In *EuroCG*, Lugano, Switzerland, Mar. 2016.

[3] G. Bradshaw and C. O'Sullivan. Adaptive medial-axis approximation for sphere-tree construction. *ACM Transactions on Graphics (TOG)*, 23(1):1–26, 2004.

[4] F. Cazals, T. Dreyfus, S. Sachdeva, and N. Shah. Greedy geometric algorithms for collection of balls, with applications to geometric approximation and molecular coarse-graining. In *Computer Graphics Forum*, volume 33, pages 1–17. Wiley Online Lib., 2014.

[5] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson. *Introduction to Algorithms.* 2001.

[6] B. Miklos, J. Giesen, and M. Pauly. Discrete scale axis representations for 3d geometry. In *ACM Transactions on Graphics (TOG)*, volume 29, page 101. ACM, 2010.

[7] H. Schwerdtfeger. *Geometry of complex numbers: circle geometry, Moebius transformation, non-euclidean geometry.* Courier Corporation, 1979.