

A Closer Cut: Computing Near-Optimal Lawn Mowing Tours

Sándor P. Fekete* Dominik Krupke* Michael Perk* Christian Rieck*
Christian Scheffer†

Abstract

For a given polygonal region P , the Lawn Mowing Problem (LMP) asks for a shortest tour T that gets within Euclidean distance 1 of every point in P ; this is equivalent to computing a shortest tour for a unit-disk cutter C that covers all of P . As a geometric optimization problem of natural practical and theoretical importance, the LMP generalizes and combines several notoriously difficult problems, including minimum covering by disks, the Traveling Salesman Problem with neighborhoods (TSPN), and the $\exists\mathbb{R}$ -complete Art Gallery Problem (AGP). So far, there have only been theoretical approximation algorithms with worst-case bounds of $2\sqrt{3}\alpha_{\text{TSP}} \approx 3.46\alpha_{\text{TSP}}$, where α_{TSP} is the approximation factor for the geometric TSP. Here, $\alpha_{\text{TSP}} = 1 + \varepsilon$ is theoretically possible by using one of the famous geometric approximation schemes; however, these methods are not practically applicable for concrete instances. Moreover, there have not been any exact methods for the LMP that compute provably near-optimal solutions for instances of interesting size, owing to the combination of geometric difficulties, such as a succinct characterization of optimal solutions, as well as the lack of useful lower bounds that provide practically small performance gaps.

In this paper, we conduct the first study of the Lawn Mowing Problem with a focus on practical computation of near-optimal solutions. To this end, we provide new theoretical insights: Optimal solutions are polygonal paths with a bounded number of vertices, i.e., they do not have any curved pieces, allowing a restriction to straight-line solutions; on the other hand, there can be relatively simple instances for which optimal solutions require a large class of irrational coordinates. On the practical side, we present a primal-dual approach with provable convergence properties based on solving a special case of the TSPN restricted to *witness sets*.

In each iteration, this establishes both a valid solution and a valid lower bound, and thereby a bound on the remaining optimality gap. As we demonstrate in an extensive computational study, this allows us to achieve provably optimal and near-optimal solutions for a large spectrum of benchmark instances with up to 2000 vertices.

1 Introduction

There are many facets of theoretical and practical difficulty of geometric optimization problems. On the theoretical side, the classic *Traveling Salesman Problem* (TSP) is NP-hard, making it unlikely that there is a polynomial-time algorithm that produces provably optimal solutions. Moreover, it is unknown how to efficiently evaluate a sum of square roots, so it is unclear whether the TSP for a set of points in the plane with Euclidean distances even belongs to NP. For the famous *Art Gallery Problem* (AGP) of finding a minimum number of guards to cover a simple polygon based on visibility, membership in NP is indeed unlikely, as it belongs to the class $\exists\mathbb{R}$. Furthermore, problems of optimal covering are also known to be prohibitively difficult from a *practical* perspective: Even the largest square that can be covered by n unit disks has only been established up to $n = 7$.

We consider a geometric optimization problem that generalizes and combines these challenges. In the *Lawn Mowing Problem* (LMP), we are given a (not necessarily simple or even connected) polygonal region P and a disk cutter C of radius r ; the task is to find a closed roundtrip of minimum Euclidean length such that the cutter “mows” all of P , i.e., a shortest tour that moves the center of C within distance r from every point in P . The LMP naturally occurs in a wide spectrum of practical applications, such as robotics, manufacturing, farming, quality control and image processing, so it is of both theoretical and practical importance.

Given that the LMP combines the Euclidean TSP, the AGP, and covering by disks, it is not surprising that it is also both theoretically and practically difficult: It is NP-hard (as it generalizes the TSP), membership in

*Department of Computer Science, TU Braunschweig, Braunschweig, Germany

†Faculty of Electrical Engineering and Computer Science, Bochum University of Applied Sciences, Bochum, Germany

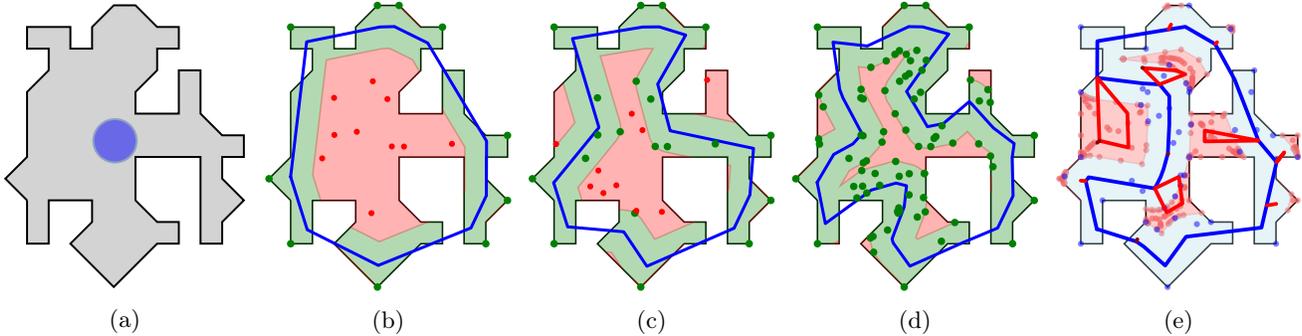


Figure 1: A sequence of iterations in the primal-dual scheme. **(a)** An example instance. The blue circle shows the tool size. **(b)** A lower bound of 32.4 with coverage of 51 %, arising in iteration 0 from an initial witness set. **(c)** A lower bound of 36.56 with coverage of 70 %, arising in iteration 1 from an enhanced witness set. **(d)** A lower bound of 42.37 with coverage of 84 %, arising in iteration 7. **(e)** An upper bound of 65.35 with full coverage, achieved in iteration 3.

NP is unclear (as it involves evaluating Euclidean distances); it is also extremely hard from a practical perspective (as it comprises covering by disks). In fact, the only known positive algorithmic result is an approximation algorithm with a factor of $2\sqrt{3}\alpha_{\text{TSP}} \approx 3.46\alpha_{\text{TSP}}$ [7], where α_{TSP} is the factor for the geometric TSP. So far, no results aiming at methods with *practically* good performance are known, in part because of the difficulties of (I) characterizing optimal solutions (due to the continuous nature of the LMP) and of (II) providing tight lower bounds: Neither the area of P nor its diameter can provide such bounds, as replacing P by a dense subset of points with zero area does not change the length of an optimal tour, which can be much longer than the diameter of P .

1.1 Our contribution We provide a number of theoretical and practical results for the Lawn Mowing Problem.

- We establish a characterization of optimal lawn mowing tours by proving that an optimal tour for a polygonal region P consists of line segments between a finite set of vertices. This allows a focus on polygonal solutions, and the ensuing primal-dual scheme.
- On the other hand, we show that even relatively simple regions P can require LMP solutions with a wide range of irrational vertices. While this does not establish $\exists\mathbb{R}$ -hardness, it gives some indication of the underlying difficulty, as it did for the AGP.
- We establish a primal-dual algorithm for the LMP by iteratively covering an expanding *witness set* of finitely many points in P . In each iteration, computing a lower bound involves solving a special case

of a TSP instance with neighborhoods, the *Close-Enough TSP* (CETSP) to provable optimality; for an upper bound, this is enhanced to provide full coverage. In each iteration, this establishes both a valid solution and a valid lower bound, and thereby a bound on the remaining optimality gap. (See Figure 1 for an illustration.)

- We prove that this discretization method leads to provably good results: both the uncovered area and the maximum distance of points from the region swept by our lower-bound tours converge to zero as we enhance the witness set.
- We present a comprehensive study to demonstrate the practical usefulness of our methods, based on a wide spectrum of benchmark instances with up to 2000 vertices. The outcomes include provably optimal solutions, limited optimality gaps, and improved lower bounds.

1.2 Related work There is a wide range of practical applications for the LMP, including manufacturing [9, 35, 36], cleaning [17], robotic coverage [19, 20, 33, 38], inspection [27], CAD [26], farming [11, 21, 45] and pest control [13]. In Computational Geometry, the Lawn Mowing Problem was first introduced by Arkin et al. [6], who later gave the currently best approximation algorithm with a performance guarantee of $2\sqrt{3}\alpha_{\text{TSP}} \approx 3.46\alpha_{\text{TSP}}$ [7], where α_{TSP} is the performance guarantee for an approximation algorithm for the TSP; while α_{TSP} may be $(1 + \varepsilon)$ in theory, based on the methods of Arora [10] or Mitchell [42], neither approach is practically useful. A variant (in which cost was incurred both for traveling and covering) was considered by Fekete et al. [30], who gave a 4-approximation for

the special case of polyominoes. Closely related is the TSP with neighborhoods (TSPN), for which it suffices to visit a neighborhood for each of a given set of discrete vertices; this was first considered by Arkin and Hassin [8], in a graph setting by Gendreau et al. [34], for “fat” neighborhoods by Mitchell [43] and heuristically by Yuan and Zhang [48]. A particularly relevant special case of the TSPN is the *Close-Enough TSP* (CETSP), in which it suffices to get within a Euclidean distance of r , i.e., for which the neighborhood is an r -disk. Dumitrescu and Tóth [24] gave an $O(1)$ approximation; they also provide a broad overview of other theoretical results for the TSPN. Practical methods were considered by Mennell [41], Behdani and Smith [14], and Coutinho et al. [22].

The Art Gallery Problem (AGP) is connected to our work by a combination of theoretical and practical issues. Also a problem of optimal geometric covering, the AGP has to deal with the theoretical difficulties of possibly irrational coordinates, as shown by Abrahamson et al. [1] (answering an open problem by Fekete [3]); subsequently, this result served as a stepping stone towards a proof of $\exists\mathbb{R}$ -completeness [2]. On the practical side, powerful methods (e.g., by Baumgartner et al. [12], Kröller et al. [39], or de Rezende et al. [23]) for computing a good set of guards for a polygonal region P are based on finding solutions for discrete *witness sets* within P , leading to a primal-dual optimization method; see [18] for an animated multimedia description. This approach is closely related to our primal-dual method for the LMP.

Optimally covering even relatively simple regions by a set of n unit disks has received a considerable attention, but is excruciatingly difficult. For covering rectangles by n unit disks, Heppes and Mellissen [37] gave optimal solutions for $n \leq 5$; Melissen and Schuur [40] extended this for $n = 6, 7$. See the website by Friedman [32] for illustrations of the best known solutions (only some of which are proven to be optimal) for $n \leq 12$. As early as 1915, Neville [44] computed the optimal arrangement for covering a disk by five unit disks, but reported a wrong optimal value; much later, Bezdek [15, 16] gave the correct value for $n = 5, 6$. As recently as 2005, Fejes Tóth [47] established optimal values for $n = 8, 9, 10$. The question of incomplete coverings was raised in 2008 by Connelly; Szalkai [46] gave an optimal solution for $n = 3$. Progress on covering by (not necessarily equal) disks has been achieved by Fekete et al. [28, 29].

1.3 Preliminaries A (simple) *polygon* P is a (non-self-intersecting) shape in the plane, bounded by a finite number n of line segments. The *boundary* of a

polygon P is denoted by ∂P . A *tour* is a closed continuous curve $T : [0, 1] \rightarrow \mathbb{R}^2$ with $T(0) = T(1)$; we denote the (Euclidean) length of T by $\ell(T)$. The *cutter* C is a disk of radius r , centered in its midpoint. The *Minkowski sum* of two sets $A, B \subset \mathbb{R}^2$ is the set $A \oplus B = \{a + b \mid a \in A, b \in B\}$. The *coverage* of a tour T with the disk cutter C is $T \oplus C$. The coverage of a point $p \in T$ is $\{p\} \oplus C$. A *lawn mowing tour* T of a polygon P with a cutter C is a tour whose coverage contains P . An *optimal* lawn mowing tour T^* is a lawn mowing tour of shortest length. For a discrete set of points P , a tour T *traverses* P if $P \subset T$, i.e., for each $p \in P$ there is a $t \in [0, 1]$ such that $T(t) = p$.

2 Optimal tours are straight

In general, a lawn mowing tour may consist of line segments and curved arcs. We show that any optimal tour mowing a polygon can only contain line segments.

THEOREM 2.1. *For every polygon P and a circular cutter, an optimal lawn mowing tour T exclusively consists of a set of line segments.*

2.1 Outline of the proof Because our proof of Theorem 2.1 is rather technical, we formulate various lemmata within its context, see Figure 2. By means of these lemmata, the proof is partitioned as follows.

Lemma 2.1. The key lemma for the proof of Theorem 2.1 is that given an optimal tour, the coverage of a curved arc cannot contain a region that is covered more than once.

Lemma 2.2. Based on Lemma 2.1, we show that each point p from the boundary of an arc t has two closest points q, r from the optimal tour.

Lemma 2.3. We show that the two closest points q, r to p lie in a line with p ; we call r the *successor* of q .

Theorem 2.1. Alternately applying Lemmas 2.2 and 2.3 leads to a sequence of successors on a common line, with a final successor r' outside the polygon. This implies that the final successor lies on an arc that covers a region outside of the polygon, leading to a contradiction by Lemma 2.1.

2.2 Preliminaries for the proof A (*curved*) *arc* is the image of a smooth function $t : [0, 1] \rightarrow \mathbb{R}^2$ that is either strongly convex or strongly concave, see Figure 3. A *segment* is the image of a linear function $t : [0, 1] \rightarrow \mathbb{R}^2$. A *component* t is an arc or a segment. We call $t(0)$ and $t(1)$ the *start* and *end point* of t . Any tour $T : [0, 1] \rightarrow \mathbb{R}^2$ can be partitioned into a sequence of

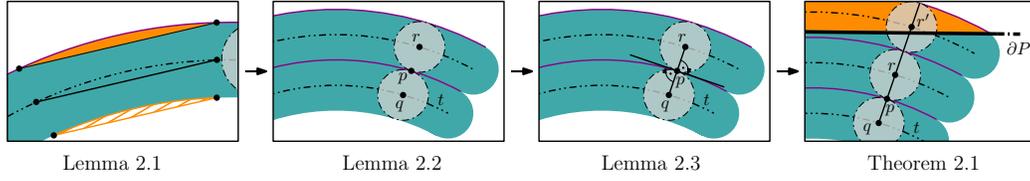


Figure 2: A symbolic overview of our proof that optimal tours only contain line segments.

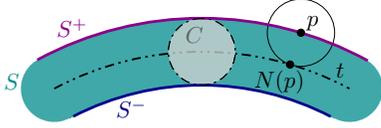


Figure 3: The coverage strip $S(t)$ (or S in short) that results from traversing t with the midpoint of the circular cutter C . S^+ and S^- denote the convex, and concave side, respectively. p is a point in the interior of S^+ , and $N(p)$ is its closest point on t .

components t_0, \dots, t_{k-1} , such that the end point of t_i equals the start point of $t_{i+1 \bmod k}$ for $i \in \{0, \dots, k-1\}$.

For any component t , the *coverage strip* $S(t)$ is the region derived as the Minkowski sum $C \oplus t$. The boundary $\partial S(t)$ can be partitioned into four curves as follows, see Figure 3: For each point p on the boundary, we consider its closest point $N(p)$ on t regarding the Euclidean metric. The closure of all points p on the boundary such that $N(p)$ is neither the start nor the endpoint of t is the union of two way-connected components. In particular, when following the boundary of $S(t)$ in clockwise orientation, the “convex” component (denoted by $S^+(t)$) has positive curvature, while the “concave” component (denoted by $S^-(t)$) has negative curvature; see Figure 3. If it is clear from the context, we denote by $S^+ := S^+(t)$, $S^- := S^-(t)$, and $S := S(t)$.

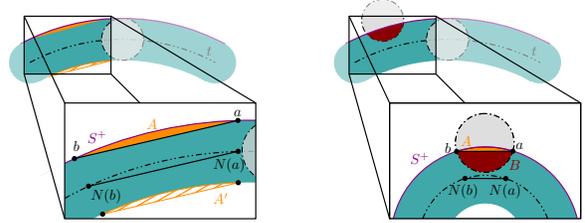
2.3 Details of the proof A simple observation is that any tour can be partitioned into a set of line segments and curved arcs. To show that the optimality of a tour excludes the existence of a curved arc, we consider for the sake of contradiction an optimal lawn mowing tour $T := \{t_0, \dots, t_{k-1}\}$ for a circular cutter C with at least one curved arc $t := t_i$. Note that all our arguments are independent from the cutter’s radius.

The key lemma for the proof is that “shortcutting” curved arcs by line segments while maintaining the coverage of the entire tour reduces its length.

LEMMA 2.1. *Let ab be a segment inside of the coverage strip S of a curved arc t , with its end points a, b on the convex side S^+ of S . Let A be the region bounded by ab*

and S^+ . If A does not have to be covered by t (because it is already covered otherwise), then T is not optimal.

Proof. Let $N(a), N(b) \in t$ be the closest points to a, b , respectively. Replacing the part of t between $N(a)$ and $N(b)$ by the segment $N(a)N(b)$ results in a shorter tour, see Figure 4(a). This tour still covers the polygon, because the arc-segment-arc sequence $(t_1, N(a)N(b), t_2)$ covers S , except for the region A that by assumption does not have to be covered by the sequence. Thus the tour containing the original arc t was not optimal. \square



(a) A covered region A . (b) Covering a point with the cutter.

Figure 4: (a) A region A that does not have to be covered by the coverage strip of a curved arc t , allowing a shortcut of t . (b) The intersection of the coverage of a point with the coverage strip of a curved arc implying the situation of Figure 4(a).

The following corollary is straightforward.

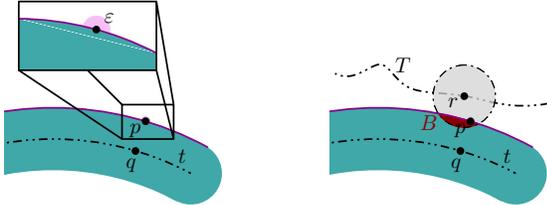
COROLLARY 2.1. *Let r be a point on the tour T , and t be an arc of T . If the coverage of r intersects the convex side S^+ of t , then T is not optimal, see Figure 4(b).*

Note that in the context of Lemma 2.1 shortcutting t by inserting the segment $N(a)N(b)$ is allowed, i.e., still maintains a covering tour, because the region A does not have to be covered by t . Simultaneously, we obtain another region A' that is now additionally covered by the coverage strip of $N(a)N(b)$, see Figure 4(b). Intuitively speaking, A' allows for another application of Lemma 2.1 to another coverage strip, and so on.

For the remaining details, denote by p an inner point of the polygon, which lies also in the *interior* of the convex side S^+ , i.e., a point p whose closest point on t is neither $t(0)$ nor $t(1)$. We show that p needs to have two closest points on T .

LEMMA 2.2. p has at least two closest points on T .

Proof. As p lies in the interior of a convex side, there is a point $N(p)$ on an arc t and within a distance of 1 to p . By assumption, p does not lie on the boundary of the polygon, so there is at least a second point r from which the cutter covers p , as otherwise there is an uncovered region, see Figure 5(a). Assume one of these points has a smaller distance to p , see Figure 5(b). Corollary 2.1 implies that T is not optimal, a contradiction. \square



(a) An uncovered portion. (b) A multi-covered region.

Figure 5: (a) An arbitrarily small area (in pink) around p and above S^+ that is not covered by the tour. (b) If the point r is closer than a distance of 1 to p , then there is a region (in dark red) that is covered multiple times.

In the context of Lemma 2.2, we call r the *successor* of $N(p)$. As t is smooth, S^+ is smooth as well. Hence, the tangent to S^+ in p is well-defined. We say that a segment *lies orthogonal* to S^+ in p when it intersects the tangent to S^+ in p orthogonally.

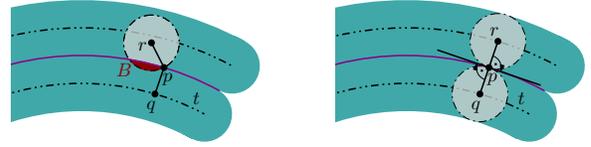
LEMMA 2.3. The segment between q and its successor r lies orthogonal to S^+ in p .

Proof. Cutter boundaries centered at q and r share at least the point p . Assume that qr does not intersect S^+ orthogonally in p . This implies that at least one of the cutters contains another point from the interior of S^+ , because p is an inner point of S^+ , see Figure 6(a). Corollary 2.1 implies that T is not optimal, a contradiction. \square

Combining Lemmas 2.2 and 2.3 yield the following.

COROLLARY 2.2. p has exactly two closest points on T .

Intuitively, the convex side S^+ of t is squeezed between two cutters centered in q and r . Applying Lemmas 2.1, 2.2, and 2.3 yields a proof of Theorem 2.1, which we restate here.



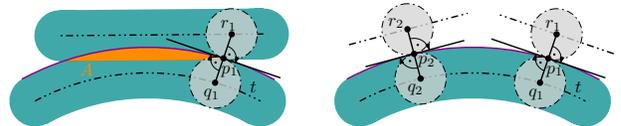
(a) Non-orthogonal intersection. (b) Orthogonal intersection.

Figure 6: (a) The situation when qr does not intersect S^+ orthogonally in p , i.e., a point from the interior of S^+ lies in the interior of a cutter. (b) The segment qr intersects S^+ orthogonally in p .

THEOREM 2.1. For every polygon P and a circular cutter, an optimal lawn mowing tour T exclusively consists of a set of line segments.

Proof. Applying Lemmas 2.2 and 2.3 yields the existence of the successor r of q , such that p , q , and r lie in a line that intersects S^+ of t orthogonally in p .

To ensure that r does not lie on a segment, we slightly perturb p in the interior of S^+ , such that r lies on an arc. In particular, let r_1 and r_2 be two successors resulting from two different points p_1 and p_2 from the interior of S^+ , such that r_1 and r_2 lie on segments s_1 and s_2 , respectively. s_1 lies orthogonal to r_1p_1 ; otherwise Lemma 2.1 implies that the entire tour is not optimal, see Figure 7(a). Analogously, we conclude that s_2 lies orthogonal to r_2p_2 . Hence, s_1 and s_2 are different segments, see Figure 7(b). As the set of segments is a discrete set, there is a point p in the interior of S^+ , such that the resulting successor does not lie on a segment, as S^+ is continuous. An analogous argument implies that there is a point p in the interior of S^+ , such that the successor r is neither a point on a segment nor a start or endpoint of an arc, i.e., r is an inner point of an arc.



(a) Non-parallel segments. (b) Locally parallel segments.

Figure 7: (a) The point r_1 lies on a segment that is not parallel to the tangent in p_1 . (b) The points r_1 and r_2 lie on two distinct segments that are parallel to the tangents in p_1 and p_2 , respectively.

This argument of the existence of a successor from the interior of an arc is repeated, resulting in a sequence $R = \{r, \dots, r'\}$ of successors lying on a line. At some point in this process, a successor r' is reached,

fulfilling exactly one of the following cases: (1) Either the segment qr' crosses the boundary of the polygon (see Figure 8(a)), or (2) the convex side of the arc of r' touches the polygon boundary in a point s (see Figure 8(b)). This means that the line qr' lies orthogonal to the boundary of P . In this situation, we denote the intersection point of qr' with the coverage of r' by s , which is equal to the intersection between qr' and ∂P .

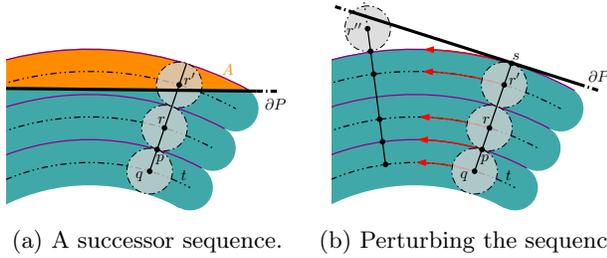


Figure 8: (a) The sequence of constructed successors $R := \{r, \dots, r'\}$ lying in a common ray intersecting the boundary of the polygon P inside the coverage of r' . (b) Perturbing the successors for avoiding a final successor whose coverage lies tangentially to ∂P .

In the first case, we obtain a region A that must not be covered by the tour, hence Lemma 2.1 implies that T cannot be optimal.

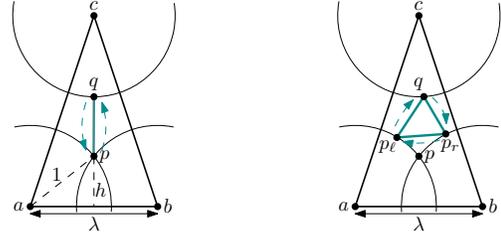
In the second case, we perturb p including the entire constructed sequence R of successors and the intersection point s of the ray pr' with ∂P . This results in the scenario that s is an inner point of P , leading to a continuation of the construction of successors until the first case is reached. This concludes the proof. \square

3 Irrational tour vertices

In the previous section we saw that optimal lawn mowing tours must be polygonal paths. This makes it critical to characterize coordinates of potential tour vertices. As we show in the following, irrational coordinates may be involved in an optimal tour, indicating possible algebraic (and hence algorithmic) difficulties.

THEOREM 3.1. *For any rational number $0 < \lambda \leq 2$, there are simple polygons P with rational vertex coordinates for which an optimal lawn mowing tour contains a vertex with a coordinate of $h := \sqrt{1 - (\lambda/2)^2}$.*

Proof. For $0 < \lambda \leq 2$, let $P = (a, b, c)$ be the polygon with $a = (0, 0)$, $b = (\lambda, 0)$, and $c = (\frac{\lambda}{2}, 2)$, see Figure 9. It is easy to see that any lawn mowing tour must have at least one point that touches the unit circles C_a, C_b , and C_c that are centered at the vertices of P .



(a) An optimal tour. (b) A suboptimal tour.

Figure 9: (a) A polygon $P = (a, b, c)$ with the optimal lawn mowing tour, connecting p and q . (b) A possible lawn mowing tour without using an irrational point.

We show that the point p on the intersection of C_a and C_b within P has to be part of the optimal lawn mowing tour, see Figure 9(a).

Assume that T is an optimal tour, and let q be some point on T with $q \in \partial C_c$. For the sake of contradiction, assume that $p \notin T$. Thus, there exist two points $p_\ell, p_r \in T$, with $p_\ell \in \partial C_a$, and $p_r \in \partial C_b$, see Figure 9(b). As T connects p_ℓ, p_r , and q , the length of the tour is $\ell(T) \geq d(p_\ell, p_r) + d(p_\ell, q) + d(p_r, q)$. Independent of the choice of the points p_ℓ, p_r , and q , we observe that $d(p_\ell, q) + d(p_r, q) > 2d(p, q)$. Moreover, $T' = pqp$ is a feasible tour with length $\ell(T') < \ell(T)$, contradicting the assumption that T is an optimal tour.

It remains to show that p has an irrational coordinate, and that T' is the unique optimal tour. The three points a, b , and p describe an isosceles triangle of base length λ , thus the y -coordinate of p is $\sqrt{1 - (\lambda/2)^2}$. The point q that is closest to p clearly lies on the vertical line $x = \frac{\lambda}{2}$. It is easy to see that the tour that connects the points $p = (\frac{\lambda}{2}, \sqrt{1 - (\lambda/2)^2})$ and $q = (\frac{\lambda}{2}, 1)$ covers P and has length $\ell(T') = 2 \left(1 - \sqrt{1 - (\lambda/2)^2}\right)$. Because p has to be a vertex of the tour, and any other choice of q results in a longer tour, T' is the unique optimal lawn mowing tour for P . \square

As a consequence, bounded accuracy in computing coordinates may become an issue, motivating the concept of ε -robust coverage, for which any point must be covered by a portion of the cutter C that is at least ε -interior to C , corresponding to a CETSP solution with neighborhood size $r - \varepsilon$. Furthermore, the algebraic difficulty motivates the following.

PROBLEM 1. *Is the LMP $\exists\mathbb{R}$ -complete?*

4 A primal-dual approach to Lawn Mowing

Now we describe our primal-dual approach to computing good tours for the LMP, based on (i) iteratively selecting finite subsets W of points from the uncovered portion of P , (ii) covering W by an optimal solution T_W for the Close-Enough TSP for W to obtain a lower bound, (iii) enhancing T_W to a feasible covering tour of P to obtain an upper bound.

In the following, we describe basic mechanisms and analytical implications; further technical details that affect the practical performance are described in the following Section 5.

4.1 Witness sets and Close-Enough TSP At each stage of our algorithm, we use a finite set W of *witnesses*, in P . It is straightforward to see that these induce lower bounds on the length of an optimal lawn mowing tour.

THEOREM 4.1. *For $W_1 \subset W_2 \subset P$, let T_{W_1} , T_{W_2} , and T_P be optimal lawn mowing tours for W_1 , W_2 , and P , respectively. Then $\ell(T_{W_1}) \leq \ell(T_{W_2}) \leq \ell(T_P)$.*

Because all points $w \in W$ are covered by a cutter C of radius r if and only if the center of C visits all disks of radius r centered at w , such a *Close-Enough TSP* tour for W corresponds to a lawn mowing tour for W , and thus a lower bound for the LMP on P . Therefore, a sequence of enhanced witness sets leads to an increasing sequence of lower bounds for $\ell(T_P)$; if some $T(W_i)$ is a feasible LMP tour for P , it is optimal.

Arkin et al. [6, 7] noted that any LMP tour may need a pseudopolynomial (but not polynomially bounded) number of tour vertices even for a square-shaped lawn with diameter d and $d \gg r$. For ε -robust coverage (as defined in Section 3), we can establish a finite upper bound on the number of necessary witnesses that is independent of n , and holds even for general sets $P \subset \mathbb{R}^2$, not necessarily just polygons or connected sets.

THEOREM 4.2. *Consider a region $P \subset \mathbb{R}^2$ of diameter d , a circular cutter C of radius r and some $0 < \varepsilon \leq \frac{r}{2}$. Then there is a witness set W of size bounded by $|W| \leq \left\lceil \frac{2d^2}{\varepsilon^2} \right\rceil \in O(d^2)$, for which any ε -robust CETSP tour is also feasible for the LMP.*

Proof. A bounding box B of P has dimensions at most $d \times d$, so we can subdivide B by an orthogonal mesh of width $\varepsilon/\sqrt{2}$, consisting of not more than $\left\lceil \frac{2d^2}{\varepsilon^2} \right\rceil$ grid cells, C_{ij} , each of diameter at most ε . Whenever $P \cap C_{ij} \neq \emptyset$, picking an arbitrary witness point from $P \cap C_{ij}$ ensures that each such cell (and hence, all of P) is covered by any ε -robust CESTP tour. \square

For general (and not just robust) witness sets, it is easy to see that all extreme points of P must be contained; if P is a convex polygon, this includes all n vertices. Thus, no such upper bound exists in terms of only d and r . For other types of regions P (such as disks), no finite bound of any kind exists, as the whole perimeter needs to be contained in a witness set.

PROBLEM 2. *For a polygon P with n vertices and diameter d , and a circular cutter of radius r , is there a witness set W of size polynomial in n , d , $1/r$, such that an optimal CETSP solution for W is a feasible lawn mowing tour of P ?*

4.2 Lower bounds: convergence As we enhance the witness set, we can ensure that the maximum distance of uncovered points from an ensuing CETSP tour converges to zero. This can be achieved by a relatively dense set of witness points (as in Theorem 4.2) or (more economically) by iteratively reducing this maximum distance when necessary. We get the following straightforward implication.

LEMMA 4.1. *Let $W \subset P$ be a finite set of witness points, such that no point of P has a distance larger than δ from a witness point, and let T_W be a feasible CETSP tour of W . Then*

$$d_{\max} := \max_{p \in P} \min_{q \in T_W \oplus C} d(p, q) \leq \delta.$$

In the remainder of this subsection, we will repeatedly apply the following theorem.

THEOREM 4.3. (FEKETE AND PULLEYBLANK [31]) *Let E be a connected planar arrangement of edges and C be a disk of radius r . If L is the total length of the edges in E , there is a closed roundtrip of length at most $2L + 2\pi r$ that visits all points of the boundary of $E \oplus C$.*

In the sequel, E will be the edges of a tour, leading to the following.

THEOREM 4.4. *Let W be a witness set, and let T_W be an optimal CETSP tour of W . Let the distance between uncovered points and covered region be bounded by the cutter radius r , i.e., $d_{\max} := \max_{p \in P} \min_{q \in T_W \oplus C} d(p, q) \leq r$. Then P has a lawn mowing tour of length at most $3\ell(T_W) + 2\pi r + 2r$.*

Proof. By assumption, any uncovered point $p \in P$ is within distance r from the boundary of $E \oplus C$. Therefore, all points that are not covered by T_W will be covered by a tour of the boundary of $E \oplus C$; by Theorem 4.3, there is such a tour of length at most $2\ell(T_W) + 2\pi r$. Concatenating this tour with T_W (at additional cost at most $2r$) achieves full coverage of P at cost at most $3\ell(T_W) + 2\pi r + 2r$, as claimed. \square

Furthermore, we can conclude the following.

LEMMA 4.2. *Let $W_0 \subset W_1 \subset \dots \subset W_i \subset \dots \subset P$ be a sequence of witness sets, with T_{W_i} being the corresponding sequence of optimal CETSP tours of W_i . If the maximum distance $d_{\max}^{(i)} := \max_{p \in P} \min_{q \in T_{W_i} \oplus C} d(p, q)$ of uncovered points to the covered region converges to zero, then the area of $P \setminus T_{W_i} \oplus C$ converges to zero.*

Proof. By assumption, the sequence $d_{\max}^{(i)}$ converges to zero, so there is an index k such that $d_{\max}^{(i)} \leq r$ for all $i \geq k$. Then by Theorem 4.1, all elements of the sequence $\ell(T_{W_i})$ are lower bounds of any feasible LMP tour, so by Theorem 4.4, they satisfy $\ell(T_{W_i}) \leq 3\ell(T_{W_k}) + 2\pi r + 2r$, i.e., remain bounded from above. Furthermore, all points in the set $P \setminus T_{W_i} \oplus C$ of uncovered points by T_{W_i} are within distance $d_{\max}^{(i)}$ of the boundary of $T_{W_i} \oplus C$; by Theorem 4.3, this boundary has length of at most $L = 2\ell(T_{W_k}) + 2\pi r$. Moreover, the area of all points within a distance $d_{\max}^{(i)}$ of any curve of length L is bounded by $A_i := (2L + 2\pi)d_{\max}^{(i)}$, so the total uncovered area is bounded by $A_i \leq (4\ell(T_{W_k}) + 4\pi r + 2\pi)d_{\max}^{(i)} \leq (12\ell(T_{W_k}) + 12\pi r + 8r + 2\pi)d_{\max}^{(i)}$. Because $(12\ell(T_{W_k}) + 12\pi r + 8r + 2\pi)$ is constant and $d_{\max}^{(i)}$ by assumption converges to zero, we conclude that A_i also converges to zero, implying the claim. \square

We summarize.

THEOREM 4.5. *By picking an appropriate sequence of witness sets W_i , we can guarantee that for the ensuing sequence of optimal CETSP solutions $T(W_i)$, the following holds.*

- *The maximum distance of uncovered points to the covered region converges to zero.*
- *The total uncovered area converges to zero.*

4.3 Solving CETSP instances Solving CETSP instances is more challenging than solving TSP instances. In 2016, Coutinho et al. [22] proposed an algorithm based on branch-and-bound and *Second-Order Cone Programming*. Each branch-and-bound node is associated with a partial tour that visits the given subset of vertices in particular order. At the root node, the algorithm chooses three vertices to generate an initial sequence. The problem of computing the exact coordinates from a predefined sequence can be formulated as a Second-Order Cone Problem (SOCP). If the solution at the root node is feasible, the solution is optimal and the algorithm terminates. Otherwise, the algorithm branches into three subproblems, one for every edge in

the current solution. At each edge a currently uncovered vertex is inserted. A node is pruned if the cost is at least equal to the best known upper bound or if its associated solution is feasible.

$$(4.1) \quad \text{minimize} \quad \sum_{k=1}^q z_k$$

subject to

$$(4.2) \quad x_{i_{k-1}} - x_{i_k} = w_k \quad k = 0, \dots, q$$

$$(4.3) \quad y_{i_{k-1}} - y_{i_k} = u_k \quad k = 0, \dots, q$$

$$(4.4) \quad \bar{x}_k - x_k = s_k \quad k = 0, \dots, q$$

$$(4.5) \quad \bar{y}_k - y_k = t_k \quad k = 0, \dots, q$$

$$(4.6) \quad w_k^2 + u_k^2 \leq z_k^2 \quad k = 0, \dots, q$$

$$(4.7) \quad s_k^2 + t_k^2 \leq r_k^2 \quad k = 0, \dots, q$$

$$z_k \geq 0; x_k, y_k, w_k, u_k, s_k, t_k \text{ free} \quad k = 0, \dots, q$$

The formulation by Coutinho et al. [22] is based on work of Mennell [41]. Let $S = \{i_0, \dots, i_q\}$, $q < n$ be a partial sequence of vertices during the execution of the branch-and-bound algorithm. The desired solution must provide exact values for the coordinates (x_{i_k}, y_{i_k}) , $k = 0, \dots, q$ such that the length of the partial tour is minimized. Denoting $i_{-1} = i_q$, the formulation is as follows.

The SOCP uses variables z_k to represent the distance between subsequent vertices i_{k-1} and i_k in the given sequence S . To compute the distance, four auxiliary variables w_k, u_k, s_k , and t_k are introduced in Constraints (4.2)–(4.5) that represent differences of coordinates that are used to calculate Euclidean distances. Constraints (4.6) define the length of the edge connecting subsequent vertices. Constraints (4.7) ensure that the hitting points will lie within their respective circles around v_i . Andersen et al. [5] showed that a SOCP can be solved in polynomial time. The branch-and-bound approach of Coutinho et al. [22] makes use of well-known optimization software that is also capable of addressing this class of problems.

4.4 Upper bounds: obtaining feasible solutions

Given is a polygon P , a circular cutter C and a maximum number of iterations. We start by initializing the witness set W_0 with some points $p \in P$.

In each iteration i , we start with generating a CETSP *start tour* T_i from W_i ; see Figure 15 for families of examples. If T_i is a feasible LMP tour, it is optimal; otherwise, we can extract the *uncovered regions* from $P \setminus (T_i \oplus C)$.

As shown in Figure 15 for each of the lower bounds, we then execute the following procedure to modify T_i until the tour becomes feasible. For every uncovered region R_j , we construct a witness set W_{R_j} and fix the closest point $p_j \in \partial T_i$ to R_j to be visited by a tour T_i computed by the CETSP solver.

After the tour T_i is computed, we add W_{R_j} to obtain W_{i+1} , so that the next start tour covers the points in W_{R_j} from the beginning. (See Section 5.1 and again Figure 15 for practical strategies for choosing W_{R_j} .) Subsequently, we extend T_i with T_{R_j} , by connecting both tours via the point p_j . Once we obtain a feasible tour T_i (such as in the examples of Figure 15), we can update the best known solution so far and continue with the next iteration for lower and upper bounds, see Algorithm A.1.

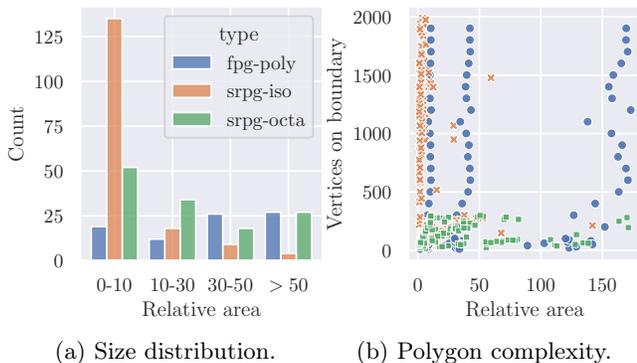


Figure 10: Distribution of instances, subdivided into different types. See Figure 17 for an enlarged version of the small instances in (b).

5 Experiments and evaluations

Now we describe our practical implementation and various algorithm engineering aspects, along with our experimental study for demonstrating the practical usefulness of our theoretical concept. All experiments were carried out on a regular desktop workstation with an AMD Ryzen 7 5800X (8×3.8 GHz) CPU and 128 GB of RAM. The code and data are publicly available¹. Consider Figure 15 for practical illustrations of the progression of lower bounds for different witness strategies, and for the corresponding upper bounds. For carrying out our experiments and evaluations, we used the *FPG*, *SRPG ISO* and *SRPF OCTA* instances from the Salzburg Database of Geometric Inputs [25]. See Figure 10 for the overall distribution, and Figure 16 for a sample of individual instances. We considered polygons

¹Source code and data: <https://github.com/tubs-alg/near-optimal-lawn-mowing-tours>

with up to $n = 2000$ vertices, in combination with cutters of varying size. Overall, this resulted in several hundred instances. As it turned out, the parameter *relative area*, i.e., the ratio between the areas of the convex hull of P and the cutter C , was more significant for the difficulty of an instance than the vertex number n , which is why we mostly focused on the variance in relative area.

5.1 Generating witness sets As we established in the previous section, enhancing the witness set leads to convergence of the resulting lower bound tours, at least in terms of covered area and distance to uncovered points. In a practical setting, solving the involved CETSP instances becomes a bottleneck, so it becomes crucial to achieve good lower bounds with witness sets of limited size.

For the initial witness set W_0 , we focus on the exterior of the polygon, as minimizing the length in further iterations will automatically pull the CETSP tour towards the interior; so we choose up to 15 witnesses either from the convex hull, or the boundary of P . If there are more than 15 candidate points, we apply a greedy max-min-dispersion to achieve a well-distributed witness set. This may already lead to optimal solutions for small instances.

For further iterations, we insert 10 more witnesses into regions that are not yet covered. For multiple uncovered regions, we do a random assignment, according to region size. Within each uncovered region R_j , we select the assigned number of new witnesses either (i) randomly, (ii) from a grid, or (iii) from the straight skeleton [4]. To avoid strongly clustered sets (which are bound to occur for the straight skeleton), we employ a dispersion technique. To prevent new witnesses from being close to the boundary but deep inside of the uncovered regions, we cluster the candidate points using k -means, and select a random witness from each cluster.

The two strategies for choosing the initial witness set and the three strategies for extending it result in six general strategies, some of which are shown in the first three columns of Figure 15.

5.2 Practical quality of lower bounds In each iteration, our algorithm computes an optimal CETSP tour on the current witness set; this results in an increasing sequence of lower bounds (due to Theorem 4.1), which are asymptotically tight, at least in terms of covered area and distance to uncovered points. Thus, computing exact CETSP solutions becomes critical. While we could still solve clustered instances with up to 1000 points to optimality, they do not necessarily yield good lower bounds for the LMP; on the other hand, well-distributed witness sets much beyond 50 points

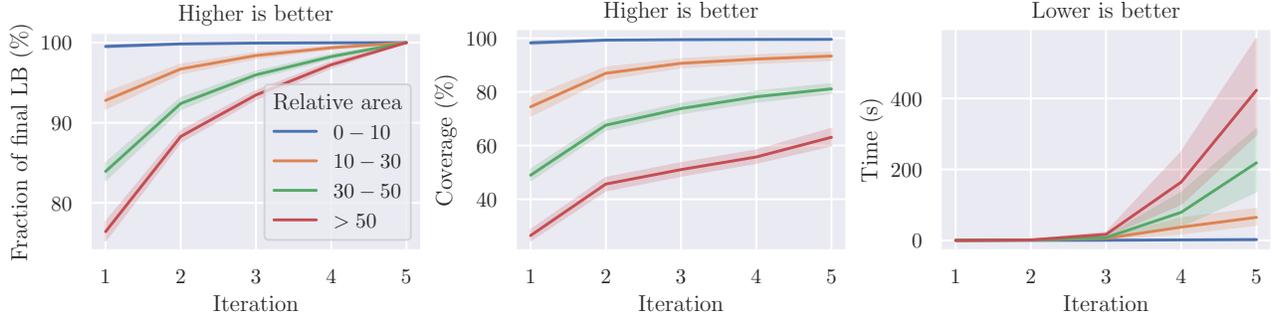


Figure 11: Improvement of lower bound and coverage over iterations, where each iteration extends the witness set. The lower bound is measured in relation to the final lower bound. All three plots share the same legend. For small instances, the first lower bound is already very close to the final upper bound, while for medium instances, it is still around 93% of the final lower bound. The larger the instances, the stronger the increase with each iteration. We get a similar result for the percentage of covered area.

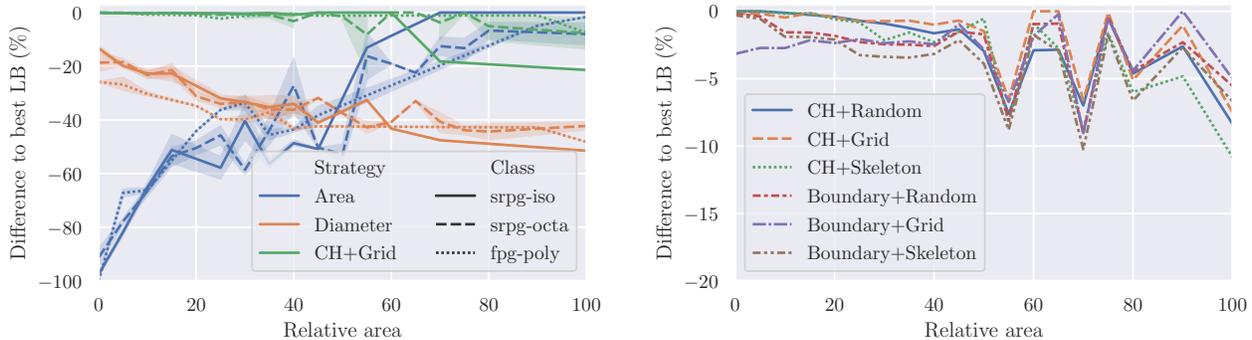


Figure 12: Comparison of lower bounds based on area or diameter with the lower bounds provided by CETSP for different witness strategies (left). Up to a relative area of 50, the simple lower bounds are more than 20% smaller than the CETSP on the convex hull and random witnesses for all classes of instances. Between the different witness strategies (right), the deviation is relatively small. Just the strategies that start with a subset of the convex hull have a small advantage.

were harder to solve optimally, but still yielded better bounds, illustrating the delicate balance between the structure of witness sets, CETSP solvability, and quality of lower bounds.

In the following, we use the best lower bound obtained within 5 iterations. The calls to the CETSP solver are limited to 30 min, after which the best lower bound is returned. Figure 11 shows the progress of lower bound and coverage. With each iteration, the CETSP gets harder to solve, and the improvements on lower bound and coverage decrease. As shown in Figure 12 (left), the resulting lower bound after the final iteration is notably better than the lower bound derived from the diameter or area of the polygon. This is true for all three classes of instances, up to a relative area of about 50, after which the CETSP strategy gets bogged down by too many witnesses. The lower bounds of the individual

witness strategies are surprisingly close, as shown in Figure 12 (right). Starting with the convex hull gives a small advantage; otherwise, the best strategies vary strongly for the instances. Example runs of some strategies over 5 iterations are shown in Figure 15. The best way to improve these (and other) lower bounds would be to add further iterations, at the expense of runtime.

5.3 Practical quality of upper bounds For small instances, the CETSP procedure frequently returns feasible (and hence, optimal) LMP tours. Otherwise, we perform up to 5 iterations in which we add 40 witnesses into each uncovered area and connect the tour to the previous tour, as described in Section 4.4. For medium-sized instances, this strategy often leads to feasible solutions that are significantly better than

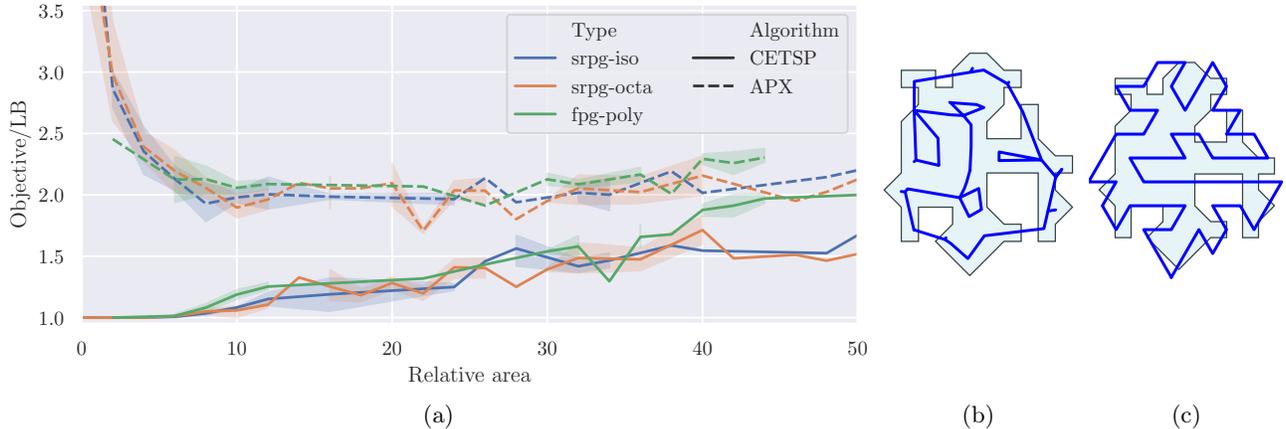


Figure 13: **(a)** Comparison of the solution quality of CETSP approach (plus heuristic) and approximation algorithm (APX). We compare the best solution of any of the witness strategies. Our approach often yields optimal solutions for instances up to a relative area of 6, see Figure 14 for examples. **(b)** An upper bound from our approach with a value of 65.35. **(c)** An upper bound from the approximation of Arkin et al. [7] with a value of 88.33.

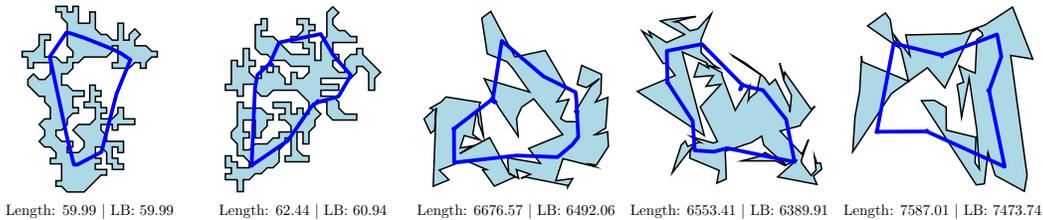


Figure 14: Examples of instances solved to (near) optimality ($\leq 3\%$ above the lower bound).

the approximation algorithm; see Figures 13(a)–13(c) for an example. For instances up to a relative area of around 6, we often obtain optimal solutions, or solutions with a negligibly small gap. For instances up to a relative area of around 30, we are at most 50% above the lower bound, and only slightly worse for 40. For larger instance sizes, Theorem 4.4 still bounds the relative gap to 200%, but the CETSP computation becomes a bottleneck. For instances up to a relative area of 40, the approximation algorithm turns out to be relatively stable at 100% above the lower bound, which is considerable better than the classical worst-case approximation guarantee by Arkin et al. [7], which is at $2\sqrt{3} + \epsilon \approx 3.46$ even when disregarding some finer technical aspects: That method is based on combining a proximity approximation for considering region boundary and interior (incurring an *asymptotic* factor of 3), a grid conversion (incurring a factor of $2/\sqrt{3}$), and a TSP approximation (incurring a theoretical factor of $1 + \epsilon$ when using a geometric PTAS, and a larger factor when using a more practical approximation method); this results in an overall factor

no better than $2\sqrt{3} + \epsilon \approx 3.46$, and even more when accounting for a non-asymptotic approximation factor with an additive term (see [7]) and PTAS limitations. For our experiments, however, the TSP instances in the approximation algorithm were solved to optimality.

6 Conclusion

The Lawn Mowing Problem generalizes a number of notoriously difficult geometric optimization problems, making it both important and highly challenging. We have provided progress on several aspects of the LMP, both on the theoretical and the practical side, demonstrating (for the first time ever) that practically useful progress may be achievable.

A spectrum of open problems remain. On the theoretical side, we already stated Problem 1 and Problem 2. On the practical side, two critical questions are the development of additional approaches for speeding up the computation of lower bounds, as well as more efficient ways to compute upper bounds, along with additional local improvement heuristics. In addition, the concept of ϵ -robustness deserves further exploration.

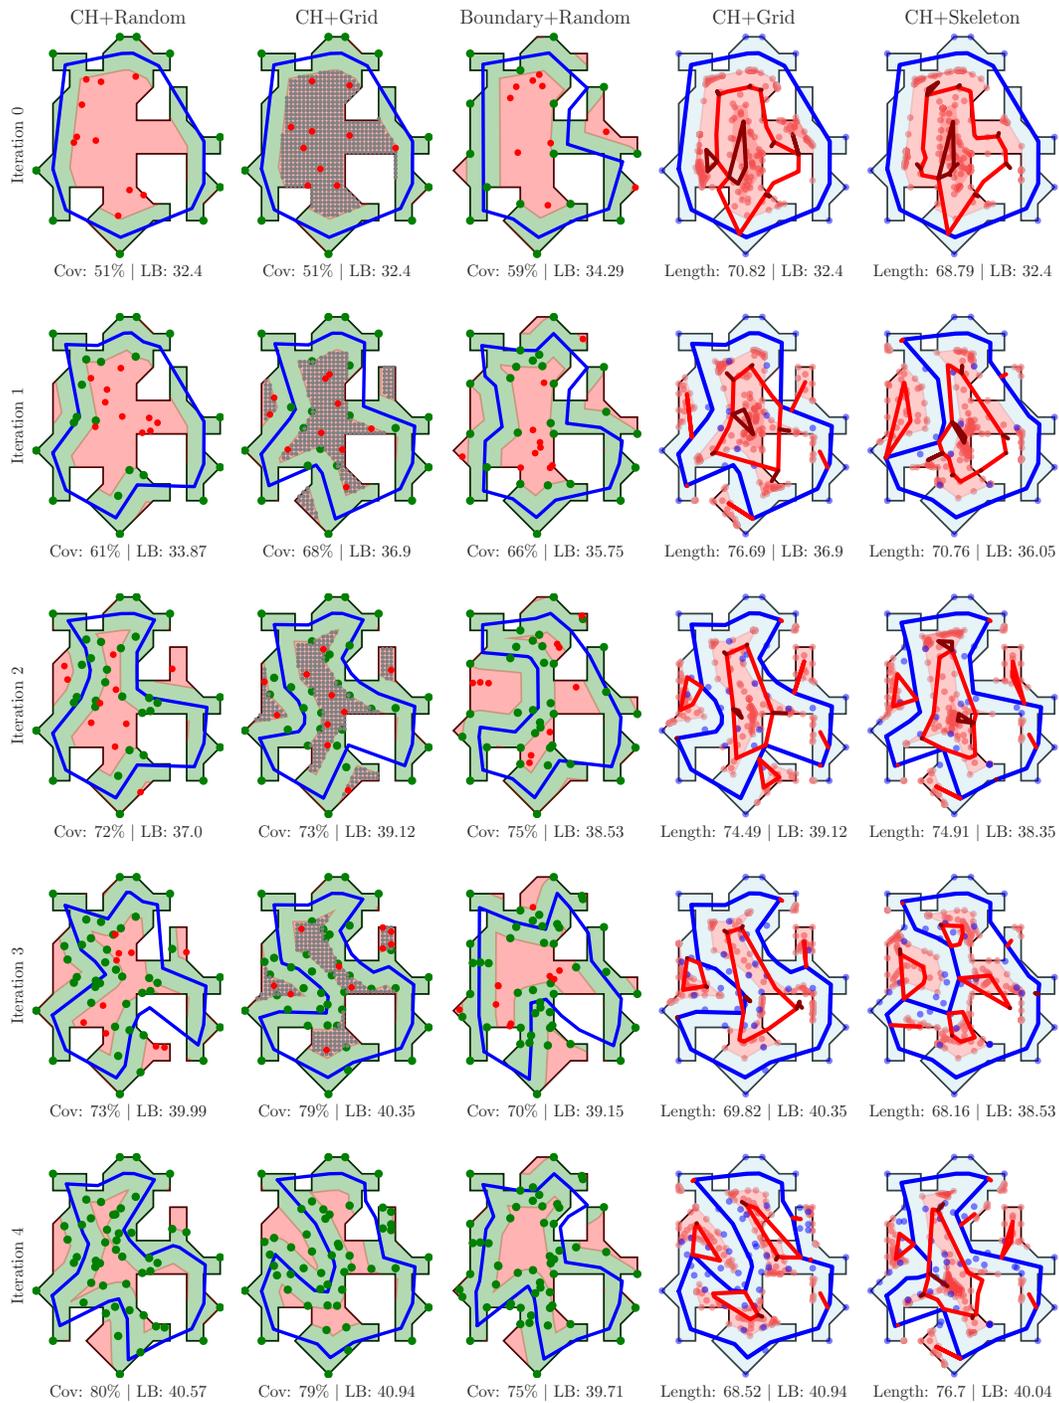


Figure 15: Examples of different witness placement strategies. The examples show just the first 5 iterations of the lower bound on the left and of the upper bound on the right. The trajectory is highlighted in blue, the used witnesses in green. The lower bounds show the covered area in green (lower bound) resp. blue (upper bound) and missed area in red. The light red tours for the upper bounds show the tour extensions after the first iteration, and the dark red tours later extensions. The new witnesses for the next iteration are highlighted in red, while the corresponding candidate set is visualized by gray points. The best lower bound achieved in these 5 iterations is 40.94 and the best upper bound 68.16.

References

- [1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. Irrational guards are sometimes needed. In *Symposium on Computational Geometry (SoCG)*, pages 3:1–3:15, 2017. doi:10.4230/LIPIcs.SoCG.2017.3.
- [2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$ -complete. *Journal of the ACM*, 69(1):4:1–4:70, 2022. doi:10.1145/3486220.
- [3] Pankaj K. Agarwal, Kurt Mehlhorn, and Monique Teillaud. Computational Geometry (Dagstuhl Seminar 11111). *Dagstuhl Reports*, 1(3), 2011. doi:10.4230/DagRep.1.3.19.
- [4] Oswin Aichholzer, Franz Aurenhammer, David Albers, and Bernd Gärtner. A novel type of skeleton for polygons. In *J. UCS The Journal of Universal Computer Science*, pages 752–761. Springer, 1996.
- [5] Erling D. Andersen, Cornelis Roos, and Tamás Terlaky. On implementing a primal-dual interior-point method for conic quadratic optimization. *Mathematical Programming*, 95(2):249–277, 2003. doi:10.1007/s10107-002-0349-3.
- [6] Esther M. Arkin, Sándor P. Fekete, and Joseph S. B. Mitchell. The lawnmower problem. In *Canadian Conference on Computational Geometry (CCCG)*, pages 461–466, 1993.
- [7] Esther M. Arkin, Sándor P. Fekete, and Joseph S. B. Mitchell. Approximation algorithms for lawn mowing and milling. *Computational Geometry*, 17:25–50, 2000. doi:10.1016/S0925-7721(00)00015-8.
- [8] Esther M. Arkin and Refael Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55(3):197–218, 1994. doi:10.1016/0166-218X(94)90008-6.
- [9] Esther M. Arkin, Martin Held, and Christopher L. Smith. Optimization problems related to zigzag pocket machining. *Algorithmica*, 26(2):197–236, 2000. doi:10.1007/s004539910010.
- [10] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998. doi:10.1145/290179.290180.
- [11] Rik Bähnemann, Nicholas Lawrance, Jen Jen Chung, Michael Pantic, Roland Siegwart, and Juan Nieto. Revisiting Boustrophedon coverage path planning as a generalized traveling salesman problem. In *Field and Service Robotics*, pages 277–290, 2021. doi:10.1007/978-981-15-9460-1_20.
- [12] Tobias Baumgartner, Sándor P. Fekete, Alexander Kröller, and Christiane Schmidt. Exact solutions and bounds for general Art Gallery Problems. In *Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 11–22, 2010. doi:10.1137/1.9781611972900.2.
- [13] Aaron T. Becker, Mustapha Debboun, Sándor P. Fekete, Dominik Krupke, and An Nguyen. Zapping Zika with a mosquito-managing drone: Computing optimal flight patterns with minimum turn cost. In *Symposium on Computational Geometry (SoCG)*, pages 62:1–62:5, 2017. doi:10.4230/LIPIcs.SoCG.2017.62.
- [14] Behnam Behdani and J. Cole Smith. An integer-programming-based approach to the close-enough Traveling Salesman Problem. *INFORMS Journal on Computing*, 26(3):415–432, 2014. doi:10.1287/ijoc.2013.0574.
- [15] Károly Bezdek. *Körök optimális fedései (Optimal Covering of Circles)*. PhD thesis, Eötvös Loránd University, 1979.
- [16] Károly Bezdek. Über einige optimale Konfigurationen von Kreisen. *Ann. Univ. Sci. Budapest Rolando Eötvös Sect. Math.*, 27:143–151, 1984.
- [17] Richard Bormann, Joshua Hampp, and Martin Hägele. New brooms sweep clean—An autonomous robotic cleaning assistant for professional office cleaning. In *International Conference on Robotics and Automation (ICRA)*, pages 4470–4477, 2015. doi:10.1109/ICRA.2015.7139818.
- [18] Dorit Borrmann, Pedro J. de Rezende, Cid C. de Souza, Sándor P. Fekete, Stephan Friedrichs, Alexander Kröller, Andreas Nüchter, Christiane Schmidt, and Davi C. Tozoni. Point guards and point clouds: Solving general Art Gallery Problems. In *Symposium on Computational Geometry (SoCG)*, pages 347–348, 2013. doi:10.1145/2462356.2462361.
- [19] Tauã M. Cabreira, Lisane B. Brisolara, and Paulo R. Ferreira Jr. Survey on coverage path planning with unmanned aerial vehicles. *Drones*, 3(1):4, 2019. doi:10.3390/drones3010004.
- [20] Howie Choset. Coverage for robotics—a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31(1):113–126, 2001. doi:10.1023/A:1016639210559.
- [21] Howie Choset and Philippe Pignon. Coverage path planning: The Boustrophedon cellular decomposition. In *Field and Service Robotics*, pages 203–209, 1998. doi:10.1007/978-1-4471-1273-0_32.
- [22] Walton P. Coutinho, Roberto Quirino do Nascimento, Artur A. Pessoa, and Anand Subramanian. A branch-and-bound algorithm for the Close-Enough Traveling Salesman Problem. *INFORMS Journal on Computing*, 28(4):752–765, 2016. doi:10.1287/ijoc.2016.0711.
- [23] Pedro J. de Rezende, Cid C. de Souza, Stephan Friedrichs, Michael Hemmer, Alexander Kröller, and Davi C. Tozoni. Engineering art galleries. In *Algorithm Engineering - Selected Results and Surveys*, pages 379–417, 2016. doi:10.1007/978-3-319-49487-6_12.
- [24] Adrian Dumitrescu and Csaba D. Tóth. Constant-factor approximation for TSP with disks. In *A Journey Through Discrete Mathematics*, pages 375–390. Springer, 2017. doi:10.1007/978-3-319-44479-6_15.
- [25] Günther Eder, Martin Held, Steinþór Jasonarson, Philipp Mayer, and Peter Palfrader. Salzburg database of polygonal data: Polygons and their generators. *Data in Brief*, 31:105984, 2020. doi:10.1016/j.dib.2020.

- 105984.
- [26] Gershon Elber and Myung-Soo Kim. Offsets, sweeps and Minkowski sums. *Computer-Aided Design*, 31(3), 1999. doi:10.1016/S0010-4485(99)00012-3.
- [27] Brendan Englot and Franz Hover. Sampling-based coverage path planning for inspection of complex structures. In *International Conference on Automated Planning and Scheduling (ICAPS)*, volume 22, pages 29–37. AAAI, 2012. URL: <http://www.aaai.org/ocs/index.php/ICAPS/ICAPS12/paper/view/4728>.
- [28] Sándor P. Fekete, Utkarsh Gupta, Phillip Keldenich, Christian Scheffer, and Sahil Shah. Worst-case optimal covering of rectangles by disks. In *Symposium on Computational Geometry (SoCG)*, pages 42:1–42:23, 2020. doi:10.4230/LIPIcs.SocG.2020.42.
- [29] Sándor P. Fekete, Phillip Keldenich, and Christian Scheffer. Covering rectangles by disks: The video. In *Symposium on Computational Geometry (SoCG)*, pages 71:1–75:5, 2020. doi:10.4230/LIPIcs.SocG.2020.75.
- [30] Sándor P. Fekete, Joseph S. B. Mitchell, and Christiane Schmidt. Minimum covering with travel cost. *Journal of Combinatorial Optimization*, 20:32–51, 2012. doi:10.1007/s10878-010-9303-0.
- [31] Sándor P. Fekete and William R. Pulleyblank. Traveling the boundary of Minkowski sums. *Inf. Process. Lett.*, 66(4):171–174, 1998. doi:10.1016/S0020-0190(98)00053-2.
- [32] Erich Friedman. Circles covering squares web page. <https://erich-friedman.github.io/packing/circovsqu>, 2014. Online, accessed December 02, 2021.
- [33] Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12):1258–1276, 2013. doi:10.1016/j.robot.2013.09.004.
- [34] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. The covering tour problem. *Operations Research*, 45(4):568–576, 1997. doi:10.1287/opre.45.4.568.
- [35] Martin Held. *On the Computational Geometry of Pocket Machining*, volume 500 of *LNCS*. Springer, 1991. doi:10.1007/3-540-54103-9.
- [36] Martin Held, Gábor Lukács, and László Andor. Pocket machining based on contour-parallel tool paths generated by means of proximity maps. *Computer-Aided Design*, 26(3):189–203, 1994. doi:10.1016/0010-4485(94)90042-6.
- [37] Aladár Heppes and Hans Melissen. Covering a rectangle with equal circles. *Periodica Mathematica Hungarica*, 34(1-2):65–81, 1997. doi:10.1023/A:1004224507766.
- [38] Katharin R. Jensen-Nau, Tucker Hermans, and Kam K. Leang. Near-optimal area-coverage path planning of energy-constrained aerial robots with application in autonomous environmental monitoring. *IEEE Transactions on Automation Science and Engineering*, 18(3):1453–1468, 2021. doi:10.1109/TASE.2020.3016276.
- [39] Alexander Kröller, Tobias Baumgartner, Sándor P. Fekete, and Christiane Schmidt. Exact solutions and bounds for general Art Gallery Problems. *Journal of Experimental Algorithmics*, 17:2.1–2.23, 2012. doi:10.1145/2133803.2184449.
- [40] Johannes B. M. Melissen and Peter C. Schuur. Covering a rectangle with six and seven circles. *Discrete Applied Mathematics*, 99(1-3):149–156, 2000. doi:10.1016/S0166-218X(99)00130-4.
- [41] William K. Menell. *Heuristics for Solving Three Routing Problems: Close-Enough Traveling Salesman Problem, Close-Enough Vehicle Routing Problem, Sequence-Dependent Team Orienteering Problem*. PhD thesis, University of Maryland, 2009. URL: <http://hdl.handle.net/1903/9822>.
- [42] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric tsp, kmst, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999. doi:10.1137/S0097539796309764.
- [43] Joseph S. B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *Symposium on Discrete Algorithms (SODA)*, pages 11–18, 2007. URL: <https://dl.acm.org/doi/10.5555/1283383.1283385>.
- [44] Eric H. Neville. On the solution of numerical functional equations. *Proceedings of the London Mathematical Society*, 2(1):308–326, 1915. doi:10.1112/plms/s2_14.1.308.
- [45] Timo Oksanen and Arto Visala. Coverage path planning algorithms for agricultural field machines. *Journal of Field Robotics*, 26(8):651–668, 2009. doi:10.1002/rob.20300.
- [46] Balázs Szalkai. Optimal cover of a disk with three smaller congruent disks. *Advances in Geometry*, 16(4):465–476, 2016. doi:10.1515/advgeom-2016-0018.
- [47] Gábor F. Tóth. Thinnest covering of a circle by eight, nine, or ten congruent circles. *Combinatorial and Computational Geometry*, 52:361–376, 2005.
- [48] Bo Yuan and Tiantian Zhang. Towards solving TSPN with arbitrary neighborhoods: A hybrid solution. In *Australasian Conference on Artificial Life and Computational Intelligence*, pages 204–215, 2017. doi:10.1007/978-3-319-51691-2_18.

A Pseudocode for the algorithm

Algorithm A.1 Pseudocode for the primal-dual algorithm.

Input: A polygon P , a cutter C with radius r , max iterations max_i

Output: A feasible tour T and a lower bound b_l

$b_l = 0$

$T_{best} = \emptyset$

Initialize W_0 with some points

for $i \leftarrow 0, 1, \dots, max_i - 1$ **do**

if $T_{best} = \emptyset \vee \ell(T_{best}) > b_l$ **then**

 Solve CETSP for centers W_i and $r_k = r$ ($0 \leq k < |W_i|$) to get T_i and a lower bound b_l^{CETSP}

$b_l = \max\{b_l^{CETSP}, b_l\}$

$W_{i+1} = W_i$

do

$uncoveredRegions = P \setminus (T_i \oplus C)$

for all $R_j \in uncoveredRegions$ **do**

 Initialize W_{R_j} from R_j

$p = \arg \min_{p \in \partial T_i} d(p, R_j)$

 Solve CETSP for centers $W_{R_j} \cup \{p\}$ and $r_k = r, r_p = 0$ ($0 \leq k < |W_{R_j}|$) to get

T_{R_j}

$W_{i+1} = W_{i+1} \cup W_{R_j}$

 Connect T_i and T_{R_j} via p

end for

while $uncoveredRegions \neq \emptyset$

if $T_{best} = \emptyset$ **then**

$T_{best} = T_i$

else

$T_{best} = \arg \min_{T \in \{T_i, T_{best}\}} \ell(T)$

end if

end if

end for

return T_{best}, b_l

B Illustrations of instances and solution progress

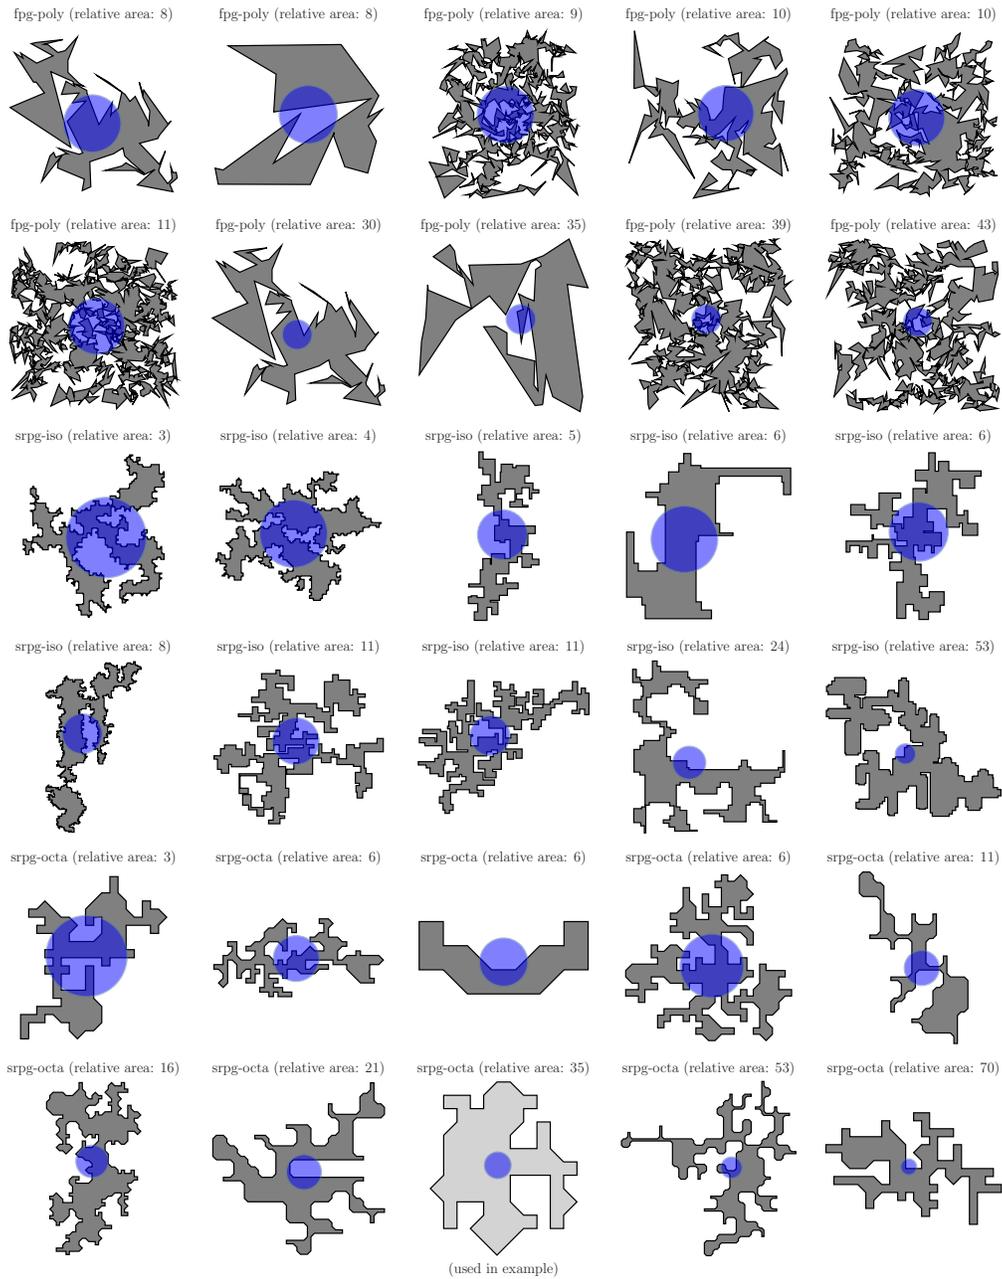


Figure 16: Examples of the used polygons. The blue circles show the respective tool size; the *relative area* denotes the ratio of convex hull area and cutter area.

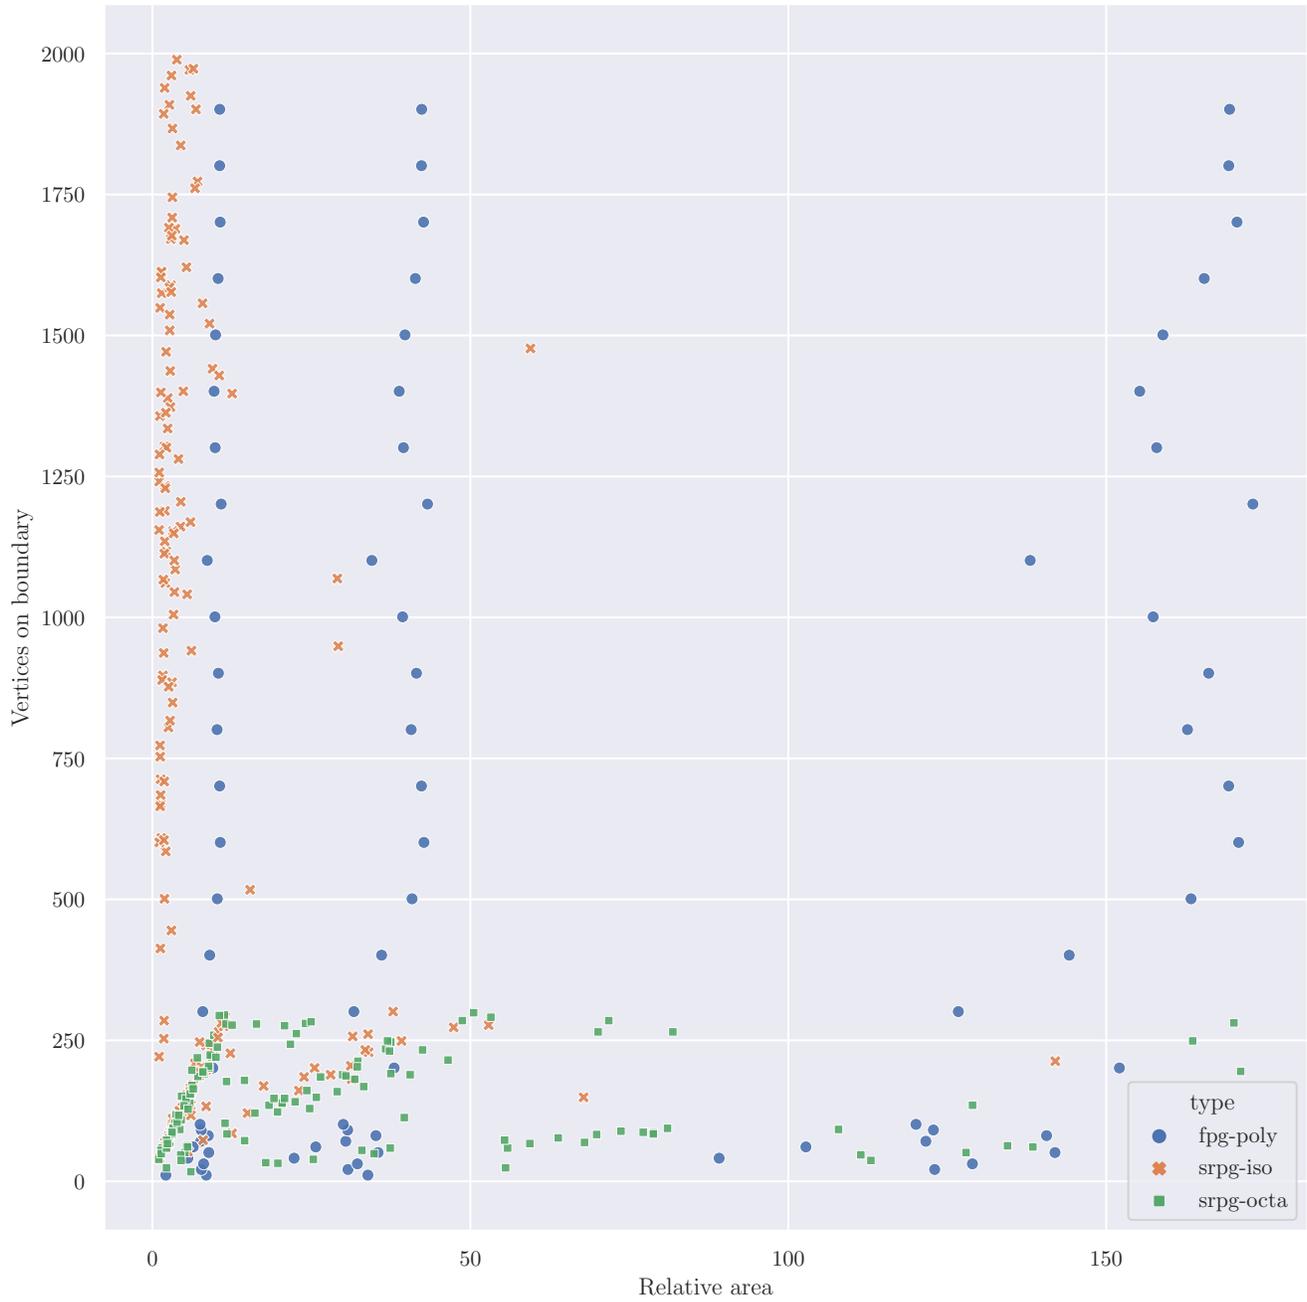


Figure 17: Enlarged distribution of instances, subdivided into different types.