An Approximation Algorithm for the Art Gallery Problem*†

Édouard Bonnet[1] and Tillmann Miltzow[2]

1 **Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI), Budapest, Hungary**
`edouard.bonnet@lamsade.dauphine.fr`

2 **Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI), Budapest, Hungary**
`t.miltzow@gmail.com`

Abstract

Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be visible to each other if the line segment between x and y is contained in $\mathcal{P}$. The Point Guard Art Gallery problem asks for a minimum-size set S such that every point in $\mathcal{P}$ is visible from a point in S. The set S is referred to as guards. Assuming integer coordinates and a specific general position on the vertices of $\mathcal{P}$, we present the first $O(\log \mathrm{OPT})$-approximation algorithm for the point guard problem. This algorithm combines ideas in papers of Efrat and Har-Peled [18] and Deshpande et al. [15, 16]. We also point out a mistake in the latter.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases computational geometry, art-gallery, approximation algorithm

Digital Object Identifier 10.4230/LIPIcs.SoCG.2017.20

1 Introduction

Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be visible to each other if the line segment between x and y is contained in $\mathcal{P}$. The point-guard art gallery problem asks for a minimum-size set S such that every point in $\mathcal{P}$ is visible from a point in S. The set S is referred to as guards.

Victor Klee introduced the art gallery problem to Václav Chvátal, who showed that $\lfloor n/3 \rfloor$ guards are always sufficient and sometimes necessary for a polygon with n vertices [11]. In 1978, Steve Fisk gave an elegant proof of the same result [21]. This constitutes the first combinatorial result related to the art gallery problem.

Related problems. A large amount of research is committed to the study of combinatorial and algorithmic aspects of the art gallery problem, as reflected by the following surveys [33, 34, 31]. This research is focused on the art gallery problem and its many variants, based on different definitions of visibility, restricted classes of polygons, different shapes and positions of guards, etc. The most natural definition of visibility is arguably the one we gave above. Other possible definitions are: x sees y if the axis-parallel rectangle spanned by x and y is contained in $\mathcal{P}$; x sees y if the line segment joining x to y intersects $\mathcal{P}$ at most c times,

* A full version of the paper is available at `http://arxiv.org/abs/1607.05527`.

† supported by the ERC grant PARAMTIGHT: "Parameterized complexity and the search for tight complexity results", no. 280152.

33rd International Symposium on Computational Geometry (SoCG 2017).
Editors: Boris Aronov and Matthew J. Katz; Article No. 20; pp. 20:1–20:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

for some value of c; x sees y if there exists a straight-line path from x to y within $\mathcal{P}$ with at most c bends. Common shapes of polygons comprise: simple polygons, polygons with holes, simple orthogonal polygons, x-monotone polygons and star-shaped polygons. Common placements of guards include: vertex guards and point guards as defined above, but also edge-guard (guards are edges of the polygon), segment guards (guards are interior segments of the polygon) and perimeter guards (guards must be placed on the boundary of $\mathcal{P}$).

On the algorithmic side, very few variants are known to be solvable in polynomial time [30, 17] and most results are on approximating the minimum number of guards [15, 16, 23, 26, 27, 18, 28]. Many of the approximation algorithms are based on the fact that the range space defined by the visibility regions has bounded VC-dimension for simple polygons [24], combined with some algorithmic ideas of Clarkson [13, 8].

On the negative side, Eidenbenz et al. [19] showed NP-hardness and inapproximability for the most principal variants. In particular, they show that getting a PTAS for those variants is very unlikely, even on simple polygons. For polygons with holes, they even show that there is no $o(\log n)$-approximation algorithm, unless P=NP. Their reduction from Set Cover also implies that the art gallery problem is W[2]-hard on polygons with holes and that there is no $n^{o(k)}$ algorithm, to determine if k guards are sufficient, under the Exponential Time Hypothesis (ETH) [19, Sec.4]. Recently, the authors of the present paper show a similar result for simple polygons (i.e., without holes) [7].

Point Guard Art Gallery Problem. Notwithstanding the large amount of research on the art gallery problem, there is only one exact algorithmic result on the point guard variant. The result is not so well-known and attributed to Micha Sharir [18]: one can find in time $n^{O(k)}$ a set of k guards for the point guard variant, if it exists. This result is quite easy to achieve with some tools from real algebraic geometry [3] and seemingly hopeless to prove without this powerful machinery (see [4] for the very restricted case $k = 2$). Although the algorithm utilizes remarkably sophisticated tools, it uses almost no problem-specific insights and no better exact algorithm is known. Moreover, we recall that the papers [19, 7] suggest that there is no significantly better exact algorithm even for simple polygons.

Regarding *approximation* algorithms for the point guard variant, the results are similarly sparse. For general polygons, Deshpande et al. gave a randomized pseudo-polynomial time $O(\log n)$-approximation algorithm [15, 16]. However, we will see that their algorithm is not correct. Efrat and Har-Peled gave a randomized polynomial time $O(\log |OPT_{\text{grid}}|)$-approximation algorithm by restricting guards to a very fine grid [18]. However, they can not prove that their grid solution is indeed an approximation of an optimal guard placement. In this paper, we develop the ideas of Deshpande et al. in combination with the algorithm of Efrat and Har-Peled. Here, OPT denotes an optimal set of guards and OPT_{grid} an optimal set of guards that is restricted to some grid. Finally, let us mention that there exist approximation algorithms for monotone and rectilinear polygons [28], when the very restrictive structure of the polygon is exploited.

Lack of progress and motivation. Note that the art gallery problem can be seen as a geometric hitting set problem. In a hitting set problem, we are given a universe U and a set of subsets $S \subseteq 2^U$ and we are asked to find a smallest set $X \subseteq U$ such that $\forall s \in S\ \exists x \in X : x \in s$. Usually the set system is given explicitly or can be at least easily restricted to a set of polynomial size. In our case, the universe is the entire polygon (not just the boundary) and the set system is the set of visibility regions (given a point $x \in \mathcal{P}$, the visibility region $\text{Vis}(x)$ is defined as the set of points visible from x). The lack of progress has come from the obvious yet crucial fact that the set system is *infinite* and that no one has found a way to restrict the universe to a finite set (see [12, 1] for some attempts). We also wish to quote a recent remark by Bhattiprolu and Har-Peled [5] both confirming that the point guard is the most principal variant and highlighting the challenge of finding an approximation algorithm: *"One of the more interesting versions of the geometric hitting set problem, is the art gallery problem, where one is given a simple polygon in the plane, and one has to select a set of points (inside or on the boundary of the polygon) that "see" the whole polygon. While much research has gone into variants of this problem [31], nothing is known as far as an approximation algorithm (for the general problem). The difficulty arises from the underlying set system being infinite, see [18] for some efforts in better understanding this problem."*

Besides theoretical considerations, there is a series of work to find efficient implementations to solve the art gallery problem in practice; see [14] for a survey on this large body of work. There as well the focus lies on the point guard variant. One of the key challenges is to find a discretization of the solution space, as was pointed out recently [22]: *"...a finite discretization whose existence in the AGP [(Art Gallery Problem)] is, to the best of our knowledge, still unknown and poses a key challenge w.r.t. software solving the AGP."* Although we cannot answer this question with respect to exact computation, we show that a fine enough grid is a sufficient discretization of the solution space with respect to constant-factor approximation; see Lemma 4. We also highlight certain fundamental problems related to solution-space discretization.

Our contribution. Recently Elbassioni showed how the framework of Brönnimann and Goodrich [8] can be extended to infinite range spaces, if one allows that some small δ-fraction of the ground set is not covered [20]. The main application of his paper is to yield an approximation algorithm for a variant of the point guard art gallery problem when one is allowed to guard only *almost all* the polygon. We show here how to achieve the same asymptotic approximation factor, while guarding the whole polygon. However, we rely on two assumptions on the gallery, which we detail below.

▶ **Assumption 1** (Integer Vertex Representation)**.** *Vertices are given by integers, represented in binary.*

An *extension* of a polygon $\mathcal{P}$ is a line that goes through two vertices of $\mathcal{P}$.

▶ **Assumption 2** (General Position Assumption)**.** *No three extensions meet in a point of $\mathcal{P}$ which is* not *a vertex and no three vertices are collinear.*

Note that we allow that three (or more) extensions meet in a vertex or outside the polygon.

No three points lie on a line is a typical assumption in computational geometry and discrete geometry. Often this assumption is a pure technicality. In some cases, however, the result might in fact be wrong without this assumption. In our case, we do believe that Lemma 4 could be proven without Assumption 2, but it seems that some new ideas would be needed. See [2] for an example where the main result is that some general position assumption can be weakened. The idea of general position assumptions is that a small random perturbation of the point set yields the assumption with probability almost 1. In case that the points are given by integers small random perturbations, destroy the integer property. But random perturbations could be performed in the following way: first multiply all coordinates by some large constant $2^C \in \mathbb{N}$ and then add a random integer x with $-C \leqslant x \leqslant C$.

The integer representation assumption (Assumption 1) seems to be very strong as it gives us useful distance bounds not just between any two different vertices of the polygon, but also between any two objects that do not share a point (see Lemma 8). On the other hand, real computers work with binary numbers and cannot compute real numbers with arbitrary precision. The real-RAM model was introduced as a convenient theoretical framework to simplify the analysis of algorithms with numerical and/or geometrical flavors, see for instance [25, page 1]. Also note that Assumption 1 can be replaced by assuming that all coordinates are represented by rational numbers with specified nominator and denominator. (There could be other potentially more compact ways to specify rational numbers.) Multiplying all numbers with the smallest common multiple of the denominators takes polynomial time, makes all numbers integers and does not change the geometry of the problem.

▶ **Theorem 3.** *Under Assumptions 1 and 2, there is a randomized approximation algorithm with approximation factor $O(\log |OPT|)$ for* Point Guard Art Gallery *for simple polygons. The running time is polynomial in the size of the input.*

The main technical idea is to show the following lemma:

▶ **Lemma 4** (Global Visibility Containment)**.** *Let $\mathcal{P}$ be some (not necessarily simple) polygon. Under Assumptions 1 and 2, it holds that there exists a grid Γ and a guard set $S_{grid} \subseteq \Gamma$, which sees the entire polygon and $|S_{grid}| = O(|S|)$, where S is an optimal guard set.*

To be a bit more precise, let M be the largest appearing integer. Then the number of points in Γ is polynomial in M. This is potentially exponential in the size of the input. Thus algorithms that rely on storing all points of Γ explicitly do not have polynomial worst case running time. The algorithm of Efrat and Har-Peled [18] does *not* store every point of Γ explicitly and, with the lemma above, the algorithm gives an $O(\log |OPT|)$-approximation on the grid Γ.

While Lemma 4 tells us that we can restrict our attention to a finite grid, when considering constant factor approximation, the same is not known for exact computation. In particular, it is not known whether the Point Guard problem lies in NP. Recently, some researchers popularized an interesting complexity class, called $\exists\mathbb{R}$, being somewhere between NP and PSPACE [10, 32, 9, 29]. Many geometric problems, for which membership in NP is uncertain, have been shown to be complete for this class. This suggests that there might be indeed no polynomial sized witness for these problems as this would imply $NP = \exists\mathbb{R}$. The history of the art gallery problem suggests the possibility that the Point Guard problem is $\exists\mathbb{R}$-complete. If $NP \neq \exists\mathbb{R}$, then this would imply that there is indeed no hope to find a witness of polynomial size for the Point Guard problem.

Given a polygon $\mathcal{P}$, we will always assume that all coordinates of its vertices are given by *positive* integers in binary. (This can be achieved in polynomial time.) We denote by M the largest appearing integer and we denote by $\mathrm{diam}(\mathcal{P})$ the largest distance between any two points in $\mathcal{P}$. Note that $\mathrm{diam}(\mathcal{P}) < 2M$. We denote $L = 20M > 10$. Note that $\log L$ is linear in the input size. We define the grid

$$\Gamma = (L^{-12} \cdot \mathbb{Z}^2) \cap \mathcal{P}.$$

In other words, we scale the integer grid by L^{-12} and take all points of the grid within the polygon $\mathcal{P}$. Note that all vertices of $\mathcal{P}$ have integer coordinates and thus are included in Γ.

▶ **Theorem 5** (Efrat, Har-Peled [18])**.** *Given a simple polygon $\mathcal{P}$ with n vertices, one can spread a grid Γ inside $\mathcal{P}$, and compute an $O(\log OPT_{\mathrm{grid}})$-approximation for the smallest* *subset of* Γ *that sees* $\mathcal{P}$*. The expected running time of the algorithm is*

$$O(n\,OPT_{\text{grid}}^2 \log OPT_{grid} \log(n\,OPT_{\text{grid}}) \log^2 \Delta),$$

where Δ *is the ratio between the diameter of the polygon and the grid size.*

The term OPT_{grid} refers to the optimum, when restricted to the grid Γ. For the solution S that is output by the algorithm of Efrat and Har-Peled the following holds $|S| = O(|OPT_{\text{grid}}| \log |OPT_{\text{grid}}|)$. However, Efrat and Har-Peled make no claim on the relation between $|S|$ and the actual optimum $|OPT|$. Note that the grid size equals $w = L^{-12}$, thus $\Delta \leqslant L^{12+1} = L^{13}$ and consequently $\log \Delta \leqslant 13 \log L$, which is polynomial in the size of the input.

Efrat and Har-Peled implicitly use the real-RAM as model of computation: elementary computations are expected to take $O(1)$ time and coordinates of points are given by real numbers. As we assume that coordinates are given by integers, the word-RAM or integer-RAM is a more appropriate model of computation. All we need to know about this model is that we can upper bound the time for elementary computations by a polynomial in the bit length of the involved numbers. Thus, going from the real-RAM to the word-RAM only adds a polynomial factor in the running time of the algorithm of Efrat and Har-Peled. Therefore, from the discussion above we see that it is sufficient to prove Lemma 4.

Organization. In Section 2, we describe the counterexample to the algorithm of Deshpande et al. [16]. This proves useful as a starting point of Section 3 in which we show Lemma 4. Due to space constraints, the detailed proofs of the lemmas can only be found in the full version [6]. Finally in Section 4, we indicate some remaining open questions.

2 Counterexample

In this section, we point out a mistake in the algorithm of Deshpande et al. [15, 16]. This mistake though constitutes an interesting starting point for our purpose.

The algorithm by Deshpande et al. can be described from a high level perspective as follows: maintain and refine a triangulation T of the polygon until every triangle $\Delta \in T$ satisfies the so-called *local visibility containment* property. The local visibility containment property of Δ certifies that every point $x \in \Delta$ can only see points that are also seen by the vertices of Δ. However, we will argue that it is impossible to attain the local visibility containment property with any finite triangulation; hence, the algorithm never stops.

Actually, we will show two lemmas, which describe fundamental issues with such an approach. Let $D \subseteq \mathcal{P}$ be a finite set of points in the polygon and $x \in \mathcal{P}$, then we denote by $D_x = \{\, d \in D : \text{dist}(d, x) \leqslant 1 \,\}$.

▶ **Lemma 6.** *There is a polygon* $\mathcal{P}$ *such that for any finite set* D*, there exists a point* x *such that* x *sees a point* p *that is not visible from* D_x*.*

Thus in case that each triangle in the triangulation by Deshpande et al. has diameter smaller than 1 Lemma 6 shows that the promised local visibility property cannot hold. We imagine that all vertices of the triangulation T form the set D. It was claimed that for each point x the triangle Δ containing x sees whatever x sees. Now, Lemma 6 says that even the larger set D_x cannot see everything that is seen by x. Thus in particular the triangle Δ cannot see everything seen by x.

The triangles of Deshpande et al. might be very large and thus not contained in D_x. The next Lemma addresses the issue of large triangles.

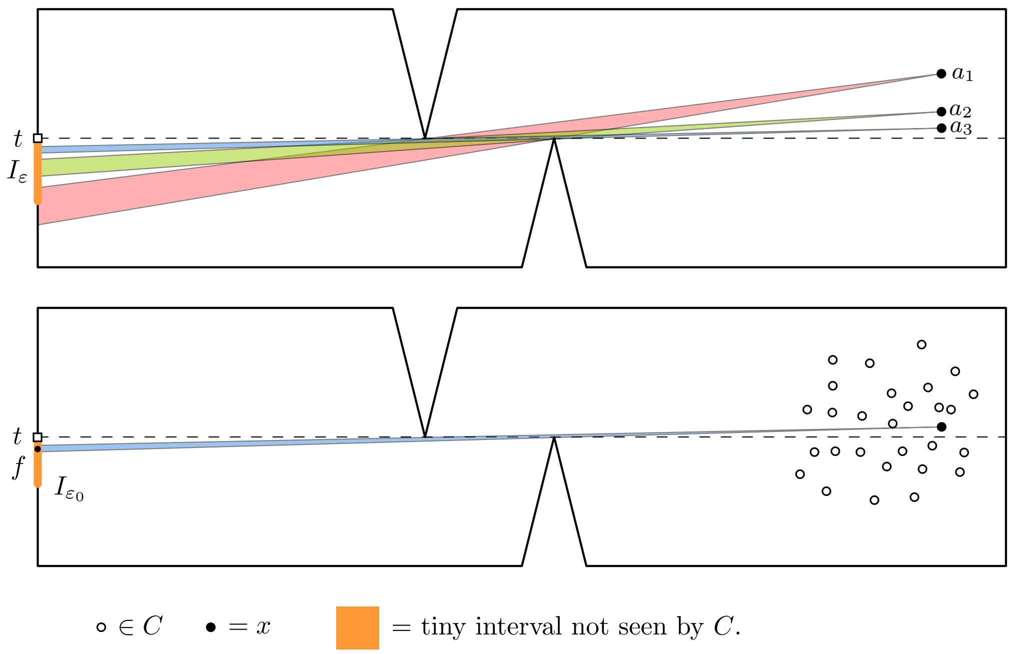

Figure 1 Illustration of the polygon with the property described in Lemma 6.

▶ **Lemma 7.** *Let $c \in \mathbb{N}$ be any constant. There exists a polygon $\mathcal{P}_c$ such that for any finite set of points D, there exists a point $x \in \mathcal{P}_c$ such that any subset $S \subseteq D$ of size $c-1$ cannot see the entire visibility region of x.*

Note that the point x depends on the set D. If we invoke Lemma 7 with $c = 4$ it refutes the algorithm of Deshpande et al. for good as follows. Consider the polygon $\mathcal{P}_4$ as in Lemma 7. For the purpose of contradiction suppose that there exists a triangulation T with the local visibility containment property. We denote by D the set of vertices of T. According to Lemma 7, there exists a point x such that any three points of D cannot see everything that x sees. In particular, the three vertices of the triangle Δ containing x cannot see everything that is seen by x. But T is supposed to have exactly this property — a contradiction.

Again, we want to mention that the paper of Deshpande et al. has ideas that helped to achieve the result of the present paper. In particular, we will show that the local visibility containment property does indeed hold most of the time.

Proof of Lemma 6. See Figure 1, for the definition of polygon $\mathcal{P}$ and the following description. We have two opposite reflex vertices with supporting line ℓ. The sequence of points $(a_i)_{i\in\mathbb{N}}$ are chosen closer and closer to ℓ on the right side of the polygon above ℓ. None of the a_i's can see t, as this would require to be actually on ℓ. Furthermore we denote by I_ε the *open* interval of length ε below t with endpoint t. The interval is indicated in orange in Figure 1. It is clear that for each $\varepsilon > 0$, there exists an i such that a_i sees at least part of I_ε. Let D be any finite set of points. Consider now any finite collection of points $C \subseteq D$ with distance at most 2 to the limit of the $(a_i)_{i\in\mathbb{N}}$. As we will choose x as one of the a_i's it holds $D_x \subseteq C$. For each point $p \in C$ exists an $\varepsilon(p)$ such that p sees nothing of the open interval $I_{\varepsilon(p)}$. Let $\varepsilon_0 = \min_{p\in C} \varepsilon(p)$. None of the points of C see anything of the open interval I_{ε_0}. Recall that the visibility of the a_i's come arbitrarily close to t. Thus, there is some a_k that sees a point f on interval I_{ε_0}. We define x to be a_k. Recall that the set D_x is contained in C. We conclude no point of D_x sees the point f, which is seen by x, as claimed. ◀

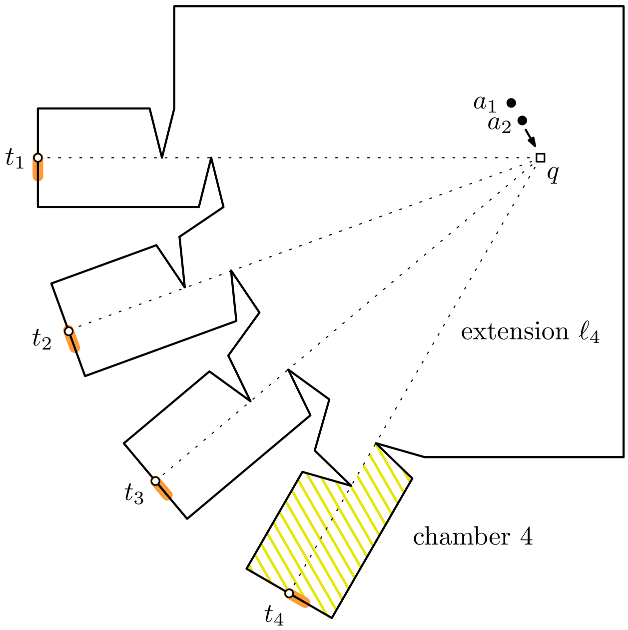

Figure 2 Illustration of the polygon with the property described in Lemma 7.

Proof of Lemmma 7. See Figure 2, for the following description of the polygon $\mathcal{P}_c$. We build $\mathcal{P}_c$ from c disjoint chambers with an entrance of opposite reflex vertices. The chambers are arranged in a way that all the extensions of the opposite reflex vertices meet in a common point q. In this way, we get c extensions $\ell_1, \ldots, \ell_c$. We denote by t_i the intersection of the extension ℓ_i with the i-th chamber. An important nuance in the construction is the fact that one can see into all the chambers *simultaneously* from points b arbitrarily close to q.

Again we construct a sequence $(a_i)_{i \in \mathbb{N}}$ such that it works as in the proof of Lemma 6, but for all chambers simultaneously. For this let a_i be any point with $\text{dist}(a_i, q) = 1/i$ and the property that it sees into each chamber, for all $i \in \mathbb{N}$. As in the proof of Lemma 6, it holds that each a_i sees a small interval close to t_j, for all $i \in N$ and $j \in \{1, \ldots, c\}$. In particular each such interval *approaches* t_j, for all $j = \{1, \ldots, c\}$.

Let $\varepsilon > 0$. We denote the *open* interval of length ε to the right of t_j on the boundary of $\mathcal{P}_c$ with $I_{\varepsilon,j}$ (indicated orange in the figure). No point $b \in \mathcal{P}$ can see two intervals $I_{\varepsilon,j}$ and $I_{\varepsilon,j'}$ *entirely*, for any $\varepsilon > 0$ and $j \neq j'$. Because to see the *whole* interval $I_{\varepsilon,j}$ requires to be in chamber j. However, no point can be in two chambers simultaneously. To avoid confusion, we want to point out that no a_i can see any interval $I_{\varepsilon,j}$ entirely. But for every $\varepsilon > 0$, there exists an i_0 such that a_{i_0} sees part of *all* $I_{\varepsilon,j}$ simultaneously.

Let $D \subseteq \mathcal{P}$ be any finite set. Then there exists some ε_0 such that any point $p \in D$ sees at most one entire interval $I_{\varepsilon_0,j}$ and nothing of any other interval $I_{\varepsilon_0,j'}$ for any $j' \neq j$. To see this consider first the case that p is contained in one of the chambers. Then the statement is clear as it cannot see any point of any other chamber. In the other case p is outside of any chamber and thus cannot see any interval $I_{\varepsilon,j}$ entirely for any j. Thus in this case there exists some $\varepsilon(p) > 0$ such that p sees nothing of $I_{\varepsilon,j}$ for any j and $0 < \varepsilon < \varepsilon(p)$. We choose $\varepsilon_0 = \min_{p \in D} \varepsilon(p)$.

By the definition of $(a_i)_{i \in \mathbb{N}}$ there exists some $x = a_k$ that sees at least one point $f_j \in I_{\varepsilon_0,j}$, for all $j = \{1, \ldots, c\}$. As no point of D sees two f_j simultaneously, we need at least c points of D to see $f_1, \ldots, f_c$. ◀

3 Detailed exposition of the proof

The details, which we skipped here due to space constraints, can be found in [6]. Nonetheless, all the ideas and crucial facts are highlighted in this conference version, and illustrated with a number of figures. In almost all figures, the proportions of some objects and distances may not reflect the reality of things. They are displayed this way to convey a message, albeit exaggeratedly.

Our high-level proof idea is that the local visibility containment property holds for every point x that is sufficiently far away from all extension lines. (We slightly tune the meaning of the local visibility containment property.) This constitutes the first step. (Recall that the extension of two vertices is the line that contains these vertices.) In a second step, we will show that a point in the gallery cannot be close to more than two extensions at the same time. We will add one vertex for each extension that x is close to. Recall that the vertices of the polygon are also in Γ.

The first step is much more tedious than the second one. A reason for that is that many observations that seem true at first sight turn out to be erroneous. Therefore, some extra care is needed for this step in the definitions of the concepts and in breaking a general situation to a distinction of more elementary cases. The crux is mainly to identify these elementary cases and to properly handle them. All the other proofs are elementary.

3.1 Benefit of Integer Coordinates

The integer coordinate assumption not only implies that the distance between any two vertices is at least 1 but it also gives useful lower bounds on distances between any two objects of interest that do not share a point. The next lemma lists all such lower bounds that we will need later. We denote by $\mathrm{dist}(u, v), \mathrm{dist}(u, \ell)$ and $\mathrm{dist}(\ell, \ell')$ the Euclidean distance between the points u and v, the point u and the line ℓ, and the lines ℓ and ℓ', respectively.

▶ **Lemma 8.** *Let $\mathcal{P}$ be a polygon with integer coordinates and L as defined above. Let v and w be vertices of $\mathcal{P}$, ℓ and ℓ' supporting lines of two vertices, and p and q intersections of supporting lines.*

1. *$dist(v, w) > 0 \Rightarrow dist(v, w) \geqslant 1$.*
2. *$dist(v, \ell) > 0 \Rightarrow dist(v, \ell) \geqslant L^{-1}$.*
3. *$dist(p, \ell) > 0 \Rightarrow dist(p, \ell) \geqslant L^{-5}$.*
4. *$dist(p, q) > 0 \Rightarrow dist(p, q) \geqslant L^{-4}$.*
5. *Let $\ell \neq \ell'$ be parallel. Then $dist(\ell, \ell') \geqslant L^{-1}$.*
6. *Let $\ell \neq \ell'$ be any two non-parallel supporting lines and α the smaller angle between them. Then holds $\tan(\alpha) \geqslant 8L^{-2}$.*
7. *Let $a \in \mathcal{P}$ be a point and ℓ_1 and ℓ_2 be some non-parallel lines with $dist(\ell_i, a) < d$, for $i = 1, 2$. Then ℓ_1 and ℓ_2 intersect in a point p with $dist(a, p) \leqslant dL^2$.*

As these bounds are important for the intuition of the forthcoming ideas, we will give an example by proving Item 2.

Proof of Item 2. The distance d can be computed as

$$d = \frac{|(v - w_1) \cdot (w_2 - w_1)^{\perp}|}{\|w_2 - w_1\|_2} \geqslant \frac{1}{\mathrm{diam}(\mathcal{P})} \geqslant \frac{1}{L}.$$

Figure 3 illustrates how to derive this elementary formula. Here, $\cdot$ denotes the scalar product, $x^{\perp}$ is the vector x rotated by $90°$ counter-clockwise, and $\|x\|_2$ is the Euclidean norm of x.

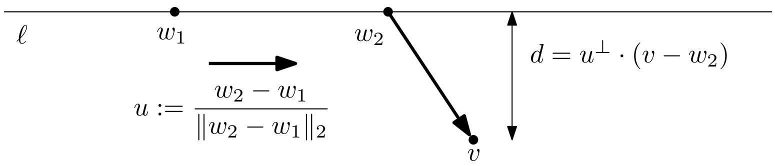

Figure 3 Computing the distance between a line and a vertex.

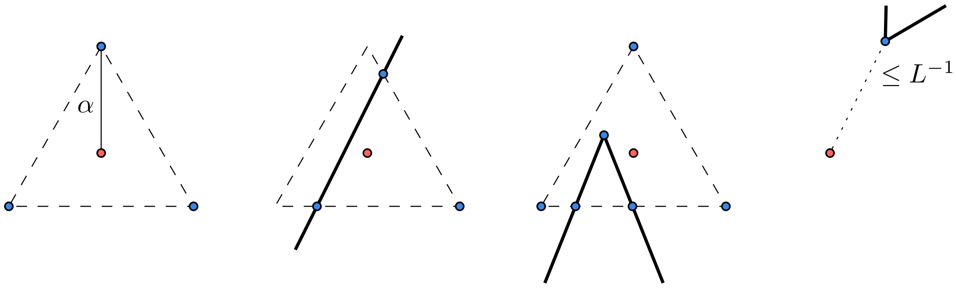

Figure 4 The red point indicates a point of the original optimal solution. The blue points indicate the surrounding grid points that we choose. The polygon is indicated by bold lines. From left to right, we have three cases: the interior case, the boundary case, and the corner case. To the very right, we indicate that in every case the vertices of $\mathcal{P}$ with distance less than L^{-1} are also included in α-grid$^*(x)$.

The numerator of this formula is at least 1 as it is a non-zero integer by assumption. The denominator is upper bounded by the diameter of $\mathcal{P}$, which is in turn upper bounded by L. ◀

3.2 Surrounding Grid Points

Given a point $x \in \mathcal{P}$ and a number α much smaller than the grid width, we will define α-grid(x) as a set of grid points around x, see Figure 4. The parameter α is an upper bound on the distance between x and α-grid(x). (We will chose later $\alpha = L^{-11}$.) In case that there exists a vertex v of $\mathcal{P}$ with distance $\mathrm{dist}(x, v) \leqslant L^{-1}$, we define α-grid$^*(x) = \alpha$-grid$(x) \cup v$. Later, we will make use of the fact that $|\alpha\text{-grid}^*(x)| \leqslant 7$.

The following precise definition depends on the position of x and the value α. It is included for the interested reader, but not strictly needed to understand the remainder of the main body. Let c be a circle with radius α and center x. Then there exists a unique equilateral triangle $\Delta(x)$ inscribed c such that the lower side of $\Delta(x)$ is horizontal. We distinguish three cases. In the *interior* case, $\Delta(x)$ and $\partial\mathcal{P}$ are disjoint. In the *boundary* case, $\Delta(x)$ and $\partial\mathcal{P}$ have a non-empty intersection, but no vertex of $\mathcal{P}$ is contained in $\Delta(x)$. In the *corner* case, one vertex of $\mathcal{P}$ is contained in Δ. It is easy to see that this covers all the cases. We also say a point x is in the interior case, and so on. In the *interior case* α-grid(x) is defined as follows. Let v_1, v_2, v_3 be the vertices of Δ. Then the grid points g_i, which are closest to v_i, for all $i = 1, 2, 3$ form the surrounding grid points. In the *boundary case* α-grid(x) is defined as follows. Let S be the set of vertices of Δ and all intersection points of $\partial\mathcal{P}$ with $\partial\Delta(x)$. For each point $v \in S$, we define the grid point g_v closest to v and accordingly we define $G_S = \{\, g_v \,:\, v \in S \,\}$. Then α-grid$(x) = G_S$. In the *corner case* α-grid(x) is defined as follows. Let S be the set of vertices of Δ and all intersection points of $\partial\mathcal{P}$ with $\partial\Delta(x)$. For each point $v \in S$, we define the grid point g_v closest to v and accordingly we define $G_S = \{\, g_v \,:\, v \in S \,\}$. Then α-grid$(x) = G_S$. In any case, if there is a reflex vertex

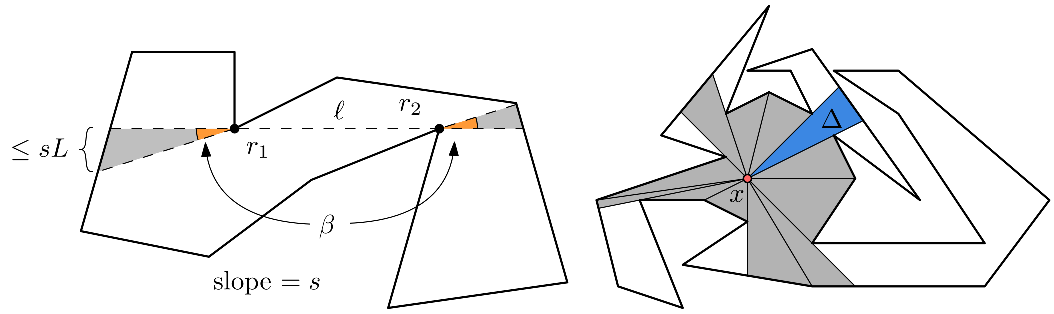

(a) A polygon with two opposite reflex vertices and their s-bad region.

(b) The star triangle decomposition of the visibility region of x.

Figure 5

r with $\text{dist}(x, r) \leqslant L^{-1}$ then we include r in the set $\alpha\text{-grid}^*(x) = r \cup \alpha\text{-grid}(x)$ as well. We will usually denote the points in $\alpha\text{-grid}(x)$ with g_1, g_2 or just g.

3.3 Local Visibility Containment

Let s be a fixed parameter to be specified later ($s = L^{-9}$). For any extension ℓ we define an s-bad region, see the gray area in Figure 5a for an illustration. Note that the bad region consists of two connected components, each being a triangle. (There can be no vertex in the interior of the triangle, because of Lemma 8 Item 2.) The parameter $s = \tan(\beta)$ is indicated in the figure. Furthermore, for each point x, the visibility region can be decomposed into triangles as indicated in Figure 5b. The region is called bad region, because Lemma 9 does not hold for points in those regions.

Let Δ be some triangle of the visibility region of x (in blue in Figure 5b). In this section, we denote the defining vertices of Δ by r_1 and r_2.

Then the main lemma asserts that $\alpha\text{-grid}^*(x)$ sees Δ except if x is in an s-bad region of the vertices defining Δ.

▶ **Lemma 9** (Special Local Visibility Containment Property). *Let r_1 and r_2 be two consecutive vertices in the clockwise order of the vertices visible from $x \in \mathcal{P}$ and let x be outside the s-bad region of the vertices r_1 and r_2 and Δ the triangle of the visibility region of x defined by r_1 and r_2. We make the following assumptions: $s \leqslant L^{-3}$, $\alpha \leqslant L^{-7}$ and $16L\alpha \leqslant s$. Then α-grid$^*(x)$ sees Δ.*

Important is the one-to-one correspondence between the triangles that cannot be seen and the extension line that we can make responsible for it.

The proof is structured in many cases. At first the triangle Δ is split into a *small triangle* (R_1) and a trapezoid (R_2), as indicated in Figure 7. We show separately, for R_1 and R_2 that $\alpha\text{-grid}^*(x)$ sees these two regions.

Another important case distinction is on whether Δ contains a point $g \in \alpha\text{-grid}^*(x)$, see Figure 6. In the first case g sees Δ as Δ is convex. In the other case, we can identify two points $g_1, g_2 \in \alpha\text{-grid}(x)$ to the left and right of Δ. For all what follows we are only concerned with the second case.

We are confronted with the situation that there might be a vertex v that is not in Δ, but obstructs the vision of g_1, g_2 in one way or another. Whenever this happens, we distinguish two cases: Either $\text{dist}(v, x) < L^{-1}$ or $\text{dist}(v, x) \geqslant L^{-1}$. In the first case $v \in \alpha\text{-grid}^*(x)$. To understand the case that v is "far" from x, one must realize that L^{-1} is huge compared to

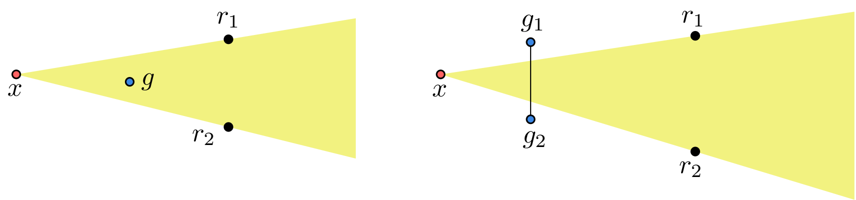

Figure 6 Left: The point g is contained in Δ and thus g sees Δ, as Δ is convex. Right: The line segment s cuts Δ.

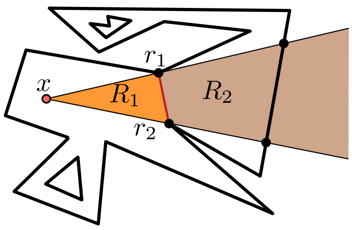

Figure 7 To show that each triangle of the visibility region is visible by α-grid$^*(x)$, we treat the small triangle R_1 and the trapezoid R_2 individually. In particular, as we do not make use of the finiteness of R_2, we just assume it is an infinite cone.

$\alpha = L^{-11}$. Thus x and $g \in \alpha$-grid(x) are affected by v in a very similar way. Unfortunately, not in exactly the same way. For instance x sees the segment seg(r_1, r_2), which has length at least one, and it is easy to show that certain points of $g_1, g_2 \in \alpha$-grid(x) see the entire segment, except a sub-segment of length at most L^{-2}, see Figure 8. This sub-segment is completely irrelevant, but we have to deal with it. These issues arise at various places. It makes forthcoming definitions more tedious and requires the proofs to be carried out with extra care.

To see that α-grid$^*(x)$ sees R_1 relies mainly on the insight, which we already mentioned above, that reflex vertices v, with dist$(x, v) \geqslant L^{-1}$ can only block a very small part of the visibility of α-grid(x) at the bottom of segment seg(r_1, r_2), as illustrated in Figure 8. For the case that there exists a reflex vertex v with dist$(x, v) < L^{-1}$, recall that v is included in α-grid$^*(x)$. Therefore, even if a reflex vertex obstructs the vision of g_2 onto seg(r_1, r_2), then g_2 can see the entire upper half of region R_1 and similarly, g_1 sees the entire lower part of R_1, as illustrated in Figure 8. Thus g_1 and g_2 see together the entire region R_1. Note that the argument does not rely on x being outside a bad region.

To prove that R_2 can be seen by α-grid$^*(x)$ is more demanding. As it seems not useful to use the boundedness of R_2, we just assume it to be an infinite cone and we show that α-grid$^*(x)$ sees this cone. Obviously, the part of $\partial\mathcal{P}$ "behind" seg(r_1, r_2) is not considered blocking. The crucial step to show that R_2 can be seen by α-grid$^*(x)$ is to show that the black region as indicated in Figure 9 does not exist. The idea is that this is implied if ray$_1$ and ray$_2$ diverge.

In other words if ray$_1$ and ray$_2$ never meet then the black region is empty. For this purpose, we make use of the fact that dist$(g_1, g_2) \approx \alpha$, for any $g_1, g_2 \in \alpha$-grid(x) by definition, while dist$(r_1, r_2) \geqslant 1$, because of integer coordinates. Thus intuitively, the distance of ray$_1$ and ray$_2$ is closer at its apex than at the segment seg(r_1, r_2). Indeed any two rays ray$_a$ and ray$_b$ will not intersect, if the following three conditions are met, see Figure 10.

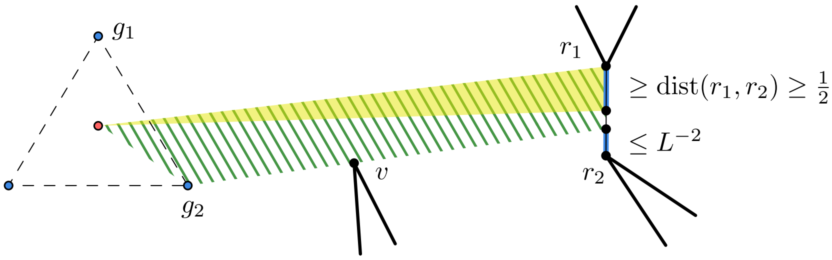

Figure 8 The point $g_2 \in \alpha\text{-grid}(x)$ sees the upper half of the Region R_1 as the green region is completely contained inside the polygon.

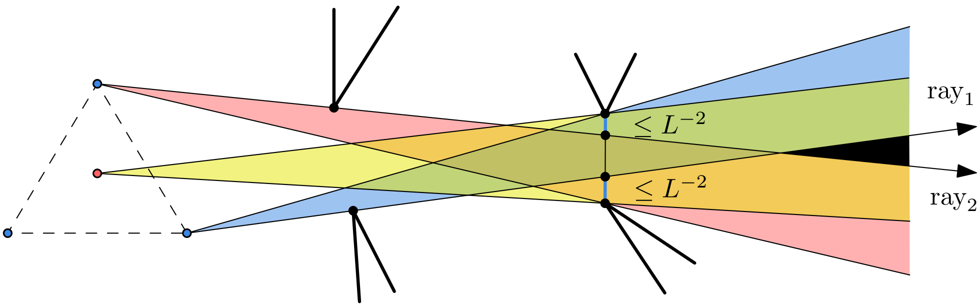

Figure 9 The visibility of the grid points $g \in \alpha\text{-grid}(x)$ can be blocked, but we can bound the amount by which it is blocked. The key idea to show that R_2 can be seen by $\alpha\text{-grid}(x)$ is to show that the region indicated in solid black is empty.

- The apex of ray_a and ray_b are "close".
- The "defining" points q_a, q_b are "far" apart.
- Both apices are outside of the s-bad region of q_a and q_b.

In order to invoke the last statement, we have to show that g_1 and g_2 are outside some appropriately defined bad regions. For this we use that x is outside the s-bad region of r_1 and r_2. The points where ray_1 and ray_2 intersect $\mathrm{seg}(r_1, r_2)$ play the role of the defining points.

3.4 Global Visibility Containment

Given a minimum solution OPT, we describe a set $G \subseteq \Gamma$ of size $O(|OPT|)$ and we show that G sees the entire polygon, see Figure 12 for an illustration. For each $x \in OPT$, G contains $\alpha\text{-grid}^*(x)$. Furthermore if x is contained in an s-bad region, G contains at least one of the vertices defining this bad region. It is clear by the previous discussion that G sees the entire polygon, as the only part that is not seen by $\alpha\text{-grid}^*(x)$ are some small regions, which are entirely seen by the vertices bounding it.

It remains to show that there is no point in three bad regions. For this, we heavily rely on the integer coordinates and the general position assumption. Note that the integer coordinate assumption implies not just that the distance between any two vertices is at least 1 but also that the distance between any extension ℓ and a vertex v not on ℓ is at least L^{-1}. Also the angle between any two extensions is at least L^{-2}. (Recall that L is an upper bound on the diameter and the largest appearing integer.) These bounds and other bounds of this kind imply that if any three bad regions meet in the interior, then their extension lines must meet in a single point, see Figure 11. We exclude this by our general position assumption.

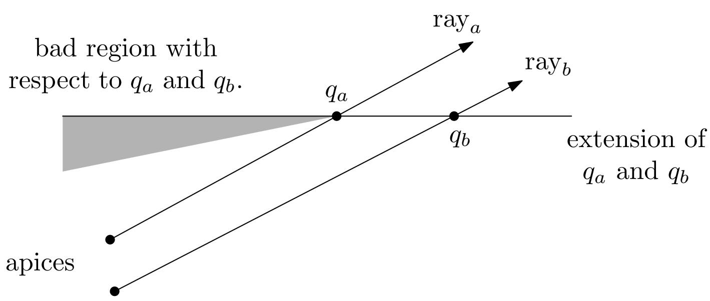

Figure 10 If the distance of the rays is closer at its apices than at q_a and q_b then we can conclude that the rays are diverging and never crossing.

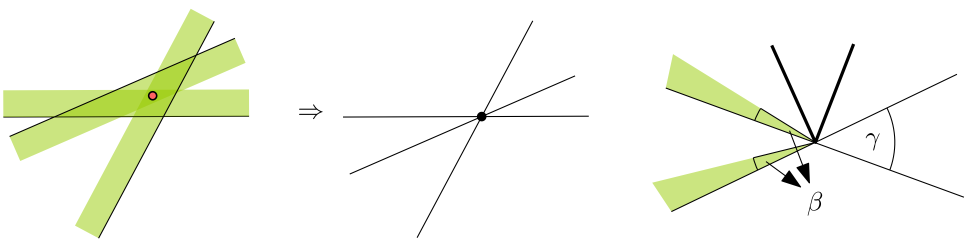

Figure 11 Three bad regions meeting in an interior point implies that the extensions must meet in a single point. No two bad regions intersect in the vicinity of a vertex, as they are defined by some angle $\beta \ll L^{-2}$. But the angle γ between any two extensions is at least L^{-2}.

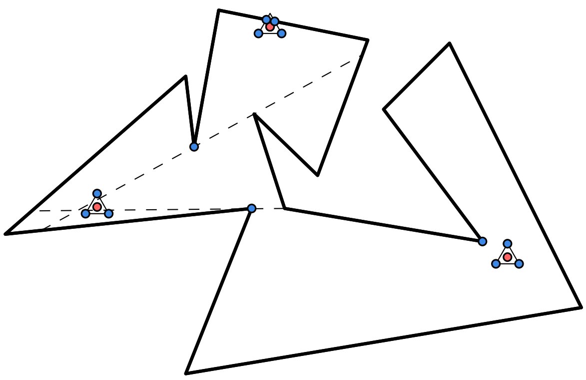

Figure 12 The red dots indicate the optimal solution. The blue dots indicate the set $G \subseteq \Gamma$ that are part of an approximate solution. The red dot on the top is in the interior case and four grid points are added around it. The red dot on the left is too close to two supporting lines and we add one of the reflex vertices of each of the supporting lines. The red dot to the right has distance less than L^{-1} to a reflex vertex, so we add that vertex to G as well.

Close to a vertex, we use a different argument: No two bad regions intersect in the vicinity of a vertex, as bad regions are defined by some angle β with $\tan(\beta) \ll L^{-2}$. But the angle γ between any two extensions is at least L^{-2}.

Recall that $|\alpha\text{-grid}^*(x)| \leqslant 7$. Together with the argument above follows that each x is in at most 2 bad regions and $|G| \leqslant (7+2)|OPT| = O(|OPT|)$.

4 Conclusion

We presented an $O(\log |OPT|)$-approximation algorithm for the Point Guard Art Gallery problem under two relatively mild assumptions. The most natural open question is whether Assumption 2 can be removed. We believe that this is possible but it will require some additional efforts and ideas. Another improvement of the result would be to achieve an approximation ratio of $O(\log n)$ for polygons with holes. This would match the currently best known algorithm for the Vertex Guard variant and the lower bound for both problems. In that respect, it is noteworthy that Lemma 4 does not require the polygon to be simple. One might also ask about the inapproximability of Point Guard Art Gallery for simple polygons. For the moment, the problem is only known to be inapproximable for a certain constant ratio (quite close to 1), unless P=NP. It would be interesting to get superconstant inapproximability under standard complexity theoretic assumptions or improved approximation algorithms.

References

1. Eyüp Serdar Ayaz and Alper Üngör. Minimal witness sets for art gallery problems. *EuroCG*, 2016.
2. János Barát, Vida Dujmovic, Gwenaël Joret, Michael S. Payne, Ludmila Scharf, Daria Schymura, Pavel Valtr, and David R. Wood. Empty pentagons in point sets with collinearities. *SIAM J. Discrete Math.*, 29(1):198–209, 2015.
3. Saugata Basu, Richard Pollack, and Marie-Francoise Roy. *Algorithms in real algebraic geometry*. Springer, 2007.
4. Patrice Belleville. Computing two-covers of simple polygons. Master's thesis, McGill University, 1991.
5. Vijay V. S. P. Bhattiprolu and Sariel Har-Peled. Separating a voronoi diagram via local search. In *SOCG*, pages 18:1–18:16, 2016.
6. Édouard Bonnet and Tillmann Miltzow. An approximation algorithm for the art gallery problem. *CoRR*, 1607.05527, 2016. URL: `http://arxiv.org/abs/1607.05527`.
7. Édouard Bonnet and Tillmann Miltzow. The parameterized hardness of the art gallery problem. In *ESA 2016*, pages 19:1–19:17, 2016. Arxiv identifier: 1603.08116.
8. Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995. `doi:10.1007/BF02570718`.
9. John Canny. Some algebraic and geometric computations in PSPACE. In *STOC*, pages 460–467. ACM, 1988.
10. Jean Cardinal. Computational geometry column 62. *SIGACT News*, 46(4):69–78, December 2015. `doi:10.1145/2852040.2852053`.
11. Václav Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, 1975.
12. Kyung-Yong Chwa, Byung-Cheol Jo, Christian Knauer, Esther Moet, René van Oostrum, and Chan-Su Shin. Guarding art galleries by guarding witnesses. *Int. J. Comput. Geometry Appl.*, 16(2-3):205–226, 2006.

13 Kenneth L. Clarkson. Algorithms for polytope covering and approximation. In *WADS 1993*, pages 246–252, 1993. doi:10.1007/3-540-57155-8_252.

14 Pedro Jussieu de Rezende, Cid C. de Souza, Stephan Friedrichs, Michael Hemmer, Alexander Kröller, and Davi C. Tozoni. Engineering art galleries. *CoRR*, abs/1410.8720, 2014. URL: http://arxiv.org/abs/1410.8720.

15 Ajay Deshpande. A pseudo-polynomial time $O(\log^2 n)$-approximation algorithm for art gallery problems. Master's thesis, Department of Mechanical Engineering, Department of Electrical Engineering and Computer Science, MIT, 2006.

16 Ajay Deshpande, Taejung Kim, Erik D. Demaine, and Sanjay E. Sarma. A pseudopolynomial time $O(\log n)$-approximation algorithm for art gallery problems. In *WADS 2007*, pages 163–174, 2007. doi:10.1007/978-3-540-73951-7_15.

17 Stephane Durocher and Saeed Mehrabi. Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In *MFCS 2013*, pages 314–324. Springer, 2013.

18 Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Inf. Process. Lett.*, 100(6):238–245, 2006. doi:10.1016/j.ipl.2006.05.014.

19 Stephan Eidenbenz, Christoph Stamm, and Peter Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.

20 Khaled Elbassioni. Finding small hitting sets in infinite range spaces of bounded VC-dimension. *CoRR*, abs/1610.03812, 2016. accepted to SoCG 2017.

21 Steve Fisk. A short proof of Chvátal's watchman theorem. *J. Comb. Theory, Ser. B*, 24(3):374, 1978. doi:10.1016/0095-8956(78)90059-X.

22 Stephan Friedrichs, Michael Hemmer, James King, and Christiane Schmidt. The continuous 1.5d terrain guarding problem: Discretization, optimal solutions, and PTAS. *JoCG*, 7, 2016.

23 Subir Kumar Ghosh. Approximation algorithms for art gallery problems in polygons. *Discrete Applied Mathematics*, 158(6):718–722, 2010.

24 Alexander Gilbers and Rolf Klein. A new upper bound for the VC-dimension of visibility regions. *Computational Geometry*, 47(1):61–74, 2014.

25 Heuna Kim and Günter Rote. Congruence testing of point sets in 4-space. In *SoCG*, volume 51 of *LIPIcs*, pages 48:1–48:16, 2016. Arxiv identifier: 1603.07269.

26 James King. Fast vertex guarding for polygons with and without holes. *Comput. Geom.*, 46(3):219–231, 2013. doi:10.1016/j.comgeo.2012.07.004.

27 David G. Kirkpatrick. An $O(\log \log OPT)$-approximation algorithm for multi-guarding galleries. *Discrete & Computational Geometry*, 53(2):327–343, 2015. doi:10.1007/s00454-014-9656-8.

28 Erik A. Krohn and Bengt J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.

29 Jirí Matousek. Intersection graphs of segments and $\exists\mathbb{R}$. *CoRR*, 1406.2636, 2014.

30 Rajeev Motwani, Arvind Raghunathan, and Huzur Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *J. Comput. Syst. Sci.*, 40(1):19–48, 1990. doi:10.1016/0022-0000(90)90017-F.

31 Joseph O'rourke. *Art gallery theorems and algorithms*, volume 57. Oxford University Press Oxford, 1987.

32 Marcus Schaefer. *Complexity of Some Geometric and Topological Problems*, pages 334–344. Springer, 2010. doi:10.1007/978-3-642-11805-0_32.

33 Thomas C. Shermer. Recent results in art galleries. *IEEE*, 80(9):1384–1399, 1992.

34 Jorge Urrutia et al. Art gallery and illumination problems. *Handbook of computational geometry*, 1(1):973–1027, 2000.