VC-dimension of Exterior Visibility

Volkan Isler and Sampath Kannan and Kostas Daniilidis*
Department of Computer and Information Science, University of Pennsylvania
Levine Hall, 3330 Walnut Street, Philadelphia, PA 19104-6389, USA
Emails: {isleri,kannan,kostas}@cis.upenn.edu
Pavel Valtr
Institute for Theoretical Computer Science and
Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
Email: valtr@kam.mff.cuni.cz

Abstract

In this paper, we study the Vapnik-Chervonenkis (VC)-dimension of set systems arising in 2D polygonal and 3D polyhedral configurations where a subset consists of all points visible from one camera. In the past, it has been shown that the VC-dimension of planar visibility systems is bounded by 23 if the cameras are allowed to be anywhere inside a polygon without holes [1]. Here, we consider the case of *exterior* visibility, where the cameras lie on a constrained area outside the polygon and have to observe the entire boundary. We present results for the cases of cameras lying on a circle containing the polygon (VC-dimension=2) or lying outside the convex hull of a polygon (VC-dimension= 5). The main result of this paper concerns the 3D case: we prove that the VC-dimension is unbounded if the cameras lie on a sphere containing the polyhedron, hence the term exterior visibility.

Keywords: VC-dimension, sensor placement, sampling, visibility.

*Corresponding Author is Kostas Daniilidis, University of Pennsylvania, Levine Hall, 3330 Walnut Street Philadelphia, PA 19104-6228, USA, Tel/Fax: +1 215 898 8549 /5732048, Email: kostas@cis.upenn.edu.

1 Introduction

Imagine a known 3D polyhedral environment where a set of cameras has to be placed in such a way that every point in the environment is visible. The 2D version is known as the art gallery problem [2, 3, 4] and sufficiency results exist for several versions of this problem. For example, $\lfloor \frac{n}{3} \rfloor$ cameras can cover any simple polygon. However, such results are inapplicable in robotic and image-based rendering applications where the environments can be very complex with millions of vertices. A further application is placing antennas for line-of-sight broadband communication [5]. Imagine that backbones end at each neighborhood and that communication inside 1km can be achieved with line-of-sight laser beams that can carry from 10Mbps up to 1.25 Gbps bandwidth. Assuming that a consumer can put a receiving antenna at the window of her studio or even on a kiosk in a street, the coverage problem becomes a visibility problem where the cameras become arrays of distributing antennas.

In this paper, we consider aspects of the problem of minimizing the number of viewpoints without sacrificing the goal of complete visibility. The particular scenarios we are addressing are surveillance, object inspection, and image based rendering. In the case of surveillance, we need a complete coverage at any time so that no event will be missed. This is the reason why coverage with one mobile guard (shortest watchman route - solvable in polynomial time [6]) is not applicable. In case of object inspection [7], we know the prior geometry of an object, and we need the minimal number of views so that the object will be checked regarding defects. In this scenario, the object might be placed on a turntable and we ask then for the minimal number of rotations. The objects might even be medical organs which have to be imaged from very few viewpoints of an endoscope guided by a robot manipulator. In the case of image based rendering, we might have a prior rough map of the environment but we need to obtain a detailed reconstruction with a range sensor. In other cases, we have a geometric model, but we do not know anything about the color or texture of an object. In all these cases, it is important that the rendered environment does not have any holes because of originally uncovered areas.

We are not going to address here the equally important problem of unknown environments or objects related to model building tasks. Several algorithms exist for exploring unknown environments [8] building upon fundamental results in the on-line traversal of graphs [9] or on

Markov processes for modeling partially known dynamic scenes [10]. Significant contributions have been also achieved in the Next Best View planning problem [11, 12, 13] for surface acquisition. However, the results proven in this paper have implications for the unknown case as well: Choosing sensor locations randomly is a method frequently used for sensor placement in unknown environments [14]. The VC-dimension theory enables us to answer the question: *How many random samples (sensors) do we need in order to cover a given region?* Our results, together with the theory of ϵ-nets, provide an upper bound to the answer to this question when omni-directional cameras are used as sensors.

It is well known [2] that the minimal guard coverage problem is NP-hard. To study the existence of approximation algorithms, we can consider minimal guard coverage as an instance of the set-cover problem. The general version of minimum set-cover cannot be approximated with a ratio better than $\log n$. However, we do not know whether any set-cover instance can be realized as a visibility system. A powerful interface between set-cover and the particular geometric setup is the Vapnik-Chervonenkis (VC) dimension which enables us to quantitatively bound how general a set system is.

In this paper, we present new bounds on the VC-dimension for three instances of the problem. In section 2, we formalize the problem statement and describe our contribution to the state of the art. Next, we deal with 2D configurations and prove new bounds on the VC-dimension for cameras on a circle (subsection 3.1). Specifically, we prove that if the cameras are restricted to lie on a circle, the VC-dimension is 2 and if they remain outside the convex hull of the object, the VC-dimension is 5 (subsection 3.2). In section 4, we prove the main result that the VC-dimension for 3D configurations is unbounded. We conclude with a summary in section 5.

2 Problem statement

2.1 VC-dimension and set-cover

A set system is a pair $(X, \mathcal{R})$ where X is a set and $\mathcal{R}$ is a collection of subsets $R \subseteq X$.

Definition 1 *Given a set system $(X, \mathcal{R})$, let A be a subset of X. We say A is shattered by $\mathcal{R}$ if $\forall Y \subseteq A$, $\exists R \in \mathcal{R}$ such that $R \cap A = Y$. The VC-dimension of $(X, \mathcal{R})$ is the cardinality* *of the largest set that can be shattered by $\mathcal{R}$ [15].*

The VC-dimension, introduced first in supervised learning for pattern classification, plays an important role also in randomized and geometric algorithms [16, 17]. As an example we mention the related notion of an ϵ-net without going into definitions: If the set system $(X, \mathcal{R})$ has a constant VC-dimension d, then with high probability, a small number $(O(\frac{d}{\epsilon}\log\frac{1}{\epsilon}))$ of points sampled from X intersects all the subsets in R whose sizes are greater than $\epsilon|X|$ (such a sample is called an ϵ-net). Another useful property is that if $(X, \mathcal{R})$ has a constant VC-dimension d, then the number of subsets in $\mathcal{R}$ is $O(n^d)$ where $n = |X|$. A related result [18] that implicitly deals with ϵ-nets is the existence of polyhedra that require arbitrary number of guards even if every point inside the polyhedron can see an area equal to ϵ fraction of the total interior area, using the notion of ϵ-good polygons introduced in path planning [19].

Given a set system, the minimum *set-cover* problem asks for a minimum cardinality set $\mathcal{S} \subseteq \mathcal{R}$ such that $\bigcup_{R\in\mathcal{S}} R = X$. The *hitting set* problem is the dual of set-cover and its decision version reads: We are given a set X and a collection of sets $\mathcal{R}$ where each $R \in \mathcal{R}$ is a subset of X. We are also given a number k. The question is whether there is a subset $H \subset X$ such that $|H| \leq k$ and for each $R \in \mathcal{R}$, $R \cap H \neq \emptyset$.

Both problems are known to be NP-complete and can be approximated to within a log factor of the maximum set sizes (in either the primal or the dual system) and not better [20]. For sets systems with finite VC-dimension d, however, Brönnimann and Goodrich presented an algorithm which returns a solution whose size is at most $O(d \cdot \log OPT \cdot OPT)$ [21]. Here, OPT is the cardinality of the optimal solution. This is a significant improvement on the previous approximation, when the cardinality of the optimal solution is smaller than the cardinality of the maximum set size in the set system.

2.2 Visibility and set-cover

In this paper we will address the problem of minimal guard coverage or camera placement. An *instance* of the camera placement problem is: Given a polygon or a polyhedron P and a specification of possible camera locations, find a minimum set-cover of the system $(P, \{V(c_i)\})$, where $V(c_i)$ is the set of points visible on P from camera c_i and the index i varies over all possible camera locations. The definition of $V(c_i)$ can capture any optical constraints on what a camera can see. We will refer to the specification of possible camera locations as a *setup*. We say a set S of cameras *cover* P if $\bigcup_{c_i \in S} V(c_i) = P$. Depending on the application, P may refer to the boundary of the object or it may be extended to include the interior of the environment as well.

Camera placement can also be seen as a particular case of the hitting set problem. The set X is the set of possible camera locations. For each point p on the boundary of the polyhedron P, there is a set R_p consisting of all camera locations that can see p. The hitting set problem assumes a finite set X and we have to implicitly deal with this issue when we attempt to pose camera placement as such a problem. Typically, we do so by using a sampling technique or by showing that a finite set of extremal points need to be considered.

Throughout this paper we will represent cameras with their projection centers c_i and say that c_i sees the point $p \in P$ if the only intersection of the line segment pc_i with P is p. We extend the notion of visibility to sets as follows: We say that a camera sees a set of points ω if it can see all the points in ω. The following notation will be useful for VC-dimension proofs: Let $P_m = \{p_1, \ldots, p_m\}$ be m points on P. We say that camera c sees the subset $\omega \subseteq P_m$ if c can see all points in ω but no point in $P_m \setminus \omega$.

By the *VC-dimension of a setup*, we will refer to the VC-dimension of the maximum number of points that can be shattered over all instances of the camera placement problem for a specific setup. For example, if there are no restrictions on cameras and we want to cover simple polygons, we would like to find the VC-dimension of the set system $(P, \{V(c_i)\})$ as P varies over the set of all simple polygons. Therefore, in order to give a lower bound m on the VC-dimension of a setup it suffices to present one instance where m points are shattered, but for an upper bound m one needs to show that there exists no instance such that $m + 1$ points can be shattered.

2.3 Related work on VC-dimension

In general, it is possible to consider visibility systems as set systems and camera placement as a set-covering problem [22]. The general version of the minimum set-cover problem cannot be approximated better than a factor of $\log n$. However, as mentioned in the first section it is not clear that general set-cover instances can be realized by visibility systems.

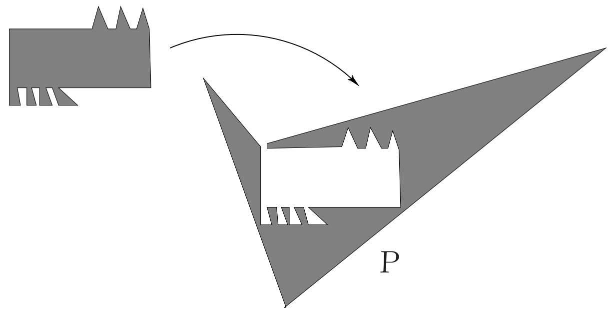

Figure 1: Interior visibility extends to exterior visibility by turning the polygon inside out.

Valtr proved that the VC-dimension of the system $(P, \{V(x) \mid \forall x \in P\})$, where P is a simple polygon and $V(x)$ is the visibility polygon of point x in P, is bounded between 6 and 23 [1]. This result applies to all 2D configurations we consider, simply by turning the polygon inside-out (see figure 1, also [3]). He also established a bound of $O(\log h)$ for polygons with h holes. A similar result for visibility under angle and distance constraints has been obtained in [23].

On the other hand, it is not clear how to exploit a bounded VC-dimension to obtain an improved approximation algorithm. Approximation algorithms for minimum guard coverage have been considered [22, 24, 23] for different versions of the problem. However there is still a gap between the inapproximability results and existing algorithms.

As mentioned before, the minimum set-cover of set systems with bounded VC-dimension can be approximated within a logarithmic factor of the optimal value [21]. However, this by itself does not improve on the existing $\log n$ approximations, as the optimum can be as big as $n/3$ [2]. Nevertheless, this algorithm was used in [23] to get rid of the dependency of the approximation factor to the number of samples (rather than the number of vertices). In fact, obtaining a constant approximation algorithm for guard placement in polygons without holes is one of main open problems in the field of art galleries.

For the setup where cameras are restricted to lie on a circle, an approximation algorithm that returns at most one more camera than the minimum necessary is presented in [25]. A similar algorithm can be found in [26] for the polygon separation problem. Placing cameras outside the convex hull of a polygon is related to hitting lines with points [27, 28]. Unfortunately there are no improved algorithms for this restricted case either.

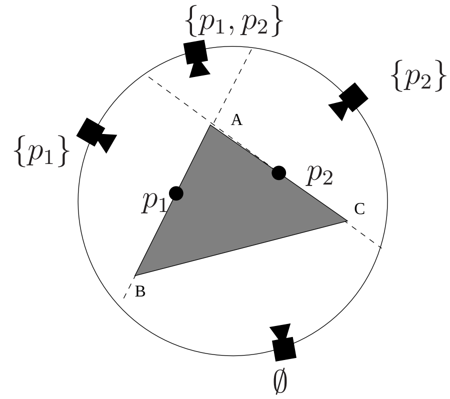

Figure 2: Points $\{p_1, p_2\}$ can be shattered by four cameras. Each camera is labeled with the subset it can see. In this figure the polygon P is $\triangle ABC$.

3 Results on planar configurations

3.1 2DCIRCLE

Consider a setup in which we want to cover a polygon P using cameras restricted to a circle C around P. Note that not all polygons are completely visible from a circle outside. In this section and the next, we restrict ourselves to the points on P that are visible from at least one point on the circle. Polygons that are entirely visible from the viewing circle are called externally visible polygons [29].

Definition 2 *We define 2DCIRCLE as a setup where a set of cameras whose locations are restricted to a circle C are to cover a simple polygon that is contained in C.*

We need the following technical lemma before proving our main proposition. Its proof is omitted because it is straightforward.

Lemma 3 *Each point p on the polygon P is visible along a continuous arc on the circle C and nowhere else.*

We now prove the following proposition.

Proposition 4 *The VC-dimension of 2DCIRCLE is exactly 2.*

Proof: For any m points on the polygon P, the m visibility arcs have $2m$ endpoints and hence there are only $2m$ combinatorially distinct camera positions. Since 2^m cameras are necessary to shatter m points, we need $2^m \geq 2m$ which is only true for $m \leq 2$.

The lower bound is proven by the example in Figure 2 where the points p_1 and p_2 are shattered by the 4 cameras shown. ∎

3.2 2DCONVEX

Let us now relax the restriction on camera locations so that we allow cameras anywhere outside the convex hull of the polygon.

Definition 5 *We define 2DCONVEX as a setup where a set of cameras located outside the convex hull of a simple polygon P are to cover P.*

The upper bound on the VC-dimension of 2DCONVEX slightly increases but it is still a small constant significantly less than the upper bound for the general case, 23.

Lemma 6 *The VC-dimension of 2DCONVEX is less than or equal to 5.*

Proof: Suppose that $Q = \{p_1, \ldots, p_6\}$ is a set of 6 points shattered on a polygon P. Let C be the boundary of conv$P$, the convex hull of P, and let E be the exterior of C, i.e. $E = R^2 \setminus (\text{conv}P)$. For each $i = 1, \ldots, 6$, a point of E sees the point p_i if and only if it lies between a certain pair of disjoint open half-lines emanating from C. We call these two open half-lines *the i-rays* and refer to the point on C from which they emanate as their endpoints. The 12 i-rays, $i = 1, \ldots, 6$, partition E into *cells*. Formally, cells are maximal connected components of the plane after the removal of *conv$P$* and of the 12 rays.

We can also obtain the cells in the following way. One by one and in an arbitrary order, we remove the 12 rays from E. After the removal of the first ray there is still one component (cell). Each other ray divides every intersected cell into two smaller cells. Thus, if the ray intersects k of the previously removed rays, its removal increases the number of cells at most by $k + 1$. Since there are at most $\binom{12}{2} - 6 = 60$ intersections between the 12 rays (the two i-rays are disjoint for each i), the number of cells in the final arrangement is at most $60 + 11 = 71$.

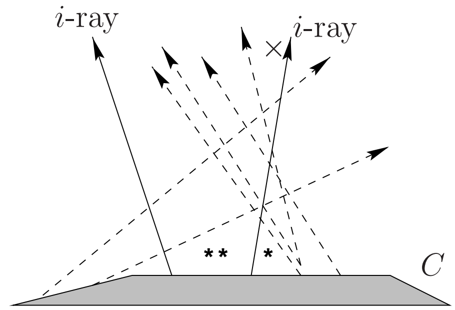

Figure 3: Case 1 of Lemma 6: The $\{p_i\}$-cell gets two marks and the $\emptyset$-cell gets one mark. The unbounded $\{p_i\}$-cell is marked $\times$.

It follows from the construction that cameras placed at different locations in the same cell see the same subset ω of $Q = \{p_1, \ldots, p_6\}$. We then call the cell an *ω-cell*. To get a contradiction it suffices to show that there are no $2^6 - 1 = 63$ cells seeing distinct nonempty subsets of Q.

First, suppose that one of the points p_i lies on C. Then the two i-rays are parts of lines tangent to C. It is easily verified that each j-ray, $j \neq i$, is disjoint from at least one of the two i-rays in this case. Whenever a j-ray is disjoint from an i-ray, $j \neq i$, this decreases the number of cells in the final arrangement by 1. Thus, the number of cells is at most $71 - 10 = 61$ in this case, which is not enough.

Thus, we may suppose that each point p_i lies inside $\mathrm{conv}P$. Whenever an i-ray is disjoint from a j-ray, $i \neq j$, for technical reasons we "create" a new, "abstract" cell and associate it with this pair of disjoint rays. The total number of "real" and "abstract" cells is exactly 71, provided no three rays intersect in a single point (otherwise it is smaller).

We suppose that every camera c_ω is placed in an unbounded cell whenever it is possible. We remove the camera $c_\emptyset$ and 63 cameras in 63 different cells remain. We say that a cell is *empty*, if it contains no camera. All "abstract" cells are empty.

We now describe a procedure which distributes 18 auxiliary marks $1, 1, 1, 2, 2, 2, \ldots, 6, 6, 6$ in some of the empty cells in such a way that at most two marks are placed in one cell. It will follow that at least $18/2 = 9$ cells are empty.

For each $i = 1, \ldots, 6$, the three marks i are distributed as follows. We say that an i-ray and a j-ray form an *i-pair*, if they are disjoint. We distinguish several cases.

Case 1: There is at most one i-pair. One of the i-rays is intersected by all j-rays for $j \neq i$.

Hence its endpoint sees no point $p_j, j \neq i$. We place two marks i in the adjacent $\{p_i\}$-cell and one mark i in the adjacent $\emptyset$-cell (see figure 3). Note that the cell with two marks i is empty because it is bounded and there is an unbounded $\{p_i\}$-cell in this case.

Case 2: There are two i-pairs.

One mark i is placed in each of the two abstract cells associated with the i-pairs. The remaining, third mark i is put in an $\emptyset$-cell chosen as follows.

Subcase 2a: No point $p_j, j \neq i$, is visible from the endpoint e of one of the i-rays. In this case we put the third mark i in the $\emptyset$-cell adjacent to e.

Subcase 2b: Condition in Subcase 2a is not valid. In this case, note that each i-ray must be responsible for exactly one i-pair. Otherwise, one of the i-rays would intersect all other j-rays and the endpoint of this i-ray would only see p_i and this is covered in Subcase 2a. Let i_1 and i_2 be the two i-rays. Suppose that the endpoint of i_1 sees some $p_j, j \neq i$. Then the wedge of visibility of p_j must intersect i_1 in a bounded interval starting at the endpoint. For all other $k \neq j$, the k-rays intersect i_1 and it is immediate that this intersection is a bounded interval again. Thus, as we go to infinity along i_1, we must have the case that only p_i is visible on one side and no point is visible on the other side. We put the third mark i to this unbounded $\emptyset$-cell adjacent to i_1. Note that this $\emptyset$-cell is not adjacent to C, as can be verified from the analogous reasoning for i_2.

Case 3: There are three or more i-pairs. We put a mark i in three abstract cells associated with distinct i-pairs.

We now check that no cell receives more than two marks. Any abstract cell is associated with a pair of disjoint rays and may receive at most two marks corresponding to these two rays. A $\{p_i\}$-cell receives at most two marks i. An $\emptyset$-cell adjacent to C may receive at most two marks corresponding to the two rays having the endpoints on the boundary of the cell. An $\emptyset$-cell non-adjacent to C may receive at most two marks corresponding to rays forming unbounded edges of the cell. All other cells receive no marks.

Hence, no cell receives more than two marks. Therefore there are at least $18/2 = 9$ cells with at least one mark. They are all empty. Thus, at most $71 - 9 = 62$ cells are nonempty. They cannot contain all the 63 cameras $c_\omega, \emptyset \subset \omega \subseteq Q$. ∎

Note that relaxing the camera locations from 2DCIRCLE to 2DCONVEX indeed increases the VC-dimension, as we see in the following lemma.

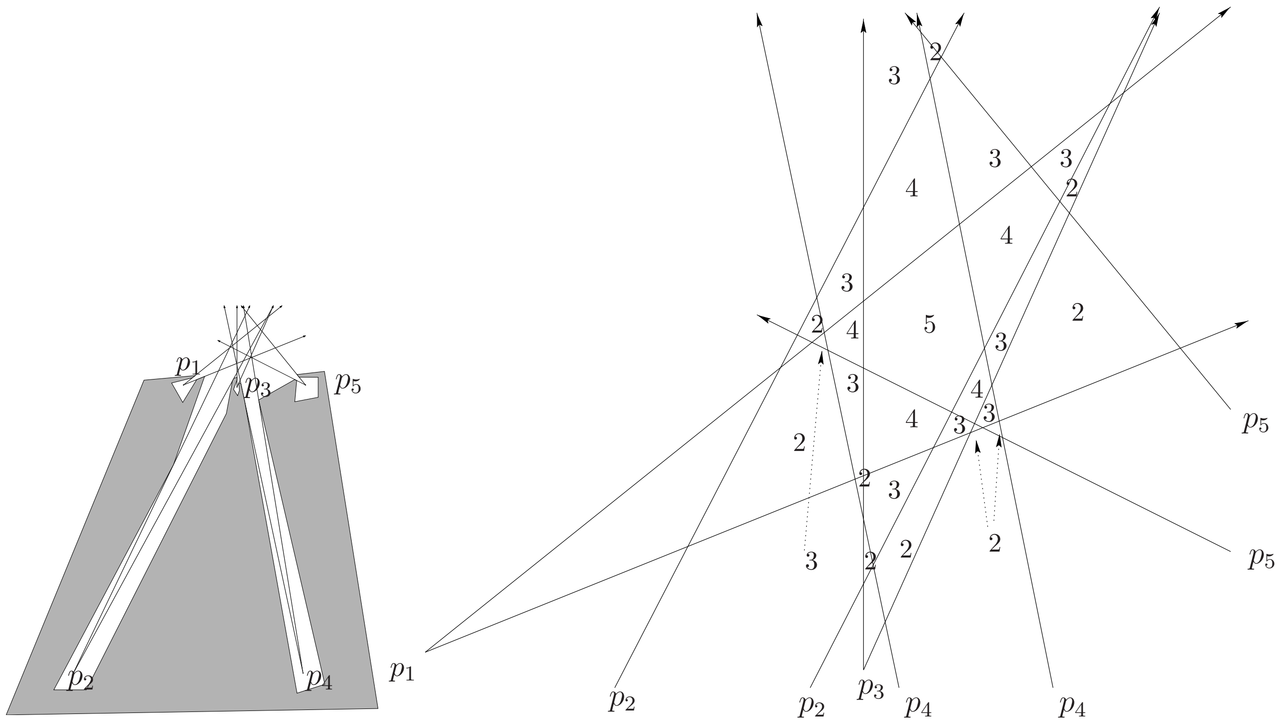

Figure 4: The small picture on the left shows a polygon with 5 shattered points $\{p_1, p_2, p_3, p_4, p_5\}$. The figure on the right is a detail showing the intersection pattern of the i-rays. Each label $2, 3, 4, 5$ indicates the cardinality of the subset seen from the region.

Lemma 7 *The VC-dimension of 2DCONVEX is greater than or equal to 5.*

Proof: An instance where five points are shattered is presented in figure 4. ∎

The result of this section is summarized in the following theorem which follows from lemma 6 and lemma 7:

Theorem 8 *The VC-dimension of 2DCONVEX is exactly 5.*

Remark: If we further remove the restriction that the cameras are outside the convex hull, then the best known bound is 23 and the reader is referred to [1].

4 Results on 3D configurations

In this section, we consider the following setup which arises in typical tele-immersive applications:

Definition 9 *We define* $3DSPHERE$ *as a setup where we are given a polyhedron* $\mathcal{P}$*, and a viewing sphere* $\mathcal{S}$ *such that* $\mathcal{P}$ *is totally contained in* $\mathcal{S}$*.*

We show that even when the viewing region is restricted to a sphere, there are polyhedra with n vertices such that $\Theta(\log n)$ points on the polyhedron can be shattered from the viewing sphere $\mathcal{S}$ that contains $\mathcal{P}$. Namely, we prove the following theorem:

Theorem 10 *The VC-dimension of 3DSPHERE is* $\Theta(\log n)$*.*

In the next two subsections, we present the upper and lower bounds for the VC-dimension of 3DSPHERE, in lemmas 11 and 12, respectively.

4.1 Upper Bound

In this section we present an upper bound on the VC-dimension of 3DSPHERE. In [30], Platinga and Dyer define *aspects* as changes in the topology of the image of a polyhedron. After presenting a catalogue of events (structural changes in the image of the polyhedron), they construct the viewing space partition, VSP, which is a partition of the viewpoint space into maximal regions of constant aspect. They also present tight bounds for the number of regions in VSP. They show that the size of the VSP for a general (i.e. non-convex) polyhedron under orthographic projection is $\Theta(n^6)$ and their model for the orthographic projection is exactly the same as 3DSPHERE with $\mathcal{S}$ at infinity.

Lemma 11 *The VC-dimension of 3DSPHERE is* $O(\log n)$*.*

Proof: Let $P_m = \{p_1, \ldots, p_m\}$ be any m points to be shattered on a polyhedron. If we define an aspect as appearance/disappearance of $p_i, i = 1, \ldots, m$, and restrict the camera locations to a sphere that contains the polyhedron, we can use the catalogue of events in [30] to show that the size of the VSP for this new notion of aspects is still $\Theta(n^6)$. However, in order to shatter m points, one needs 2^m distinct partitions. Since we must have $2^m = O(n^6)$, we have $m = O(\log n)$ which gives us the desired upper bound. ∎

4.2 Lower Bound

In this section we show that the upper bound $\log n$ on the VC-dimension of 3DSPHERE is indeed tight, our main result is stated in the following lemma for theorem 10.

Lemma 12 *For any given m, there exists a polyhedron $\mathcal{P}$ with $\Theta(2^m)$ vertices such that there are m marked points on $\mathcal{P}$ that can be shattered from 2^m cameras on the viewing sphere $\mathcal{S}$.*

Proof: We take a regular simplex (tetrahedron) T inside a viewing sphere S. Let F be a facet (2-dimensional face) of T. We take a set M of m points slightly above the center of F. Further, we place 2^m cameras $c_\omega, \omega \subseteq M$, on S visible from each point of M. We choose them so that no camera lies on a line determined by two points of M.

Then, for any $p \in \omega \subseteq M$, we choose a segment $s(p, \omega)$ having one endpoint on the face F and intersecting the segment pc_ω and no other segments $p'c_{\omega'}, p' \in M, \omega' \subseteq M$. Clearly, we may choose the segments $s(p, \omega)$ pairwise disjoint. Then we replace each $s(p, \omega)$ by a very thin simplex $S(p, \omega)$ growing out of the face F and intersecting the segment pc_ω and no other segments $p'c_{\omega'}, p' \in M, \omega' \subseteq M$. If the simplices $S(p, \omega)$ are thin enough, then they are pairwise disjoint.

Let P be the union of T with the $m \cdot 2^{m-1}$ simplices $S(p, \omega)$. Then P shatters the set M and has $m \cdot 2^{m+1} + 4$ vertices. It has a facet F' obtained from F by the removal of $m \cdot 2^{m-1}$ triangular holes. We arbitrarily triangulate F'. If we now perturb the vertices of F', then P may be changed to a simplicial polyhedron $\mathcal{P}$ such that F' is replaced by triangular facets corresponding to the chosen triangulation of F'. ∎

Proof of theorem 10: Theorem 10 is a direct consequence of lemmas 11 and 12. ∎

5 Conclusion

Visibility of a polygonal or polyhedral configuration can be cast as a set system with subsets defined by the visibility region of each camera. The minimal guard coverage problem can then be formulated as a minimum set-cover problem. The constraints imposed by the geometry of the setup can be captured with the VC-dimension of the visibility set system. It was known [1] that the upper bound of the VC-dimension for polygons is 23. In this paper, we improved this bound for two cases of exterior visibility: cameras on a circle containing the polygon and cameras outside the convex hull of a polygon. The circle case has significant practical implications because it minimizes the number of stations of a turn-table or a laser-scanner pedestal in 3D-modeling and object inspection. The placement of cameras outside the convex hull is significant in surveillance and image based rendering. For this case as well as the case of arbitrary placement inside a polygon the existence of approximation algorithms with constant ratio is still an open problem.

In the 3D case, we showed that for any n, a polyhedron with n vertices can be constructed such that the VC-dimension of its exterior visibility is $\Theta(\log n)$.

Acknowledgments: We highly appreciate the comments by the anonymous reviewers which substantially improved the readability of the paper. The first three authors are grateful for support through the following grants: NSF-IIS-0083209, NSF-IIS-0121293, NSF-CCR-98-20885, NSF-CCR-01-05337, MURI DAAH-19-02-1-03-83, and a Penn Research Foundation grant. Research by P. Valtr was supported by project LN00A056 of The Ministry of Education of the Czech Republic.

References

[1] P. Valtr, "Guarding galleries where no point sees a small area," *Israel Journal of Mathematics*, vol. 104, pp. 1–16, 1998.

[2] J. O'Rourke, *Art Gallery Theorems And Algorithms.* Oxford University Press, 1987.

[3] T. Shermer, "Recent results in art galleries," *Proceedings of the IEEE*, vol. 80, pp. 1384–1399, 1992.

[4] J. O'Rourke, "Visibility," in *Handbook of Discrete and Computational Geometry* (J. E. Goodman and J. O'Rourke, eds.), ch. 25, pp. 467–480, CRC Press LLC, 1997.

[5] A. Acampora, "Last mile by laser," *Scientific American*, vol. 287, pp. 48–53, July 2002.

[6] W. Chin and S. Ntafos, "Shortest watchman routes in simple polygons," *Discrete and Computational Geometry*, vol. 6, no. 1, pp. 9–31, 1991.

[7] R. Sharma and J. Molineros, "Computer vision-based augmented reality for guiding manual assembly," *PRESENCE: Teleoperators and Virtual Environments, MIT Press*, vol. 6, pp. 292–317, 1997.

[8] F. Hoffmann, C. Icking, R. Klein, and K. Kriegel, "The polygon exploration problem," *SIAM Journal on Computing*, vol. 31, no. 2, pp. 577–600, 2002.

[9] X. Deng, T. Kameda, and C. Papadimitriou, "How to learn an unknown environment i: The rectilinear case," *Journal of the ACM*, vol. 45, pp. 215–245, 1998.

[10] S. M. LaValle and R. Sharma, "On motion planning in changing, partially-predictable environments," *International Journal of Robotics Research*, vol. 16, pp. 775–805, 1997.

[11] K. Tarabanis, P. Allen, and R. Tsai, "A survey of sensor planning in computer vision," *IEEE Trans. Robotics and Automation*, vol. 11, pp. 86–104, 1995.

[12] K. Kutulakos and C. Dyer, "Recovering shape by purposive viewpoint adjustment," *International Journal of Computer Vision*, vol. 12, pp. 113–136, 1994.

[13] V. Hlavac, A. Leonardis, and T. Werner, "Automatic selection of reference views for image-based scene representations," *European Conference on Computer Vision*, pp. 526–535, 1996.

[14] N. Bulusu, J. Heidemann, and D. Estrin, "Adaptive beacon placement," in *Int. Conf. on Distributed Computing Systems*, 2001.

[15] V. N. Vapnik and A. Y. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Theory of Probability and its Applications*, vol. 16, no. 2, pp. 264–280, 1971.

[16] J. Matoušek, "Geometric computation and the art of sampling," in *IEEE Symposium on Foundations of Computer Science – Tutorial*, 1998.

[17] P. K. Agarwal and S. Sen, "Randomized algorithms for geometric optimization problems," *Handbook of Randomization (P. Pardalos et al., eds.), Kluwer Academic Publishers*, 2003.

[18] P. Valtr, "On galleries with no bad points," *Discrete & Computational Geometry*, vol. 21, pp. 193–200, 1999.

[19] L. E. Kavraki, J.-C. Latombe, and R. Motwani, "Randomized query processing in robot path planning," *Journal of Computer and System Sciences*, vol. 57, pp. 50–60, August 1998.

[20] P. Slavik, "A tight analysis of the greedy algorithm for set cover," *Journal of Algorithms*, vol. 25, pp. 237–254, 1997.

[21] H. Brönnimann and M. Goodrich, "Almost optimal set covers in finite VC-dimension," *Discrete & Computational Geometry*, vol. 14, no. 4, pp. 463–479, 1995.

[22] S. Ghosh, "Approximation algorithms for art gallery problems," *Proc. of the Canadian Information Processing Society Congress*, 1987.

[23] H. Gonzalez-Banos and J. Latombe, "A randomized art-gallery algorithm for sensor placement.," *Proc. ACM Symp. on Computational Geometry (SoCG'01)*, pp. 232–240, 2001.

[24] S. Eidenbenz, C. Stamm, and P. Widmayer, "Inapproximability results for guarding polygons and terrains," *Algorithmica*, vol. 31, pp. 79–113, 2001.

[25] V. Isler, S. Kannan, and K. Daniilidis, "VC-dimension of exterior visibility of polyhedra," Tech. Rep. MS-CIS-01-34, Computer and Information Science, University of Pennsylvania, 2001.

[26] J. S. B. Mitchell and S. Suri, "Separation and approximation of polyhedral surfaces," in *Proceedings of the 3rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 296–306, 1992.

[27] V. S. A. Kumar, S. Arya, and H. Ramesh, "Hardness of set cover with intersection 1," in *Automata, Languages and Programming*, pp. 624–635, 2000.

[28] B. Brodén, M. Hammar, and B. J. Nilsson, "Guarding lines and 2-link polygons is apx-hard," in *13th Canadian Conf. on Computational Geometry*, pp. 45–48, 2001.

[29] C. Hernando, M. E. Houle, and F. Hurtado, "On local transformation of polygons with visibility properties," *Theoretical Computer Science*, vol. 289, pp. 919–937, October 2002.

[30] W. Plantinga and C. Dyer, "Visibility, occlusion, and the aspect graph," *International Journal of Computer Vision*, vol. 5, no. 2, pp. 137–160, 1990.