Randomness and Computation
Joint Workshop "New Horizons in Computing" and "Statistical Mechanical Approach to Probabilistic Information Processing" (18-21 July, 2005, Sendai, Japan)
Lecture Note in Tutorial Sessions

On finding a guard that sees most and a shop that sells most[1]

Otfried Cheong[2]
Dept. of Computer Science, KAIST, 373-1 Guseong-dong, Yuseong-gu, Daejeon 305-701, South Korea

1 Introduction

We consider two problems where our goal is to find a point x such that the area of the region $V(x)$ "controlled" by x is as large as possible. In the first problem, we are given a simple polygon P, and $V(x)$ is the *visibility polygon* of x, that is, the region of points y inside P such that the segment xy does not intersect the boundary of P. In the second problem, we are given a set of points T, and $V(x)$ is the *Voronoi cell* of x in the Voronoi diagram of the set $T \cup \{x\}$, that is, the set of points that are closer to x than to any point in T.

In both problems, it is straightforward (but tedious) to write a closed formula describing the area of the region controlled by a point x. This area function (inside a region where $V(x)$ has the same combinatorial structure) is the sum of the areas of triangles that depend on the location of x. The function domain consists of a polynomial number of regions, and the function has a different closed form in each region: it is the sum of $\Theta(n)$ low-degree rational functions in two variables, which do not have common denominator. It seems difficult to solve the problem of finding the maximum of this function analytically and efficiently, and we resort to approximation.

In this paper we address the question of efficiently finding a point x that approximately maximizes the area of $V(x)$. More precisely, let $\mu(x)$ be the area of $V(x)$, and let $\mu_{opt} = \max_x \mu(x)$ be the area for the optimal solution. Given $\delta > 0$, we show how to find x_{app} such that $\mu(x_{app}) \geq (1-\delta)\mu_{opt}$.

The main motivation for our first problem arises from *art-gallery* or *sensor placement* problems. In a typical problem of this type, we are given a simple polygon P, and wish to find a set of points (*guards*) so that each point of P is seen by least one guard. This problem is NP-hard. Art-gallery problems have attracted a lot of research in the last thirty years [16, 17]. A natural heuristic for solving art-gallery problems is to use a greedy approach based on area: We first find a guard that maximizes the area seen, next find a guard that sees the maximal area not seen by the first guard, and so on until each point of P is seen by some guard.

Ntafos and Tsoukalas [14] show how to find, for any $\delta > 0$, a guard that sees an area of size $(1-\delta)\mu_{opt}$. Their algorithm requires $O(n^5/\delta^2)$ time in the worst case. We give a probabilistic algorithm that finds a $(1-\delta)$-approximation in time $O((n^2/\delta^4)\log^3(n/\delta))$ with high probability. We also show that approximating the largest visible polygon up to a constant factor is 3SUM-hard [11], implying that our algorithm is probably close to optimal as far as the dependency on n is concerned.

Our second problem is motivated by the task of placing a new supermarket such that it takes over as many customers as possible from the existing competition. If we assume that customers are uniformly distributed and shop at the nearest supermarket, then our task is indeed to find a point x such that the Voronoi region of x is as large as possible. The area of Voronoi regions has been considered before in the context of games, such as the Voronoi game [1, 5] or the Hotelling game [15]. As far as we know, the only previous paper discussing maximizing the Voronoi region of a new point is by Dehne et al. [9], who show that the area function has only a single local maximum inside a region where the set of Voronoi

[1] Joint work with Alon Efrat and Sariel Har-Peled. An earlier version of this work has appeared in the proceedings of SODA 2004.
[2] E-mail: otfried@cs.kaist.ac.kr

neighbors does not change and is in convex position. They give an algorithm for finding (approximately) the optimal new point numerically based on Newton approximation.

We show that given a set T of n points and a $\delta > 0$, we can find a point x_{app} such that $\mu(x_{app}) \geq (1-\delta)\mu_{opt}$, where $\mu(x)$ is the area of the Voronoi region of x in the Voronoi diagram of $T \cup \{x\}$, and $\mu_{opt} = \max_x \mu(x)$. The (deterministic) running time of the algorithm is $O(n/\delta^4 + n \log n)$.

Our framework captures a variety of other problems, where the goal is to maximize the area of some region which depends on a multi-dimensional parameter. As an example of such a further application, we consider the problem of matching two planar shapes P and Q under translations or rigid motions. The area of overlap (or the area of the symmetric difference) of two planar regions is a natural measure of their similarity that is insensitive to noise [3, 7]. Mount et al. [13] first studied the function mapping a translation vector to the area of overlap of a translated simple polygon P with another simple polygon Q, showing that it is continuous and piecewise polynomial of degree at most two. If m and n are the number of vertices of P and Q, respectively, then the function has $O((nm)^2)$ pieces, and can be computed within the same time bound. No algorithm is known that computes the translation maximizing the area of overlap that does not essentially construct the whole function graph. De Berg et al. [7] gave an $O((n+m)\log(n+m))$ time algorithm to solve the problem in the case of convex polygons, and gave a constant-factor approximation. Alt et al. [3] gave a constant-factor approximation for the minimum area of the symmetric difference of two convex polygons. For rigid motions, Brass recently showed how to find the rigid motion maximizing the area of overlap of two collections of triangles in time $O(n^8)$. A $(1-\varepsilon)$-approximation algorithm for the case of convex polygons with running time $O((\log n)/\varepsilon + (1/\varepsilon)\log(1/\varepsilon))$ was given by Ahn et al. [2]. Finally, de Berg et al. [8] consider the case where P and Q are disjoint unions of m and n unit disks, with $m \leq n$. They compute a $(1-\varepsilon)$-approximation for the maximal area of overlap of P and Q under translations in time $O((nm/\varepsilon^2)\log(n/\varepsilon))$, and under rigid motions in time $O((n^2m^2/\varepsilon^3)\log n)$.

Our framework applies immediately to this problem: for a translation vector x, let $P(x)$ denote the translation of P by x, and let $\mu(x)$ be the ratio of the areas of $P(x) \cap Q$ and P. Clearly, $0 \leq \mu(x) \leq 1$, where $\mu(x) = 1$ for a perfect match (this model allows to search for P appearing as a subpattern in Q). Let $\mu_{opt} := \max_x \mu(x)$. We show how to find a translation x_{app} such that $\mu(x_{app}) \geq \mu_{opt} - \varepsilon$. Note that the error is *absolute* here. This makes sense in shape matching: if μ_{opt} is small (say less than 1/10), then P and Q cannot be matched well. In many applications it will be sufficient to know that no decent match is possible, rather than a precise estimate on how poor the match is.

If P and Q are polygonal regions of complexity m and n, the running time of our procedure is $O(m + (n^2/\varepsilon^4)\log^2 n)$ for translations, and $O(m + (n^3/\varepsilon^4)\log^5 n)$ for rigid motions.

2 The framework

Consider a function V that assigns to $x \in \mathbb{R}^d$ a region $V(x) \subset \mathbb{R}^m$. Our goal is to compute the point $x_{opt} \in \mathbb{R}^d$ maximizing the volume of $V(x_{opt})$, or rather to compute a point $x_{app} \in \mathbb{R}^d$ such that the volume of $V(x_{app})$ is at least $(1-\varepsilon)$ times the volume of $V(x_{opt})$.

In all our applications, the region $V(x)$ can be computed efficiently for any given $x \in \mathbb{R}^d$. One can therefore use any optimization package to compute an approximation to the maximum value of $V(x)$. However, such techniques may get stuck on local maxima, whose value can be much smaller than the true global maximum.

Our framework is based on the idea of replacing the continuous volume measure by a discrete measure: we generate a set $S \subset \mathbb{R}^m$ of suitable sample points, and use the cardinality of $V(x) \cap S$ as an approximation of the volume of $V(x)$. This then turns the optimization problem into a discrete problem: for each $s \in S$, let $W(s) = \{x \in \mathbb{R}^d \mid s \in V(x)\}$. For each point $x \in \mathbb{R}^d$, the cardinality of $V(x) \cap S$ is the number of regions $W(s)$ containing x, and so we can find the point $x \in \mathbb{R}^d$ maximizing this quantity by constructing the arrangement of all the regions $W(s)$. The sample-point based estimate is sufficiently tight such that the point that maximizes it is a good approximation for the optimal solution.

This framework can be applied to a large number of problems directly, but often the direct application does not result in an efficient algorithm. We demonstrate on the three applications mentioned in the introduction how insight into the geometry of the problem can be used to fine-tune the application of our framework.

3 Maximizing the Voronoi region

Let T be a given fixed set of n points in the plane. For a point x not necessarily in T, let $V_T(x)$ denote the Voronoi region of x in the Voronoi diagram of $T \cup \{x\}$, and let $\mu(x)$ denote the area of $V_T(x)$. We are looking for a point x_{opt} maximizing $\mu_{opt} = \mu(x_{opt})$. For points x outside the convex hull of T, $\mu(x)$ would be infinite. There are quite a few ways of avoiding these boundary situations: using torus topology, restricting the point (i.e., supermarket) to lie within a polygon (i.e., city limits), or by adding a boundary that acts as an additional site. In the following we choose the first option, and assume the input is a set of points in a unit square with torus topology. The reader can easily modify the arguments to handle the boundary in a different way.

The *reach* of a Voronoi region $V_T(x)$ is the distance between the site x and the furthest point inside $V_T(x)$, or, in other words, the radius of the smallest disc centered at x containing $V_T(x)$. We can estimate μ_{opt} as follows.

Lemma 1 *Let ℓ be the largest reach of any Voronoi region $V_T(t)$, for $t \in T$. Then*

$$\pi\ell^2/4 \leq \mu_{opt} \leq \pi\ell^2.$$

Proof. Let p be a point realizing the reach ℓ, that is, its distance to the nearest site is ℓ. It follows that $V_T(p)$ contains the disc with center p and radius $\ell/2$, and so $\mu(p) \geq \pi\ell^2/4$. The lower bound follows.

Let now x be the point realizing the optimal solution, that is $\mu(x) = \mu_{opt}$, and let $y \in V_T(x)$ be the point furthest from x. Its distance to the nearest site in T is at most ℓ, and so its distance to x is at most ℓ. It follows that $V_T(x)$ is contained in the disc with radius ℓ and center x, implying the upper bound. □

Note that the largest reach ℓ is also the radius of the largest empty circle. It can be computed in $O(n \log n)$ time by computing the Voronoi diagram V of T and inspecting every vertex of V.

Our goal is to find a point x_{app} such that $\mu(x_{app}) \geq (1-\delta)\mu_{opt}$, for some parameter $\delta > 0$. We partition the unit square (containing T) into a grid of squares with side length ℓ. For each grid cell $\mathcal{Q}$, we define an estimate function e_S, such that for any $x \in \mathcal{Q}$ we have

$$|e_S(x) - \mu(x)| \leq \frac{\delta\pi\ell^2}{8},$$

and we pick a point $x_\mathcal{Q} \in \mathcal{Q}$ maximizing $e_S(x_\mathcal{Q})$. Let's first argue that this solves the problem: Let x_{app} be the point $x_\mathcal{Q}$ that maximizes $e_S(x_\mathcal{Q})$. Then

$$\begin{aligned}\mu(x_{app}) &\geq e_S(x_{app}) - \frac{\delta\pi\ell^2}{8} \geq e_S(x_{opt}) - \frac{\delta\pi\ell^2}{8} \geq \mu_{opt} - \frac{\delta\pi\ell^2}{4} \geq \mu_{opt} - \delta\mu_{opt} \\ &= (1-\delta)\mu_{opt},\end{aligned}$$

and so x_{app} is the desired approximate solution.

It remains to show how to define e_S and how to find the point $x_\mathcal{Q}$, for each grid cell $\mathcal{Q}$. Let's fix a grid cell $\mathcal{Q}$, and let x be a point in $\mathcal{Q}$. The reach of $V_T(x)$ is at most ℓ, and so $V_T(x)$ can intersect only $\mathcal{Q}$ itself and its eight neighboring grid cells. Consequently, all points of T participating in the definition of $V_T(x)$ lie in $\mathcal{Q}$ and the 24 grid cells at distance at most 2ℓ. Let $\mathcal{Q}'$ denote the union of these 25 grid cells, and let $T_\mathcal{Q} = T \cap \mathcal{Q}'$.

We make use of the following simple lemma.

Lemma 2 *Let S be a square grid of density ε in the plane, that is, the distance between neighboring grid points is ε, and let C be a convex body of diameter at most D. Then*

$$\left|\mu(C) - \varepsilon^2 \left|C \cap S\right|\right| \leq 4D\varepsilon.$$

Proof. Consider the tessellation of the plane into little squares of side length ε, where each point of S is the center of one little square. The boundary of C intersects at most $4D/\varepsilon$ little squares, which implies the bound. □

We set $\varepsilon = \delta\pi\ell/64$ and let S be a square grid of density ε, covering $\mathcal{Q}'$. For a point $x \in \mathcal{Q}$, let

$$e_S(x) = \varepsilon^2 \, |V_T(x) \cap S|$$

be the *estimate of the Voronoi region of* x. Making use of the fact that the diameter of $V_T(x)$ is at most 2ℓ, we then have by Lemma 2

$$|e_S(x) - \mu(x)| \leq 8\ell\varepsilon \leq \delta\pi\ell^2/8,$$

and by what we observed above, it remains to find the point $x_{\mathcal{Q}} \in \mathcal{Q}$ maximizing $e_S(x_{\mathcal{Q}})$. To this end, we define

$$W(s) = \{x \in \mathcal{Q} \mid s \in V_T(x)\}.$$

Note that $W(s)$ is simply the largest disc with center s that contains no point of T in its interior, clipped to $\mathcal{Q}$. Let $\mathcal{W}_S = \{W(s) \mid s \in S\}$ and consider the arrangement $\mathcal{A}(\mathcal{W}_S)$. Our problem has reduced to finding a point in $\mathcal{Q}$ that is contained in the largest number of clipped discs in $\mathcal{W}_S$.

Theorem 1 *Given a set T of n points in the plane and a parameter $\delta > 0$, one can deterministically compute, in time $O(n/\delta^4 + n \log n)$, a point x_{app} such that $\mu(x_{app}) \geq (1-\delta)\mu_{opt}$.*

Proof. We start by computing the Voronoi diagram of T and inspecting its vertices to determine the largest reach ℓ. We then define the square grid, and determine the set of points $T_{\mathcal{Q}}$ relevant in each grid cell. Since a point of T is relevant in at most 25 grid cells, the total size of the sets $T_{\mathcal{Q}}$ is $O(n)$.

For each grid cell $\mathcal{Q}$ we take a square grid S of density $\varepsilon = \delta\pi\ell/64$. It consists of $M = 25\ell^2/\varepsilon^2 = O(1/\delta^2)$ points. For $s \in S$, the clipped disc $W(s)$ can be determined by finding the nearest neighbor to s in $T_{\mathcal{Q}}$. We do this by simply comparing the distance from s to each point in $T_{\mathcal{Q}}$. The arrangement $\mathcal{W}_S$ is computed by a sweep-line algorithm in time $O(M^2)$. The number of discs containing each face of the arrangement can again be determined by a simple transversal. We pick a point $x_{\mathcal{Q}}$ from the face maximizing the estimate $e_S(x_{\mathcal{Q}})$.

By the choice of ℓ, every grid cell is within distance at most 2ℓ from a point of T. The number of grid cells handled is therefore at most $O(n)$. Each point of T appears at most $25M$ times in a nearest-neighbor computation, and so the overall running time is $O(n \log n + nM + nM^2) = O(n/\delta^4 + n \log n)$. ◻

4 Shape matching

We now briefly discuss the application of our framework to the shape matching problem. Let P and Q be polygonal regions of complexity m and n, respectively. For a translation vector x, let $V_{PQ}(x)$ denote $P(x) \cap Q$, where $P(x) = \{p + x \mid p \in P\}$ is the region obtained by translating P by x. Let $\mu(x) = \mu(V_{PQ}(x))/\mu(P)$ denote the ratio of the areas of $V_{PQ}(x)$ and P. Our goal is to find the translation x maximizing $\mu(x)$. As before, let μ_{opt} be $\max_x \mu(x)$, and let x_{opt} be such that $\mu_{opt} = x_{opt}$.

We normalize such that $\mu(P) = 1$, and so $\mu(x)$ becomes simply $\mu(V_{PQ}(x))$. We sample a set S of M points in P, and for a translation x (identified with a point x in the plane), we count the fraction of sample points that is translated into Q to obtain the estimate

$$e_S(x) = \frac{|S(x) \cap Q|}{|S|},$$

where $S(x) = \{s + x \mid s \in S\}$.

We will use the following Chernoff-bound for the absolute error.

Lemma 3 *Let X_i, $i = 1, ..., r$, be independent random variables with values 0 and 1, let $X = \sum_{i=1}^{r} X_i$, and let ε be a parameter with $0 < \varepsilon < 1$. Then*

$$\mathbf{Pr}[|X - E[X]| > \varepsilon r] < 2\exp(-\varepsilon^2 r/2).$$

We now define, for $s \in S$, $W(s) = \{x \mid s + x \in Q\}$. Obviously, $W(s)$ is a translated copy of Q. Let $\mathcal{A}_Q(S)$ be the arrangement of all regions $W(s)$. As before, we choose the vertex x_{app} of $\mathcal{A}_Q(S)$ that maximizes $e_S(x_{app})$. We have the following lemma.

Lemma 4 *Let $0 < \varepsilon < 1$ be a parameter, let S be a uniform sample from P of size $M \geq \frac{c \log n}{\varepsilon^2} + 2$, for a suitable constant c, and let x be a vertex of the arrangement $\mathcal{A}_Q(S)$. Then,*

$$\mathbf{Pr}[|e_S(x) - \mu(x)| > \varepsilon] \leq \frac{1}{n^6}.$$

Proof. Consider a fixed vertex x of $\mathcal{A}_Q(S)$. We observe that x is the intersection of the boundaries of $W(s_1)$ and $W(s_2)$, for two points $s_1, s_2 \in S$. Let $T = S \setminus \{s_1, s_2\}$. The random sample T is independent of x. We find that

$$|e_T(x) - e_S(x)| \leq 2/(M-2) \leq \varepsilon/10,$$

for c large enough.

We now define random variables X_s, for $s \in T$, as follows: $X_s = 1$ if $s + x \in Q$, else $X_s = 0$. Let $X = \sum_{s \in T} X_s$ and let $r = |T| = M - 2$. We then have $X = e_T(x) \cdot r$, and $E[X] = \mu(x) \cdot r$. By Lemma 3, we have

$$\mathbf{Pr}[|e_T(x) - \mu(x)| > \varepsilon/2] = \mathbf{Pr}[|X - E[X]| > \varepsilon r/2] < 2\exp(-\varepsilon^2 r/8) \leq \frac{2}{n^{c/8}}.$$

The arrangement $\mathcal{A}_Q(S)$ consists of M translated copies of Q, and has at most $M^2 n^2$ vertices. This implies that the probability that for *any* vertex x of $\mathcal{A}_Q(S)$ we have $|e_S(x) - \mu(x)| > \varepsilon$ is at most $2M^2 n^2 / n^{c/8}$, which is less than $1/n^6$ for large enough c. ◻

Lemma 4 implies that with high probability the vertex x_{app} of $\mathcal{A}_Q(S)$ maximizing $e_S(x_{app})$ fulfills $\mu(x_{app}) \geq \mu(x_{opt}) - 2\varepsilon$. We therefore have the following main theorem.

Theorem 2 *Given polygonal shapes P and Q of complexity m and n in the plane. In time $O(m + (n^2/\varepsilon^4) \log^2 n)$ we can compute a translation x_{app} such that $\mu(x_{app}) \geq \mu_{opt} - \varepsilon$ with high probability.*

If P and Q are disjoint unions of m and n unit disks, then the running time reduces to $O((n/\varepsilon^4) \log^3 n)$.

Proof. We start by triangulating P, so that we can take the sample S uniformly at random from P. The arrangement $\mathcal{A}_Q(S)$ can be computed in time $O(n^2 M^2) = O((n^2/\varepsilon^4) \log^2 n)$. The vertex x_{app} maximizing $e_S(x_{app})$ can then be found by a simple traversal of the arrangement.

If P and Q are disjoint unions of unit disks (the case studied by de Berg et al. [8]), the arrangement $\mathcal{A}_Q(S)$ consists of M translated copies of a set of n disjoint unit disks. Each disk can intersect only a constant number of disks in each $W(s)$, and so the total number of vertices of $\mathcal{A}_Q(S)$ is at most $O(nM^2)$. It can be computed in time $O(nM^2 \log n) = O((n/\varepsilon^4) \log^3 n)$, for instance by a plane sweep, and again x_{app} can be found using a traversal. ◻

Finally, we consider the problem of maximizing $\mu(V_{PQ}(x))$ under *rigid motions* x. The probabilistic analysis above goes through nearly unchanged, using a three-dimensional configuration space for x. It is, however, no longer attractive to explicitly compute the arrangement $\mathcal{A}_Q(S)$, as the regions $W(s)$ are now bounded by curved surfaces. Fortunately, we do not need the arrangement $\mathcal{A}_Q(S)$ as long as we can somehow enumerate its vertices. A vertex of $\mathcal{A}_Q(S)$ is defined by a rigid motion that moves up to three points of S onto the boundary of Q. For each triple of sample points from S and each triple of edges of Q, this can happen only a constant number of times, and so we can enumerate all possible vertices of $\mathcal{A}_Q(S)$ in time $O(M^3 n^3) = O((n^3/\varepsilon^6) \log^3 n)$. For each candidate motion x, we test whether it maps each $s \in S$ into Q. Using a point location data structure for Q, this can be done in time $O(M \log n)$, and so the total running time is $O(m + M^4 n^3 \log n) = O(m + (n^3/\varepsilon^4) \log^5 n)$.

References

[1] H.-K. Ahn, S.-W. Cheng, O. Cheong, M. Golin, and R. van Oostrum. Competitive facility location along a highway. In *7th Annual International Computing and Combinatorics Conference*, volume 2108 of *LNCS*, pages 237–246, 2001.

[2] H.-K. Ahn, O. Cheong, C.-D. Park, C.-S. Shin, and A. Vigneron. Maximizing the overlap of two planar convex sets under rigid motions. In *Proc. 20th Ann. Symp. on Computational Geometry*, pages 356–363, 2005.

[3] H. Alt, U. Fuchs, G. Rote, and G. Weber. Matching convex shapes with respect to the symmetric difference. *Algorithmica*, 21:89–103, 1998.

[4] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley Inter-Science, 2nd edition, 2000.

[5] O. Cheong, S. Har-Peled, N. Linial, and J. Matousek. The one-round Voronoi game. *Discrete Comput. Geom.*, 31(1):125–138, 2004.

[6] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[7] M. de Berg, O. Cheong, O. Devillers, M. van Kreveld, and M. Teillaud. Computing the maximum overlap of two convex polygons under translations. *Theo. Comp. Sci.*, 31:613–628, 1998.

[8] M. de Berg, P. Giannopoulos, C. Knauer, R. van Oostrum, and R. C. Veltkamp. Maximizing the area of overlap of two unions of convex homothets under translation. In *Scandinavian Workshop on Algorithms*, volume 3111 of *LNCS*, pages 138–149, 2004.

[9] F. Dehne, R. Klein, and R. Seidel. Maximizing a Voronoi region: The convex case. In *Proc. 13th Annu. Internat. Sympos. Algorithms Comput.*, pages 624–634, 2002.

[10] L. Gewali, A. Meng, J. S. B. Mitchell, and S. Ntafos. Path planning in $0/1/\infty$ weighted regions with applications. *ORSA J. Computing*, 2:253–272, 1990.

[11] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5:165–185, 1995.

[12] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.

[13] D. M. Mount, R. Silverman, and A. Y. Wu. On the area of overlap of translated polygons. *Computer Vision and Image Understanding: CVIU*, 64(1):53–61, 1996.

[14] S. C. Ntafos and M. Z. Tsoukalas. Optimum placement of guards. *Info. Sciences*, 76:141–150, 1994.

[15] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial tessellations: Concepts and applications of Voronoi diagrams*. Probability and Statistics. Wiley, 2nd edition edition, 2000.

[16] J. O'Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[17] J. Urrutia. Art gallery and illumination problems. In J. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North-Holland, 2000.

[18] P. Valtr. Guarding galleries where no point sees a small area. *Israel J. Math*, 104:1–16, 1998.