# Solving The Art Gallery Problem Using Gradient Descent

Master's Thesis



Georgiana Juglan

Supervisor: Tillmann Miltzow Second Examiner: Frank Staals

Computing Science November 9, 2022 Student Nr. 2996307 Utrecht University The Netherlands



Universiteit Utrecht

## Abstract

The Art Gallery Problem is one of the central problems in computational geometry. It was recently shown that the problem is  $\exists \mathbb{R}$ -complete. Thus it is unlikely to be solvable by methods that are usually applied to NP-complete problems. These methods include heuristics with provable guarantees to come as close as possible to the optimal solution in polynomial time. Nonetheless, one of the most important methods to solve  $\exists \mathbb{R}$ -complete problems is called gradient descent. Curiously, there was no attempt yet to solve the Art Gallery Problem with gradient descent. By aiming to maximise the area seen by the guards in the Art Gallery Problem, we get a continuous cost function, which allows us to compute the gradient. Using the gradient and other heuristics inspired from neural networks, we try to solve the Art Gallery Problem practically using gradient descent. Specifically, we aim to use the library CGAL. State-of-the-art implementations show how feasible this approach is. We visualise some preliminary results and discuss the performance of the algorithm on different input shapes and sizes.

## Contents

1 Introduction							
	1.1	The Art Gallery Problem	5				
	1.2	Gradient Descent	6				
	1.3	Thesis Goal	6				
2 Literature Summaries							
	2.1	Efficient Computation of Visibility Polygons	7				
		2.1.1 Algorithm of Joe and Simpson	7				
		2.1.2 Algorithm of Asano	8				
		2.1.3 New Algorithm: Triangular Expansion	10				
		2.1.4 Experiments	10				
	2.2	Irrational Guards Are Sometimes Needed	11				
		2.2.1 Intuition for Triangular Pockets	12				
		2.2.2 Intuition for Quadrilaterial Pockets	13				
		2.2.3 Intuition for Rectangular Pockets	13				
		2.2.4 Monotone Polygon Construction	13				
		2.2.5 Rectilinear Polygon Construction	14				
	2.3	Implementation of Guarding Algorithms	15				
		2.3.1 Algorithm of Ghosh	15				
		2.3.2 Algorithm of Bhattacharya, Ghosh and Roy	16				
		2.3.3 New Algorithm	17				
		2.3.4 Experiments	17				
2.4 An Algorithm with Performance Guarantees		An Algorithm with Performance Guarantees	19				
		2.4.1 One-Shot Vision-Stable Algorithm	20				
		2.4.2 Iterative Vision-Stable Algorithm	21				
		2.4.3 Experimental Results	21				
3	The	Theory 23					
3.1 Guarding the Polygon with One Point		Guarding the Polygon with One Point	23				
		3.1.1 Gradient Descent	23				
		3.1.2 Computing the Gradient	24				
		3.1.3 Computing the New Guard's Position	26				
3.2 (		Guarding the Polygon with Multiple Guards	27				
		3.2.1 Computing the Gradient for Two Guards	27				
		3.2.2 Computing the Gradient for Multiple Guards	28				
	3.3	Momentum	30				
3.4 Reflex Vertex Pull		Reflex Vertex Pull	31				
		3.4.1 Pull onto the Reflex Vertex	33				
		3.4.2 Pull Capping	35				
	3.5	Reflex Area	35				
	3.6	Line Search	37				
	3.7 Angle Behind Reflex Vertex		37				
3.8 Hide		Hidden Movement	38				
	3.9 Greedy Initialisation		40				

<b>4</b>	Practice					
	4.1	Heuris	tics	41		
		4.1.1	Without Momentum	41		
		4.1.2	Without Pulling To and Onto the Reflex Vertex	43		
			4.1.2.1 Evaluation and Results	44		
		4.1.3	Without Pull Capping	47		
		4.1.4	Without Reflex Area	49		
		4.1.5	Without Line Search	51		
		4.1.6	Without Angle Behind Reflex Vertex	52		
		4.1.7	No Hidden Movement	54		
		4.1.8	Greedy Initialisation	55		
4.2		Scaling	g for the Comb Polygon	56		
	4.3	Hyper	parameter Sensitivity	59		
5	Problems Encountered					
6	3 Discussion					
	6.1	Future	Work	64		

## 1 Introduction

#### 1.1 The Art Gallery Problem

The Art Gallery Problem [18] is a central problem in computational geometry. It can be introduced as follows: given a simple polygon P with n vertices, we are interested in finding the minimum number of guards that are able to see the whole polygon. A simple polygon is a polygon that has no holes. Thus, we can define the visibility of a guard g in the polygon  $P \subset \mathbb{R}^2$ . The guard  $g = (x, y) \in P$ sees another point  $q \in P$  if the line segment  $\overline{gq} \subseteq P$ . The points that are visible from g form the visibility polygon (region) Vis(g). In the Art Gallery Problem, we are looking for a minimum size set of guards S that can see the whole polygon P.

Figure 1.1 displays an example of the Art Gallery Problem with polygon P guarded by 2 guards (|S| = 2). The visibility region Vis of each guard is marked with a different colour. For guard  $g \in S$ , its visibility region Vis(g) is emphasised with the pink contour. The vertex r blocks part of the view of g. The vertex r is thus called "reflex", because its interior angle is larger than 180°. Reflex vertices are only found in concave polygons, so convex polygons can be guarded by only one guard.



Figure 1.1: Example of an Art Gallery Problem instance with polygon P guarded by 2 guards. The visibility area Vis(g) is emphasised in pink.

The Art Gallery Problem is  $\exists \mathbb{R}$ -complete [1], which means it is even harder to solve than NPcomplete problems [19]. For this reason, approximation algorithms have been extensively used to address it ([5], [11], [13]). Nonetheless, to the best of our knowledge there is no related work on approaching the Art Gallery Problem using gradient descent. As such, we approach the Art Gallery Problem from a new perspective using gradient descent.

#### 1.2 Gradient Descent

Gradient descent is an iterative optimisation algorithm for finding the minimum of a continuous differentiable function. The core idea of gradient descent is to repeatedly move in the opposite direction of the gradient of the function at the current point using a specific step size (learning rate). High learning rates result in approaching a local optimum faster, but risk overshooting it. Conversely, small learning rates are more precise, with the compromise of a longer computation and convergence time. When there is no more change in the gradient, then an optimum has been reached. Gradient descent does not guarantee that the found optimum is global. For this reason, it can remain stuck in local optima.

Figure 1.2 illustrates an example of applying gradient descent. The optimisation function takes the shape of a curve. Starting from an arbitrary point on the curve, the goal is to reach the minimum of the function (the bottom of the curve). This is done by computing the gradient (derivative) of the function at the current point. The gradient is displayed in red and tells that the largest increase in value for the function is going up the curve. Because we are interested in finding the minimum of the function, we move in the opposite direction of the gradient. So, we move down the curve with the given step size. Continuing from the new point on the curve, the process is repeated until the minimum is reached. If we were interested in finding the maximum of the function, we would move in the same direction with the gradient.



Figure 1.2: Illustrative example of applying gradient descent.

#### 1.3 Thesis Goal

This thesis is an algorithm engineering paper. The main goal is to create and implement an algorithm that uses gradient descent to approximate the solution to the Art Gallery Problem. The explored research question is whether the Art Gallery Problem can be efficiently solved using gradient descent. Additional to gradient descent, the algorithm deploys different heuristics to address various short-comings and edge-cases. These are discussed based on polygon shapes and sizes. The algorithm is implemented in C++ using the CGAL library [7].

This paper is organised as follows: Section 2 offers an overview of existing work. Section 3 explains how gradient descent is used to solve the Art Gallery Problem and what heuristics we can make use of to improve the performance of our algorithm. Section 4 offers an overview of our algorithm's performance. Due to the time constraints of the thesis, extensive experiments have only been performed for one of the heuristics. Namely, Section 4.1.2.1 shows a preliminary evaluation for the experiments, based on the number of iterations required to solve a problem instance. For the rest, only an intuition about the advantages and use-cases for each of the heuristics is given. Lastly, Section 5 discusses issues this project has encountered.

## 2 Literature Summaries

This section offers an overview of previous works addressing the Art Gallery Problem [18]. Working with the visibility region is a basic tool in computational geometry, specifically in the context of the Art Gallery Problem. The Art Gallery Problem is an  $\exists \mathbb{R}$ -complete problem [1]. The problem class  $\exists \mathbb{R}$  consists of instances that can be reduced in polynomial time to a decision problem of whether a system of polynomial equations with integer coefficients and any number of real variables has a solution. Since NP  $\subseteq \exists \mathbb{R}$ , the  $\exists \mathbb{R}$  class is an at least as hard to solve complexity class. That is due to the fact that  $\exists \mathbb{R}$ -complete problems are not know to have solutions of polynomial space.

Building on the aforementioned concepts, this section will further inspect how visibility in polygons can be efficiently computed [6], how it is not always possible to place guards at rational coordinates [1], in addition to two improved polygon guarding algorithms [16], [13].

#### 2.1 Efficient Computation of Visibility Polygons

Bungiu, Hemmer, Hershberger, Huang and Kröller [6] introduce the implementations and their experimental evaluations for two existing algorithms ([14], [3]) and a newly developed one for computing visibility in polygons. These implementations are available in the CGAL library [7], starting with version 4.5.

Therefore, Bungiu et al. present three algorithms and their practical performance.

#### 2.1.1 Algorithm of Joe and Simpson

The algorithm of Joe and Simpson [14] runs in O(n) time and space.

Let  $v_i$ , for i = 1, 2, ..., n, be the boundary vertices of the polygon P. Let g be a guard in P, and let s be a stack data structure. The stack s will be used to keep track of the vertices determining Vis(g).

The algorithm begins by scanning the boundaries of P. The scanning is done through shooting rays  $p\vec{v}_i$ , for i = 1, 2, ..., n in this order. The endpoints  $v_i$  and  $v_{i+1}$  of each ray form a boundary edge  $\overline{v_i v_{i+1}}$ . In this way, the processing of  $v_i$  and  $v_{i+1}$  is done by checking whether the points are in Vis(g). This means that the position of every  $v_{i+1}$  with respect to the ray  $p\vec{v}_i$  is checked. If  $v_{i+1}$  is found in front of the ray  $p\vec{v}_i$  (if  $v_{i+1}$  is seen from g), then both  $v_i$  and  $v_{i+1}$  are added to the processing stack s. For every newly pushed vertex on s, the algorithm checks whether the segment  $\overline{v_i v_{i+1}}$ obscures any of the previously added vertices. If that is the case, then the endpoints of the obscured line segment are declared obsolete and deleted. The polygon comprised of the vertices from s forms at the end the visibility polygon Vis(g).

Figure 2.1 displays an example run of the Algorithm of Joe and Simpson [14] for polygon P and guard g. First, the ray from g is shot through vertex  $v_1$ , and  $v_1$  is added to s. Then, the ray from g is shot through  $v_2$ . Since ray  $p\vec{v}_2$  takes a right turn from  $p\vec{v}_1$ , this means that  $v_2$  is still in the visibility region of g. For this reason,  $v_2$  is also added to s. The ray passing through  $v_2$  also intersects the boundary of P in a point  $v'_2$ . To account for the fact that g can see "behind"  $v_2$  and is still inside P, the boundary vertex  $v'_2$  is hence added to s. Next, the ray  $p\vec{v}_3$  takes a left turn from  $v_2$ , which means that  $v_3$  is not seen from g.



Figure 2.1: An example run of the algorithm of Joe and Simpson [14] for polygon P guarded by g and boundary vertices  $v_i$ , for  $i = \{1, 2, ..., n\}$ .

Similarly,  $v_4$  and  $v_5$  are added to s. However, because ray  $p\vec{v}_6$  takes a left turn from  $p\vec{v}_5$ , segment  $\overline{v_5v_6}$  obscures  $\overline{v_4v_5}$ . So,  $v_4$  and  $v_5$  are removed from s. Ray  $p\vec{v}_6$  then intersects the boundary of P in  $v'_6$ . At the end,  $Vis(g) = \{v_1, v_2, v'_2, v_6, v'_6, v_7, v_8, v_9, v_{10}\}$ , as shown on the boundary of the green area.

#### 2.1.2 Algorithm of Asano

The algorithm of Asano [3] runs in  $O(n \log n)$  time and O(n) space. It uses a plane sweep approach with event line L.

Let P be a polygon determined by vertices  $\{a_1, a_2, a_3, b_1, b_2, b_3, b_4\}$ . P may have holes. Suppose we want to guard it by point g. The algorithm of Asano begins by efficiently sorting all the vertices of P based on their polar angles with respect to the guard g. Figure 2.2 displays an example run of the algorithm. The points will be treated in the order of  $a_2, a_1, a_3, b_4, b_3, b_2, b_1$  with respect to g and their angular comparisons

$$\measuredangle a_2 Op < \measuredangle a_1 Op < \measuredangle a_3 Op < \measuredangle b_4 Op < \measuredangle b_3 Op < \measuredangle b_2 Op < \measuredangle b_1 Op,$$

where O = (0, 0).

Then, the event line L starts sweeping around g as shown in Subfigures 2.2a - 2.2g. Every line segment that L intersects is stored in a balanced binary tree T in the order of their intersection angle. As T is updated, a new vertex of Vis(g) is stored each time the segment closest to g in T changes. It is important to mention that the intersection between L and the line segments is not explicit, but is instead determined by comparisons between the endpoints' coordinates. For instance, in Figure 2.2, the endpoint  $b_2$  of line segment  $\overline{b_2a_2}$  is the first one L intersects. Point  $b_2$  is thus added to T. Then, Lcontinues sweeping and adds  $b_1, a_2$  and  $a_1$  to T. Although  $\overline{b_2a_2}$  and  $\overline{b_1a_1}$  represent line segments  $s_2$ and  $s_1$ , respectively, the intersection of L with them is not explicitly computed, but is determined based solely on the positions of their endpoints:  $s_1$  is farther away from g because  $q, a_2$  and  $b_2$  are on the same side of  $s_2$ .

8

P

P



O = (0, 0)

(a)  $a_2$  is added to the empty Vis(g) such that  $Vis(g) = \{a_2\}.$ 





(c)  $a_3$  is added to Vis(g) such that  $Vis(g) = \{a_1, a_2, a_3\}$ , and line segment  $\overline{a_1 a_3}$  is added to T such that  $T = \{\overline{b_1 a_1}, \overline{a_1 a_3}\}.$ 



(e)  $b_3$  is added to Vis(g)such that  $Vis(g) = \{a_1, a_2, a_3, b_4, b_3\}$ , and line segment  $\overline{b_4 b_3}$  is added to Tsuch that  $T = \{\overline{b_1 a_1}, \overline{a_1 a_3}, \overline{a_3 b_4}, \overline{b_4 b_3}\}$ .



(b)  $a_1$  is added to Vis(g) such that  $Vis(g) = \{a_1, a_2\}$ , and line segment  $\overline{b_1 a_1}$  is added to the empty binary tree  $T = \{\overline{b_1 a_1}\}$ .



(d)  $b_4$  is added to Vis(g) such that  $Vis(g) = \{a_1, a_2, a_3, b_4\}$ , and line segment  $\overline{a_3b_4}$  is added to T such that  $T = \{\overline{b_1a_1}, \overline{a_1a_3}, \overline{a_3b_4}\}.$ 



(f)  $b_2$  is added to Vis(g) such that  $Vis(g) = \{a_1, a_2, a_3, b_3, b_4, b_2\}$ , and line segment  $\overline{b_2a_2}$  is added to T such that  $T = \{\overline{b_1a_1}, \overline{a_1a_3}, \overline{a_3b_4}, \overline{b_4b_3}, \overline{b_2a_2}\}.$ 

Figure 2.2: Example run of the Algorithm of Asano [3] on polygon P and guard g. The vertices of P are added to the binary tree T in the order of their angle between g and the origin O = (0,0). The result of the algorithm is visibility region  $Vis(g) = \{a_1, a_2, a_3, b_3, b_4, b_2\}$ .



(g)  $b_1$  is not added to Vis(g) because it is obstructed by the line segment  $\overline{b_2a_2}$  which is already in T. For the same reason, line segments  $\overline{b_3b_1}$  and  $\overline{b_1a_1}$  are also not added in T.

Figure 2.2: Example run of the Algorithm of Asano [3] on polygon P and guard g. The vertices of P are added to the binary tree T in the order of their angle between g and the origin O = (0,0). The result of the algorithm is visibility region  $Vis(g) = \{a_1, a_2, a_3, b_3, b_4, b_2\}$ .

#### 2.1.3 New Algorithm: Triangular Expansion

The algorithm introduced by Bungiu et al. [6] is named Triangular Expansion and runs in  $\Omega(n^2)$  time and O(n) space. It begins by triangulating P in  $O(n \log n)$  time if P has holes, and O(n) otherwise. The implementation runtime is constrained by CGAL, which makes use of the Delaunay triangulation algorithm [9] with  $O(n^2)$  time for the worst case, but with better performance in practice.

Taken from [6] and annotated to suit the explanations in these summaries, Figure 2.3 depicts an example run of the algorithm on a polygon with holes P. Starting from the viewpoint g, the triangle containing g is located by performing a simple walk. Trivially, g sees the entire triangle it is contained in. The algorithm continues by recursively expanding the view of g from one triangle into the next, until there are no more triangles to expand into. The view of g becomes restricted by the reflex vertices l and r of the third triangle entered by the recursive step. Since l and r are reflex vertices, the view past them is further restricted until the boundaries l' and r' of P, respectively, are reached. Line segments  $\overline{ll'}$  and  $\overline{rr'}$  are added to Vis(g) in their angular order around g. At the end, Vis(g) will contain the segments delimiting the visibility polygon of g.

#### 2.1.4 Experiments

Bungiu et al. do not report on benchmarks with query points on edges in the interior polygon. This is because they claim that their implementations perform similarly to other already implemented algorithms. Instead, they use two real-world scenarios (a simple polygon of Norway with 20981 vertices, and a cathedral polygon with 1209 vertices) and a worst-case polygon for the Triangular Expansion algorithm.

In terms of results on the real-world polygons, the Triangular Expansion algorithm has a 2-factor improved performance when compared to Asano's algorithm, and performs "one order of magnitude" faster than Joe and Simpson's algorithm. For the worst-case scenario, Asano's algorithm outperforms



Figure 2.3: The Triangular Expansion Algorithm Example - recursion entering triangle  $\Delta$  through edge e [6].

the Triangular Expansion algorithm with increasing input complexity.

Thus, despite the Triangular Expansion algorithm being outperformed in the worst-case scenario, Bungiu et al. add efficient implementations for 3 different polygon visibility algorithms in the CGAL library. The choice of algorithms when using the library can be adapted based on the input polygons.

#### 2.2 Irrational Guards Are Sometimes Needed

Abrahamsen, Adamaszek, and Miltzow [1] study the placement of the guards in the context of the Art Gallery Problem [18]. Namely, it focusses on and confirms that there are polygons with integer coordinates that require guards placed at points with irrational coordinates. Generalising, Abrahamsen et al. show that  $\forall n \in \mathbb{N}$ , there is a family of simple monotone polygons that can either be guarded by 3n guards with irrational coordinates, or requires at least 4n guards with rational coordinates. The result is also extended to a rectilinear polygon that can be guarded by 9 guards with irrational coordinates, or requires at least 10 guards with rational coordinates. The family of simple monotone polygons that can be guarded by 3n guards with irrational coordinates, as well as the rectilinear polygon that can be guarded by 9 guards with irrational coordinates are also discussed.

First, we will address the question regarding whether polygons given by integer coordinates require guard positions with irrational coordinates in any optimal solution by introducing a crafted monotone polygon P. A polygon is monotone if there exists a line l such that every line orthogonal to l intersects its boundary at most twice. Polygon P is depicted in Figure 2.4 and is constructed using a rectangle, six triangular pockets (in green), three rectangular pockets (in blue) and four quadrilaterial pockets (in red) are added. The intuition behind the polygon will be explained in the upcoming subsections. The figure is accredited to Abraham et al. [15].



Figure 2.4: The polygon P developed by Abraham et al. [15] to emphasise the need for irrational guards sometimes.

#### 2.2.1 Intuition for Triangular Pockets

The six triangular pockets are created in order to force a guard on a line segment, as depicted in Figure 2.5. The figure is taken from [1]. As such, they are grouped into three pairs: top pockets are paired with bottom pockets corresponding to their position in P. For every pair (leftmost, middle, rightmost) of triangular pockets, there is a corresponding line  $l_{\ell}, l_m, l_r$ , respectively, that joins the peaks of each triangular pocket in a pair. A guard can see both the peaks t, b of the pockets only if it is placed on the line segment  $\overline{tb}$  (Figure 2.5a). Hence, one guard is needed per pair, resulting in 3 guards in total.





(a) The only way that one guard can see both t and b is when the guard is on the blue line.

(b) The only way to guard the polygon with three guards requires one guard on each of the green lines  $l_{\ell}, l_m, l_r$ .

Figure 2.5: Forcing guards to lie on specific line segments [15].

#### 2.2.2 Intuition for Quadrilaterial Pockets

The four quadrilateral pockets are created in order to force a guard in a region bounded by a curve. Figure 2.6 was taken and adapted from [15], and depicts the visualisation for this case. Consider some position of  $g_{\ell}$  on  $l_{\ell}$ . As  $g_{\ell}$  moves up on  $l_{\ell}$ , the point  $p_{t}^{\ell}$  slides towards the right on segment  $e_{t}^{\ell}$ , and point  $p_{b}^{\ell}$  slides to the left on segment  $e_{b}^{\ell}$ . In this way, the lines  $p_{t}^{\ell}b$  and  $p_{b}^{\ell}d$ , where b, d are reflex vertices, intersect and form the curve  $c_{\ell}$ . Considering thus the fixed position of a guard  $g_{m}$  either on the left or on the right of the curve  $c_{\ell}$ , there exists a position of guard  $g_{\ell}$  on  $l_{\ell}$  such that edge  $e_{t}^{\ell}$  and line  $p_{b}^{\ell}g_{m}$  are seen by  $g_{\ell}$ . Then, only if  $g_{m}$  is on the left or on  $c_{\ell}$  can it see the remaining parts of the pockets that are not seen by  $g_{\ell}$  (edge  $e_{b}^{\ell}$ , line  $p_{t}^{\ell}g_{\ell}$ ). Analogously, in order for the right side of Pand the right pair of quadrilateral pockets to be seen by both guards  $g_{m}$  and  $g_{r}$ ,  $g_{m}$  has to be on the curve  $c_{r}$  or on its right. Since  $g_{m}$  has to also satisfy its position on  $l_{m}$ , its only feasible position is as the intersection point between  $c_{\ell}, c_{r}$  and  $l_{m}$ .



Figure 2.6: Restricting a guard to a region bounded by a curve [15].

#### 2.2.3 Intuition for Rectangular Pockets

The three rectangular pockets are created in order to force a guard to a single irrational point. Further building onto the previously discussed instance from Figure 2.6, the three rectangular pockets on the laterals and top of P are added. They allow for additional constraints for the positions of the guards  $g_{\ell}, g_m, g_r$ . Based on the curve equations of  $c_{\ell}$  and  $c_r$ , we can thus calculate the irrational position of  $g_m = (3.5 + 5\sqrt{2}, 1.5\sqrt{2})$ . Subsequently, we can fix the positions of  $g_{\ell}$  and  $g_r$  on lines  $l_{\ell}$ and  $l_r$ .

#### 2.2.4 Monotone Polygon Construction

The polygon P in question was found through experimentation in GeoGebra [10]. The triangular and rectangular pockets were fixed. Finding then the position of the rectangular pockets required the existence of a rational line that contained the irrational coordinates of guard  $g_m$  from Figure 2.6. The irrational coordinates of the two other guards  $g_\ell, g_r$  were chosen such that they would be able to see the rest of the polygon that is unseen by  $g_m$ . The rectangular pockets were added based on their coordinates.

In this way, by placing each guard placed on the lines  $l_{\ell}, l_m, l_r$  at unique irrational positions, respectively, dependencies between guard positions were created. The resulting guard set becomes thus  $S = \{g_{\ell}, g_m, g_r\}$ . Given the unique and irrational positions of the guards, the only method for seeing the whole polygon P using guards with rational positions would be to use at least 4. This statement can also be generalised such that a family of polygons  $(\mathcal{P}_n)_{n \in \mathbb{Z}_+}, \forall n$  can be guarded by 3n guards with irrational coordinates, or requires 4n guards with rational coordinates. The coordinates determining the polygons are hence polynomial in n.

#### 2.2.5 Rectilinear Polygon Construction

Similarly, a rectilinear polygon  $P_R$  can be created given integer coordinates. Polygon  $P_R$  would require guards with irrational coordinates in any optimal solution. It can be guarded by 9 guards if guards can be placed at points with irrational coordinates. Otherwise, the smallest optimal guard set S with rational coordinates would be required to have size 10.

Polygon  $P_R$  can be constructed by extending P with the 6 non-rectiliniar grey parts  $(Q_1, Q_2, Q_3, Q_4, T_1, T_2)$  in Figure 2.7, such that each of them requires at least 1 guard in the interior. The figure was taken from [15]. When placing 6 guards in the grey parts, the white areas of  $P_R$  remain unseen. As such, 3 guards must still be similarly placed at the 3 irrational points on lines  $l_\ell, l_m, l_r$  as in the case of P, in order to guard the rest of  $P_R$ .



Figure 2.7: Rectilinear Polygon  $P_R$  [15].

## 2.3 Implementation of Polygon Guarding Algorithms for Art Gallery Problems

Maleki and Mohades [16] implement efficiently two existing approximation algorithms ([11], [4]) for computing visibility in simple polygons. Namely, they introduce practical implementations for visibility algorithms that already offer theoretical guarantees. Additionally, they develop their own visibility algorithm that aims to offer performance guarantees for polygon visibility computation. They lastly evaluate experimentally their implementations for the three algorithms in the context of the Art Gallery Problem [18].

To begin with, we will introduce some terminology used throughout this summary. The algorithms distinguish in their implementation between vertex guards and point guards in a polygon P. Vertex guards can be placed only on the vertices of P. Point guards can be placed without restriction inside P. Lastly, the algorithms are tested on weak visibility polygons. A polygon P has weak visibility if all boundary vertices of P can see all the points in P.

#### 2.3.1 Algorithm of Ghosh

The algorithm of Ghosh [11] runs in  $O(n^4)$  time and yields an  $O(\log n)$ -approximation algorithm for computing the minimum vertex guard for simple polygons.

One of the most important concepts the algorithm works with is that of a convex component. Given a polygon P, we can form subsets of the vertices in P such that all the subsets form convex subpolygons of P. We call these subsets of vertices the convex components of P.

The algorithm begins by computing the set of all the convex components C of a given polygon P. Then, each vertex  $j \in P$  creates a set  $F_j$  with the convex components from C that it is able to fully see. Let  $F := \{F_j \mid \forall j \in P, F_j \subseteq C \text{ seen by } j\}$ . Then, the algorithm checks for overlaps between the sets of F. For every vertex i and its corresponding convex components set  $F_i$ , the algorithm checks whether any of the other sets  $F \setminus F_i$  sets are included in  $F_i$ . That is, if  $F_j \subseteq F_i, F_j \in F, i \neq j$ . If that is the case, then vertex i sees at least as little as j. The vertex j is thus not needed as a guard, so  $F_j$ is removed from F. Vertex i is added in the final guard set S. The algorithm repeats until the set Fbecomes empty. When that happens, it means that the algorithm found a set S of guards who see all the convex components of P without overlap.

An example run of the algorithm is depicted in Figure 2.8. Let P be the polygon in question, and its boundary vertices  $\mathcal{V} = \{1, 2, 3, 4, 5, 6\}$ . Polygon P is divided into two convex components  $C_1$  and  $C_2$ , such that  $C = \{C_1, C_2\}$ . The sets  $F_j, \forall j \in P$  are also displayed. Since  $F_1 \subseteq F_2 \subseteq F_3 \subseteq F_4 \subseteq F_6$ and  $F_5 \subseteq F_3 \subseteq F_4 \subseteq F_6$ , then  $F_1, F_2, F_3, F_4$ , and  $F_5$  are removed. The remaining set  $F_6$  yields the final guard set  $S = \{6\}$ , which can see both convex regions of P, and hence the whole polygon.



Figure 2.8: Example run of the Algorithm of Ghosh [11] with polygon P divided into two convex components  $C_1$  and  $C_2$ . The resulting guard set is  $S = \{6\}$ .

#### 2.3.2 Algorithm of Bhattacharya, Ghosh and Roy

The algorithm of Bhattacharya et al. [4] runs in  $O(n^2)$  time and yields a 6-approximation for computing the minimum vertex guard for weak visibility polygons without holes.

It begins by choosing two neighbours u and v as parents for every vertex in P. It then computes the Shortest Path from each pair of parents (u, v) to every other vertex in P. The Shortest Path is a path between two vertices such that the distance between them is minimal. If all distances between two adjacent vertices are the same, then the length of the path corresponds to the number of edges between the vertices. The Shortest Path from a pair of vertices (u, v) to any other vertex w is the length of the minimum path between u and w, and v and w.

Then, all vertices that can be reached from u and v are unmarked. In increasing angular order from  $\overline{uv}$ , every vertex  $w \in P$  is checked for visibility from u, and v. If all vertices w are visible, then u, and v are added to S. Otherwise, w is added to S and the procedure continues with w as the starting node. All vertices that become seen from S are marked. At the end, the algorithm checks whether the vertices in S have overlapping visibility regions, and duplicates are removed.

An example run of the algorithm is depicted in Figure 2.9. Let P be the polygon in question, and its boundary vertices  $\mathcal{V} = \{a, b, c, d, e, f\}$ . The algorithm starts with vertices a and c as parents of b. Clockwise, vertex d is visible from  $\overline{ac}$ , but e is not. For this reason, e is added to the guard set S, and vertices d and f are chosen as the new parents. In increasing angular order, all vertices a, b, c, e are visible from  $\overline{ef}$ , so d and f are added to S. Since the visibility regions of d and f coincide (Vis(d) = Vis(f)), and the visibility region of e is included in that of d and f ( $Vis(e) \subseteq Vis(d) \subseteq Vis(f)$ ), eand d can be removed from S. As such, the final guard set becomes  $S = \{f\}$ .



Figure 2.9: Example run of the Algorithm of Bhattacharya et al. [4] for polygon P. The final guard set is  $S = \{f\}$ .

#### 2.3.3 New Algorithm

The algorithm introduced by Maleki and Mohades is focussed on polygons with large number of vertices n and different amounts of reflex vertices r. If the number of reflex vertices is significantly lower than the total number of vertices  $(r \leq \log \log n)$ , guards are placed at all reflex vertices. Otherwise, they are placed according to the algorithm of Ghosh [11].

#### 2.3.4 Experiments

Algorithms of Ghosh [11] and Bhattacharya et al. [4] are tested on weak visibility polygons, and the newly introduced algorithm is tested on simple polygons.

Let Procedure 1 be a procedure for generating weak visibility polygons that is illustrated in Figure 2.10: given two points p = (k, 0), q = (-k, 0), n random points  $\{x_1, ..., x_n\}$  sorted on the distance from p on  $\overline{pq}$ , n sorted random angles  $\{\alpha_1, ..., \alpha_n\} \in (0, \pi)$ , and n vertices  $\{y_1, ..., y_n\}$  are created by shooting n rays at the corresponding angle  $\alpha_i, \forall x_i$ . Then, n reflex vertices  $\{z_1, ..., z_n\}$  are added in the quadrilateral formed by vertices  $x_i y_i y_{i+1} x_{i+1}$ . The figure is accredited to [16].



Figure 2.10: Generated weak visibility polygon for n = 3 [16] through Procedure 1.

The algorithms of Ghosh [11] and Bhattacharya et al. [4] were tested using polygons generated by Procedure 1 with  $n \in \{10, 11, ..., 15\}$  vertices and  $r \in \{2, 3\}$  reflex vertices. The results suggest that for low values of n and r, the algorithm of Ghosh [11] performs better when using the number of guards as evaluation criteria. Similarly, the algorithm of Ghosh [11] performed better both when tested on a weak visibility polygon with low  $n = \overline{10, 15}$  and  $\frac{n}{2} \leq r$ , and when tested with large  $n \in \{100, 400\}$ .

Secondly, let Procedure 2 be a procedure for generating simple polygons with custom number of reflex vertices r is devised in Figure 2.11. The figure was taken from [16] and annotated to suit the explanations in these summaries. Starting from a simple convex polygon P with n vertices u, v, x1, x2, ..., x7, polygon P is triangulated such that every triangle has a joint edge with its boundaries; r reflex vertices r1, r2, r3 are randomly added inside P, and the boundaries outside of the reflex vertices are moved such that all r points are now forming boundaries.



Figure 2.11: Generated simple polygon P for n = 12, r = 3 [16] through Procedure 2.

The new algorithm is tested on simple polygons constructed as mentioned by starting from low r and gradually increasing it. The results are reported positively in the sense that the |S| always remains close to the optimal, as a 2-approximation solution.

Therefore, through the newly implemented algorithm, Maleki and Mohades [16] testify that the algorithm of Ghosh [11] performs like a constant approximation in practice, and often better than its theoretical bound when tested on complex simple polygons.

## 2.4 A Practical Algorithm with Performance Guarantees for the Art Gallery Problem

Hengeveld and Miltzow [13] introduce the idea of vision-stability, and describe an iterative algorithm that guarantees to find the optimal guard set for every vision-stable polygon in order to address the Art Gallery Problem [18]. Additionally, it proves that the set of vision-stable polygons admits an FPT algorithm when parametrised by the number of vertices visible from a chord of the polygon, and gives a one-shot algorithm for that matter.

Hengeveld and Miltzow [13] start by introducing the notion of vision-stability (see Figure 2.12). The figure was taken from [13] and annotated to facilitate the understanding of this summary. The authors define guards as "enhanced" if they can "see around the corner" by an angle of  $\delta$  ( $vis_{\delta}(q)$ ), or "diminished" if their vision is "blocked" by reflex vertices by an angle of  $\delta$  ( $vis_{-\delta}(q)$ ). As such, the size of the minimum  $\delta$ -guarding set is defined as  $opt(P, \delta)$ . A polygon P has vision-stability  $\delta$  if the optimal number of enhanced guards is the same as the optimal number of diminished guards for that matter ( $opt(P, -\delta) = opt(P, \delta)$ ).



Figure 2.12: Enhanced  $(vis_{\delta}(q))$  and diminished  $(vis_{-\delta}(q))$  visibility regions of the one point q [13].

The core idea of Hengeveld and Miltzow's approach is to discretise the problem in order to improve the computation time. In this way, only a subset of potential guards and points to be guarded from Pwould be considered. In order to do so, they introduce the notions of a candidate guard set C and a witness set W. The role of the candidate guard set C is to discretise the guard searching space. As such, we can find an upper bound opt[C] on the final guard set. Thus C is "pretty close" on the actual optimum guard set opt. The candidate guard set C is then initialised with a subset of all the potential guards. The guards that are included should have their visibility regions overlap as little as possible. Similarly, the witness set W discretises the space of the guards' visibility regions. At initialisation time, W contains a subset of all vertices and faces of P. It is important to also include faces in W. In this way, we account for the fact that if a point q is included in a face f, then f sees at least as much as q. So, we can determine an approximation of how much more f sees when compared to g. Thus, by adding the face as a witness, we are creating a discretisation of the space containing a superset of possible guard positions.

Firstly, rays are shot from every reflex vertex such that the angle between any two rays is at most  $\delta$ , as observed in Figure 2.13, as taken from [13]. All the intersection points of the rays within P define the candidate guard set C, together with the faces they are part of. A witness guard set W is then defined by picking an arbitrary interior point of every face of P, together with all faces of P. In this way, the problem can be discretised in a way in which the solution to the discretised problem are also solutions to the continuous problem.



Figure 2.13: Shooting rays from all reflex Vertices with angles between them at most  $\delta$  [13].

#### 2.4.1 One-Shot Vision-Stable Algorithm

At first, an Integer Program (IP) is built and named the One-Shot Vision-Stable Algorithm. It computes the minimum number of point guards in polynomial time in a vision-stable polygon P with vision-stability  $\delta$  and r reflex vertices. The guards are encoded by binary integers, and linear equations and inequalities are used to encode visibility. Assuming that  $\delta$  is part of the input, the algorithm returns the optimal solution only if the vision-stability of P is larger or equal to  $\delta$ . The one-shot vision-stable algorithm is reliable, as it also verifies that its result is an optimal solution.

Let  $f_1$  be the optimisation function of the IP. The function  $f_1$  uses variables [[c]] for every candidate  $c \in C$ , and [[w]] for every witness  $w \in W$ . Its role is to count the number of guards that are used with the purpose of minimising this number. Additionally, it makes use of parameter factor  $1 + \varepsilon$  to ensure that vertex candidates are preferred over face candidates. By choosing  $\frac{1}{\varepsilon} = |C| + |W| + 1$ , we ensure that  $\varepsilon$  is sufficiently small. As such, the optimisation function  $f_1$  of the IP is

$$f_1 = \sum_{c \in \operatorname{vertex}(C)} [[c]] + (1 + \varepsilon) \sum_{c \in \operatorname{face}(C)} [[c]] + \varepsilon \sum_{w \in \operatorname{face}(W)} [[w]].$$

For every witness w, its visible set of candidates are denoted as vis(w), where *face* and *vertex* refer to whether the candidate is a face or a vertex, respectively.

Additionally, let vis(w) be the set of candidates that see witness  $w \in W$  completely. For every vertex-witness  $w \in vertex(W)$ , we formulate a constraint such that all  $w \in vertex(W)$  is seen:

$$\sum_{[[c]] \in vis(w)} c \ge 1, \forall w \in \operatorname{vertex}(W).$$

We similarly formulate a constraint for every face-witness  $w \in \text{face}(W)$ . However, we add [[w]] as a variable to relax it, so that the focus is placed on first seeing all vertex-witnesses

$$[[w]] + \sum_{[[c]] \in vis(w)} c \ge 1, \forall w \in face(W).$$

Nonetheless, the one-shot vision-stable algorithm proves to be too slow for practical reasons. The bottleneck is however not solving the IP in itself, but computing visibilities. Hence, an iterative vision-stable algorithm is devised, that is faster and retains similar performance guarantees.

#### 2.4.2 Iterative Vision-Stable Algorithm

Starting with the smaller sets of candidates C and witnesses W, another IP is deployed in order to find a minimum guard set  $S \subseteq C$  that sees all vertex-witnesses. This algorithm is then called the Iterative Vision Stable Algorithm, as it changes C and W in each iteration.

As the Iterative Vision Stable Algorithm is also named the Big IP, it also has an objective function  $f_2$ . The function  $f_2$  minimises the number of face-guards used in S and the number of unseen face-witnesses. If S contains only point guards and sees all face-witnesses, the algorithm reports the optimal solution by using the One-Shot IP. As such,  $f_2$  assures that all witnesses w are seen by their set of candidates vis(w), and enforces that only the number of guards s resulted from the One-Shot IP is used. The optimisation function with its constraints of the Big IP becomes

$$f_2 = \sum_{x \in splittable(W \cup C)} [[x]] \tag{1}$$

$$\sum_{[[c]]\in C} c = s, s \in \mathbb{Z}, s \text{ the number of used guards in the One-Shot IP}$$
(2)

$$\sum_{[[c]]\in vis(w)}^{n} c \ge 1 \tag{3}$$

$$1 - \left(\varepsilon \sum_{[[c]] \in vis(w)} c\right) \ge [[w]]. \tag{4}$$

If not all point guards are used, or not all face-witnesses are seen, the faces of P are split (discretised). Then the algorithm continues to the next iteration. The One-Shot IP is then used as soon as all faces are deemed unsplittable (they have reached a certain granularity  $\lambda$ ).

Inspired by the Iterative Algorithm, an FPT algorithm for the Art Gallery Problem can be created. In this case, the FPT algorithm would have as parameter of chord-visibility width. Given a chord c of P, we count the number n(c) of vertices visible from c. The chord-visibility width of P(cw(P)) is the maximum n(c) over all possible chords c. By using this measure, it is possible to describe the local complexity of a polygon: many synthetic arbitrary polygons have much smaller chord-visibility width than they have reflex vertices, and polygons constructed with hardness reductions have the chord-visibility width proportional to the total number of vertices.

#### 2.4.3 Experimental Results

The previously introduced algorithms are tested experimentally in terms of their running times and solution quality with respect to the input size, chord-visibility width and vision-stability of the input polygons P. The tests were run on 30 instances of arbitrary simple polygons of sizes 60, 100, 200 and 500 vertices. The algorithms have been implemented both with the theoretical performance guarantees  $\mathcal{T}$  (if no solution is found in  $\mathcal{T}$ , abort) and with the practical optimisation of critical witnesses.

When running the Iterative Algorithm with optimisations, reasonable practical results were found: all tested instances were solved to optimality, although the overall running time was not improved when compared to state-of-the-art algorithms. The median running time was also lower than the average running time. This could be accounted for by the sensitivity of the algorithm to the visionstability of a polygon. Thus, a few polygons that are hard to solve (have low vision-stability) outweigh the rest of the running times.

The Iterative Algorithm without optimisations was on the 60 vertices instance and with the theoretical limit  $\mathcal{T}$  as a safe guard. A solution was found in an efficient manner (within an hour) only for 25 out of 30 instances. For the rest of the 5 instances, the Big IP performed unnecessary splits, which negatively influenced the running times.

Additionally, the correlation between the granularity  $\lambda$  of the input polygon was tested for the iterative algorithm with performance guarantees  $\mathcal{T}$ . There appeared to be a strong correlation between the running times and  $\lambda$ . Thus, lower granularity implied a shorter running time.

In terms of the convergence to the optimal solution, the iterative algorithm with optimisations has been used to show the fast convergence in the first few iterations. The speed of the convergence slows down as the algorithm nears its end due to the increasing number of candidates and witnesses that are added to C and  $W^*$ , respectively, in the later iterations. Nonetheless, it appears that computing the visibility area of the polygon is still the bottleneck in the algorithm's performance.

### 3 Theory

This section investigates the theory aspects of solving the Art Gallery Problem using gradient descent. Namely, we explain what the optimisation function is in the context of the Art Gallery problem, and how it can be solved with gradient descent. We initially apply gradient descent to only one guard, then extend it to multiple guards.

#### 3.1 Guarding the Polygon with One Point

First, we explore how gradient descent is applied for the case that we want to guard the polygon using only one guard. As such, we tackle the main challenge of this section: defining the optimisation function of the Art Gallery Problem to be continuous. Then, we compute its gradient descent. The gradient descent computation is used then for defining how the new positions of a guard is computed.

#### 3.1.1 Gradient Descent

Let P be a polygon and  $g = (x, y) \in P$  a guard. We are interested in computing the best direction for moving g inside P such that the visibility area Vis(g) increases. That is, exploring what would be a better position g' to move g to such that g "sees more" of the polygon P.

We define  $f(g) = \operatorname{Area}(g)$  as the area seen by a guard g. Let  $\bigtriangledown \operatorname{Area}_r(g)$  be the local change in the area guarded by point g around a reflex vertex r seen by guard g, thus  $r \in Vis(g)$ . Given all reflex vertices i, the total (global) change in the area seen by g is summed up to  $\operatorname{Area}(g) = \sum_i \operatorname{Area}_i(g)$ . Figure 3.1 offers an example for this case for a polygon P and its reflex vertices  $r_1$  and  $r_2$ . The polygon P is guarded by g, and its position is modified to g' by a small change  $\partial y$  in its y-coordinate. The visibility areas of g are  $\operatorname{Area}_{r_1}$  and  $\operatorname{Area}_{r_2}$  around reflex vertices  $r_1$  and  $r_2$ , respectively. In this way, the total change in the visibility area of g is computed as  $\bigtriangledown \operatorname{Area}(g) = \bigtriangledown \operatorname{Area}_{r_1}(g) + \bigtriangledown \operatorname{Area}_{r_2}(g)$ .

Thus, we consider f(g) as the continuous objective function of the Art Gallery Problem. We then use gradient descent as a method to optimise the objective function f. We define below what the methodology of gradient descent consists of.



Figure 3.1: Global change in the area seen by g when moved by  $\partial y$  to a new position g'.

Let  $\nabla f$  be the gradient of f. The gradient then indicates the direction of the steepest descent for the objective function f(g). The learning rate (step size)  $\alpha$  is the size of the steps taken to reach the optimum. It is typically a small value, and it is evaluated and chosen based on the behaviour of the optimisation function.

After the gradient  $\nabla f$  is computed, we use it to calculate the new optimised position g' of guard g:

$$g' = g + \alpha \bigtriangledown f(g)$$

In later sections we experiment with various learning rates. As such, we explore how they influence the performance of our algorithm in relation to different test polygons.

#### 3.1.2 Computing the Gradient

Given that f is a function that describes the visibility area of a point g, we first need to define how its gradient is computed. We simplify the gradient computation without losing generality. As such, we rotate the plane with rotation matrix R, so that g and any reflex vertex r have the same x-coordinate. In this way, we only need to compute the gradient when we vary the y-coordinate. The computation of the gradient remains the same regardless of the rotation applied to the plane.

We use the notation  $\frac{\partial f}{\partial y}$  to denote the change in the visibility area f(g) when the plane is rotated and then the y-coordinate is modified by a small amount  $\partial y \to 0$ . In this way, we define

$$\nabla f(g) = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)^{\mathsf{T}} \tag{5}$$

to be the gradient of f given that P is guarded by a point g.

We now create a canonical geometric construction that allows us to further define and compute  $\nabla f$ . In this case, we consider the normalised length of the gradient as  $|| \nabla f(g)|| = 1$ . This canonical construction is displayed in Figure 3.2.



Figure 3.2: Canonical gradient construction for when the position of g is varied by a small amount  $\partial y$  around reflex vertex r to the new position g'.

Take a boundary line segment of P, r a reflex vertex inside P and g a guard whose optimal position we are interested in. The reflex vertex r is seen by g. Let  $\overline{gr} = a$  be known, and let  $\triangle rgg' = \triangle_2$ . Similarly, let b be the known distance between r and the polygon boundary in question.

Let  $\partial y$  be a tiny change in the *y*-coordinate of *g*. Let *g'* be the new position of *g* given the change  $\partial y$ . The point *g'* sees more around *r* on the polygon boundary. Let *L* be the new segment that *g'* sees on the polygon boundary. As such, let  $\Delta_1$  denote the increase in the visibility area of *g* when it moves to position *g'*:

$$\nabla \operatorname{Area}_r(g) = \Delta_1.$$
 (6)

Given that triangles  $\triangle_1$  and  $\triangle_2$  are square triangles, their areas is calculated as:

$$\operatorname{Area}_{\Delta_1}(g) = \frac{bL}{2},$$
$$\operatorname{Area}_{\Delta_2}(g) = \frac{a\partial y}{2}$$

Given that  $\overline{gg'}$  is parallel to polygon's boundary, we use Thales's Theorem [2] in triangles  $\triangle_1$ and  $\triangle_2$  to compute the length L:

$$\frac{\partial y}{L} = \frac{a}{b}$$
$$L = \frac{b\partial y}{a}.$$
(7)

So, the area of  $\triangle_1$  is computed:

$$\operatorname{Area}_{\Delta_1}(g) = \frac{Lb}{2}$$

$$\stackrel{(7)}{=} \frac{\frac{b\partial y}{a}b}{2}$$

$$\operatorname{Area}_{\Delta_1}(g) = \frac{b^2\partial y}{2a}.$$
(8)

We compute the gradient  $\bigtriangledown Area_r$  in a plane rotated by R for a point g and a reflex guard r seen by g as

$$R \bigtriangledown \operatorname{Area}_{r}(g) \stackrel{(5)}{=} \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)^{\mathsf{T}}$$
$$\stackrel{(6)}{=} \left(0, \frac{\operatorname{Area}_{\bigtriangleup_{1}}}{\partial y}\right)^{\mathsf{T}}$$
$$R \bigtriangledown \operatorname{Area}_{r}(g) \stackrel{(8)}{=} \left(0, \frac{b^{2}}{2a}\right)^{\mathsf{T}}.$$
(9)

Therefore, for all the reflex vertices r guard g sees, the total gradient  $\nabla f$  becomes the sum of all the partial gradients  $\nabla f_r$  as

$$\nabla f(g) = \sum_{i \in R(g)} \nabla \operatorname{Area}_i(g), \tag{10}$$
$$R(g) = \{ \text{reflex vertices of } P \text{ seen by } g \}.$$

#### 3.1.3 Computing the New Guard's Position

We use the coordinates of the gradient  $\nabla f$  to compute the movement direction of the guard g given all the reflex vertices from P seen by g. In order to do so, we use the construction depicted in Figure 3.3.



Figure 3.3: Computing the new position g' of guard g around reflex vertex r based on the gradient  $\nabla f$ .

Let  $\vec{v}$  be the vector corresponding to the direction of movement from guard g to a reflex vertex r, such that  $\vec{v} = (r - g) = (x, y)^{\intercal}$ , with norm  $||\vec{v}|| = a$ . So,  $||\frac{\vec{v}}{a}|| = 1$ .

Let  $\vec{v}^{\perp} = (g'-g)$  be the vector corresponding to the direction of movement from guard g to its new position g'. Vector  $\vec{v}^{\perp}$  is orthogonal to  $\vec{v}$ , in the same direction as  $\nabla f$ , such that  $||\vec{v}|| = ||\vec{v}^{\perp}|| = a$ . We use the coordinates of  $\vec{v}^{\perp}$  to compute the coordinates of  $\nabla f$  and thus the direction in which gneeds to move.

The coordinates of  $\vec{v}^{\perp}$  are computed using the construction from Figure 3.4. Since  $\vec{v}^{\perp} \perp \vec{v}$ , and g and r are on the right-hand side of  $\vec{v}^{\perp}$ , the coordinates of  $\vec{v}^{\perp}$  are rotated by  $-90^{\circ}$  so that  $\vec{v}^{\perp} = (-y, x)^{\mathsf{T}}$ . Analogously for the case when g and r are rotated by  $180^{\circ}$  to the left-hand side of  $\vec{v}^{\perp}$ , the coordinates of  $\vec{v}^{\perp}$  are rotated by  $90^{\circ}$  to  $\vec{v}^{\perp} = (-x, y)^{\mathsf{T}}$ .



Figure 3.4: Computing the coordinates of  $\vec{v}^{\perp}$  given the guard g and the reflex vertex r = (x, y).

We know that the norm of the total gradient  $\nabla f$  is

$$|| \nabla f(g)|| \stackrel{(10)}{=} || \sum \left(0, \frac{b^2}{2a}\right)^{\mathsf{T}} ||$$
$$\stackrel{(9)}{=} \sqrt{\left(\frac{b^2}{2a}\right)^2}$$
$$|| \nabla f(g)|| = \frac{b^2}{2a}.$$
(11)

Since  $\nabla f$  has the same direction as  $\vec{v}^{\perp}$ , we wish to normalise it with  $\frac{1}{a}$  (the norm of  $\vec{v}^{\perp}$ ). Therefore, the gradient for guard g and one reflex vertex  $r \in Vis(g)$  is computed as

$$\nabla f(g) = \vec{v}^{\perp} \frac{b^2}{2a} \frac{1}{a}.$$
(12)

As mentioned before, the total gradient for guard g and all the reflex vertices r the guard sees is

$$\nabla f(g) = \sum_{r \in R(g)} \nabla \operatorname{Area}_r,$$
$$R(g) = \{ \text{reflex vertices of } P \text{ seen by } g \}.$$

The new position g' of guard g based on all the reflex vertices it sees is:

$$g' = g + \alpha \bigtriangledown f(g). \tag{13}$$

#### 3.2 Guarding the Polygon with Multiple Guards

In this section we investigate how to generalise the computation of gradient descent to multiple guards. Previously we defined the gradient descent of one guard in order to compute its movement towards optimality. We extend this computation to multiple guards. Namely, we define how the gradient descent changes when the visibility region of one guard is seen by other guards as well.

#### 3.2.1 Computing the Gradient for Two Guards

Let point  $g_1$  be the guard we have previously computed the gradient for. Let point  $g_2$  be another guard in the polygon. The position of  $g_2$  is optimised around reflex vertex  $r_2$ , as seen in Figure 3.5. Guard  $g_2$  sees a part of the visibility region already seen by  $g_1$ . Let the shared seen region be  $\Delta_2$ , and the region that is only visible by  $g_1$  be  $\Delta_1$ .



Figure 3.5: Canonical gradient construction for when the visible area of  $g'_1$  is also seen by guard  $g_2$ .

The visibility regions of guards  $g_1$  and  $g_2$  are computed. Let p be the intersection point between them. Let d be the point seen by  $g_1$  on the polygon boundary. Point p divides the segment seen by  $g_1$ behind the reflex vertex  $r_1$  into two subsegments  $\overline{r_1p_1}$  and  $\overline{p_1p_2}$ . The lengths of the subsegments are  $b_1$ and  $b_2$ , respectively, with  $b_1 + b_2 = b$ . Recall that b is the distance between the reflex vertex  $r_1$  and the polygon boundary. Thus,  $b_2$  corresponds to the length of the shared seen segment.

As before in equation 8, the area seen by  $g_1$  is

$$\operatorname{Area}_{\triangle_1 + \triangle_2}(g_1) = (b_1 + b_2)^2 \frac{\partial y}{2a}$$

$$\operatorname{Area}_{\Delta_1}(g_1) = b_1^2 \frac{\partial y}{2a}$$

So, we do not need to take  $\triangle_2$  into account when computing the gradient of  $g_1$ .

Now we can also compute the shared area seen  $\triangle_2$  as:

$$\operatorname{Area}_{\Delta_2}(g_1) = \operatorname{Area}_{\Delta_1 + \Delta_2}(g_1) - \operatorname{Area}_{\Delta_1}(g_1)$$
$$= (b_1 + b_2)^2 \frac{\partial y}{2a} - b_1^2 \frac{\partial y}{2a}$$
$$\operatorname{Area}_{\Delta_2}(g_1) = \left[ (b_1 + b_2)^2 - b_1^2 \right] \frac{\partial y}{2a}.$$
(14)

We use thus equation 14 in the next section to introduce a generalisation for the computation for the exclusive area seen by  $g_1$ .

#### 3.2.2 Computing the Gradient for Multiple Guards

We generalise the gradient computation to m guards and m intersection points. Let g be the guard whose position we wish to optimise around reflex vertex r as before. Let the areas highlighted in blue in Figure 3.6 be the parts of the visibility region of g that are also seen by the other m-1 guards. Let the visibility region of g intersect other guards' visibility regions in m intersection points  $p_1, p_2, ..., p_m$ . Let the distance from the reflex vertex r to the intersection points be  $b_1, b_2, ..., b_m, b$ .



Figure 3.6: Canonical gradient construction for where the pink parts of the visible area of g are also seen by other guards. The grey polygons  $\triangle_1, \triangle_3, ..., \triangle_{m-1}$  are the areas that are exclusively seen by g.

We generalise equation 14. Namely, we need to subtract the areas that are seen by other guards from the total area seen by guard g.

We start to move in the direction from the polygon boundary to reflex vertex r. We know the intersection points  $p_{m-1}..., p_3, p_1$  of the shared seen regions with the visibility region of g. Thus, we compute the distances  $\overline{rp_{m-1}}, ..., \overline{rp_3}, \overline{rp_1}$  between the reflex vertex r and the beginning of the shared

visibility regions. Analogously, we compute the distances  $\overline{rp_m}, ..., \overline{rp_4}, \overline{rp_2}$  between the reflex vertex r and the end of the shared visibility regions.

The areas of polygons  $\triangle_1, \triangle_2, ..., \triangle_m$  do not grow linearly. For this reason, we cannot simply subtract the sum of the shared areas from the total area seen by g.

An explanation about why this is the case is given in Figure 3.7. Take overlapping polygons  $s_1, s_2, s_3, s_4$  with areas  $\text{Area}_{s_1} = 2$ ,  $\text{Area}_{s_2} = 3$ ,  $\text{Area}_{s_3} = 4$ ,  $\text{Area}_{s_4} = 5$ . We want to compute the area of polygons  $s_1$  and  $s_2$ . However, it is incorrect to simply add their areas together as

$$Area_{s_1+s_3} = 2 + 4 = 6.$$

as  $s_2$  is overlapping in between them. Instead, from the total area

$$\operatorname{Area}_{s_1+s_2+s_3+s_4} = \operatorname{Area}_{s_4} = 5$$

we sequentially subtract the areas we are not interested in  $(Area_{s_2}, Area_{s_4})$  and add the ones we are interested in  $(Area_{s_1}, Area_{s_3})$ . Namely,

 $Area_{s_1+s_3} = Area_{s_1+s_2+s_3+s_4} - Area_{s_4} + Area_{s_3} - Area_{s_2} + Area_{s_1} = 5 - 5 + 4 - 3 + 2 = 3.$ 



Figure 3.7: Example for computing the area of overlapping polygons with areas  $\operatorname{Area}_{s_1} = 2$ ,  $\operatorname{Area}_{s_2} = 3$ ,  $\operatorname{Area}_{s_3} = 4$ ,  $\operatorname{Area}_{s_4} = 5$ . It is incorrect to compute  $\operatorname{Area}_{s_1+s_3} = 2 + 4 = 6$ . Instead,  $\operatorname{Area}_{s_1+s_3} = \operatorname{Area}_{s_1+s_2+s_3+s_4} - \operatorname{Area}_{s_4} + \operatorname{Area}_{s_3} - \operatorname{Area}_{s_2} + \operatorname{Area}_{s_1} = 5 - 5 + 4 - 3 + 2 = 3$ .

Similarly to Figure 3.7, we compute the area seen exclusively by g. From the total Area $_{\Delta_1+\Delta_2,...,\Delta_m}$  seen by g we need to alternatively subtract the shared areas Area $_{\Delta_m},...,$  Area $_{\Delta_4}$ , Area $_{\Delta_2}$ , then add the exclusively seen areas Area $_{\Delta_{m-1}},...,$  Area $_{\Delta_3}$ , Area $_{\Delta_1}$ .

This results in the fact that the area seen by g exclusively is computed as

$$\begin{aligned} \operatorname{Area}_{\triangle_1 + \triangle_3 + \ldots + \triangle_{m-1}}(g) &= \operatorname{Area}_{\triangle_1 + \triangle_2 + \ldots + \triangle_m}(g) \\ &- \operatorname{Area}_{\triangle_{m-1}}(g) + \operatorname{Area}_{\triangle_{m-2}}(g) \\ &- \ldots \\ &- \operatorname{Area}_{\triangle_2}(g) + \operatorname{Area}_{\triangle_1}(g) \\ &= \left(b^2 - b_m^2 + b_{(m-1)}^2 - \ldots - b_2^2 + b_1^2\right) \frac{\partial y}{2a} \end{aligned}$$

Alternatively, if  $\triangle_1$  is first seen by other guards (so  $b_{11} = 0$ ), then all the signs are flipped. This claim is also supported by the intuition of Figure 3.7. If we are interested in the areas of  $s_2$  and  $s_4$ , then it is incorrect to compute them as

$$Area_{s_2+s_4} = 3 + 5 = 8$$

This is because  $s_1$  and  $s_3$  are overlapping in between them. Instead,

$$Area_{s_2+s_4} = Area_{s_1+s_2+s_3+s_4} + Area_{s_4} - Area_{s_3} + Area_{s_2} - Area_{s_1} = 5 + 5 - 4 + 3 - 2 = 7.5$$

#### 3.3 Momentum

This section introduces an extension to the regular gradient descent algorithm: momentum. Momentum builds upon the idea of considering the past states of the gradient descent computation. In this way, the past states create "inertia" to the newly computed gradient state. This results in the movement trajectory to be smoother. This is especially advantageous for the gradient descent computation. The reason is that the gradient descent computations can be noisy. So some of them are extremely large, while others are insignificant. Momentum addresses this non-uniformity [12].

An example for the benefits of using momentum are found in Figure 3.8. Subfigure 3.8a displays a gradient descent computation without momentum. Initially, the jumps are large and become smaller as the minimum is approached. In Subfigure 3.8b these movements become smoothened out. The gradient descent steps become thus more even. This allows the overall trajectory to be less noisy and chaotic.



Figure 3.8: Illustrative examples of gradient descent trajectories with and without momentum.

So, the position  $g_i$  at iteration *i* of a guard *g* is not calculated anymore based only on the current computation of gradient descent. Instead, it also takes into account past values of gradient descent. In order to decide how much past states influence the value of the current position, we take each gradient value with a weight.

Let  $M_i(g_i)$  be the momentum for a guard at iteration *i*. Let  $\nabla f_i(g_i)$  be the computation of the gradient descent in iteration *i*. Let the weight of a gradient descent value be a hyperparameter  $\gamma$ . As such, we take into account the value of the previous gradient  $\nabla f_{i-1}(g_{i-1})$  with weight  $\gamma$ . However, we still want the current gradient  $\nabla f_i(g_i)$  to influence the guard's movement. This can happen with a weight of  $1 - \gamma$ . So, the computation for the momentum at iteration *i* becomes

$$M_i(g_i) = \gamma M_{i-1}(g_{i-1}) + (1-\gamma) \bigtriangledown f_i(g_i).$$

We start with

$$M_0(g_0) = (1 - \gamma) \bigtriangledown f_0(g_0)$$

We then check for correctness how the next iterations of the momentum are carried out:

$$\begin{split} M_{i}(g) &= \gamma M_{i-1}(g_{i-1}) + (1-\gamma) \bigtriangledown f_{i}(g_{i}) \\ &= \gamma (\gamma M_{i-2}(g_{i-2}) + (1-\gamma) \bigtriangledown f_{i-1}(g_{i-1})) + (1-\gamma \bigtriangledown f_{i}(g_{i})) \\ &= \gamma^{2} M_{i-2}(g_{i-2}) + (1-\gamma)(\gamma \bigtriangledown f_{i-1}(g_{i-1}) + \bigtriangledown f_{i}(g_{i})) \\ &= \gamma^{3} M_{i-3}(g_{i-3}) + (1-\gamma)(\gamma^{2} \bigtriangledown f_{i-2}(g_{i-2}) + \gamma \bigtriangledown f_{i-1}(g_{i-1}) + \bigtriangledown f_{i}(g_{i})) \\ &\dots \end{split}$$

In this way, our momentum computation exponentially decreases the influence of a past state over the guard's current movement. So, the previous value of the momentum exerts its inertia with  $\gamma$ influence over the new value of the momentum.

Similarly, the new position  $g_i$  of guard g with learning rate  $\alpha$  is now calculated as

$$g_i = g_{i-1} + \alpha M_i(g_{i-1}).$$

#### 3.4 Reflex Vertex Pull

This section presents an additional idea to our gradient descent algorithm: pulling a guard towards reflex vertices. The pull strategy makes use of the second derivative of our optimisation function f(g) = Area(g).

As mentioned previously, we are using the first derivative to compute the gradient of a guard's movement. That is, we use the first derivative of f(g) to find the direction of a guard's movement that increases its seen area. Additionally, we also make use of the second derivative of f(g) to explore what the rate of change in the area increase is.

We can achieve a more rapid increase in the seen area if we move a guard closer to a reflex vertex. The closer a guard is moved to a reflex vertex, the larger the increase in the area seen past the vertex. If a guard is moved directly on a reflex vertex, then the area seen past the vertex is maximised.

We explore how the pull towards a reflex vertex r can be computed, with an example in Figure 3.9. Let the guard g = (0,0), the reflex vertex r = (2,0) and the polygon boundary intersection point be (5,0). So, the distances a = 2 and b = 3 are known. Let  $h_r$  be the pull in question. Let  $g'_y$ be the new coordinates of guard g when only the gradient is taken into account. In that case, we are only interested in the small movement  $\partial y$ . Analogously, let  $g'_x$  be the new coordinate of guard gwhen g moves by a small amount  $\partial x$  towards the reflex vertex. As previously, and without losing generality, assume that g and r have the same x-coordinate by rotating the plane R. Combining the two movements together results in g' being the final position of g.

The gradient is computed as

$$\nabla f(g) \stackrel{(12)}{=} (0, \frac{b^2}{2a^2})^{\mathsf{T}} = (0, \frac{9}{8})^{\mathsf{T}} = (0, 1.125)^{\mathsf{T}}.$$

The pull is computed as

$$h_r(g) \stackrel{(15)}{=} (\frac{b^2}{2a^3}, 0)^{\mathsf{T}} = (\frac{9}{16}, 0)^{\mathsf{T}} = (0.625, 0)^{\mathsf{T}}.$$

So, the new position g' of the guard g given the learning rate  $\alpha = 0.3$  becomes

 $g' = g + \alpha(\bigtriangledown f(g) + h(g)) = 0.3((0, 1.125)^{\mathsf{T}} + (0.625, 0)^{\mathsf{T}}) = (0.1875, 0.3375).$ 



Figure 3.9: Computing the movements of the guard g based on both the gradient and the pull towards reflex vertex r. The new position of the guard with learning rate  $\alpha = 0.3$  becomes g' = (0.1875, 0.3375).

We explain how we arrived at the pull computation. The pull is the second derivative of the norm of the gradient, so

$$h_r(g) = \bigtriangledown || \bigtriangledown f_r(g)|| = \left(\frac{\partial \bigtriangledown f(g)}{\partial x}, \frac{\partial \bigtriangledown f(g)}{\partial y}\right)^{\mathsf{T}} = \left(\frac{\partial \bigtriangledown f(g)}{\partial x}, 0\right)^{\mathsf{T}}$$

We calculate  $h_r(g)$  as follows: the Euclidean norm of the gradient is  $|| \bigtriangledown f(g) || = \frac{b^2}{2a}(11)$ , so the norm of the second gradient is

$$||h_r(g)|| = || \bigtriangledown || \bigtriangledown f(g)||||$$
$$= || \bigtriangledown (\frac{b^2}{2a})||$$
$$= \frac{b^2}{2a}\frac{d}{da}$$
$$= b^2\frac{1}{2a}\frac{d}{da}$$
$$||h_r(g)|| = -b^2\frac{1}{2a^2}.$$

Note that we are using the norm approximation heuristic

$$\nabla || \bigtriangledown f(g)|| \approx \bigtriangledown \sum_{i \in R(g)} ||f_i(g)||,$$
$$R(g) = \{ \text{reflex vertices of } P \text{ seen by } g \}.$$

The reason for this choice is that computing the norm of the partial gradients, summing them and then computing their gradient is much easier than computing the gradient of their sum

$$\nabla || \nabla f|| = \nabla || \sum_{i \in R(g)} f_i||,$$

$$R(g) = \{ \text{reflex vertices of } P \text{ seen by } g \}.$$

We are aware of the fallacies of this approximation. Take the opposing unit vectors  $a_1 = (1,0)^{\intercal}$ ,  $a_2 = (-1,0)^{\intercal}$ . The sum of their norms  $\sum_i ||a_i|| = 2$ , whereas the norm of their sum is  $||\sum_i a_i|| = 0$ . Clearly, they are not the same. Nonetheless, we still consider this approximation due to computation speed improvements and easiness of use.

The pull  $h_r(g)$  is directed towards the reflex vertex r. So, it has the same orientation as vector  $\vec{v} = (r-g)$ . We normalise  $h_r(g)$  with the norm of vector  $\vec{v}$ . Its norm is  $||\vec{v}|| = \frac{1}{q}$ . So,

$$h_r(g) = \vec{v} \frac{-b^2}{2a^2} \frac{1}{a} = \vec{v} \frac{-b^2}{2a^3}.$$
(15)

Let h(g) be the total pull for guard g. As for the gradient, the total pull for guard g and all reflex vertices r the guard sees is

$$\begin{split} h(g) &= \sum_{r \in R(g)} h_r(g), \\ R(g) &= \{ \text{reflex vertices of } P \text{ seen by } g \}. \end{split}$$

Therefore, the movement of a guard g to the new position g' takes both the gradient and the pull into account:

$$g' = g + \alpha(\nabla f(g) + h(g)).$$

Additionally, we choose how much influence the pull itself has in the movement of the guard by adding a hyperparameter  $\beta$ :

$$g' = g + \alpha(\nabla f(g) + \beta h(g)).$$

#### 3.4.1 Pull onto the Reflex Vertex

We have now created a heuristic for pulling a guard closer to a reflex vertex based on the increase in the seen area behind the reflex vertex. It could happen however that the pull towards the reflex vertex is forceful. In this case, the guard could be moved past the reflex vertex, in between the reflex vertex and the polygon boundary. Although the area seen behind the reflex vertex would be maximised, the guard would "unsee" (parts of) the area seen from its initial position g. In order to address this particular edge-case, the guard is placed on the reflex vertex when the pull is strong enough.

This section thus expands on a procedure that decides when a guard is best placed on a reflex vertex based on its pull towards it.

Firstly, we must define the condition of what a "strong enough" pull means. Such a pull would move the guard g "close enough" to the reflex vertex. Hence, we need to define what a minimum the distance between the new guard position g' and reflex vertex r would be. The most straightforward way to do so would be to assign a hyperparameter to the distance between g and r. Recall that  $||\overline{gr}|| = a$ . For now, let  $\frac{2}{3}$  be the factor of closeness of g' to r. This means, that when the guard is more than two thirds of the distance close to the reflex vertex, then the guard is moved onto the reflex vertex:

$$||h_r(g)|| > \frac{2}{3}a.$$
 (16)

Next, we need to account for the case where the guard is close to multiple reflex vertices. Namely, we consider the specific edge-case when a guard is in between two reflex vertices, illustrated in Figure 3.10. Let D be the minimum distance in between any two reflex vertices of polygon P:

$$\begin{split} D &= \min_{q \neq r} \text{ distance}(q, r), \\ &\forall q, r \in P \text{ reflex vertices.} \end{split}$$

Let g be a guard in between two reflex vertices  $r_1$  and  $r_2$  at distance D from each other. Let the pulls  $h_{r_1}(g)$  and  $h_{r_2}(g)$  be strong towards  $r_1$  and  $r_2$ , respectively. However, because they are opposites in directions, they cancel each other out, resulting in g possibly not changing its position at all.



Figure 3.10: Guard g is in between two reflex vertices  $r_1$  and  $r_2$ . Because they are collinear, g is equidistant from them and the pulls are equally strong  $||h_{r_1}(g) = h_{r_2}(g)$ , they cancel each other out.

We now introduce another condition for pulling the guard towards the closer reflex vertex when pulls are cancelling each other out. If a guard g is not equidistantly placed in between two reflex vertices, we are interested in computing what the closer reflex vertex is. So, if g is closer to a reflex vertex than half of the minimum distance in between any two reflex vertices  $\frac{D}{2}$ , then we consider it close enough if:

$$a < \frac{D}{2}.$$

In this case we choose to still move towards one of the reflex vertices and make progress.

An example of moving on top of a reflex vertex is found in Figure 3.11. Let the guard g = (1, 0), the reflex vertex r = (2, 0) and the polygon boundary intersection point be (5, 0). So, the distances a = 1 and b = 3 are known. Let  $h_r(g)$  be the pull in question. The gradient is computed as

$$\nabla f(g) \stackrel{(12)}{=} (0, \frac{b^2}{2a^2})^{\mathsf{T}} = (0, \frac{9}{2})^{\mathsf{T}} = (0, 4.5)^{\mathsf{T}}.$$

The pull is computed as

$$h_r(g) \stackrel{(15)}{=} (\frac{b^2}{2a^3}, 0)^{\mathsf{T}} = (\frac{9}{2}, 0)^{\mathsf{T}} = (4.5, 0)^{\mathsf{T}}$$

The guard is close enough to r to be moved onto it (16):

$$||h_r(g)|| > \frac{2}{3}a \iff 4.5 > \frac{2}{3}.$$

So the new coordinate of the guard is g' = r = (2, 0). they Note that Figure 3.11 displays how the closer the guard to a reflex vertex is, the stronger its pull. This comes in contrast with Figure 3.9. Given the same reflex vertex r and polygon boundary coordinates, the guard g = (0, 0) was farther away, and its pull  $h_r(g) = (0.625, 0)^{\mathsf{T}}$  was thus not as strong. Hence, we only want to place guards on reflex vertices when their pull is strong enough.



Figure 3.11: Computing the movements of the guard g based on both the gradient and the pull towards reflex vertex r. The guard g is close enough to r (16), so g is placed on top of the reflex vertex r.

#### 3.4.2 Pull Capping

Nonetheless, it can happen that the pull towards a reflex vertex is significantly larger in comparison to the gradient and the momentum. In that case, if the guard is not pulled onto the reflex vertex, it would at least have a huge jump towards the reflex vertex.

We want to smoothen out possibly erratic movement jumps. Just like in the case of momentum, we want our pull to be smoothened out when it is "suddenly too large". We define what "suddenly too large" is based on the momentum. If the pull is larger than a factor of  $\mu$  than the momentum and a constant c, then we cap it at that value. So, at step i, the momentum  $h_i(g_i)$  is capped as:

if 
$$||h_i(g_i)|| > \mu||M_i(g_i)|| + c$$
, then  
 $h_i(g_i) = h_i(g_i) \frac{\mu||M_i(g_i)||}{||h_i(g_i)||}.$ 

The factor  $\mu$  becomes thus a hyperparameter to experiment with.

#### 3.5 Reflex Area

An additional heuristic is introduced in this section: the concept of *reflex area*. The reflex area heuristic counteracts a specific edge-case: a guard placed on a reflex vertex would move behind the reflex vertex. As such, it would stop seeing the area that it was seeing before. We want to prevent this loss in area seen. So, we are restricting the movement of the guard away from the reflex vertex such that it continues seeing the same areas it was seeing from the reflex vertex.

This case is illustrated in Figure 3.12. In Subfigure 3.12a, guard g starts moving with pull  $h_r(g)$  towards reflex vertex r. The pull is strong enough to place g on r in Subfigure 3.12b. In this case, g sees everywhere around r. However, the new gradient  $\nabla f(g)$  of g moves it past r in Subfigure 3.12c. The initially seen area before r is now not completely seen anymore by g.

Let rr' and rr'' be the extensions to the polygon boundary segments whose intersection is the reflex vertex r. We call *reflex area* the area between line segments rr' and rr'' that is inside the polygon. Subfigure 3.12d draws the reflex area more closely. If a guard g has to move outside of it, we project its new position g' onto the closest reflex line (in this case, rr'). If g has to move inside the reflex area, it can do so unaffectedly. In this way, we maintain the gained property of a guard seeing everything around a reflex vertex while still allowing it to move away from the reflex vertex.



(a) Guard g starts moving with pull  $h_r(g)$  towards reflex vertex r. Guard g cannot see past r.



(c) Guard g has moved past reflex vertex r. Guard g cannot see anymore before r.



(b) Guard g is placed on reflex vertex r. Guard g sees both before r and past it.



(d) The movement of g is restricted inside the reflex area.

Figure 3.12: Guard g moves towards and on the reflex vertex r. Eventually, its newly computed position is away from r and the reflex area. This results in g not seeing the initial area anymore. So, its movement is restricted to the reflex area  $\triangle rr'r''$ . If g needs to move outside the reflex area, its new position is projected on the closest reflex line.

#### 3.6 Line Search

We introduce in this section another extension to the regular gradient descent algorithm: line search [21]. Given a step size, its aim is to search for the best solution on the gradient descent line.

Figure 3.13 illustrates an example for this extension. Take guard g and reflex vertex r. Let  $\frac{1}{x}$  be the starting search factor, and s the step size factor. Recall that  $M_i(g)$  is the optimal movement direction for a guard at iteration i. As such, line search computes the optimal guard position  $g_{i+1} = g_i + \frac{s^i}{x}M_i(g_i)$ , with t the best step power for the guard from  $\{\frac{s^0}{xM_i}, \frac{s}{x}M_i(g_i), \frac{s^2}{x}M_i(g_i), \ldots\}$ . The dashed line represents the direction of gradient descent, as computed based on the gradient and reflex vertex pull computations. Given a step size s, 3 different possible new guard positions are identified:  $g'_{i_1} = g_i + \frac{1}{x}M_i(g_i), g'_{i_2} = g_i + \frac{s}{x}M_i(g_i), g'_{i_3} = g_i + \frac{s^2}{x}M_i(g_i)$ . The optimal guard is the one who has the largest increase in the area seen. If the global area is not increased, a best solution is also chosen if a guard is gaining visibility to a previously unseen part of the polygon, regardless of the area increase.

It is worth mentioning that Line Search deems the use of a learning rate obsolete. Because the search factor finds the optimal position on the direction line, it acts thus as a learning rate.

#### P boundary



Figure 3.13: The gradient descent line gives the direction of searching for the optimal guard position. Given step sizes  $\frac{1}{x}, \frac{s}{x}, \frac{s^2}{x}$ , there are 3 possible guard positioning on the dashed gradient descent line:  $g'_{i_1}, g'_{i_2}, g'_{i_3}$ . The best of them is chosen based on the size of the newly visible area seen.

#### 3.7 Angle Behind Reflex Vertex

We propose in this section a heuristic that further fine-tunes the factor with which a guard's movement is influenced by a reflex vertex. Currently, we only take into account the distance b between the reflex vertex and the polygon boundary. Intuitively, the unseen area behind the reflex vertex should also play a role in the computation of the gradient: guards should be drawn faster to larger areas. In order to do so, we take into account the normalised value of angle  $\theta$  behind the reflex vertex.

A visualisation for this heuristic is found in Figure 3.14. The pull of the guard g towards the reflex vertex r is influenced by the normalised value of  $\theta$  as follows:

$$g' = g + \left(\frac{\theta}{2\pi} + c\right)(\bigtriangledown f(g) + \beta h(g)).$$

We additionally add a constant c to account for small angles. For example, if the normalised value of  $\theta$  is close to  $10^{-5}$ , then  $c = 10^{-2}$ . Then the guard still has a significant move towards the unseen area, no matter how small it is. In this way, smaller areas are not overlooked.



Figure 3.14: The normalised angle  $\mu$  behind the reflex vertex r takes into account the factor with which the guard g is drawn to r.

#### 3.8 Hidden Movement

This section tackles a computation speed-up for the case in which guards have a movement vector of 0 (they don't move). This can happen when the area seen by a guard is already completely seen by other guards. We consider that having a movement vector of 0 is detrimental to the progress of the algorithm. Note that we call *movement vector* the sum of all heuristics that apply to computing the new position of a guard. The reason behind this is that it is unlikely that a guard's optimal position has been found when its movement vector is 0. Hence, we would like every guard to move, no matter how little.

In order to allow guards to still move when their movement vector is 0, we deployed the *hidden* gradient heuristic. This heuristic is based on the fact that we allow guards whose movement vector is 0 to still move with a newly computed "hidden" movement vector. So, if there are guards whose movement vector is 0, we recompute their movement vector by not taking into account the area seen by the guards who have a non-zero movement vector.

An example of this approach is found in Figure 3.15. Let  $G = \{g_1, g_2, g_3\}$  be the complete set of guards. Initially in Subfigure 3.15a, only  $g_1$  has a non-zero movement vector. The visibility regions of  $g_2$  and  $g_3$  are overlapping with that of  $g_1$ , so their movement vectors are 0. Let  $G_0 = \{g_1\}$  be the set of guards who at step 0 have a non-zero movement vector. In this case, only  $g_1$  Then, Subfigure 3.15b displays the remaining set  $G \setminus G_0$  of guards with a zeroed movement vector. The visibility area of guard  $g_1$  overlaps with  $g_2$ , so only guard  $g_2$  has a non-zero movement vector. So,  $g_2$  is part of the set  $G_1 = \{g_2\}$  of guards who at step 1 have a non-zero movement vector. Lastly, we compute the non-zero movement vector for guard  $g_3$  in Subfigure 3.15c. The guard  $g_3$  can now be part of the set  $G_2 = \{g_3\}$  which contains the guards who at step 2 have a non-zero movement vector. In Subfigure 3.15d all the guards have been moved to their new positions  $g'_1, g'_2, g'_3$ , respectively. So the movement vectors have been computed as  $G = G_0 \cup G_1 \cup G_2$ , where  $G_1$  and  $G_2$  contain the guards with hidden movements.



(a) Guard  $g_2$  and  $g_3$  have a movement vector of 0, because guard  $g_1$  sees all the areas they see. So, the set of guards with a non-zero movement vector is  $G_0 = \{g_1\}$ .



(c) Guard  $g_2$  had a non-zero movement vector, so we compute the movement vector of guard  $g_3$ without  $g_2$ . Guard  $g_3$  has a non-zero movement vector. So, the set of guards with a non-zero movement vector is  $G_2 = \{g_3\}$ .



(b) Guard  $g_1$  had a non-zero movement vector, so we compute the movement vectors of guards  $g_2$  and  $g_3$  without  $g_1$ . Guard  $g_2$  sees everything that  $g_3$  sees, so  $g_3$  has movement vector 0. So, the set of guards with a non-zero movement vector is  $G_1 = \{g_2\}.$ 



(d) The new positions  $g'_1, g'_2, g'_3$  of the guards such that  $G = G_0 \cup G_1 \cup G_2$ .

Figure 3.15: Example of hidden movement computation for guard set  $G = \{g_1, g_2, g_3\}$  in a corridorlike polygon. The visibility areas of each of the guards  $g_1, g_2$ , and  $g_3$  are shown in purple, green and orange, respectively.

The hidden movement heuristic can be generalised. Let G be the complete set of guards. For each step i, let  $G_i = G \setminus G_{i-1}$  be the set of guards with a non-zero movement vector after removing the guards with a movement vector from step i - 1. The set  $G_0$  is the set of guards with a non-zero movement vector before removal of any guards. At the end, the union of all subsets  $G_i$  comprises the set of all guards  $G = G_0 \cup G_1 \cup \dots$  In this way, all guards have a non-zero movement vector and make progress.

#### 3.9 Greedy Initialisation

In this section we introduce another heuristic for our algorithm: *greedy initialisation*. This heuristic sequentially places guards at starting positions in areas that are unseen by other guards. In this way, the algorithm has a head start with a larger covered area than when guards are placed arbitrarily anywhere in the polygon.

Figure 3.16 offers an example of a greedy initialisation. The first guard  $g_1$  is arbitrarily placed inside polygon P as shown in Subfigure 3.16a. Then, guard  $g_2$  is arbitrarily placed in an unseen part of P, as displayed in Subfigure 3.16b. The algorithm continues then with the optimisation of the guards' positions.



(a) Guard  $g_1$  has been arbitrarily placed inside the polygon P.



(b) Guard  $g_2$  has been arbitrarily placed inside the polygon P, outside the visibility region of guard  $g_1$ .

Figure 3.16: Greedy Initialisation for guards  $g_1$  and  $g_2$  inside polygon P. The visibility regions of the guards are displayed in orange and purple, respectively. In this way, the algorithm gains a head start for optimising the positions of the guards.

Nonetheless, placing guards arbitrarily inside unseen areas of the polygon deemed to be too difficult in CGAL given the time constraints of this thesis. For this reason, we are placing the guards in the middle of the leftmost segment of every unseen area. We are aware of the limitations of this deterministic approach.

## 4 Practice

This section provides the implementation details of the theory part (Section 3). Then, we introduce our experimental setup. The experiments include basic test polygons to showcase the execution of the program. In the experimental setup we discover the benefits of each heuristic. We then analyse how the program scales in the context of the comb polygon. Lastly, we mention the importance of the hyperparameter values used.

The algorithm is implemented in C++ and makes use of the CGAL library [7]. The project is publicly found at https://github.com/geo-j/master-thesis. It contains 22 code files and 1873 lines of code (excluding comments and empty lines). From the CGAL library we are using the Triangular Expansion Visibility algorithm [6] to compute the visibility polygons.

The program takes the segments of a boundary of a simple polygon, and the initial coordinates of the guards as input. For each iteration, the new position of the guards, the area seen by each guard and their movement direction details (momentum, pull, events i.e. pulling on top of reflex vertices) are logged in a file. Using the logfile, the program creates visualisations for the position of each guard, their visibility polygon area and their movement direction vector for each iteration. The program also plots the area seen both by each guard and in total. The visualisations are created using the Python 3 libraries matplotlib [17] and the scikit-geometry [20]. The plots displaying guard movements inside polygons show, for each guard, the total gradient vector in red. For each reflex vertex, the partial gradient vectors are green and the pull vectors are purple.

#### 4.1 Heuristics

In this section we observe the role played by each of the heuristics used. We additionally notice how different heuristics are relevant for different types of polygons. In order to do so, we run the program with all the heuristics but one, for each of the heuristics. By analysing the difference in movement for each of the guards, we assess the influence every heuristic has on different types of polygons.

It is worth mentioning that the experiments are not run with greedy initialisation. When initialising guards greedily, some polygons already become solved from the beginning. This does not allow us to properly showcase the other heuristics. Nonetheless, we still discuss the greedy initialisation heuristic at the end of this section.

We use fixed hyperparameters for all the runs. This allows us to focus on the differences between the heuristics. As such, we use the momentum hyperparameter  $\gamma = 0.8$  and pull attraction  $\beta = 1$ . The hyperparameter values were chosen through experimentation. The algorithm is sensitive to the hyperparameter choices, so for other values it often becomes stuck in local optima.

#### 4.1.1 Without Momentum

In this section we discuss the impact momentum has on the overall behaviour of the algorithm. As introduced in Section 3.3, momentum takes into account the position history of the guards. In this way, the overall trajectory of a guard is smoothened out.

We observe this behaviour with an arbitrary polygon in Figure 4.1. The polygon requires a minimum of 3 guards to be fully seen. They start at the same fixed position in both cases. We compare how the guards move with all the heuristics to when not using momentum.



Figure 4.1: Arbitrarily shaped polygon.

Figure 4.2 displays the area seen per iteration for the arbitrary polygon. Both the total and the individual areas seen by each guard are shown. Starting with almost the whole polygon seen, the guards are eventually optimally placed. Nonetheless, using momentum clearly makes a difference in Subfigure 4.2a, than when not using it in Subfigure 4.2b. Momentum allows the overall seen area to keep a steady trajectory towards its maximum. Additionally, guards quickly find their optimum in only 4 iterations, without many oscillations. In Subfigure 4.2b however we observe how the total area fluctuates. The guards display large jumps close to iterations 5 and 20. These jumps cause the overall progress towards the optimum to be less stable. For example, when guard 2 ( $g_2$ ) has a sudden drop in the area it sees around iteration 5, the total area seen naturally drops as well. The algorithm only recovers after iteration 20, when guard 0 ( $g_0$ ) makes another large jump. This behaviour also emphasises how the movement of one guard heavily influences the other guards' trajectories. As a result, the guards' trajectories to optimality become noisier and slower (more iterations are needed).



Figure 4.2: Total area seen per iteration for an arbitrary polygon guarded by 3 guards.

Thus, it becomes clearer how momentum allows the smoothening of noisy guard movements. We reckon that because guards are holding a steadier trajectory, they are more likely to achieve the optimum in less iterations. When not using momentum, the number of iterations increases substantially. Momentum appears to be a crucial improving heuristic to our whole algorithm, both in terms of speed-up and in the smoothness of the process.

#### 4.1.2 Without Pulling To and Onto the Reflex Vertex

In this section we discuss the impact that the pull to and onto the reflex vertex heuristic has upon the behaviour of the algorithm. Sections 3.4 and 3.4.1 introduced the pull to and onto reflex vertices heuristic, respectively. These heuristic tackles the idea that the closer a guard is to a reflex vertex, the stronger it is pulled towards it. If a guard is "close enough" to a reflex vertex, then the locally maximum seen area would be achieved by placing the guard on top of the reflex vertex. Thus, we believe that studying the impact of these two heuristics together is the most natural.

Figure 4.3 displays the guards' movement both when we are using all the heuristics and when we are not pulling them to and onto reflex vertices. The polygon is arbitrarily shaped and requires three guards for full guarding. Subfigures 4.3a and 4.3c show how the green guard's movement changes when it is placed onto the reflex vertex. Subfigures 4.3b and 4.3d show how the guard's movement changes when it is not pulled towards the reflex vertex. We observe how in Subfigure 4.3c the green guard is placed onto the reflex vertex. It then strives to reach the unseen polygon part in the upper left corner. On the other hand, in Subfigure 4.3d the guard moves away from the reflex vertex.



Gradient Computation for Iteration #0

(a) All heuristics. The green guard is pulled towards the reflex vertex.



(c) All heuristics. The green guard is onto the reflex vertex.

(b) The green guard has a pull towards the reflex vertex. However, we are not taking it into account.



(d) The green guard is placed according to the momentum computation, without being pulled towards or onto the reflex vertex.

Figure 4.3: Example of different movements of guards with and without pulling them to and onto reflex vertices in an arbitrarily shaped polygon guarded by three guards.

**4.1.2.1** Evaluation and Results We compare the performance of our algorithm in two settings: with all heuristics and without the pull to and onto reflex vertices. We observe this behaviour on the three polygons displayed in Figure 4.4. For each polygon, the algorithm is run 20 times for each heuristic setting, each time with randomised starting guards coordinates. Our evaluation consists of comparing the average number of iterations required for the algorithm to finish in the two heuristic cases. We additionally take note of the average total area seen per iteration.



(a) **Polygon 1:** arbitrarily shaped polygon that requires 3 guards for complete guarding.



(b) **Polygon 2:** comb-shaped polygon with four teeth that requires 4 guards for complete guarding.



(c) **Polygon 3:** arbitrarily shaped polygon that requires 2 guards for complete guarding.

Figure 4.4: The polygons used for the experiments.

Figure 4.5 displays how the total area seen on average per iteration for the three polygons develops in the two heuristics cases. There is a quite clear distinction between using all heuristics and not using the pull for all the polygons. Overall, the algorithm requires more iterations when the pull heuristic is not used.

The initial total area seen is higher on average for the all heuristics case than for the no pull case. Nonetheless, the algorithm in the no pull case compensates this disadvantaged start within the first 6 iterations for all polygons. In these first 6 iterations, reach pass 96% total polygon visibility. When using all heuristics, the algorithm manages to achieve the full visibility in less than 10 iteration after the initial spike. In the case of the no pull algorithm, the convergence towards the optimum is much slower in terms of number of iterations. Subfigure 4.5c displays the worst such case (polygon 3), where the no pull algorithm took thrice as many iterations than in the all heuristics case. Subfigure 4.5b displays a less dramatic case for polygon 2. There not using the pull results in more smoothened out convergence curve, that is only a quarter longer than the all heuristics case.

The difference in the average number of iterations happens due to the polygon shape. Polygons 1 and 3 are polygons where guards are required to travel past reflex vertices. So, the absence of the pull influences the speed and the efficiency of the movement much more. This is because without a pull, guards only travel around reflex vertices. Subfigure 4.5a displays such a case for polygon 1. In the first 10 iterations without a pull, the guards appear to oscillate between different positions. When using the pull, guards have a more stable trajectory. We reckon that this happens due to the fact that placing guards on reflex vertices maximises the seen area around the vertex. Otherwise, guards are required to find their optimal position around the reflex vertex. Conversely, polygon 2 (Subfigure 4.5b) shows a less drastic difference between the two heuristic cases. This is due of its shape: its reflex vertices do not have such a strong influence. Namely, an optimal positioning in this case can also be achieved by moving around reflex vertices, and not necessarily towards them.



Figure 4.5: Average total area seen per iteration for all the experiment polygons.

Figure 4.6 displays the average number of iterations for all polygons in the two heuristic cases. As previously observed, it takes more iterations on average to solve a polygon when not using the pull heuristic. Nonetheless, not using this heuristic also results in a less predictable guard movement, with a larger variability. This is especially noticeable for polygons 1 and 3, in Subfigures 4.6a and 4.6c, respectively. In these two cases, it is unlikely that there is a statistical difference between the two approaches. This is because the median values of both cases lie within the first and third quantile of each other's boxes. However, the initial guard placement plays an important role in determining how well our algorithm performs. When no pull is used, some outliers appear. Those correspond to unfortunate initial guard coordinates. Using all heuristics mitigates such unlucky cases by swiftly moving to and placing guards onto reflex vertices and maximising the area seen around the reflex vertex. This results in a more predictable guard movement that faster converges towards an optimum.

A few outliers appear also when using all heuristics on polygon 3. An optimal solution is still achieved in less iteration on average than when not using the pull heuristic. Nonetheless, the starting guard positioning in the case of all heuristics used on polygon 3 appears to be positively skewed. This means that on average the starting positioning of the guards gave the algorithm a head start. It is statistically unclear how this influences our algorithm's performance. In the case of polygon 2 in Subfigure 4.6b, the improved performance of the all heuristic algorithm becomes clearer. The median number of iterations lays underneath the first quantile when no pull is used. This indicates that a statistical difference between them could be likely. It is worth discussing what the implications of this difference are. We have discussed that the shape of polygon 2 allows guards to find an optimum solution similarly fast to when all heuristics are used. Still, the pull appears to play a statistically important role in the faster convergence towards the optimum.



Figure 4.6: Average number of iterations for all the experiment polygons.

Therefore, we believe that pulling guards towards reflex vertices results in reaching the optimum in less iterations. This appears to be an improvement for all polygon shapes, regardless of whether guards need to move only around or also towards and onto reflex vertices. This difference is less likely to be supported statistically for polygons with stronger pulls towards reflex vertices. It is unclear whether this statistical uncertainty is due to the heuristic themselves, or the skewed starting guard positions. Nonetheless, on average, the algorithm can solve polygons up to three times as fast when using all heuristics than when not using the pull.

#### 4.1.3 Without Pull Capping

In this section we discuss how capping the pull towards a reflex vertex influences the progress of the algorithm. Section 3.4.2 introduced the pull capping notion. The reason for this choice was to not allow the pull to overpower the value of the momentum. So, if the pull is larger than  $\mu$  times the momentum, it is capped at the momentum value multiplied by  $\mu$ .

We compare how the guards move when we are using all the heuristics to when we are not capping the pull towards reflex vertices. We use the comb polygon with four teeth as example. The guards have a pull cap hyperparameter  $\mu = 1$ . Thus, if the pull is larger than the momentum, we cap it at the size of the momentum. This value was chosen experimentally in order to clearly show the benefits of using this heuristic. So, if the pull is as large as the momentum, we prioritise pulling guards onto reflex vertices if they are "close enough".



(a) All heuristics. The guard's pull towards the reflex vertex is capped.



(c) No pull onto the reflex vertex. The guard's pull towards the reflex vertex is not capped.



(b) All heuristics. The guard does not act upon the pull towards the reflex vertex.



(d) No pull onto the reflex vertex. The guard is drawn closer to the reflex vertex.

Figure 4.7: Example of different movements of guards with and without pull capping in a comb polygon with four teeth.

Figure 4.7 displays the two reflex vertex pull cases: Subfigures 4.7a and 4.7b show how the blue guard movement changes when its pull is capped; Subfigures 4.7c and 4.7d show the blue guard's movement without pull capping. We observe how in Subfigure 4.7a the blue guard has its pull towards the reflex vertex capped. In that case, the movement vector is larger than the pull, so the guard doesn't move towards the reflex vertex. So, the blue guard moves as its movement vector dictates to the position in Subfigure 4.7b. This encourages the guard to explore the polygon more rather than directly moving towards a reflex vertex. Conversely, Subfigure 4.7d displays the blue

guard moving closer to the reflex vertex it is pulled to. In subsequent iterations, it is placed on the reflex vertex.

Figure 4.8 shows how the global behaviour of the algorithm is influenced by the capping. Interestingly enough, when capping the pull, the guards' behaviour is more erratic. This results in more iterations before the whole polygon is seen (Subfigure 4.8a). On the other hand, using no capping allows a steadier increase in the total area seen (Subfigure 4.8b). Eventually, the blue guard (g0) is drawn and placed onto the reflex vertex in iteration 2. It is worth mentioning that it does not move in the last iteration because the other guards have already covered the whole polygon. Nonetheless, the reflex vertex placement does not improve the locally seen area. It however reinforces the previously discussed idea that placing guards on reflex vertices is globally beneficial for the algorithm. As such, tuning the hyperparameter  $\mu$  is a crucial point of exploration for deciding how and when pull capping would always prove most favourable.



Figure 4.8: Total and individual areas seen per iteration for a comb polygon with four teeth.

In this way, we claim that capping the guards' pull towards a reflex vertex highly depends on the hyperparameter  $\mu$ . When not capped, it emphasises the importance of placing guards on reflex vertices. However, when the pull of guards is capped, we are using a more exploratory approach to discovering the polygon. Thus, we can experimentally decide whether to use this heuristic.

#### 4.1.4 Without Reflex Area

In this section we discuss the importance of keeping guards in the reflex area of a reflex vertex they had been placed on. Section 3.5 introduced this heuristic. The idea counteracts the edge-case of guards moving away from reflex vertices they had been placed on. When they move away from the reflex vertex, they might unsee a large previously seen area. It is possible that afterwards their movement vector directs them to move back onto the reflex vertex. We want to prevent them from unnecessarily moving away and then back on the vertex in a loop. For this reason, we keep them in the reflex area. The reflex area is then the area determined by the two segment lines meeting in the reflex vertex. By staying in this area, the guard would be able to move away from the reflex vertex in a direction that does not undo its progress.

We compare how the guards move when we are using all the heuristics to when we are not keeping guards into a reflex area. We use another arbitrary polygon as an example, that needs two guards to be fully seen.

Figures 4.9 and 4.10 compare the algorithm progress when using and not using the reflex area heuristic, respectively. Subfigures 4.9a - 4.9e show the guards' movement when one of them is bound to the reflex area. In this case, the blue guard wants to move away from the reflex vertex. The reflex area does not allow it to move away from the reflex vertex, so it stays there. Subfigures 4.9b - 4.9h display the guards' movement when the reflex area heuristic is not applied. In this case, the blue guard moves away from the reflex vertex (Subfigure 4.9d), only to be drawn back to it in the next iteration (Subfigure 4.9f). This repetitive behaviour continues without eventually reaching the optimum state.

Figure 4.10 displays the difference in the areas seen between the two approaches. When using all heuristics (Subfigure 4.10a), the algorithm terminates in 7 iterations. The continuous line marks the guard who did not move away from the reflex vertex. Conversely, when not using the reflex area heuristic (Subfigure 4.10b), the algorithm appears not to terminate. One of the guards loops away and back on the reflex vertex. The other guard loops so that it sees the area that is being seen and unseen by the other guard. In this way, they do not synchronise in order to see the polygon at the same time. This shows in the repetitive plot.

In this way, the reflex area heuristic becomes important to the correct run of the algorithm. It addresses the issue of guards moving repeatedly away and back on reflex vertices. This behaviour causes the other guards to move towards the area that becomes unseen by the first guard. The guards continue to move out of sync with each other, resulting in the polygon never being fully seen.



(a) All heuristics. The blue guard is on the reflex vertex.



(c) All heuristics. The blue guard is still on the reflex vertex (in the reflex area).



(e) All heuristics. The blue guard is still on the reflex vertex (in the reflex area).



(g) All heuristics. The blue guard is still on the reflex vertex (in the reflex area).



(b) No reflex area. The blue guard is on the reflex vertex.



(d) No reflex area. The blue guard moved away from the reflex vertex.



(f) No reflex area. The blue guard is drawn back to the same reflex vertex.



(h) No reflex area. The blue guard is back on the reflex vertex.

Figure 4.9: Comparison between using and not using the reflex area heuristic on an arbitrary polygon guarded by two guards.



(a) Area seen for all heuristics on another arbitrary polygon.

(b) Area seen for no reflex area heuristic on another arbitrary polygon.

Figure 4.10: Area comparison between using and not using the reflex area heuristic on an arbitrary polygon guarded by two guards.

#### 4.1.5 Without Line Search

In this section we discuss the impact line search has on the overall behaviour of the algorithm. As introduced in Section 3.6, line search determines how far a guard should move towards the direction movement vector. In this way, it computes the optimal position of a guard on the movement direction given a step size.

Line Search has step size s, maximum movement factor x and starting factor  $\frac{s^t}{x}$ , t = 0 as hyperparameters. For our experiments, we choose s = 2 and x = 32. We start with a factor of  $\frac{1}{3}2$  for the movement vector. We increase it by a step size of 2 up to factor 32. In this way we generate 10 solutions in total, one for every step  $\frac{1}{32}$ ,  $\frac{1}{16}$ ,  $\frac{1}{8}$ , ..., 1, 2, 4, ..., 32. This allows us to search a larger space of position possibilities knowing the direction of the movement vector. We pick the best solution along all the step size based on the largest area increase.



Figure 4.11: Polygon in the shape of a comb with four teeth.

A suggestive way to observe how well line search works is with the comb polygon with four teeth from Figure 4.11. Comb polygons with t teeth require t guards to be guarded, one guard for each tooth. This is because the space in between the teeth is large enough so that guards cannot fully see two teeth simultaneously. In our case, t = 4. We compare how the guards move when we are using all the heuristics to when we are not using line search. They start at the same fixed position in both cases, so we focus on observing their movements.

Figure 4.12 displays the area seen per iteration for the comb polygon with four teeth. Both the total area seen and the individual area seen by each guard are shown. Starting with around 82.5% total area seen, the global optimum is eventually found. Nonetheless, using line search clearly makes a difference between Subfigures 4.12a and 4.12b. The first noticeable difference is the number of iterations. Using line search allows the guards to find their optimal positions in 3 iterations, with a steady increase in the total area seen. Note that not all guards were moved in iteration 3. This is because the polygon was fully seen after moving only two of the guards. On the other hand, in subfigure 4.12b not using line search results in the optimal position to be found in more than 80 iterations. Additionally, 3 of the guards found their optimal position after the 30<sup>th</sup> iteration, whereas the last 50 iterations are spent on only one guard finding its own.



Figure 4.12: Total area seen per iteration for the comb polygon with four teeth, guarded by four guards.

Therefore, we reckon that line search significantly and more efficiently speeds up the process of finding the optimal position for each guard. In this way, each guard moves faster to its optimal position. The situation where multiple guards that have found their optimal position have to wait for only one guard to find its own is also avoided.

#### 4.1.6 Without Angle Behind Reflex Vertex

In this section we discuss the benefits of using the angle behind reflex vertices. Section 3.7 introduced this heuristic. The idea behind this technique is that guards should be drawn faster to larger unseen areas behind reflex vertices. In this way, we prioritise unseen areas based on their size, while still accounting for the small unseen areas.

We compare how the guards move when we are using all the heuristics to when we are not taking into account the angle behind the reflex vertex. We use an arbitrary polygon as an example, that needs three guards to be fully seen. Figure 4.13 shows a comparison between using and not using the angle behind the reflex vertex heuristic. Subfigures 4.13a and 4.13b display the first two iterations of the algorithm when all heuristics are used. Subfigures 4.13c and 4.13d focus on the first two iterations when not using the heuristic. We observe a major difference in the way the final movement vector is computed for the orange guard in the two cases. When all heuristics are used, the movements of the guards are much smaller and smoother. For example, in Subfigure 4.13c, the orange guard has a large movement vector outside the polygon due to the unseen part in the upper pocket. When we also take into account the angles behind the reflex vertices in the pocket, its movement vectors become much smaller (Subfigure 4.13a). In fact, the orange guard is also drawn slightly to the right part of the polygon, as it has a larger angle behind the reflex vertex. That part of the polygon is already seen by the blue guard, so it starts moving towards the upper pocket in Subfigure 4.13b.



(a) All heuristics. The orange guard's movement is towards upper right.



(c) No angle behind the reflex vertex. The orange guard's movement is all the way to the left boundary of the polygon.







(d) No angle behind the reflex vertex. The orange guard's movement is towards the left-side outside the polygon.

Figure 4.13: Example of different movements of guards with and without the angle behind the reflex vertex in an arbitrarily shaped polygon guarded by three guards. The total gradient vector is displayed in red, partial gradient vectors are green and the pull vector is purple.

In terms of efficiency, the optimal solution is achieved in more iterations when the angle heuristic is used. This is observed in Figure 4.14. When using all heuristics, the optimal solution is achieved in 5 iterations (Subfigure 4.14a). The movement of the guards is smooth. Two of the guards have an increasing seen area, whereas the orange guard moves slowly towards the upper pocket of the polygon. Not using the angle heuristic allows us to achieve the same goal in only 3 iterations, in a less smooth manner (Subfigure 4.14b).

Therefore, computing the angle behind the reflex vertex is a heuristic that allows us to fine-tune and smoothen the movement of the guards. In this way, we focus on moving fast towards the bigger unseen areas, while not neglecting the smaller ones. We could say that this heuristic offers a trade-



Figure 4.14: Total and individual areas seen per iteration for an arbitrarily shaped polygon guarded by three guards.

off between the number of iterations and path smoothness. For this reason, the performance of the algorithm is influenced by the type of polygons it is applied to: polygons with sharper, narrower turns could benefit the most from this heuristic.

#### 4.1.7 No Hidden Movement

In this section we discuss the importance of using the hidden movement heuristic. Section 3.8 introduced it. The idea is based on the fact that we want guards to make progress, no matter how little. If a guard's visibility region is already seen by other guards (so its movement vector is zero), it is still unlikely that its position is optimal. Thus, we want the guard to still progress.

We compare how the guards move when we are using all the heuristics to when we are not using the hidden movement heuristic. We use the comb polygon with four teeth as an example.

Figure 4.15 displays a comparison between using and not using the hidden movement in our algorithm. In Subfigure 4.15b we notice that the execution of the algorithm takes twice as many iterations than in Subfigure 4.15a. The reason behind this is the fact that for half of the iterations two of the guards do not move when no hidden gradient is used. Additionally, the total area seen in that case fluctuates. This is in contrast with the case where the hidden gradient heuristic is used. In that case, the total area seen is continuously increasing, and all guards who can make progress move.

Therefore, we reckon that the hidden gradient heuristic results in a significant improvement to our algorithm in terms of efficiency. In this way, all guards are ensured to make some progress towards their optimum and are not stalled.



Figure 4.15: Seen area for the comb polygon with four teeth guarded by four guards.

#### 4.1.8 Greedy Initialisation

In this section we discuss the benefits of using a greedy initialisation technique for our algorithm. Section 3.9 introduces it as beneficial for giving a head start to our algorithm. In our experiments we greedily initialise the guards in the middle of the leftmost segment of each of their visibility regions. The reason behind this choice is that CGAL did not offer a quick way to pick a randomised position inside the visibility area. Due to the time constraints of this thesis, we decided to use a deterministic positioning.

So far we have not used greedy initialisation in any of our examples. The reason is that this technique already solves some polygons at guard placement time. An example in this case is the comb polygon with four teeth in Figure 4.16.



Figure 4.16: The comb polygon with four teeth is already completely seen at greedy initialisation time.

In the case of polygons who are not completely seen at initialisation time, it is unclear whether the greedy placement improves the overall progress. How well the algorithm behaves is highly dependent on the initial placement of the guards. Namely, the algorithm often does not finish with different starting positions (it is stuck in local optima). Because of this drawback, we cannot create extensive experiments with different starting guard positions.

Therefore, we believe that the greedy initialisation technique would benefit most from a deterministic placement if it were tailored to the shape of the polygon. Otherwise, an arbitrary placement would be most suitable for general cases. However, due to time constraints, this heuristic's performance is limited. So, it is to be explored with a more robust implementation that places guards arbitrarily.

#### 4.2 Scaling for the Comb Polygon

In this section we observe how the algorithm scales on the comb polygon. We take the comb polygons with 2, 3, ..., 10, 15, 20, 50, 100 teeth. We note how the algorithm progresses in terms of the total area seen per iteration. As such, we run the algorithm with all heuristics (but greedy initialisation) for all the mentioned comb polygons. All guards start at the same *y*-coordinate, 0.1 units apart from each other. An example of such a starting point is found in Figure 4.17. In this way, we ensure that the performance of the algorithm can be compared using the same starting position of the guards. We deem a timeout of one hour for declaring an algorithm infeasible for larger polygons. We then analyse how the algorithm performed under each of the circumstances. Lastly, we offer some points of improvement.



Figure 4.17: In the comb polygon with 6 teeth, all guards start at the same y-coordinate, with 0.1 units in between them.

We expect that the number of iterations and the execution time needed would scale with the number of teeth t of the polygon. With every t + 1 tooth added, we would need the same time as before for the t guards to reach their respective teeth, and some extra time for the new t + 1<sup>th</sup> guard.

Figure 4.18 displays the progression of our algorithm for the comb polygons with 2, 3, ..., 10, 15, 20 teeth within one hour. The comb polygons with 50 and 100 teeth did not complete any iteration within that time. For comb polygons with 5 teeth and more than 6 teeth, the timeout was not enough

to find a feasible solution. Nonetheless, it is worth discussing their behaviour, as it highlights the different issues the algorithm faces. For comb polygons with 7, 9 and 10 teeth, the guards appear to be stuck in a local plateau that they did not manage to escape. This is traced to an edge-case where the guards are stuck in their position because their movement vector does not improve the total area. Another reason why the guards do not move is that they are still trying to maximise their local area in the case where the total area seen cannot be increased within that step. This could result in them all moving to the same part of the polygon. Since that part of the polygon maximises their local area, they not move anymore. For the comb polygon with 5 teeth, the guards appear to be stuck in a local maxima that they manage to escape. Unfortunately, they later return to it and restart the cycle. It is unclear whether the comb polygon with 8 teeth would display a similar looping behaviour. Given the timeout, the guards trying to solve the comb polygon with 8 teeth display a less predictable behaviour. Hyperparameter tuning could address this issue.



Figure 4.18: The percentage of the area seen in terms of iterations for comb polygons with 2, 3, ..., 10, 15, and 20 teeth.

Conversely, it is also crucial to address in Figure 4.19 the comb polygons which the algorithm manage to solve: 2, 3, 4 and 6 teeth. In this case, we observe a clearer scaling: the more teeth a comb polygon has, the longer it takes to be solved. What is more, the solution also scales with the number of teeth. The comb polygon with 4 and 6 teeth take twice as many iterations to be solved than the ones with 2 and 3 teeth, respectively. We would expect a similar behaviour for comb polygons with more teeth, if the algorithms would finish.



Figure 4.19: The percentage of the area seen in terms of iterations for comb polygons with 2, 3, 4, and 6 teeth.

It is also worth discussing how the initial guard placement and comb polygon shapes could influence the performance of the algorithm. Firstly, it is important to note the initial placement strategy. Guards are placed in the same lower left part of the polygon. However, the larger a polygon, the lower the initially seen area. For example, we observe how the comb polygon with 2 teeth starts at more than 80% seen area. On the other hand, the comb polygon with 20 teeth has a less than 65% seen area at start. We could argue that this type of placement gives larger polygons a start disadvantage. A way to mitigate this issue could be to find a way of placement that scales, while still keeping the initial similar positioning among different combs.

Moreover, the algorithm is sensitive to the shape of the polygons. The bigger a comb polygon grows, the sharper and narrower its teeth are. A good example in this case is Figure 4.20, which displays a comb polygon with 20 teeth. Clearly, the teeth are much narrower than in the case of a comb with 4 teeth. This raises the question about the performance of the algorithm given different angles in the polygon boundary. A possible solution to this aspect would be to scale up the polygon. However, a similar question would remain: does the algorithm necessarily perform better or worse because of the polygon shape? What would be the optimal shape?



Figure 4.20: Comb polygon with 20 teeth.

#### 4.3 Hyperparameter Sensitivity

The algorithm is highly sensitive to the hyperparameter choices, polygon shapes and sizes, and guard starting points. We already emphasised how the hyperparameter values have been chosen: other values than the ones used would result in the program crashing or not terminating. Similarly, in the context of the comb polygon, we have discussed that the polygon shape can have a crucial effect on the performance of the algorithm.

In this section we additionally explain how the guard starting points influence the outcome of the program. We use the irrational guards polygon [1] as an example. The polygon can be guarded by 3 guards with irrational coordinates, or 4 guards with rational coordinates. We thus compare these placements. We position 3 guards at an approximation of the optimal irrational coordinates, and 4 guards at arbitrary rational coordinates. The approximated irrational guards have coordinates (2, 0.59), (19, 1.71), (10.57, 2.12). The arbitrary guards have coordinates (5, 0), (5.5, 0), (6, 0), (6.5, 0).



Figure 4.21: Comparison for the irrational guards polygon between guards with approximately irrational coordinates and rational guard coordinates.

Figure 4.21 displays the difference in runs between the previously mentioned guard placements. The orange line displays the algorithm run for the guard irrational approximation coordinates. As expected, the algorithm deems that the whole polygon is seen within the first iteration. On the other side, the blue line shows the algorithm performance for the guards with rational coordinates. Unfortunately, the program in this case crashes after the 8<sup>th</sup> iteration. The error was unclear and the time constraints of this thesis did not allow us to solve it.

The irrational guards polygon becomes thus a suggestive example to the sensitivity of the algorithm to initial guard placements. The starting position of the guards can determine whether the algorithm is able to escape local maxima or not crash. Nonetheless, we are not able to provide a generalisation of this statement. Namely, we were not able to determine how the starting position of the guards influences the progress of the algorithm, while it is also dependent on the polygon shape.

In this way, we observed how the different heuristics influence the algorithm. Some heuristics, like the pull, are likely to be statistically significant and fundamental to the faster convergence of the algorithm, regardless of the polygon shape. Others, such as the pull capping, require case-specific tuning based on the start positions of the guards and the shape of the polygon. In the next section we further discuss why the algorithm is sensitive to these hyperparameters and how it could be improved.

60

## 5 Problems Encountered

The development of this algorithm has encountered a number of important problems that affected its progress. This section emphasises the most influential such issues and, if known, suggest a solution for resolving them.

#### **Edge-Cases**

To begin with, some edge-cases remained unaddressed. For example, Figure 5.1 shows a frequently appearing such edge-case: the blue guard is stuck on the vertex because its movement direction points outside the polygon. When this happens, the program would project the movement direction vector on the closest polygon boundary. In our example we would expect that the blue guard moves downwards on the polygon boundary. However, the projection of the movement direction vector inside the polygon is the vertex. The grey dashed lines display the extensions of the segments intersecting in the vertex the guard is stuck on. So, the guard cannot move anywhere else from the vertex.



Figure 5.1: Edge-case for when the blue guard cannot move because its movement direction cannot be projected towards the inside the polygon.

In order to solve this issue, we would need to devise a technique that allows the guard to move regardless of its gradient. One such heuristic could be to mirror the movement direction vector inside the polygon. The black dotted line in the figure displays the mirroring line. The mirrored movement direction vector along the mirroring line is displayed in dark red. In this way, the guard would be able to make progress along the polygon's boundaries as expected.

However, we did not encounter this edge-case frequently enough to give it a high priority on the issues list. Due to the time constraints of the thesis, it remained hence unsolved.

#### CGAL

Moreover, numerous issues were posed by the CGAL library itself. These problems mainly related to the library's non-detailed errors and brief documentation. Having to reverse engineer and delve into the source-code of CGAL slowed down the debugging process. Because some errors were not explicit at all, some crashes remained unsolved (for example, the program crashes for certain starting positions for different polygons). We encountered these issues only half-way through the thesis. Because of a lack of more robust known computational geometry libraries and time constraints, we decided to heroically tackle them.

Some issues were posed by converting between CGAL's number types and the native C++ types (double). CGAL's proprietary numeric values were greatly inefficient and effort was required to convert these to C++ native double values. Doing this greatly increased efficiency. Nonetheless, such conversion resulted in approximation errors. For instance, some guards "close enough" to reflex vertices were wrongly considered to be on the reflex vertex. Conversely, comparing point coordinates would not always work, as sometimes guards with the same coordinates would not be considered to be the same.

#### Input Polygons

Additionally, we were especially interested in the irrational guards polygon [1]. However, our program can only solve the irrational guards polygon if the guards are already very close to the optimum. For other starting positions, the program crashes with a CGAL error about the finiteness of the computed numbers. More concretely, the assertion of is\_finite(d) in the Interval\_nt.h was failing. The error gives no further information about the line of the code it is encountered at. Because of the vagueness of the error and the time constraints of the thesis, we could not fix this issue.

Lastly, we also wanted to test the program on polygons from the APGlib library [8]. The APGlib library offers an extensive testbed for polygons. Unfortunately, we encountered the same error with the APGlib polygons we used for testing. Interestingly, the program only got stuck in a local minimum for one of the polygons with 20 vertices (polygon 3). Nonetheless, the fact that we could not address the CGAL error in this thesis' time constraints was quite dissatisfying.

#### Scalability

Another problem worth mentioning is scalability. The program does not scale. As mentioned in Section 4, for comb polygons with more than 6 teeth, the waiting time already exceeds an hour to finish. We believe that this waiting time is inadequate for the size of the polygon and number of guards inside it.

Some of the largest bottlenecks that work against scalability are the visibility area and the hidden movement computations. Currently, the visibility area of each guard is updated at every iteration. We are using the Triangular Expansion Visibility [6] which runs in O(g) time, with g the number of guards. Thus, the visibility computation per iteration runs in  $O(g^2)$  time. In terms of the hidden movement computation, each guard has its gradient recomputed at most g times. This happens in the case when only one guard out of the g guards has a non-zero movement vector, and thus it gets removed from the set. The remaining g - 1 guards have their movement vector recomputed. Again, in the worst case only one guard has a non-zero gradient. Thus, a guard gets its movement vector recomputed in  $O(g^2)$  time. Other poorly scalable parts of the algorithm are based on the number of reflex vertices. The movement vector computation depends on the number of reflex vertices r. If  $r \gg g$ , then the algorithm performs poorly for polygons with significantly fewer guards than reflex vertices.

For these reasons, the program is especially sensitive to the values of the parameters. For instance, searching between fewer steps for line search can help the program find a solution faster. In this way, we decrease the number of steps taken for computing a new guard position. The program is additionally sensitive to the heuristics used and the shape of the polygons. In Section 4 we have mentioned some shapes of polygons that the algorithm can solve, with which hyperparameters and initial guard positions.

Therefore, we believe that if these edge-cases and errors get solved, the program should perform correctly and without crashing. Additionally, should the naive (brute-force) approaches be implemented more efficiently, we would expect the program to scale. We would still predict that the program gets stuck in local maxima. However, it should be possible to escape them with case-specific hyperparameter tuning and heuristic choice.

## 6 Discussion

This thesis focussed on implementing and evaluating a gradient descent algorithm to find a solution to the Art Gallery Problem. This goal has been achieved for the few discussed polygons. Nonetheless, this method raised quite some issues. For this reason, we discussed more in-depth the development process and the performance of the algorithm. There were numerous edge-cases to be taken into account. For many of them, we created various hyperparameterised heuristics. Other heuristics (such as momentum) were created as an improvement to the shortcomings of gradient descent.

The resulting program is sensitive to hyperparameter choices, polygon shapes and starting guard positions. For this reason, we only provided qualitative evaluations for the heuristics used to extend gradient descent. Namely, we discussed both specific added benefits of each of the heuristics, and a comparison based on the average number of iterations required to fully see a polygon. The way this was done is by running the program with all heuristics but the one whose practicality we are discussing. In this way, we assess the added advantage of each heuristic in its absence. The reason behind this experimental setup is that combinations of all heuristics are both verbose and unnecessary. Given that the program is still limited to a few examples, we cannot use extensive statistical tools to test its overall performance. With the given examples, we managed to assess that using the pull heuristic is likely to be statistically significant to improve the algorithm's performance in terms of the number of iterations. In this way, we provided a pragmatic overview to the usefulness of each of the heuristics and hyperparameters used. Namely, we analysed on what the type of polygons different heuristics and hyperparameter values provide concrete results.

Section 6.1 further discusses how the program could be improved in the future.

#### 6.1 Future Work

Currently, the program offers a state-of-the-art view about its practical possibilities. In the future, it would be suitable to extend it with more robust functionalities.

One of the most crucial aspects of improvement for the program would be to solve the existing errors and bugs. As of now, the program crashes for specific input polygons (see Section 5 for an extensive overview of its shortcomings). Another important element to consider would be to code it more efficiently, as the program does not scale. This is because some of the data structures and techniques used are naive (brute-forced). Implementing them in a more efficient way should improve the scalability of the program.

Additionally, some features and heuristics are not complete. For example, the greedy initialisation of the guards' position is deterministic. In order to quantitatively assess the performance of the algorithm, a truly randomised greedy initialisation would be required for statistical significance.

Lastly, it would be worth exploring how the algorithm would benefit from a different implementation. Namely, using another programming language like Python (or CGAL Python bindings) and other geometry libraries which are better documented and more reliable to use.

## Acknowledgements

I would like to thank Till Miltzow and Frank Staals for supervising the thesis and always offering creative and constructive feedback. I would also like to thank Simon Hengeveld for the helpful advice and explanation of his implementation for solving the Art Gallery Problem. Lastly, I would like to thank my partners Tim and Teun for always supporting me during both the rough and the happy times of this thesis.

## References

- Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. "Irrational Guards are Sometimes Needed". In: CoRR abs/1701.05475 (2017). arXiv: 1701.05475. URL: http://arxiv. org/abs/1701.05475.
- [2] George Johnston Allman. Greek geometry from Thales to Euclid. Hodges, Figgis, & Company, 1889.
- [3] Tetsuo Asano. "An efficient algorithm for finding the visibility polygon for a polygonal region with holes". In: *IEICE TRANSACTIONS (1976-1990)* 68.9 (1985), pp. 557–559.
- Pritam Bhattacharya, Subir Kumar Ghosh, and Bodhayan Roy. "Approximability of guarding weak visibility polygons". In: *Discrete Applied Mathematics* 228 (2017). CALDAM 2015, pp. 109-129. ISSN: 0166-218X. DOI: https://doi.org/10.1016/j.dam.2016.12.015. URL: https://www.sciencedirect.com/science/article/pii/S0166218X16306333.
- [5] Édouard Bonnet and Tillmann Miltzow. "An Approximation Algorithm for the Art Gallery Problem". In: CoRR abs/1607.05527 (2016). arXiv: 1607.05527. URL: http://arxiv.org/ abs/1607.05527.
- [6] Francisc Bungiu et al. "Efficient Computation of Visibility Polygons". In: CoRR abs/1403.3905 (2014). arXiv: 1403.3905. URL: http://arxiv.org/abs/1403.3905.
- [7] Computational Geometry Algorithms Library. https://www.cgal.org.
- [8] Marcelo C. Couto, Pedro J. de Rezende, and Cid C. de Souza. Instances for the Art Gallery Problem. www.ic.unicamp.br/~cid/Problem-instances/Art-Gallery. 2009.
- Boris Delaunay et al. "Sur la sphere vide". In: Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk 7.793-800 (1934), pp. 1–2.
- [10] Geometry, Algebra and Computer Algebra apps. https://www.geogebra.org/.
- [11] Subir Kumar Ghosh. "Approximation algorithms for Art Gallery Problems in polygons". In: Discrete Applied Mathematics 158.6 (2010), pp. 718-722. ISSN: 0166-218X. DOI: https://doi. org/10.1016/j.dam.2009.12.004. URL: https://www.sciencedirect.com/science/ article/pii/S0166218X09004855.
- [12] I Goodfelow, Y Bengio, and A Courville. "Deep learning (adaptive computation and machine learning series)". In: (2016), p. 296.
- Simon B. Hengeveld and Tillmann Miltzow. "A Practical Algorithm with Performance Guarantees for the Art Gallery Problem". In: CoRR, SoCG abs/2007.06920 (2020). arXiv: 2007.06920. URL: https://arxiv.org/abs/2007.06920.
- [14] Barry Joe and Richard B Simpson. "Corrections to Lee's visibility polygon algorithm". In: BIT Numerical Mathematics 27.4 (1987), pp. 458–473.
- [15] D. Lee and A. Lin. "Computational complexity of Art Gallery Problem". In: *IEEE Transactions on Information Theory* 32.2 (1986), pp. 276–282. DOI: 10.1109/TIT.1986.1057165.
- [16] Shiva Maleki and Ali Mohades. "Implementation of polygon guarding algorithms for art gallery problems". In: CoRR abs/2201.03535 (2022). arXiv: 2201.03535. URL: https://arxiv.org/ abs/2201.03535.
- [17] Matplotlib: Visualisation with Python. https://matplotlib.org/.

- [18] Joseph O'rourke et al. Art gallery theorems and algorithms. Vol. 57. Oxford University Press Oxford, 1987.
- [19] Marcus Schaefer. "Complexity of some geometric and topological problems". In: International Symposium on Graph Drawing. Springer. 2009, pp. 334–344.
- [20] Scientific Python Geometric Algorithms Library. https://github.com/scikit-geometry/ scikit-geometry.
- [21] WH Swann. "A survey of non-linear optimization techniques". In: *FEBS letters* 2.S1 (1969), S39–S55.