Pinwheels and Polygons: Symmetric Realizations of Polygon-Free Point Placements via SAT

Ethan Mackey

CMU-CS-25-108

May 2025

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Marijn Heule, Chair
Ruben Martins

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Science

Keywords: SAT, Symmetry, Point Placements, The Erdős–Szekeres Conjecture

For my family and loved ones

Abstract

This research explores the discovery of symmetrical realizations of convex-hexagon-free point placements on 16 points using satisfiability solving techniques. The central focus is on identifying point configurations that exhibit 3-, 4-, and 5-fold rotational symmetry. These 16 point configurations correspond to the maximal number of points that can be placed in the Euclidean plane in general position without forming any convex hexagons, a central case in the study of the Erdős–Szekeres conjecture , a foundational problem in combinatorial geometry.

Building on previous work in combinatorial geometry and SAT-based combinatorial methods, this research extends existing Boolean satisfiability encodings by incorporating symmetry constraints and structural conditions specific to the hexagon-free problem. Using these ideas, new conjunctive normal form formulas are developed to represent the search space of symmetric hexagon-free point placements.

To interpret and visualize solutions, satisfying assignments to these CNFs are passed through a point realization tool that reconstructs geometric configurations from orientation triple data. This enables the conversion of logical encodings into concrete point placements that can be analyzed and compared. Structural analysis of these placements includes examining the frequency and distribution of smaller convex polygons, such as 4-gons and 5-gons, to better understand the local geometric implications of hexagon avoidance.

The resulting symmetric configurations, especially those with four-fold and five-fold symmetry, represent some of the first structured, realizable examples of 16-point hexagon-free sets. These findings contribute new insight into the Erdős–Szekeres conjecture and offer a stepping stone toward understanding larger generalizations, such as the existence of 32-point configurations that avoid convex 7-gons.

Acknowledgments

I would like to express my deepest gratitude to my advisor, Professor Marijn Heule, for his invaluable guidance, support, and encouragement throughout the course of this research. His expertise, patience, and insightful feedback were instrumental at every stage of this work. Working under his mentorship has been a tremendously rewarding experience, and I feel truly fortunate to have had the opportunity to learn from him.

I would also like to sincerely thank Bernardo Subercaseaux for his generous help in providing and modifying the point realizer used in this project. His technical assistance and responsiveness were greatly appreciated and had a significant impact on the progress and results of my work.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This research concerns the search for symmetrical realizations of hexagon-free point placements on 16 points using SAT solving techniques. In order to understand the objective, it is first necessary to introduce the Erdős–Szekeres conjecture as well as motivate the idea of SAT solving, as the remaining sections will rely heavily on both.

1.1 The Erdős–Szekeres Conjecture

The Erdős–Szekeres conjecture is a well-known open problem in the field of combinatorial geometry, originally posed in 1935 by Paul Erdős and George Szekeres in their seminal paper "A Combinatorial Problem in Geometry" [2]. The conjecture asserts that for every integer $k > 3$, there exists a minimal number $n = h(k)$ such that any set of at least n points when placed in general position in the plane (no three points collinear) must contain a subset of k points that form the vertices of a convex k-gon.

Erdős and Szekeres proved that such a number n exists for every k and provided an upper bound on its size. They also conjectured that the exact bound is given by the following equation:

$$n = 2^{k-2} + 1 \tag{1.1}$$

The values of $h(3)$ through $h(6)$ have been verified to match the proposed bound, a simple case split was used to prove that $h(4) = 5$ in the original paper. The result for $k = 6$ was first shown by a computer proof in the 2006 paper "Computer Solution to the 17-point Erdős–Szekeres problem" by George Szekeres and Lindsay Peters [6]. Despite extensive work since the origins of the conjecture, it remains unresolved for general k, and the values of $h(7)$ and above remain unknown.

From the proof that $h(6) = 17$ it follows that every placement of 17 points in general position in the plane contains at least one convex 6-gon. It also follows that 17 is the minimal number of points such that this is true. A natural second result that follows from these two facts is that there must exist at least one way to place 16 points in the plane in general position such that no convex 6-gons are formed. These 16-point solutions are the study of this research.

1.2 Propositional Logic and SAT Solvers

Propositional logic, also known as Boolean logic, is a formal system in which logical formulas are constructed from propositional variables using logical connectives such as AND ($\wedge$), OR ($\vee$), NOT ($\neg$), IMPLIES ($\implies$), and EQUIVALENT ($\iff$). Each propositional variable represents a statement that can be true or false, and the logic of the connectives defines the overall truth evaluation of each formula based on its components.

A central problem in propositional logic is the satisfiability problem (SAT), which concerns whether, given a propositional formula, there exists an assignment of true or false to that formula's variables such that the entire formula evaluates to true. It is a well-known fact that SAT is NP-complete; however, SAT solving has emerged as a highly practical tool for a wide range of applications.

This emergence is in large part due to the ability of diverse problem domains to be represented as propositional formulas, allowing SAT solving to have applications for many real-world problems. SAT solvers are specialized tools that are designed to efficiently decide the satisfiability of propositional formulas. Most often, these solvers require that their input formulas be in conjunctive normal form (CNF), which simply means that the formula is structured as a conjunction of disjunctions of variables.

There are a great number of modern SAT solvers to choose from, each slightly different and more specialized for different applications. The solver of choice for this research was CaDiCaL, designed by Armin Biere [1]. CaDiCaL is a type of conflict driven clause learning solver that works great for general-purpose applications.

1.3 Objective

The objective of this research was to discover symmetrical realizations of hexagon-free point placements on 16 points using SAT solving techniques. Specifically, the study focuses on identifying configurations exhibiting 3-fold, 4-fold, and 5-fold rotational symmetry, which lend themselves naturally to SAT encodings on 16 points. These 16-point configurations represent the largest possible point placements that can avoid the formation of convex 6-gons, as implied by the Erdős–Szekeres conjecture.

This research builds on previous work in the field, mainly from two papers. These papers create the foundation for how to efficiently encode combinatorial geometry instances like this problem in propositional logic. Building on the base CNFs provided by these papers and through the addition of new clauses that encode the desired structure and symmetries, new formulas can be created that encode symmetric 6-gon free point placements on 16 points.

An analysis of the underlying properties and structure of these solutions would be performed to gain further insight into the Erdős–Szekeres conjecture as well as the solutions themselves. Finally, leveraging a point realizer designed by Bernardo Subercaseaux [5], satisfying assignments to these CNF formulas would be converted into realized images.

Chapter 2

Literature Review

This research builds on a broad body of prior work focused on encoding problems in combinatorial geometry, such as the Erdős–Szekeres Conjecture, into SAT frameworks. Among these, two papers had a particularly direct influence on the present work, "Axioms and Hulls" by Donald Knuth [4] and "Happy Ending: An Empty Hexagon in Every Set of 30 Points" by Marijn Heule and Manfred Scheucher [3], as they provide the structural basis for the CNF encodings used in both the analytical and constructive phases of this thesis.

2.1 CC Systems

In "Axioms and Hulls" [4], Knuth defines a cc system (short for counterclockwise system) as a combinatorial abstraction used to encode orientation information among triples of points in the plane. These systems are foundational in SAT encodings of point configuration problems, including those addressed in this thesis.

2.1.1 Orientation Triples

In the context of planar geometry, an orientation triple captures the relative orientation of three distinct points. The orientation is considered positive if the points in the triple are encountered in a counterclockwise manner, and negative if the points in the triple are encountered in a clockwise manner. These directional relationships can be abstractly represented as propositional variables. For example, the orientation variable for the triple of points (a, b, c) would be written as:

$$o_{a,b,c}$$

For these orientation variables, a true value indicates a counterclockwise turn and a false value indicates a clockwise turn. This relationship is illustrated in Figure 2.1.

In SAT-based geometric reasoning, orientation triples allow point configurations to be represented without reference to explicit coordinates. By expressing orientation relationships directly in Boolean logic, it becomes possible to encode high-level geometric properties, such as convexity, ordering, and certain substructures, into satisfiability problems. This abstraction is crucial for transforming geometric constraints into CNF formulas and was the main method used for this research.

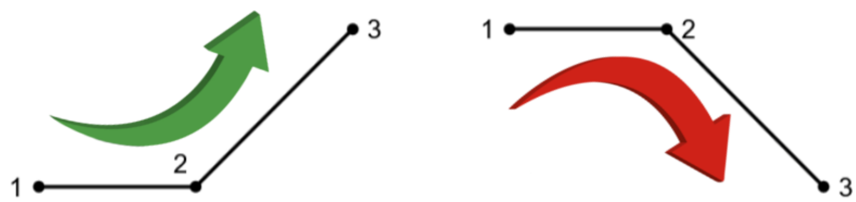

Figure 2.1: Two examples of orientation triples for points 1, 2, and 3 are illustrated. In the first case, the orientation of $o_{1,2,3}$ is counterclockwise and therefore evaluates to true. In the second case, the orientation is clockwise and consequently evaluates to false.

2.1.2 CC Axioms

Although orientation triples provide a compact and expressive way to encode local relationships between points, not all combinations of orientation assignments correspond to realizable point sets in the Euclidean plane. To ensure local logical consistency, Knuth introduces a set of axioms that describe how orientation values must behave to be more consistent with the underlying planar geometry.

The first class of axioms enforces cyclic symmetry, ensuring that the orientation of a triple remains consistent under cyclic permutations. It is encoded as:

$$o_{p,q,r} \implies o_{q,r,p} \tag{2.1}$$

The second class of axioms enforces antisymmetry, ensuring that reversing the order of the last two points in a triple negates its orientation. It is encoded as:

$$o_{p,q,r} \implies \neg o_{p,r,q} \tag{2.2}$$

The third class of axioms enforces nondegeneracy, ensuring that no three distinct points are collinear. This condition requires that at least one of the two possible orientations for a triple must hold. It is encoded as:

$$o_{p,q,r} \vee o_{p,r,q} \tag{2.3}$$

The fourth class of axioms enforces interiority, a condition that ensures consistency between the orientation of a triangle (p, q, r) and the position of a point t lying within it. It is encoded as:

$$o_{t,q,r} \wedge o_{p,t,r} \wedge o_{p,q,t} \implies o_{p,q,r} \tag{2.4}$$

The fifth class of axioms enforces transitivity, a condition that ensures the compatibility of orientation triples across five distinct points. Specifically, if p, q, and r all lie to the left of ts (the directed line formed by t and s), q lies to the left of tp, and r lies to the left of tq, then it must follow that r lies to the left of tp. It is encoded as:

$$o_{t,s,p} \wedge o_{t,s,q} \wedge o_{t,s,r} \wedge o_{t,p,q} \wedge o_{t,q,r} \implies o_{t,p,r} \tag{2.5}$$

Together, these five classes of axioms serve to constrain the orientation variables in a manner that approximates the behavior of realizable point sets in the Euclidean plane. However, they do not fully encode the complete set of geometric constraints necessary to ensure that every satisfying assignment corresponds to an actual geometric realization. This limitation arises from the fact that the realizability of a point set is fundamentally an algebraic condition, rooted in the sign of the determinant of certain coordinate matrices, which cannot be easily captured within the confines of propositional logic. Despite this, the axioms do enforce local logical consistency among the orientation variables, thereby significantly narrowing the search space and guiding the solver toward solutions that are more likely to correspond to valid realizations.

2.2 An $O(n^4)$ Encoding

In "Happy Ending: An Empty Hexagon in Every Set of 30 Points" [3], Heule and Scheucher introduce a SAT-based approach to forbidding convex 6-gons among point placements utilizing a novel encoding that improves upon the efficiency of earlier formulations. Central to their method is the assumption of a fixed ordering of the points 1 through 16 from left to right. This assumption enables a redefinition of the cc axioms that reduces the number of required clauses to $O(n^4)$, where n is the total number of points, an improvement over the $O(n^5)$ axioms found in Knuth's earlier work.

In addition to the optimized axiom encoding, their approach leverages auxiliary variables to forbid the presence of convex k-gons using only $O(n^4)$ clauses, a substantial improvement over the naive $O(n^k)$ formulation. Furthermore, preliminary symmetry-breaking constraints are introduced to reduce redundancies in the search space, further enhancing the efficiency of the encoding. Together, these techniques contribute to an efficient encoding that will form the basis of the CNF used during the analysis portion of this work.

2.2.1 $O(n^4)$ Signotope Axioms

The encoding developed by Heule and Scheucher to forbid convex 6-gons among point placements utilizes orientation triples to represent the combinatorial geometry problem as a SAT instance. Consequently, new axioms (referred to as signotope axioms) are required to ensure that every satisfying assignment maintains local logical consistency.

The following describes the method employed to forbid non-realizable patterns. Consider the set of 4 points (a, b, c, d). Using the assumption that the points are placed in this order from left to right, the directed lines ab, bc, and cd partition the region in which point d can be placed into 4 subregions. These subregions are illustrated in Figure 2.2. The corresponding orientation triples for each subregion are also illustrated in Table 2.1. When the orientation triple $o_{a,b,c}$ is positive, the third orientation variable being true implies that the second is also true. Similarly, the fourth orientation triple being true implies that the third orientation triple is true. This relationship can be encoded using the following clauses:

$$o_{a,b,c} \wedge o_{a,c,d} \implies o_{a,b,d} \tag{2.6}$$

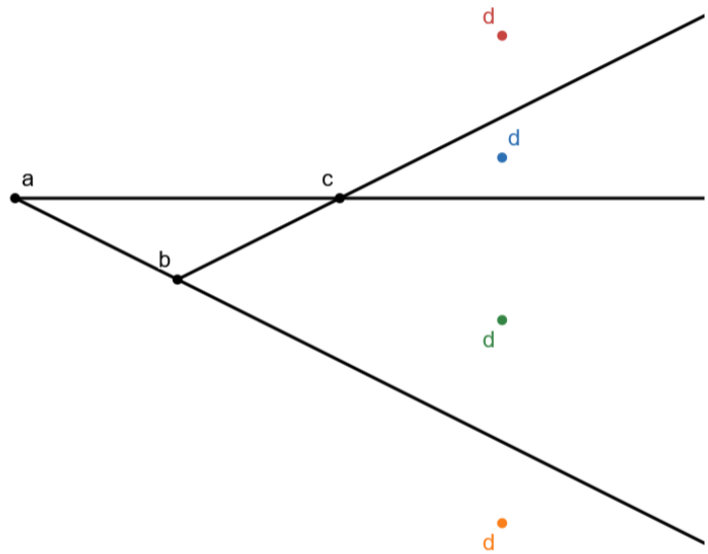

Figure 2.2: A visual depiction of the four possible subregions in which the point d can lie with respect to the points a, b, and c. The mirror case in which point b lies above the line ac is not shown.

$$o_{a,b,c} \land o_{b,c,d} \implies o_{a,c,d} \tag{2.7}$$

The mirror case in which the triple $o_{a,b,c}$ is negative can be represented using the same logic, just inverted. This relationship can be encoded using the following clauses:

$$\neg o_{a,b,c} \land \neg o_{a,c,d} \implies \neg o_{a,b,d} \tag{2.8}$$

$$\neg o_{a,b,c} \land \neg o_{b,c,d} \implies \neg o_{a,c,d} \tag{2.9}$$

For every four-element subset of the 16 points, the presence of all four sets of these clauses helps to ensure local logical consistency and makes each satisfying assignment closer to being a realizable instance in the plane. There are $\binom{n}{4}$ of these subsets, resulting in a total of $O(n^4)$ clauses added to the CNF.

2.2.2 Forbidding 6-gons

The method used by Heule and Scheucher to forbid the formation of convex 6-gons relies heavily on the use of auxiliary variables. Specifically, two types of auxiliary variables are introduced: u-variables, which encode convex polygons whose interior points lie entirely above the line formed by their endpoints, and v-variables, which encode those whose interior points lie below this line. These variables are employed in a case split to systematically eliminate all possible configurations that could result in a convex 6-gon.

$o_{a,b,c}$	$o_{a,b,d}$	$o_{a,c,d}$	$o_{b,c,d}$
+	+	+	+
+	+	+	-
+	+	-	-
+	-	-	-
-	-	-	-
-	-	-	+
-	-	+	+
-	+	+	+

Table 2.1: A table showcasing the corresponding orientation triples for each of the four regions in Figure 2.2 (the mirror case is shown but not explicitly colored).

U-Variables

As noted above, each u-variable represents a convex polygon whose interior points lie entirely above the line formed by its endpoints. Each variable includes a superscript denoting the size of the polygon and a subscript consisting of three points: the left endpoint, the rightmost non-endpoint, and the right endpoint of the polygon. For example, the variable corresponding to a convex 4-gon with endpoints a and d, and c as the rightmost non-endpoint, is written as:

$$u^4_{a,c,d}$$

In each u-variable, the remaining interior points of the polygon are not explicitly named; instead, they are implicitly encoded. As a result, each u-variable represents all possible convex polygons of the specified size that can be formed using the given three defining points. When using these variables to forbid the formation of all convex k-gons, the specific identities of the interior points are not relevant; what matters is that the variable collectively encodes all such formations. An illustration that more clearly conveys the geometric interpretation of a u-variable is shown in Figure Figure 2.3.

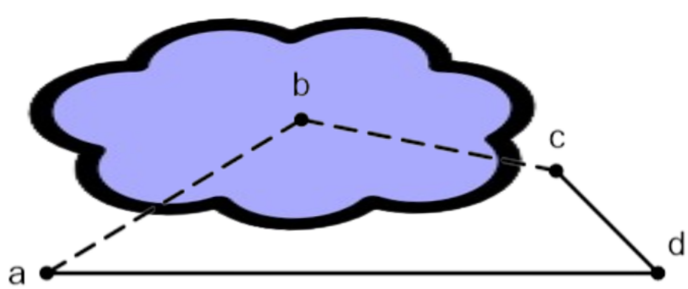

Figure 2.3: A visual representation of the auxiliary variable $u^4_{a,c,d}$. The points a, c, and d are explicitly defined while it is assumed that there exists some point b that completes the convex 4-gon.

U-variables are encoded using a recursive approach, with the base case corresponding to convex polygons of size 4. For each four-element subset of the sixteen points (a, b, c, d), the base case is defined by the following implication: If the orientation triples indicate that the points b and c lie above the line ad, specifically, if $o_{a,b,c}$ and $o_{b,c,d}$ are both false, then a convex 4-gon is formed and the corresponding u-variable $u^4_{a,c,d}$ is set to true:

$$\neg o_{a,b,c} \wedge \neg o_{b,c,d} \implies u^4_{a,c,d} \tag{2.10}$$

The recursive case builds upon the base case by extending smaller u-variables to larger ones. For a given set of points (a, b, c, d), the u-variable $u^x_{a,c,d}$ represents a convex polygon of size x with endpoints a and d, and c as the rightmost non-endpoint. This variable is encoded recursively in terms of the previously constructed u-variable $u^{x-1}_{a,b,c}$. If $u^{x-1}_{a,b,c}$ is true, this indicates the existence of a convex polygon of size $x-1$ with rightmost non-endpoint b and endpoint c. If the orientation triple $o_{b,c,d}$ is false, placing the point d below the line bc, then the polygon $u^{x-1}_{a,b,c}$ can be extended to include a new endpoint d such that all interior points lie above the line ad. This relationship is encoded by the following implication:

$$u^{x-1}_{a,b,c} \wedge \neg o_{b,c,d} \implies u^x_{a,c,d} \tag{2.11}$$

V-Variables

V-variables function analogously to u-variables, with the key distinction that they represent convex polygons whose interior points lie entirely below the line formed by their endpoints. As with u-variables, each v-variable includes a superscript indicating the size of the polygon, and a subscript consisting of three points: the left endpoint, the rightmost non-endpoint, and the right endpoint.

The recursive encoding scheme for v-variables mirrors that of the u-variables, with the exception that the relevant orientation triples must be true rather than false. In the base case, for a convex 4-gon formed by points (a, b, c, d), if the orientation triples $o_{a,b,c}$ and $o_{b,c,d}$ are true, indicating that the points b and c lie below the line ad, then the corresponding v-variable is set to true:

$$o_{a,b,c} \wedge o_{b,c,d} \implies v^4_{a,c,d} \tag{2.12}$$

In the recursive case, for a given set of points (a, b, c, d), the variable $v^x_{a,b,c}$ is defined in terms of the previously constructed variable $v^{x-1}_{a,b,c}$. If $v^{x-1}_{a,b,c}$ is true, this indicates the existence of a convex polygon of size $x - 1$ with rightmost non-endpoint b and endpoint c. If the orientation triple $o_{b,c,d}$ is true, placing the point d above the line bc, then the polygon $v^{x-1}_{a,b,c}$ can be extended to include a new endpoint d such that all interior points lie below the line ad. This relationship is encoded by the following implication:

$$v^{x-1}_{a,b,c} \wedge o_{b,c,d} \implies v^x_{a,c,d} \tag{2.13}$$

Forbidding all 6-gons

To eliminate all possible convex 6-gons from a given point configuration, a case-based approach is employed. Any convex 6-gon must have two endpoints, with the remaining four points lying strictly between them. These interior points may be located above or below the line defined by the endpoints. This results in five distinct cases, based on the distribution of the interior points: all four above the line, three above and one below, two on each side, one above and three below, and all four below the line.

Each of these cases is handled using the previously defined u- and v-variables, which represent convex polygons with interior points strictly above or below the line formed by their endpoints, respectively. To minimize the number of variables and clauses in the SAT encoding, only u- and v-variables up to size five are needed. Then, for each four-element subset (a, b, c, d) drawn from the sixteen points, five sets of clauses are added, one for each case, ensuring that no satisfying assignment corresponds to a configuration containing a convex 6-gon.

In the first case, all four interior points lie below the line that connects the end points. This configuration can be ruled out by asserting that there cannot exist a convex 5-gon encoded by $v^5_{a,b,c}$ while the orientation triple $o_{b,c,d}$ is false. This would imply that point d lies above line bc and a convex 6-gon could be created by extending the convex 5-gon with point d. To forbid this, the following clauses are added:

$$\neg(v^5_{a,b,c} \land \neg o_{b,c,d}) \tag{2.14}$$

In the second case, three interior points lie below the line that connects the end points and one interior point lies above. This configuration can be ruled out by asserting that there cannot exist a convex 5-gon encoded by $v^5_{a,b,d}$ while the orientation triple $o_{a,c,d}$ is true, as well as asserting that there cannot exist a convex 5-gon encoded by $v^5_{a,c,d}$ while the orientation triple $o_{a,b,d}$ is true. In the first case, this would imply that a convex 6-gon could be created by extending the convex 5-gon with point c, and in the second case this would imply that a convex 6-gon could be created by extending the convex 5-gon with point b. To forbid both, the following clauses are added:

$$\neg(v^5_{a,b,d} \land o_{a,c,d}) \tag{2.15}$$

$$\neg(v^5_{a,c,d} \land o_{a,b,d}) \tag{2.16}$$

In the third case, two interior points lie below the line that connects the end points and two interior points lie above. This configuration can be ruled out by asserting that there cannot exist a convex 4-gon encoded by $v^4_{a,b,d}$ and a convex 4-gon encoded by $u^4_{a,c,d}$, as well as by asserting that there cannot exist a convex 4-gon encoded by $v^4_{a,c,d}$ and a convex 4-gon encoded by $u^4_{a,b,d}$. In both cases, this would imply that a convex 6-gon could be created by the extremal points that result from combining the two convex 4-gons. To forbid both, the following clauses are added:

$$\neg(v^4_{a,b,d} \land u^4_{a,c,d}) \tag{2.17}$$

$$\neg(v^4_{a,c,d} \land u^4_{a,b,d}) \tag{2.18}$$

In the fourth case, one interior point lies below the line that connects the end points and three interior points lie above. This configuration can be ruled out by asserting that there cannot exist a convex 5-gon encoded by $u^5_{a,b,d}$ while the orientation triple $o_{a,c,d}$ is false, as well as asserting that there cannot exist a convex 5-gon encoded by $u^5_{a,c,d}$ while the orientation triple $o_{a,b,d}$ is false.

In the first case, this would imply that a convex 6-gon could be created by extending the convex 5-gon with point c, and in the second case this would imply that a convex 6-gon could be created by extending the convex 5-gon with point b. To forbid both, the following clauses are added:

$$\neg(u^5_{a,b,d} \wedge \neg o_{a,c,d}) \tag{2.19}$$

$$\neg(u^5_{a,c,d} \wedge \neg o_{a,b,d}) \tag{2.20}$$

In the fifth and final case, all four interior points lie above the line that connects the end points. This configuration can be ruled out by asserting that there cannot exist a convex 5-gon encoded by $u^5_{a,b,c}$ while the orientation triple $o_{b,c,d}$ is true. This would imply that point d lies below line bc and a convex 6-gon could be created by extending the convex 5-gon with point d. To forbid this, the following clauses are added:

$$\neg(u^5_{a,b,c} \wedge o_{b,c,d}) \tag{2.21}$$

Symmetry Breaking

Heule and Scheucher also introduced symmetry-breaking clauses that take advantage of the fixed left-to-right ordering of the points. These clauses help reduce the search space without eliminating any valid satisfying assignments. Specifically, a lemma provided in the original paper shows that the points labeled 2 through 16 can be assumed, without loss of generality, to appear around point 1 in counterclockwise order, while still preserving the left-to-right ordering. This additional constraint can be encoded using the following clauses for all $a < b < 16$:

$$o_{1,a,b} \tag{2.22}$$

Together, the components described above, including the signotope axioms, the recursive construction of auxiliary u- and v-variables, the exhaustive case-based constraints to eliminate all convex 6-gons, and the symmetry-breaking clauses, combine to form a CNF formula that encodes hexagon-free planar point placements in $O(n^4)$ clauses. This encoding will serve as a foundational component in the analysis and experimental investigations presented in the remainder of this thesis.

Chapter 3

Methodology

The core contributions of this thesis can be broadly divided into two main areas: enforcing structure and enforcing symmetry. The enforcing structure phase aimed to develop a deeper understanding of the structural properties underlying hexagon-free point placements on 16 points. This was achieved using a computationally efficient SAT encoding, which prioritized speed and allowed for a detailed examination of structural patterns that occur in placements with explicitly layered convex hull formations through the analysis of millions of candidate point placements.

In contrast, the enforcing symmetry phase focused on the concrete construction of point sets. This part of the work employed a more flexible encoding in combination with a powerful geometric realization tool. The goal of this phase was to discover explicit coordinate realizations of hexagon-free symmetric configurations, thus bridging the gap between abstract combinatorial encodings and concrete geometric embeddings.

3.1 Enforcing Structure

The structural component of this thesis begins by introducing an efficient SAT encoding designed to represent layered convex hulls within hexagon-free planar point sets on sixteen points. Using this encoding, the full space of candidate point placements consistent with each convex hull formation was exhaustively explored. In addition to identifying valid hexagon-free configurations, the internal structure of each candidate was further examined by counting the number of convex 4- and 5-gons present. This analysis provides insight into the combinatorial geometry of these configurations and highlights their inherent tendency to suppress higher-order convex substructures.

3.1.1 General Structure

Simply forbidding all convex 6-gons over a set of 16 points produces solutions that appear nearly random and are difficult to interpret or analyze. For reference, consider the solution currently displayed on Wikipedia for the Erdős–Szekeres conjecture [7], shown in Figure 3.1. By deliberately introducing structural constraints into the encoding, the search space is not only substantially reduced but also restricted to a well-defined and geometrically meaningful subset of configurations.

To guide this structured exploration, clauses were added to enforce specific layered convex hull formations with uniform sizes, such as $(3,3,3,3,3,1)$, $(4,4,4,4)$, and $(5,5,5,1)$. These formations naturally introduce both radial and rotational symmetry into the solution space. Observe the parallels between the rotationally symmetric figures and the layered convex hull figures shown in Figure 3.2. Additional clauses partition the interior of each convex layer into sectors, also shown in Figure 3.2, with each sector required to contain an equal number of interior points. To support exhaustive analysis, the encoding systematically accounts for all possibilities of how points can be distributed between layers and how they can be placed within each shell relative to the line formed by the shell's endpoints, illustrated in Figure 3.3. Together, these structural constraints expose the deeper geometric restrictions imposed by the forbidding of convex 6-gons and provide a more focused foundation for analyzing the structural properties of the resulting configurations.

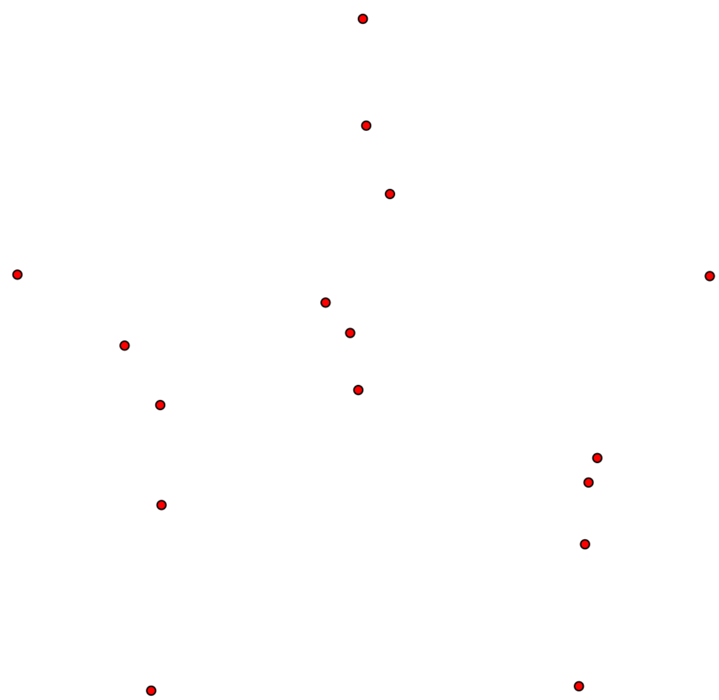

Figure 3.1: A set of 16 points in general position that avoids the formation of any convex 6-gons.

3.1.2 Analysis

Using each of the encodings, the complete space of candidate point placements consistent with each convex hull structure was exhaustively explored. This process involved evaluating all possible combinations of point assignments across the designated convex layers, along with every feasible orientation of points within each hull: specifically, determining whether each point lies above or below the central axis of the convex hull.

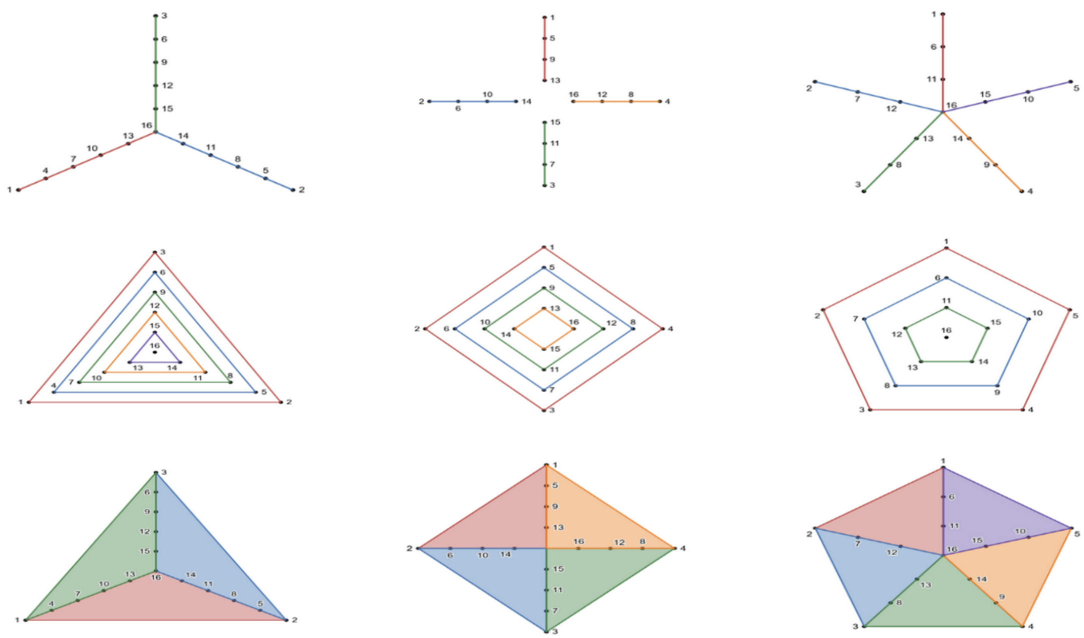

Figure 3.2: A visual comparison of point placements exhibiting rotational symmetry and their corresponding layered structures. The first row shows realizations of 16-point configurations with 3-fold, 4-fold, and 5-fold rotational symmetry, respectively. The second row reinterprets these configurations as layered convex hulls, each with layers of equal size corresponding to the symmetry order. The third row illustrates a sector-based decomposition of each configuration, where the interior of each layered hull is evenly partitioned into sectors to reflect the imposed symmetry.

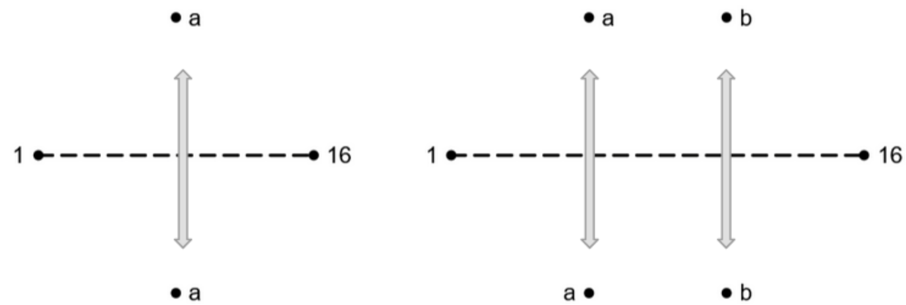

Figure 3.3: This figure illustrates the first layer configurations for the 3-fold and 4-fold symmetry cases. In each case, the endpoints (Points 1 and 16) remain fixed due to the sorted ordering of points, while the interior points are allowed to vary. These interior points can assume any of the remaining values and may shift in position relative to the central axis, enabling an exhaustive exploration of structurally distinct configurations.

For each candidate configuration, a distinct CNF encoding was generated, resulting in potentially millions of unique formulas. These encodings were then fed to the SAT solver CaDiCaL. Due to the efficiency of the encoding, CaDiCaL was able to quickly determine the satisfiability of each formula. The total number of satisfiable encodings was recorded, along with the total number of candidate configurations, to calculate a satisfiability ratio for each hull structure.

In addition to satisfiability testing, each satisfying assignment was subjected to a secondary analysis to investigate the presence of smaller convex polygon formations. This analysis was carried out using a specialized tool developed by Marijn Heule, designed specifically for CNF formulas of this kind. For each satisfying assignment, the tool reported the number of convex 4- and 5-gons present. These values were aggregated across all realizations, and the corresponding minimum, maximum, and average counts were calculated.

3.1.3 Base CNF

The base CNF encoding used for this phase of the investigation was the $O(n^4)$ formulation developed by Heule and Scheucher, as described in detail in Chapter 2. This encoding was selected for its computational efficiency, which was critical given the exhaustive nature of the structural analysis, which required potentially millions of SAT solver instances to be executed sequentially. Since the focus at this stage was not on producing explicit planar realizations, the additional ordering constraints imposed by the encoding were not a limitation. In contrast, these constraints preserved satisfiability while substantially reducing the size of the search space, enabling more tractable large-scale analysis.

3.1.4 Layered 3-Gons

The first layered convex hull formation considered was $(3, 3, 3, 3, 3, 1)$, comprising of five nested convex layers of three points each, arranged around a single central point. To enforce additional structure, the interior of each convex hull was partitioned into three sectors, with each sector defined by two adjacent points on the hull and the central point. The distribution of interior points within these sectors followed a uniform pattern: four points per sector in the outermost layer, three in the second layer, two in the third, and one in the fourth. To comprehensively explore the space of candidate point placements, all combinations of point assignments to each convex layer were examined. Additionally, for each convex hull, the orientation of its middle point, whether positioned above or below the line formed by the hull's endpoints, was allowed to vary.

Convexity

To enforce convexity within each of the three-point convex hull layers, orientation constraints were encoded based on the relative position of the middle point. Let the convex hull be composed of three ordered points (a, b, c) with $a < b < c$. If the middle point b is intended to lie above the line formed by a and c, convexity is enforced by adding the following clause, illustrated in Figure 3.4:

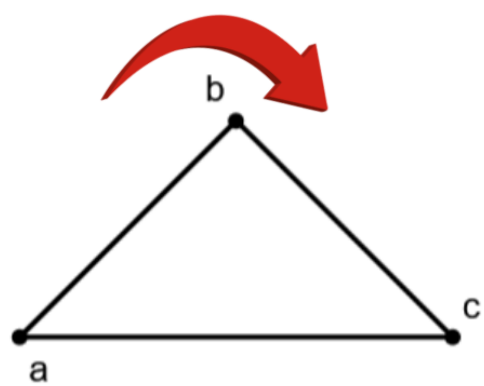

Figure 3.4: An illustration of how convexity in the first case is enforced by adding clause $\neg o_{a,b,c}$.

$$\neg o_{a,b,c} \tag{3.1}$$

Conversely, if b is intended to lie below the line ac, the opposite orientation clause is added:

$$o_{a,b,c} \tag{3.2}$$

Containment

To enforce the containment of the remaining points within each sequential convex hull layer, orientation clauses are added for each point p that is meant to lie inside a triangular hull formed by points (a, b, c). Since point indices are assumed to be ordered from left to right, $a < p < c$ is assumed as interior placement is otherwise not possible. Within this range, two distinct orderings must be considered: $a < p < b < c$ and $a < b < p < c$. For each order, two geometric cases are handled depending on whether the middle point b lies above or below the line formed by a and c.

If b lies above line ac and $p < b$ the following clauses are added to force p to lie below line ab and above line ac, illustrated in Figure 3.5:

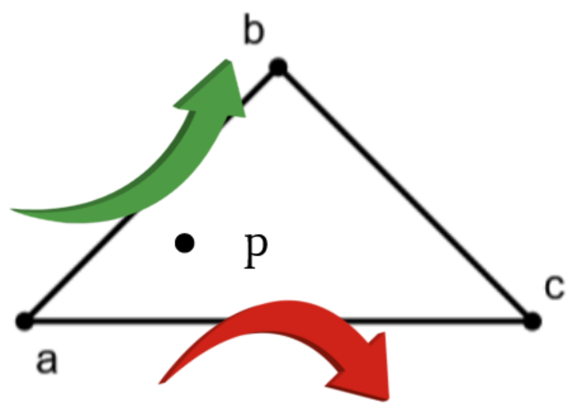

Figure 3.5: An illustration of how containment in the first case is enforced by adding $o_{a,p,b} \wedge \neg o_{a,p,c}$.

$$o_{a,p,b} \wedge \neg o_{a,p,c} \tag{3.3}$$

If b lies above line ac and $b < p$ the following clauses are added to force p to lie below line bc and above line ac:

$$o_{b,p,c} \wedge \neg o_{a,p,c} \tag{3.4}$$

If b lies below line ac and $p < b$ the following clauses are added to force p to lie below line ac and above line ab:

$$o_{a,p,c} \wedge \neg o_{a,p,b} \tag{3.5}$$

If b lies below line ac and $b < p$ the following clauses are added to force p to lie below line ac and above line bc:

$$o_{a,p,c} \wedge \neg o_{b,p,c} \tag{3.6}$$

Sectors

To ensure that within the same convex hull layer, every sector, defined by two adjacent points on the convex hull and a central point, contains an equal number of interior points, a careful case analysis was required. Consider a convex hull formed by three ordered points (a, b, c) enclosing a central point x. This configuration partitions the interior into three sectors, defined by the triangles (a, b, x), (a, c, x), and (b, c, x), which we refer to as sectors 1, 2, and 3, respectively. To ensure that each sector contains exactly m interior points, auxiliary variables were introduced. For example, the variable $s_1^2 p$ indicates whether a point p lies within sector 1 of the second convex hull layer. In this notation, the superscript denotes the convex hull layer and the subscript denotes the sector index.

Focusing on sector 1 of the first convex layer, defined by the triangle (a, b, x), $s_1^1 p$ must determine whether a point p lies within this triangle. It can be assumed that a is the leftmost of the three convex hull points, resulting in $a < b$, $a < x$, and $a < p$ in the order from left to right. The relative orders among p, x, and b yield six distinct permutations:

1. $a < p < x < b$
2. $a < p < b < x$
3. $a < x < p < b$
4. $a < x < b < p$
5. $a < b < p < x$
6. $a < b < x < p$

Among these, cases 4 and 6 can be eliminated immediately, as they place point p to the right of both b and x, making it impossible for p to lie within the triangle (a, b, x). For the remaining valid orderings, the orientations of the point p with respect to the boundary lines are used to determine whether p lies within the sector. Assume first that the point b lies below the line ac. In case 1, the inclusion of p in sector 1 can be characterized by the following logic:

$$s_1^1 p \iff (o_{a,p,x} \land \neg o_{p,x,b}) \tag{3.7}$$

If instead b lies above the line ac, the orientation conditions are reversed, as illustrated in Figure 3.6:

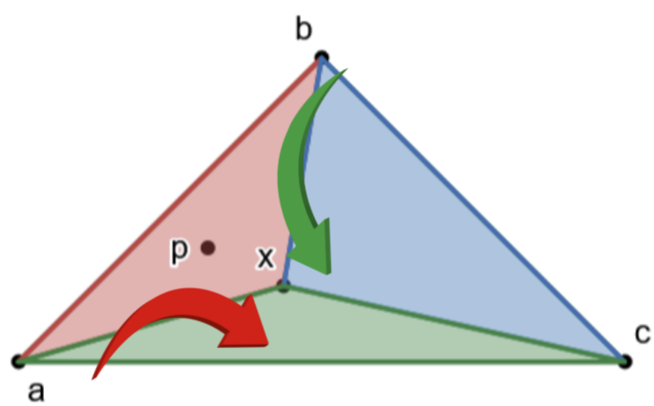

Figure 3.6: An illustration of how the auxiliary variable for point p lying in sector 1 can be enforced with respect to the lines ax and bx.

$$s_1^1 p \iff (\neg o_{a,p,x} \land o_{p,x,b}) \tag{3.8}$$

The expressions for the other valid cases follow similar logic and are not explicitly shown. To ensure that no more than m points are assigned to any sector, a cardinality constraint is encoded using the auxiliary variables. Specifically, for each subset of $m+1$ candidate points $\{p_1, p_2, \ldots, p_{m+1}\}$, the following clause is added to prevent all of them from being assigned to the same sector:

$$\left(\neg s_1^1 p_1 \vee \neg s_1^1 p_2 \vee \cdots \vee \neg s_1^1 p_{m+1}\right) \tag{3.9}$$

This enforces that at least one of the $m+1$ points does not lie in sector 1. The same approach is applied to all sectors in every convex hull layer. For brevity, the full enumeration of all the case distinctions and clauses is omitted.

Results and Discussion

In total, 14,080 candidate point placements were consistent with this formation. Of these, 4,984 were satisfiable, resulting in a satisfiability ratio of 35%.

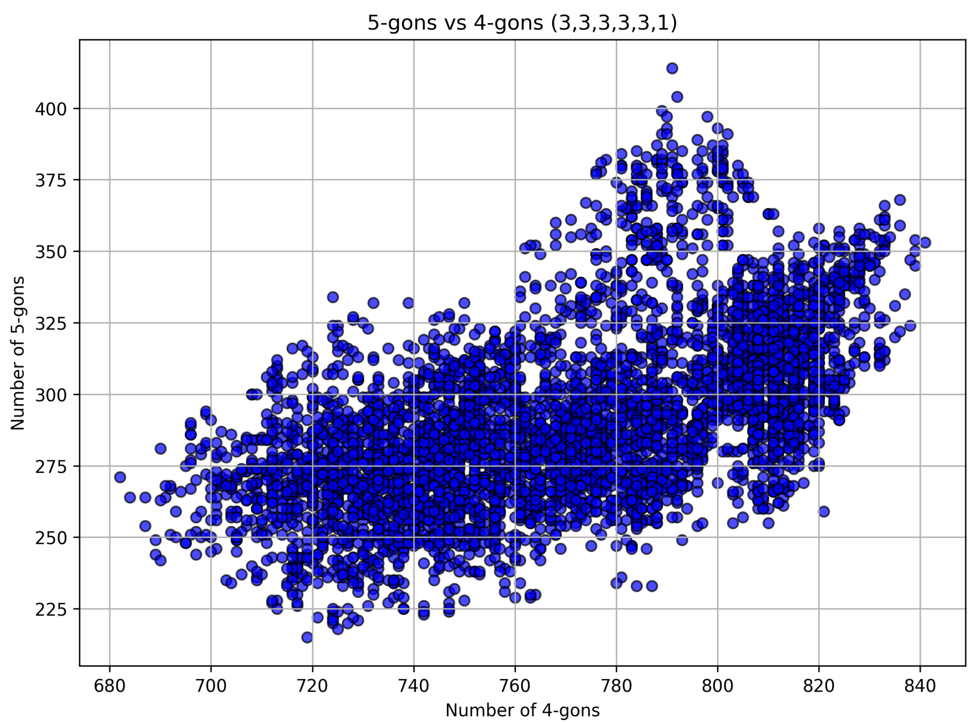

Figure 3.7: The number of convex 4-gons plotted on the x axis vs the number of convex 5-gons plotted on the y axis in the $(3,3,3,3,3,1)$ case. Each single data point represents one of the 4,984 satisfiable instances.

Among the three layered convex hull formations examined, the layered 3-gon configuration yielded the fewest total candidate point placements, with 14,080 possibilities. This reduction in search space is expected, as each convex hull in this formation includes only one central point whose orientation must be explicitly constrained, resulting in fewer combinations to evaluate.

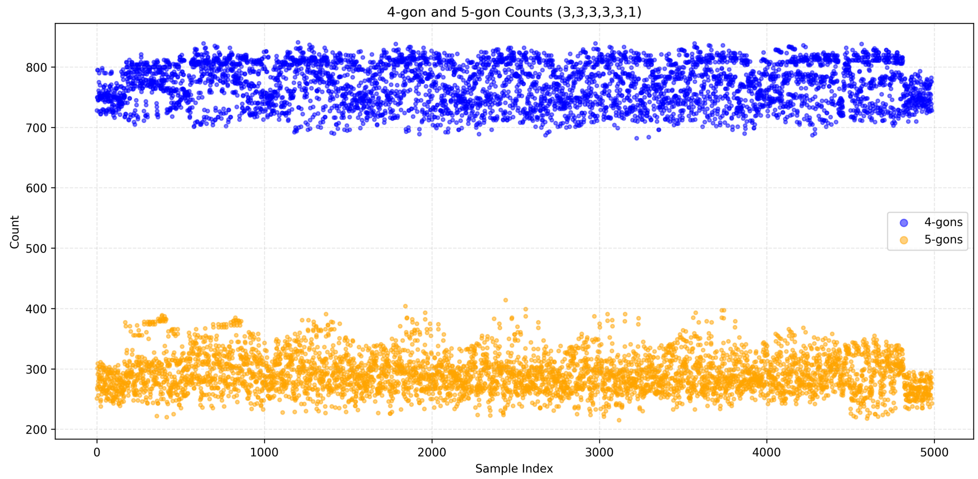

Figure 3.8: The number of convex 4-gons (blue) and convex 5-gons (orange) for each satisfying assignment in the $(3,3,3,3,3,1)$ case. The x axis sorts the instances from left to right. The leftmost possible instance in this case is shell 1 = [1,2,16], shell 2 = [3,4,15], shell 3 = [5,6,14], shell 4 = [7,8,13], shell 5 = [9,10,12], and center point = 11. The rightmost possible instance in this case is shell shell 1 = [1,15,16], shell 2 = [2,13,14], shell 3 = [3,11,12], shell 4 = [4,9,10], shell 5 = [5,7,8], and center point = 6. The y axis represents the counts of each of the different amounts of the specified convex substructures in each satisfiable instance.

Of these 14,080 candidates, 4,984 were found to be satisfiable by the SAT solver, corresponding to a satisfiability ratio of approximately 35%. This relatively high ratio suggests that a significant proportion of the candidate placements are geometrically valid under the imposed symmetry and convexity constraints. One possible explanation for this observation is that 3-gon layers, having the smallest size among the three formations, are less likely to form convex 6-gons, which aligns with the goal of identifying 6-gon-free configurations.

When examining the occurrence of convex 4-gons and 5-gons, Figure 3.7 indicates that there is no clear trend or correlation between their quantities—points appear relatively evenly scattered throughout the space. In contrast, Figure 3.8 reveals a consistent pattern: the number of convex 4-gons and 5-gons remains relatively stable across the satisfiable instances, averaging approximately 750 4-gons and 300 5-gons. This suggests a structural regularity within the realizable configurations that may be a byproduct of the underlying rotational symmetry and layered construction.

3.1.5 Layered 4-Gons

The second layered convex hull formation considered was $(4,4,4,4)$, comprising of four nested convex layers of four points each. To enforce additional structure, the interior of each convex hull was partitioned into four sectors, with each sector defined by the intersecting diagonal lines of opposite points on the hull. The distribution of interior points within these sectors followed a uniform pattern: three points per sector in the outermost layer, two in the second layer, and one in the third. To comprehensively explore the space of candidate point placements, all combinations of point assignments to each convex layer were examined. Additionally, for each convex hull, the orientations of its two middle points, whether positioned above or below the line formed by the hull's endpoints, were allowed to vary.

Convexity

To enforce convexity for each of the four-point convex hull layers, orientation constraints were encoded based on the relative position of the two middle points. Let the convex hull be composed of four ordered points (a, b, c, d) with $a < b < c < d$. If the middle points b and c are intended to lie above the line formed by a and d, convexity is enforced by adding the following clauses:

$$\neg o_{a,b,c} \wedge \neg o_{b,c,d} \tag{3.10}$$

If b is intended to lie above the line ad and c is intended to lie below the line ad, convexity is enforced by adding the following clauses and is depicted in Figure 3.9:

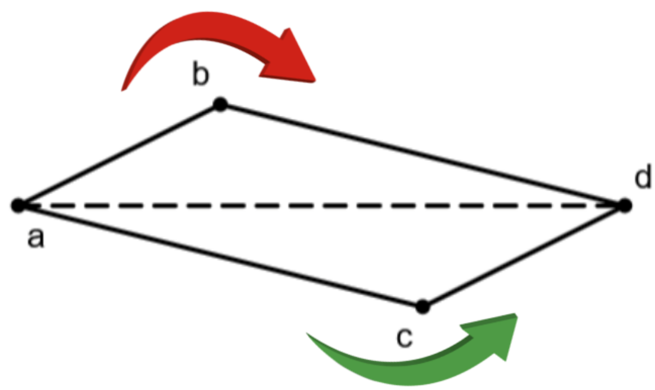

Figure 3.9: An illustration of how convexity is enforced in the case where b lies above the line ad and c lies below it using the clauses $\neg o_{a,b,d} \wedge o_{a,c,d}$.

$$\neg o_{a,b,d} \wedge o_{a,c,d} \tag{3.11}$$

The remaining two cases are just mirror images of the first two and can be encoded using the opposite orientation triples.

Containment

To enforce the containment of the remaining points within each sequential convex hull layer, orientation clauses are added for each point p that is meant to lie inside the convex hull formed by the points (a, b, c, d). Since point indices are assumed to be ordered from left to right, $a < p < d$ is assumed as interior placement is otherwise not possible. Within this range, three distinct orderings must be considered: $p < b < c$, $b < p < c$, and $b < c < p$. For each ordering, each geometric case is handled depending on whether the middle points b and c lie above or below the line formed by a and d.

The first three convexity cases are characterized by the intermediate points b and c lying above the line formed by the points a and d. In these configurations, to ensure that a point p lies within the convex quadrilateral, it must also lie above the line ad. This condition is encoded using the following clause:

$$\neg o_{a,p,d} \tag{3.12}$$

The specific segment for which point p must lie below to remain within the bounds of the convex hull depends on its relative ordering with respect to points b and c:

- If $p < b$, then p must lie below the line segment ab. This is enforced by the clause:

$$o_{a,p,b} \tag{3.13}$$

- If $b < p < c$, then p must lie below the line segment bc. This is enforced by the clause:

$$o_{b,p,c} \tag{3.14}$$

- If $p > c$, then p must lie below the line segment cd. This is enforced by the clause:

$$o_{c,p,d} \tag{3.15}$$

The second group of three convexity cases is characterized by the intermediate point b lying above the line formed by the points a and d, and the point c lying below it.

- If $p < b$, then p must lie below the line segment ab and above the line segment ac. This is enforced with the following clauses and depicted in Figure 3.10:

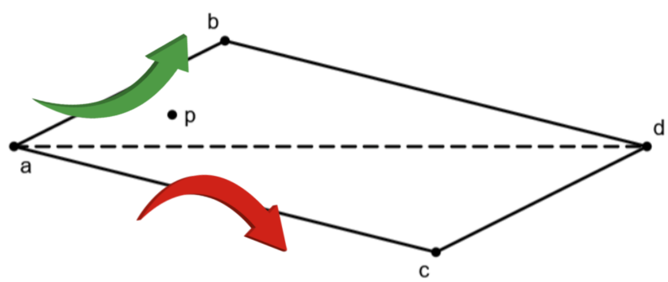

Figure 3.10: An illustration of how containment is enforced in the case where b lies above the line ad, c lies below it, and $p < b$ using the clauses $o_{a,p,b} \land \neg o_{a,p,c}$.

$$o_{a,p,b} \land \neg o_{a,p,c} \tag{3.16}$$

- If $b < p < c$, then p must lie below the line segment bd and above the line segment ac. This is enforced with the clauses:

$$o_{b,p,d} \land \neg o_{a,p,c} \tag{3.17}$$

- If $c < p$, then p must lie below the line segment bd and above the line segment cd. This is enforced with the clauses:

$$o_{b,p,d} \wedge \neg o_{c,p,d} \tag{3.18}$$

The remaining two sets of cases are just mirror images of the first two sets and can be encoded using the opposite orientation triples.

Sectors

To ensure that within the same convex hull layer, each sector contains an equal number of interior points, a careful case analysis was required. Each sector is defined by its orientation with respect to the diagonal lines formed by the first and third vertices versus the second and fourth vertices of the convex hull. The method used to enforce an equal distribution of points across sectors mirrors the approach taken in the layered 3-gon construction. As before, this requires a detailed case split based on the relative ordering and orientation of the underlying points.

Consider a convex hull formed by four ordered points (a, b, c, d). The interior of this hull is partitioned into four sectors, each determined by its position relative to the diagonals defined by the lines ac and bd. As in the 3-gon case, a case split is needed based on whether the central points b and c lie above or below the baseline ad. Furthermore, the relative ordering of the interior point p with respect to b and c must be considered. Auxiliary variables of the form $s_1^1 p$ are once again utilized to represent whether the point p is contained within sector 1 of the first convex hull layer.

Consider the case where points b and c both lie above the line ad. In this scenario, without loss of generality, we define sector 1 as the region above the line ac and below the line bd. Two subcases must be considered: the first is when $a < p < b$, and the second is when $b < p < c$. Note that the configurations $p < a$ and $b < c < p$ are excluded, since a is assumed to be the leftmost point of the convex hull, and a point p to the right of both b and c cannot lie within the defined sector 1.

In the first case, where $p < b$, the following logic is added to equate $s_1^1 p$ with p lying above the line ac and below the line bd, illustrated in Figure 3.11:

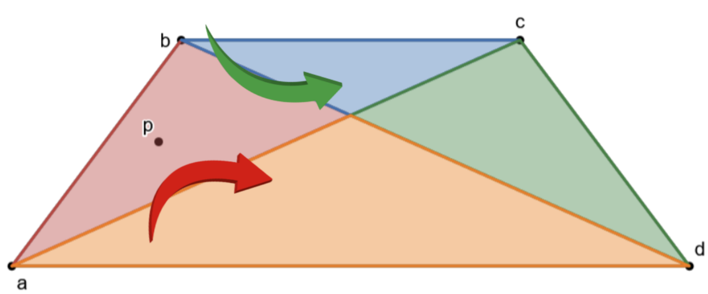

Figure 3.11: An illustration of how the auxiliary variable for point p lying in sector 1 can be enforced with respect to the lines ac and bd.

$$s_1^1 p \iff (\neg o_{a,p,c} \wedge \neg o_{p,b,d}) \tag{3.19}$$

In the second case, where $b < p < c$, the definition becomes:

$$s_1^1 p \iff (\neg o_{a,p,c} \land o_{b,p,d}) \tag{3.20}$$

For brevity, not all possible case splits are shown explicitly, but each follows the same structure to determine the sector membership across all convex hull layers. Once all auxiliary variables have been defined, the same cardinality constraint used for the 3-gon case is applied to ensure that no more than m points are assigned to any given sector.

Results and Discussion

In total, 561,600 candidate point placements were consistent with this formation. Of these, 112,142 were satisfiable, resulting in a satisfiability ratio of 20%.

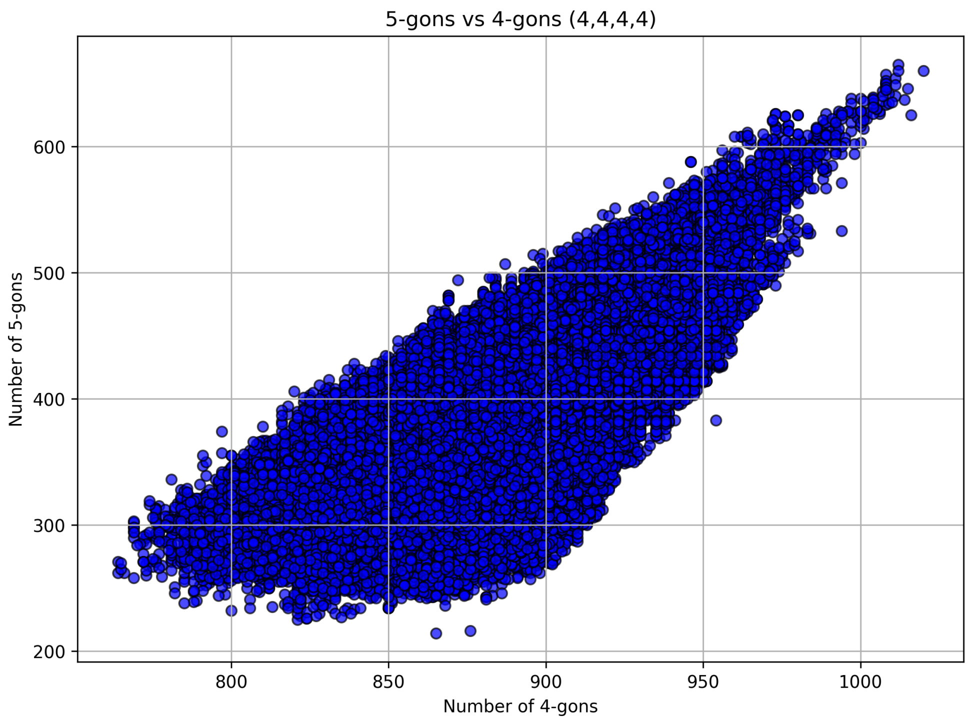

Figure 3.12: The number of convex 4-gons plotted on the x axis vs the number of convex 5-gons plotted on the y axis in the $(4, 4, 4, 4)$ case. Each single data point represents one of the 112,142 satisfiable instances.

The layered 4-gon formation yielded the second-highest number of candidate point placements, totaling 561,600. This increase is expected, as each convex layer now consists of four points, resulting in two internal points per layer that must be considered during orientation casing. Despite the larger search space, 112,142 of these candidate configurations were found to be satisfiable, producing a satisfiability ratio of approximately 20%.

While this ratio represents a decline from the 35% observed in the 3-gon case, it nonetheless constitutes a significant fraction of the total candidates. This reduction in realizability aligns

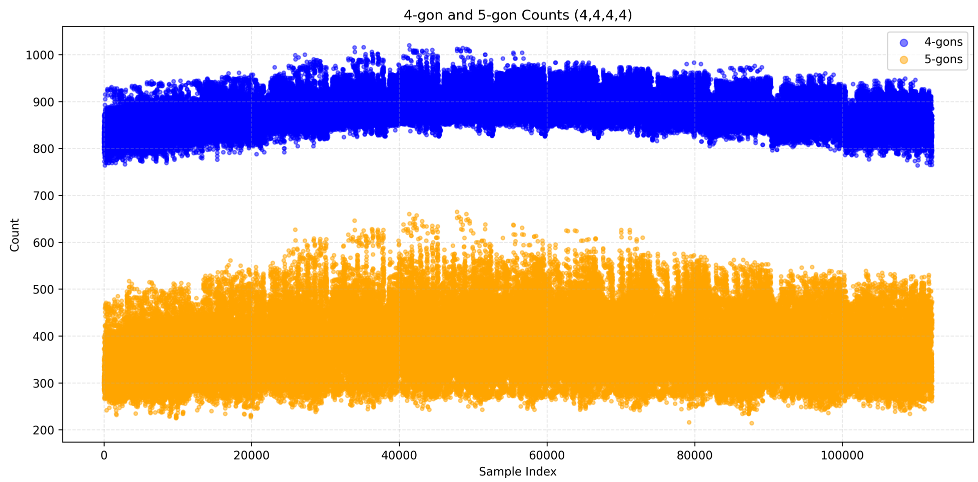

Figure 3.13: The number of convex 4-gons (blue) and convex 5-gons (orange) for each satisfying assignment in the $(4, 4, 4, 4)$ case. The x axis sorts the instances from left to right. The leftmost possible instance in this case is shell 1 = [1,2,3,16], shell 2 = [4,5,6,15], shell 3 = [7,8,9,14], and shell 4 = [10,11,12,13]. The rightmost possible instance in this case is shell 1 = [1,14,15,16], shell 2 = [2,11,12,13], shell 3 = [3,8,9,10], and shell 4 = [4,5,6,7]. The y axis represents the counts of each of the different amounts of the specified convex substructures in each satisfiable instance.

with the increased likelihood of encountering convex 6-gons: layered convex 4-gons are only two points away from forming a convex 6-gon, making the avoidance of such configurations more difficult under the imposed constraints.

The relationship between convex 4-gons and convex 5-gons becomes more evident in this formation. Figure 3.12 reveals a roughly linear distribution, beginning with a broad base and narrowing as the counts of both types of polygons increase. This pattern indicates that configurations with more convex 5-gons tend to also contain proportionally more convex 4-gons—a natural consequence of the fact that every convex 5-gon contains multiple convex 4-gon subsets.

As in the 3-gon case, Figure 3.13 displays a consistent and well-formed structure, with solutions tending to contain around 850 convex 4-gons and 400 convex 5-gons. These counts are notably higher than those found in the 3-gon configuration, which is consistent with the lower satisfiability ratio: fewer configurations can be satisfied, but those that are tend to be denser in convex substructures.

3.1.6 Layered 5-Gons

The final layered convex hull formation considered was $(5, 5, 5, 1)$, comprising of three nested convex layers of five points each, arranged around a single central point. To enforce additional structure, the interior of each convex hull was partitioned into five sectors, with each sector defined by two adjacent points on the hull and the central point. The distribution of interior points within these sectors followed a uniform pattern: two points per sector in the outermost layer, and one point per sector in the second layer. To comprehensively explore the space of candidate point placements, all combinations of point assignments to each convex layer were examined. Additionally, for each convex hull, the orientation of its middle three points, whether positioned above or below the line formed by the hull's endpoints, was allowed to vary.

Convexity

To enforce convexity for each of the five-point convex hull layers, orientation constraints were encoded based on the relative position of the three middle points. Let the convex hull be composed of five ordered points (a, b, c, d, e) with $a < b < c < d < e$. If the middle points b, c, and d are intended to lie above the line formed by a and e, convexity is enforced by adding the following clauses:

$$\neg o_{a,b,c} \wedge \neg o_{b,c,d} \wedge \neg o_{c,d,e} \tag{3.21}$$

If b and c are intended to lie above the line ae and d is intended to lie below the line ae, convexity is enforced by adding the following clauses:

$$\neg o_{a,b,c} \wedge \neg o_{b,c,e} \wedge o_{a,d,e} \tag{3.22}$$

If b and d are intended to lie above the line ae and c is intended to lie below the line ae, convexity is enforced by adding the following clauses and this is illustrated in Figure 3.14:

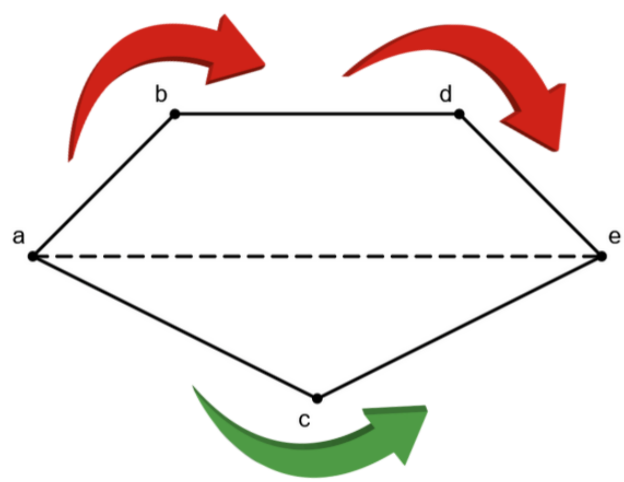

Figure 3.14: An illustration of how convexity is enforced in the case where points b and d lie above the line ae and point c lies below it using the clauses $\neg o_{a,b,d} \wedge \neg o_{b,d,e} \wedge o_{a,c,e}$.

$$\neg o_{a,b,d} \wedge \neg o_{b,d,e} \wedge o_{a,c,e} \tag{3.23}$$

If c and d are intended to lie above the line ae and b is intended to lie below the line ae, convexity is enforced by adding the following clauses:

$$\neg o_{a,c,d} \wedge \neg o_{c,d,e} \wedge o_{a,b,e} \tag{3.24}$$

The remaining four cases are just mirror images of the first four and can be encoded using the opposite orientation triples.

Containment

To enforce the containment of the remaining points within each sequential convex hull layer, orientation clauses are added for each point p that is meant to lie inside the convex hull formed by the points (a, b, c, d, e). Since point indices are assumed to be ordered from left to right, $a < p < e$ is assumed as interior placement is otherwise not possible. Within this range, four distinct orderings must be considered: $p < b < c < d$, $b < p < c < d$, $b < c < p < d$, and $b < c < d < p$. For each ordering, each geometric case is handled depending on whether the middle points b, c, and d lie above or below the line formed by a and e.

Since this process closely parallels the methods used for the 3- and 4-gon constructions, only the first representative case is detailed here. In this case, the points b, c, and d all lie above the line ae, and the point p satisfies $p < b$. To restrict p to lie within the convex hull, it must be bounded above and below by the relevant line segments that define its region within the hull. Specifically, in this configuration, the point p must be below the line ab and above the line ae. This constraint is encoded as and is illustrated in Figure 3.15:

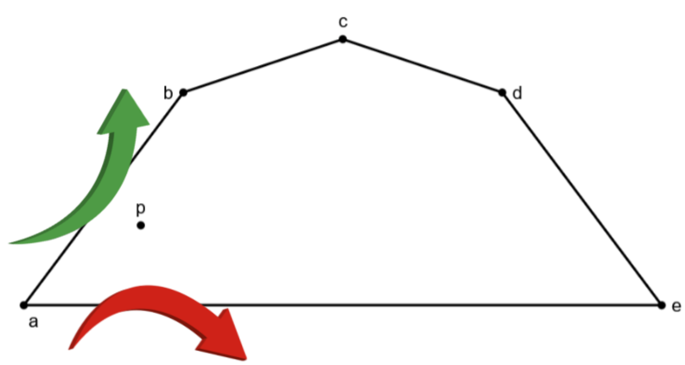

Figure 3.15: An illustration of how containment is enforced when points b, c, and d lie above the line ae and $p < b$ using the clauses $o_{a,p,b} \wedge \neg o_{a,p,e}$.

$$o_{a,p,b} \wedge \neg o_{a,p,e} \tag{3.25}$$

Sectors

As in 3- and 4-gon constructions, ensuring that each sector within a convex hull layer contains an equal number of interior points requires a careful case analysis. Consider a convex hull formed by five ordered points (a, b, c, d, e) that enclose a central point x. In the case where points b, c, and d all lie above the line ae, this configuration partitions the interior into five sectors, defined by the triangles (a, b, x), (b, c, x), (c, d, x), (d, e, x), and (a, e, x), which we refer to as sectors 1 through 5, respectively.

To enforce the constraint that each sector contains exactly m interior points, auxiliary variables of the form $s_1^1 p$ are introduced. The method for assigning these variables follows the same logic established in the 3- and 4-gon constructions.

Within the context described, four subcases arise depending on the relative position of the point p with respect to b, c and d. The first subcase occurs when $p < b$. In this situation, to indicate that p is within sector 1, the following equivalence is used to encode p lying above the line ax and below the line bx and is illustrated in Figure 3.16:

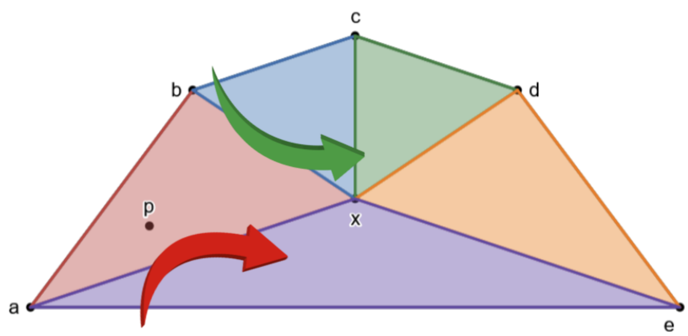

Figure 3.16: An illustration of how the auxiliary variable for point p lying in sector 1 is enforced with respect to lines ax and bx.

$$s_1^1 p \iff (\neg o_{a,p,x} \wedge \neg o_{p,b,x}) \tag{3.26}$$

For brevity, and because this procedure mirrors the approach already described in the 3-gon and 4-gon cases, the remaining subcases are not shown. Once all auxiliary variables have been defined, the same cardinality constraint used for the 3- and 4-gon cases is applied to ensure that no more than m points are assigned to any given sector.

Results and Discussion

In total, 1,677,312 candidate point placements were consistent with this formation. Of these, 9,806 were satisfiable, resulting in a satisfiability ratio of 0.58%.

The layered 5-gon formation produced the highest number of candidate point placements by a considerable margin, totaling 1,677,312 configurations. This result is expected, as each convex layer now contains five points, yielding three interior points per layer whose orientations must be individually cased. The increased degrees of freedom naturally lead to a combinatorial explosion in the number of possible placements.

However, this richness in candidate configurations is met with a dramatic drop in realizability. Only 9,806 of these candidates were found to be satisfiable by the SAT solver, resulting in a satisfiability ratio of merely 0.58%. This stark decrease from the 35% and 20% observed in the 3-gon and 4-gon formations, respectively, is theoretically sound: when convex layers of size 5 are enforced, the addition of even a single point can complete a convex 6-gon. Avoiding such completions under rigid rotational symmetry and layering constraints significantly restricts the solution space.

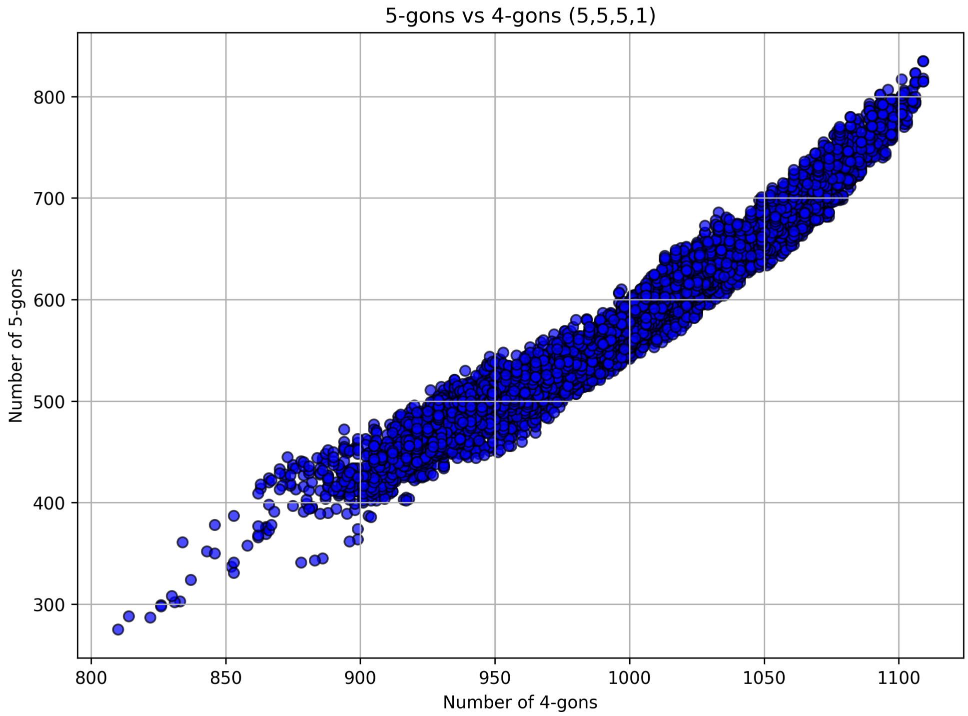

Figure 3.17: The number of convex 4-gons plotted on the x axis vs the number of convex 5-gons plotted on the y axis in the $(5, 5, 5, 1)$ case. Each single data point represents one of the 9,806 satisfiable instances.

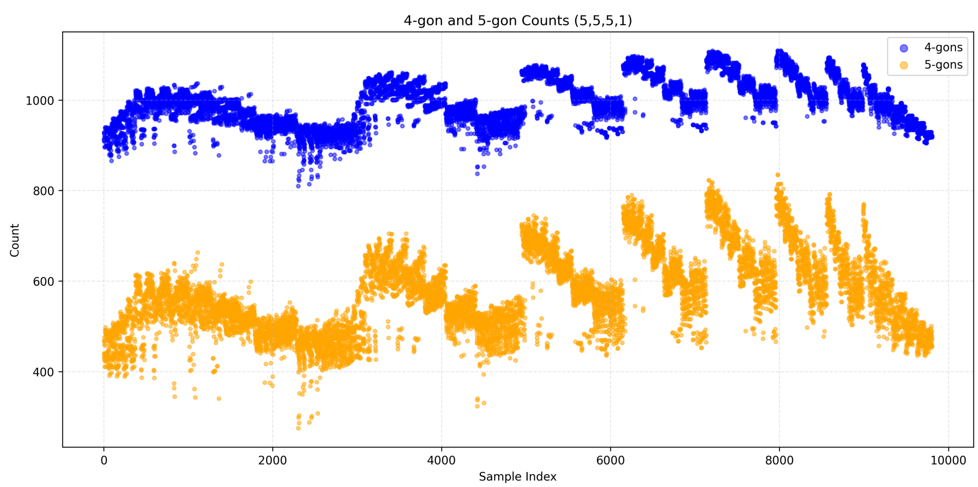

Figure 3.18: The number of convex 4-gons (blue) and convex 5-gons (orange) for each satisfying assignment in the $(5, 5, 5, 1)$ case. The x axis sorts the instances from left to right. The leftmost possible instance in this case is shell 1 = [1,2,3,4,16], shell 2 = [5,6,7,8,15], shell 3 = [6,7,8,9,14], and center point = 13. The rightmost possible instance in this case is shell 1 = [1,13,14,15,16], shell 2 = [2,9,10,11,12], shell 3 = [3,5,6,7,8], and center point = 4. The y axis represents the counts of each of the different amounts of the specified convex substructures in each satisfiable instance.

The structure of the satisfiable instances reveals further insight. In Figure 3.17, plotting the number of convex 4-gons against the number of convex 5-gons, the relationship becomes tightly linear. This suggests that realizable configurations exhibit a very specific structural rigidity—one in which the presence of convex 5-gons strongly dictates the count of convex 4-gons, likely due to the hierarchical nature of their containment.

In contrast, Figure 3.18 deviates from the uniformity observed in the previous formations. Instead of the consistent shapes seen for the 3-gon and 4-gon cases, a wave-like or rippled pattern emerges. This unexpected structure invites further exploration and could be indicative of subtle underlying constraints or periodic behaviors within the realizable 5-gon configurations. The average polygon counts are also the highest of the three, with approximately 1,000 convex 4-gons and over 500 convex 5-gons per solution. These elevated figures are consistent with the greater number of points per hull and further underscore the dense combinatorial environment in which these configurations reside.

This irregular yet pronounced structure in the results opens new directions for future work, particularly in understanding the nature of these ripples and what they reveal about the geometry and symmetry of near-critical configurations.

3.2 Enforcing Symmetry

The enforcing symmetry component of this thesis begins by introducing a more flexible, albeit computationally less efficient, SAT encoding designed to capture rotational symmetry within hexagon-free planar point sets on sixteen points. This encoding allows for the precise specification of rotationally symmetric configurations. The resulting CNF formulas were then evaluated using a SAT solver, and satisfiable instances were subsequently passed to a realization tool specifically designed for orientation-triple-based encodings. This tool generated explicit coordinate placements of planar point sets that satisfied both the imposed symmetry and geometric constraints.

3.2.1 General Structure

This section extends the methodology introduced in the structural component of the thesis. As in the previous phase, specific layered convex hull configurations with uniform sizes, such as $(3, 3, 3, 3, 3, 1)$, $(4, 4, 4, 4)$, and $(5, 5, 5, 1)$ were selected as target structures. In addition to the layered convex hull clauses, the previously defined sector-based constraints were incorporated to ensure that the points are evenly distributed radially. Building upon the earlier structure clauses, this stage also introduces new constraints that explicitly enforce rotational symmetry of orders 3, 4, and 5. Together, these structural and symmetry constraints serve a dual purpose: they narrow the search space, thereby greatly enhancing the efficiency of the realization tool, and they promote the emergence of highly regular, symmetric geometric configurations.

In contrast to the exhaustive search strategy employed earlier, this phase adopts a more intentional construction. Each point is carefully placed in the plane to satisfy the required rotational mappings. The resulting encodings, shaped by the interplay between structural and symmetry-enforcing clauses, yield solutions that visually spiral outward from a central origin in evenly spaced, symmetric branches. This striking pinwheel-shaped pattern is the conceptual and visual basis for the title of the thesis.

3.2.2 Base CNF

The base CNF encoding used to restrict the formation of convex 6-gons in this section adopts a different approach from previous constructions by eliminating the requirement that points be sorted from left to right. This modification provides greater flexibility, enabling explicit placement of points into specific convex layers without enforcing global ordering. To ensure that each satisfying assignment corresponds to a valid geometric realization, the encoding relies on the Knuth CC axioms introduced in Chapter 1. These axioms are inherently order-independent, making them well-suited for this generalized point placement strategy.

To prevent the formation of convex 6-gons, a new approach is employed based on the following geometric insight: a set of six points forms a convex 6-gon if and only if every 4-point subset of the six forms a convex 4-gon. In other words, the overall 6-gon is convex precisely when all of its 4-point subsets are convex; conversely, the presence of a single concave 4-point subset implies that the 6-gon itself must be concave. This observation follows from the fact that a concave 4-gon necessarily contains an interior point that does not lie on its convex hull. If such a point exists, it cannot lie on the convex hull of the full 6-point configuration either, thereby ruling out the possibility of a convex 6-gon. Figure 3.19 illustrates this property, and it will be used to encode the avoidance of 6 gons by ensuring that, for every 6-point subset, at least one of its 4-point subsets is concave.

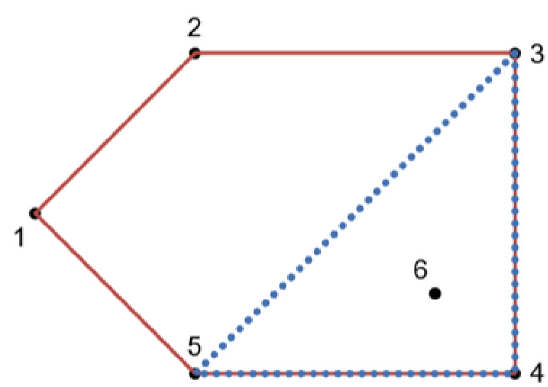

Figure 3.19: An illustration of a set of six points where the four element subset $(3, 4, 5, 6)$ is concave, thus preventing the point 6 from lying on the overall convex hull.

To encode this convexity logic using orientation triples, auxiliary variables of the form $c_{a,b,c,d}$ are introduced. Each such variable represents whether the four points a, b, c, d form a convex 4-gon, evaluating to true if the configuration is convex, and false otherwise. These variables are instantiated by analyzing the orientation of each 3-point subset of the 4-point set.

It is a well-known fact that a 4-point set in general position forms a convex quadrilateral if and only if an even number of its orientation triples are positively oriented. Although this equivalence can be verified through exhaustive case analysis (there are sixteen possible orientation assignments), it is not proven here explicitly. This insight allows us to define the truth value of $c_{a,b,c,d}$ using logical equivalences that reflect the parity of the positive orientations.

For instance, in one case where the number of positively oriented triples is even, the convexity variable is defined as follows:

$$\left(o_{a,b,c} \land \neg o_{a,b,d} \land o_{a,c,d} \land \neg o_{b,c,d}\right) \iff c_{a,b,c,d} \tag{3.27}$$

Conversely, in a case where the number of positive orientations is odd, the variable reflects a non-convex configuration:

$$\left(o_{a,b,c} \land o_{a,b,d} \land \neg o_{a,c,d} \land o_{b,c,d}\right) \iff \neg c_{a,b,c,d} \tag{3.28}$$

Clauses like these are constructed to cover all possible parity patterns, ensuring that each instance of $c_{a,b,c,d}$ correctly encodes the convexity status of its corresponding 4-point subset.

Once each convexity auxiliary variable has been initialized, a simple cardinality constraint is used to enforce that, for every subset of six points (a, b, c, d, e, f), at least one subset of size four in concave:

$$\left(\neg c_{a,b,c,d} \lor \neg c_{a,b,d,e} \lor \ldots \lor \neg c_{c,d,e,f}\right) \tag{3.29}$$

3.2.3 Realization

To obtain explicit realizations of rotationally symmetric point sets that are free of convex 6-gons, the satisfiability of the corresponding logical encodings under each of the three forms of rotational symmetry was evaluated using a SAT solver.

To extract concrete geometric realizations from these abstract encodings, a realization tool developed by Bernardo Subercaseaux [5] was employed. This solver accepts a set of orientation triples and performs local search to iteratively adjust point placements in order to eliminate conflicts and achieve a valid geometric embedding. In collaboration with the author of the tool, the solver was extended to support partial fixing of point positions during execution. This enhancement enabled a guided search process in which layers of points could be anchored incrementally. By fixing subsets of points according to previously successful placements, the solver could focus on resolving the remaining conflicts while preserving the intended rotational symmetry. This iterative strategy proved essential for discovering structured and symmetric configurations that aligned with the desired combinatorial constraints.

3.2.4 3-Fold Symmetry

Similarly to the analysis phase, the first class of structures considered in this section consists of layered convex hull formations of size $(3, 3, 3, 3, 3, 1)$, comprising five nested convex layers of three points each, arranged symmetrically around a single central point. However, unlike the previous phase, the points constituting each layer are now explicitly defined: The first convex layer consists of points 1, 2, and 3; the second layer consists of points 4, 5, and 6; and so on, with the central point explicitly assigned as point 16.

Under this encoding, the sector constraints are no longer required to enforce an equal distribution of points. Instead, each sector is associated with a specific and predetermined set of points. This simplification significantly reduces the complexity of the construction, as the need for extensive case analysis is eliminated: the sector boundaries and the points they must contain are fixed in advance.

To extend this encoding beyond the structured hull configurations used in the analysis portion, additional clauses are introduced to explicitly enforce 3-fold rotational symmetry on the point placement. Since the identities of all points are now fixed, the mapping of each point under a rotation of 120° is also explicitly known. This enables the orientation triples to be equated across each rotational instance, thereby ensuring that the underlying geometric structure remains consistently oriented under every rotation. These symmetry-preserving constraints play a crucial role in maintaining uniformity across the realization and are central to the pinwheel-like configurations studied in this section.

Convexity

To enforce convexity within each of the three-point convex hull layers, a single clause is introduced per layer. In order to maintain consistent rotational symmetry across layers, we assume, without loss of generality, that the points in each layer appear in counterclockwise order. Under this assumption, convexity can be encoded by asserting that the corresponding orientation triple for each layer is positive. Specifically, the following conjunction of orientation literals is added:

$$o_{1,2,3} \wedge o_{4,5,6} \wedge o_{7,8,9} \wedge o_{10,11,12} \wedge o_{13,14,15} \tag{3.30}$$

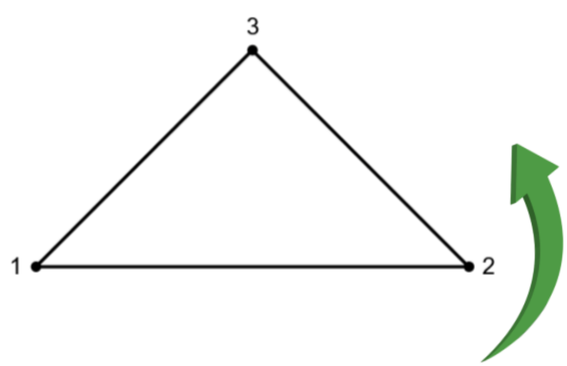

Figure 3.20: An illustration of the first convex layer for the 3-fold symmetry case, and the corresponding clause $o_{1,2,3}$ causing it to be counterclockwise.

The illustration for the first layer is shown in Figure 3.20. These clauses ensure that each trio of points forms a counterclockwise triangle, thereby preserving both the convexity and the layered symmetry necessary for the intended geometric hierarchy.

Containment

To enforce the containment of points within each successive convex hull layer, orientation clauses are added for every point p that is intended to lie strictly inside the triangle formed by a given convex hull (a, b, c). For the first convex layer, this involves ensuring that all points labeled 4 through 16 lie within the triangle defined by points $(1, 2, 3)$.

This is achieved by requiring that each such point p lies on the internal side of all three edges of the triangle. Since the points $(1, 2, 3)$ are assumed to be ordered counterclockwise, the containment corresponds to the requirement that p lies to the left of the directed edges $\overline{12}$, $\overline{23}$, and $\overline{31}$. This condition is encoded by asserting that the corresponding orientation triples are positive. This is illustrated in Figure 3.21. Specifically, for all $p \in \{4, 5, \ldots, 16\}$, the following conjunction is added:

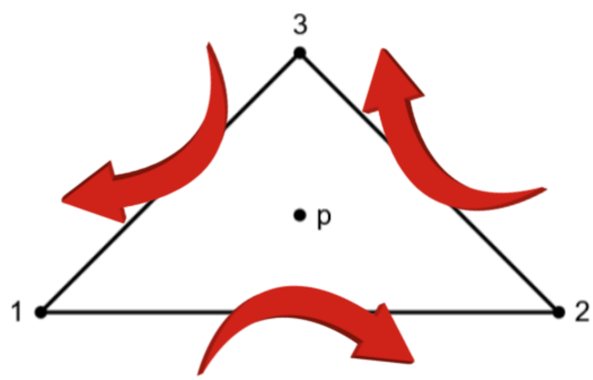

Figure 3.21: An illustration of the clauses necessary to enforce that point p is contained within the first layer convex hull.

$$o_{1,2,p} \wedge o_{2,3,p} \wedge o_{3,1,p} \tag{3.31}$$

The same logic is applied recursively to ensure that each subsequent convex layer contains all remaining points that have not yet been assigned to an outer hull. In each case, the orientation conditions are tailored to the specific triangle that defines that layer.

Sectors

As in the analysis phase, the convex hull layers are subdivided into sectors to facilitate an even radial distribution of points around the central point. In this section, where each convex hull consists of three points, the interior region of each convex layer is partitioned into three distinct sectors. Each sector is defined by a triangle formed by two adjacent points on the convex hull together with the central point. For the first layer, this results in three sectors defined by the triangles $(1, 2, 16)$, $(2, 3, 16)$, and $(3, 1, 16)$, which are treated as Sectors 1 through 3, respectively.

To create a consistent and well-defined mapping between the convex hull vertices and their corresponding sectors, we enforce the following assignment: The first point of each hull is associated with Sector 1, the second point with Sector 2, and the third point with Sector 3. For the hull $(1, 2, 3)$, this yields the following desired sector assignments for the inner points:

Sector 1: 4, 7, 10, 13
Sector 2: 5, 8, 11, 14
Sector 3: 6, 9, 12, 15

To enforce that a point p lies within its designated sector, orientation constraints are imposed relative to the triangle defining that sector. Specifically, a point is constrained to lie between the two directed boundary edges of the sector triangle. For Sector 1, defined by the triangle $(1, 2, 16)$, a point must lie below the line formed by points 1 and 16 and above the line formed by points 2 and 16, which is encoded as shown below and illustrated in Figure 3.22

$$\neg o_{1,16,p} \wedge o_{2,16,p} \tag{3.32}$$

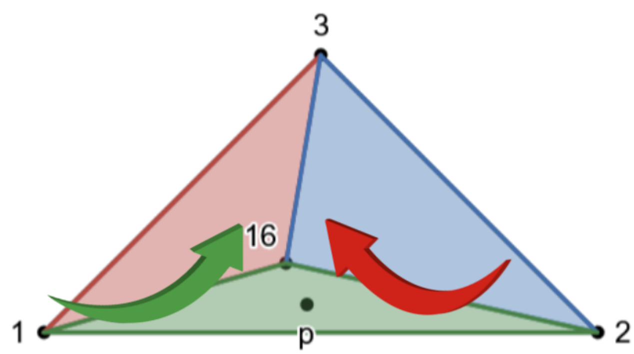

Figure 3.22: An illustration of the clauses necessary to enforce that point p is contained in the sector below the line formed by points 1 and 16 and above the line formed by points 2 and 16.

For Sector 2, defined by $(2, 3, 16)$, the encoding becomes:

$$\neg o_{2,16,p} \wedge o_{3,16,p} \tag{3.33}$$

And for Sector 3, defined by $(3, 1, 16)$, the encoding becomes:

$$\neg o_{3,16,p} \wedge o_{1,16,p} \tag{3.34}$$

These constraints are applied in the same manner across all convex hull layers, ensuring that each point lies within its designated sector. This consistent enforcement preserves the intended rotational symmetry and radial organization throughout the layered structure of the hulls.

Symmetry

To enforce 3-fold rotational symmetry over the orientation triples, the encoding introduces clauses that equate each orientation triple with its image under a $120°$ rotation. Since the structure and placement of the points are fixed, this rotational mapping is deterministic and can be computed explicitly.

Each orientation triple $o_{a,b,c}$ is mapped to a corresponding triple $o_{a',b',c'}$, which results from applying a $120°$ counterclockwise rotation to each of the points involved. The mapping from a point p to its rotated image p' is defined based on the radial sector structure established earlier and is illustrated in Figure 3.23:

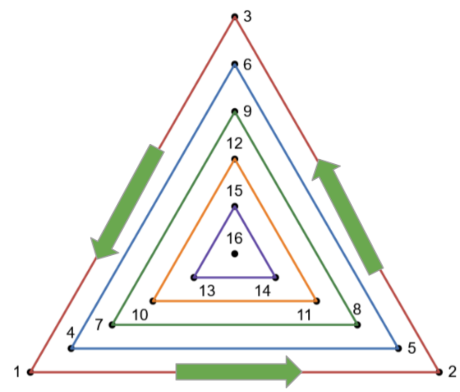

Figure 3.23: An illustration of the mapping of points in the 3-fold symmetry case. $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 1$, etc.

- If $p = 16$, the central point remains fixed under rotation, so $p' = 16$.
- Otherwise, each point in a sector is mapped to the next point in the same convex hull layer under 3-fold rotation. This mapping follows a cyclic increment, so each third point will wrap around the first point in the same layer (e.g. $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 1$, etc.).

For each orientation triple, clauses are added to enforce equivalence between the original triple and its rotated counterpart. Specifically, given variables $o_{a,b,c}$ and $o_{a',b',c'}$, the following biimplication is encoded:

$$o_{a,b,c} \iff o_{a',b',c'} \tag{3.35}$$

This logic ensures that the relative orientation of any three points remains consistent across all $120°$ rotations. When applied to all orientation triples over the point set, these constraints collectively guarantee that the entire configuration exhibits global 3-fold rotational symmetry.

Results and Discussion

Unfortunately, no realizations exhibiting 3-fold rotational symmetry were found among the 16-point configurations explored. Even under a relaxed encoding—where only the sector variables for the first convex layer were initialized—the SAT solver consistently returned unsatisfiable. This outcome appears to stem directly from the constraints introduced by the rotational symmetry clauses. While the encoding was successful in isolating point configurations that satisfied the structural criteria, the additional symmetry requirements imposed by the 3-fold rotation proved too restrictive to permit a viable geometric realization. This suggests that 3-fold symmetry, when combined with the structural framework used in this study, imposes conditions that are incompatible with the avoidance of convex 6-gons within this point limit.

3.2.5 4-Fold Symmetry

The second class of structures considered in this section consists of layered convex hull formations of size $(4, 4, 4, 4)$, consisting of four nested convex layers with four points each. The points in each layer are also now explicitly defined: the first convex layer consists of points $1, 2, 3$, and 4; the second layer consists of points $5, 6, 7$, and 8; and so on for the remaining layers.

Similarly to the 3-fold symmetric case, each sector is associated with a specific and predetermined subset of points, eliminating the need for the involved case split. Clauses analogous to those used in the 3-fold case are utilized to explicitly enforce 4-fold rotational symmetry on the point placement. By computing and equating orientation triples under the corresponding $90°$ rotation mappings provided by the sectors, the encoding ensures that the configuration maintains its symmetry across all four rotational axes.

Convexity

To enforce convexity within each of the four-point convex hull layers, a group of four clauses is introduced per layer. In order to maintain consistent rotational symmetry across layers, we assume, without loss of generality, that the points in each layer appear in counterclockwise order. Under this assumption, convexity can be encoded by asserting that the orientation triples corresponding to consecutive triples of points within each layer are positive. Specifically, for the first layer, the following conjunction of orientation literals is added and the corresponding logic is depicted in Figure 3.24:

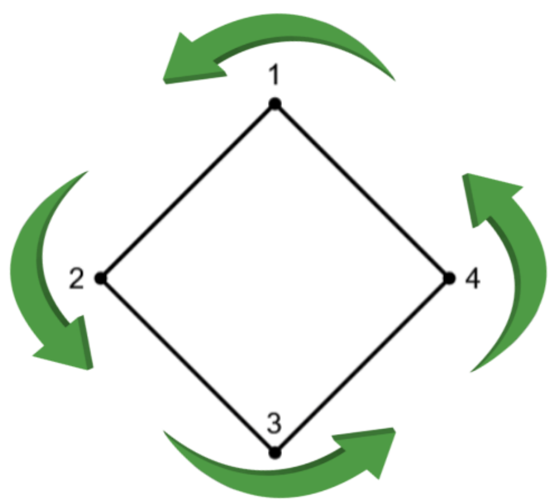

Figure 3.24: An illustration of the logic needed to enforce convexity over the first layer in the 4-fold symmetry case.

$$o_{1,2,3} \wedge o_{2,3,4} \wedge o_{3,4,1} \wedge o_{4,1,2} \tag{3.36}$$

The same approach is applied to layers 2 through 4. This guarantees that each set of four points forms a counterclockwise quadrilateral, thereby preserving both the convexity and the layered symmetry required for the desired geometric structure.

Containment

To enforce the containment of points within each successive convex hull layer of size four, we add orientation clauses for every point p that is intended to lie strictly inside the quadrilateral formed by four counterclockwise ordered points (a, b, c, d).

For the first convex hull layer, defined by points $1, 2, 3$, and 4, we require that all points labeled 5 through 16 lie within this shell. This is encoded by adding four positive orientation literals for each point p, one for each boundary edge, and is depicted in Figure 3.25:

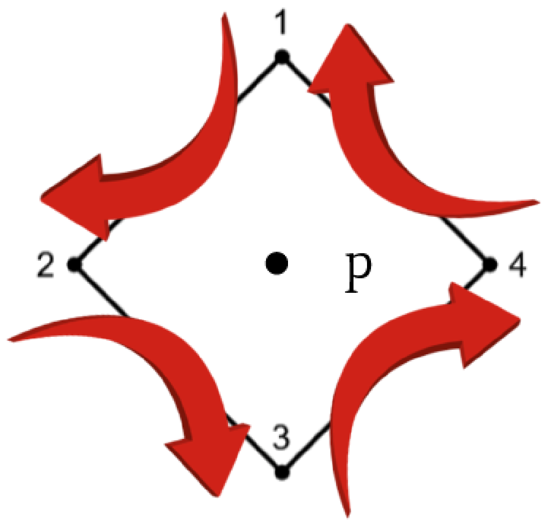

Figure 3.25: An illustration of the logic needed to enforce the containment of each point p on the interior side of each line segment of the hull.

$$o_{1,2,p} \wedge o_{2,3,p} \wedge o_{3,4,p} \wedge o_{4,1,p} \tag{3.37}$$

These four conditions ensure that point p lies within the convex 4-gon by being on the interior side of every edge. The same logic is recursively applied to each inner layer of four points. For each such layer, the boundary edges are defined by iterating cyclically through the four points in counterclockwise order, and the same set of orientation constraints is imposed on each point that must lie within the layer.

Sectors

As in previous sections, the convex hull layers are subdivided into sectors to facilitate a uniform radial distribution of interior points across the geometric structure. In this setting, where each convex hull consists of four vertices, the interior region of each convex layer is partitioned into four distinct sectors. These sectors are defined by the relative position of a point with respect to the two diagonals of the convex hull, formed by connecting opposite vertices.

For the first convex hull layer, composed of points $(1, 2, 3, 4)$, the diagonals $\overline{13}$ and $\overline{24}$ serve as the reference lines for partitioning the interior into each sector. Each sector corresponds to one of the four possible combinations of the orientation of a point, above or below, with respect to each diagonal. Specifically, the four sectors for the first layer are defined as follows:

Sector 1: above $\overline{13}$ and above $\overline{24}$
Sector 2: above $\overline{13}$ and below $\overline{24}$
Sector 3: below $\overline{13}$ and above $\overline{24}$
Sector 4: below $\overline{13}$ and below $\overline{24}$

To create a consistent and well-defined mapping between the convex hull vertices and their corresponding sectors, we enforce the following assignment: The first point of each hull is associated with Sector 1, the second point with Sector 2, the third point with Sector 3, and the fourth point with Sector 4. For the hull $(1, 2, 3, 4)$, this yields the following desired sector assignments for the inner points:

Sector 1: 5, 9, 13
Sector 2: 6, 10, 14
Sector 3: 7, 11, 15
Sector 4: 8, 12, 16

To enforce that a point p lies within its designated sector, orientation constraints are imposed relative to the diagonal boundaries defining that sector. For Sector 1, a point must lie above the lines formed by points 1 and 3 and points 2 and 4, which is depicted in Figure 3.26 and encoded as

$$o_{1,3,p} \wedge o_{2,4,p} \tag{3.38}$$

For Sector 2, the encoding becomes:

$$o_{1,3,p} \wedge \neg o_{2,4,p} \tag{3.39}$$

For Sector 3, the encoding becomes:

$$\neg o_{1,3,p} \wedge o_{2,4,p} \tag{3.40}$$

And for Sector 4, the encoding becomes:

$$\neg o_{1,3,p} \wedge \neg o_{2,4,p} \tag{3.41}$$

These constraints are applied in the same manner across all convex hull layers, ensuring that each point lies within its designated sector. This consistent enforcement preserves the intended rotational symmetry and radial organization throughout the layered structure of the hulls.

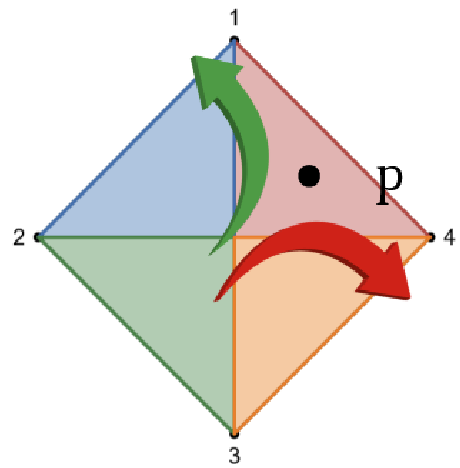

Figure 3.26: An illustration of the logic needed to enforce that point p lies in the sector above the lines 13 and 24.

Symmetry

To enforce 4-fold rotational symmetry over the orientation triples, the encoding introduces clauses that equate each orientation triple with its image under a $90°$ rotation. Since the structure and placement of the points are fixed, this rotational mapping is deterministic and can be computed explicitly.

Each orientation triple $o_{a,b,c}$ is mapped to a corresponding triple $o_{a',b',c'}$, which results from applying a $90°$ counterclockwise rotation to each of the points involved. The mapping from a point p to its rotated image p' is defined based on the radial sector structure established earlier: Each point in a sector is mapped to the next point in the same convex hull layer under 4-fold rotation. This mapping follows a cyclic increment, so each fourth point will wrap around the first point in the same layer (e.g. $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 4$, $4 \mapsto 1$, etc.) This mapping is shown in Figure 3.27.

Just like in the 3-fold case, for each orientation triple, clauses are added to enforce equivalence between the original triple and its rotated counterpart. Specifically, given variables $o_{a,b,c}$ and $o_{a',b',c'}$, the following biimplication is encoded:

$$o_{a,b,c} \iff o_{a',b',c'} \tag{3.42}$$

This logic ensures that the relative orientation of any three points remains consistent across all $90°$ rotations. When applied to all orientation triples over the point set, these constraints collectively guarantee that the entire configuration exhibits global 4-fold rotational symmetry.

Results and Discussion

In contrast to the 3-fold case, full 4-fold rotational symmetry was successfully achieved. An explicit realization satisfying all structural and symmetry constraints was produced, marking the first known 16-point configuration that avoids convex 6-gons while maintaining complete

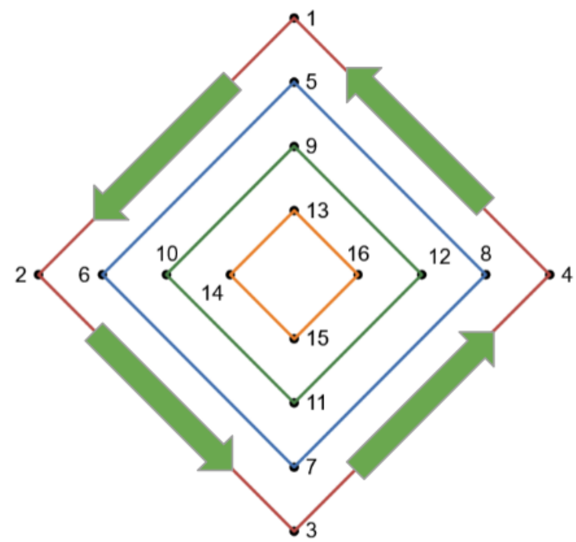

Figure 3.27: An illustration depicting the mapping of points in the 4-fold symmetry case, $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 4$, $4 \mapsto 1$, etc.

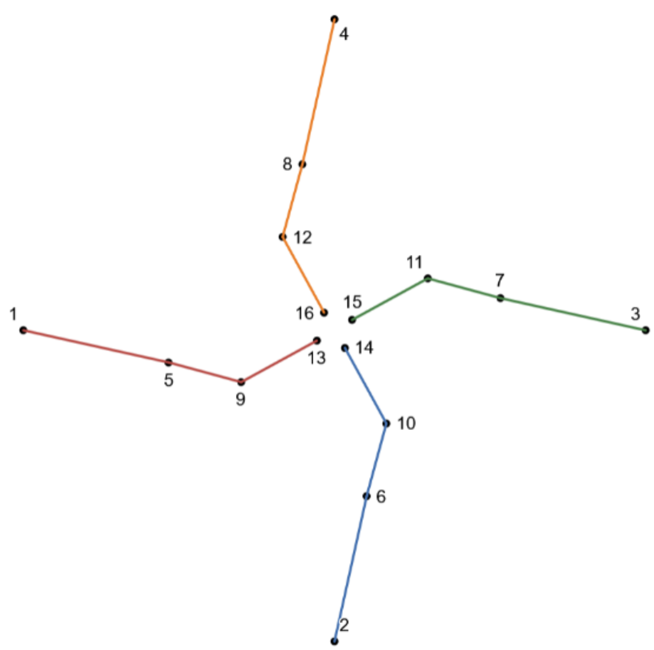

Figure 3.28: A 16 point realization exhibiting 4-fold symmetry.

4-fold rotational symmetry. The resulting point arrangement, shown in Figure 3.28, exhibits a distinctive pinwheel-like structure, a natural consequence of the sector constraints which guide the inner points into spiral-like branches emanating from the center. This visually compelling formation not only satisfies the geometric and combinatorial constraints but also highlights the power of encoding sector variables to induce symmetry in highly constrained environments. The coordinates of this realization are listed below:

Point 1:	$(-30.000000,\ 0.000000)$
Point 2:	$(0.000000,\ -30.000000)$
Point 3:	$(30.000000,\ 0.000000)$
Point 4:	$(0.000000,\ 30.000000)$
Point 5:	$(-16.000000,\ -3.100000)$
Point 6:	$(3.100000,\ -16.000000)$
Point 7:	$(16.000000,\ 3.100000)$
Point 8:	$(-3.100000,\ 16.000000)$
Point 9:	$(-9.000000,\ -5.000000)$
Point 10:	$(5.000000,\ -9.000000)$
Point 11:	$(9.000000,\ 5.000000)$
Point 12:	$(-5.000000,\ 9.000000)$
Point 13:	$(-1.700000,\ -1.000000)$
Point 14:	$(1.000000,\ -1.700000)$
Point 15:	$(1.700000,\ 1.000000)$
Point 16:	$(-1.000000,\ 1.700000)$

This solution stands as a significant benchmark for symmetric extremal configurations, and opens up new avenues for further exploration of symmetry-enforced constraints in geometric Ramsey-type problems.

3.2.6 5-Fold Symmetry

The third and final class of structures considered in this section consists of layered convex hull formations of size $(5, 5, 5, 1)$, consisting of three nested convex layers with five points each, arranged around a single central point. The points in each layer are also now explicitly defined: the first convex layer consists of points $1, 2, 3, 4$ and 5; the second layer consists of points $6, 7, 8, 9$ and 10; and so on for the remaining layers.

Similarly to the $3-$ and 4-fold symmetric cases, each sector is associated with a specific and predetermined subset of points, eliminating the need for the involved case split. Clauses analogous to those used in the 3- and 4-fold case are utilized to explicitly enforce 5-fold rotational symmetry on the point placement. By computing and equating orientation triples under the corresponding $72°$ rotation mappings provided by the sectors, the encoding ensures that the configuration maintains its symmetry across all five rotational axes.

Convexity

To enforce convexity within each of the five-point convex hull layers, a group of five clauses is introduced per layer. In order to maintain consistent rotational symmetry across layers, we assume, without loss of generality, that the points in each layer appear in counterclockwise order. Under this assumption, convexity can be encoded by asserting that the orientation triples corresponding to consecutive triples of points within each layer are positive. Specifically, for the first layer, the following conjunction of orientation literals is added and is depicted in Figure 3.29:

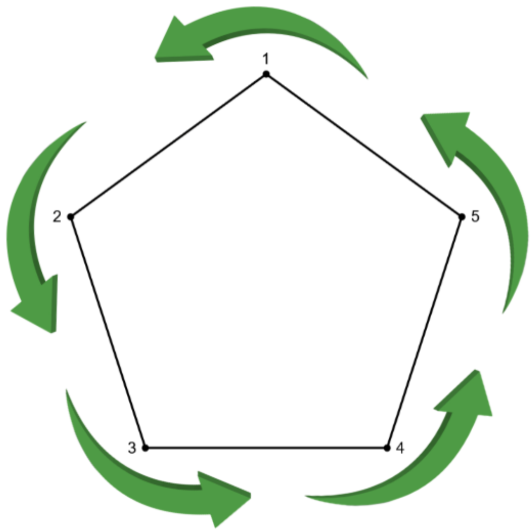

Figure 3.29: A depiction of the clauses necessary to enforce counterclockwise convexity in the 5-fold symmetry case.

$$o_{1,2,3} \wedge o_{2,3,4} \wedge o_{3,4,5} \wedge o_{4,5,1} \wedge o_{5,1,2} \tag{3.43}$$

The same approach is applied to layers 2 and 3. This guarantees that each set of five points forms a counterclockwise 5-gon, thereby preserving both the convexity and the layered symmetry required for the desired geometric structure.

Containment

To enforce the containment of points within each successive convex hull layer of size five, we add orientation clauses for every point p that is intended to lie strictly inside the 5-gon formed by five counterclockwise ordered points (a, b, c, d, e).

For the first convex hull layer, defined by points $1, 2, 3, 4$ and 5, we require that all points labeled 6 through 16 lie within this shell. This is encoded by adding five positive orientation literals for each point p, one for each boundary edge and is depicted in Figure 3.30:

$$o_{1,2,p} \wedge o_{2,3,p} \wedge o_{3,4,p} \wedge o_{4,5,p} \wedge o_{5,1,p} \tag{3.44}$$

These five conditions ensure that point p lies within the convex 5-gon by being on the interior side of every edge. The same logic is recursively applied to each inner layer of five points. For each such layer, the boundary edges are defined by iterating cyclically through the five points in

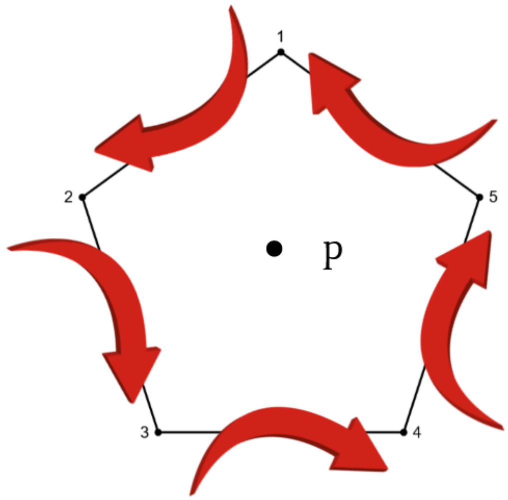

Figure 3.30: A depiction of the clauses necessary to enforce that point p is on the interior side of each of the boundary line segments of the first convex hull

counterclockwise order, and the same set of orientation constraints is imposed on each point that must lie within the layer.

Sectors

As in the previous phases, the convex hull layers are subdivided into sectors to facilitate an even radial distribution of points around the central point. In this section, where each convex hull consists of five points, the interior region of each convex layer is partitioned into five distinct sectors. Each sector is defined by a triangle formed by two adjacent points on the convex hull together with the central point. For the first layer, this results in five sectors defined by the triangles $(1, 2, 16)$, $(2, 3, 16)$, $(3, 4, 16)$, $(4, 5, 16)$, and $(5, 1, 16)$, which are treated as Sectors 1 through 5, respectively.

To create a consistent and well-defined mapping between the convex hull vertices and their corresponding sectors, we enforce the following assignment: The first point of each hull is associated with Sector 1, the second point with Sector 2, and the third point with Sector 3, etc. For the hull $(1, 2, 3, 4, 5)$, this yields the following desired sector assignments for the inner points:

Sector 1: 6, 11
Sector 2: 7, 12
Sector 3: 8, 13
Sector 4: 9, 14
Sector 5: 10, 15

To enforce that a point p lies within its designated sector, orientation constraints are imposed relative to the triangle defining that sector. Specifically, a point is constrained to lie between the two directed boundary edges of the sector triangle. For Sector 1, defined by the triangle $(1, 2, 16)$, a point must lie below the line formed by points 1 and 16 and above the line formed by points 2 and 16, this is depicted in Figure 3.31 and is encoded as

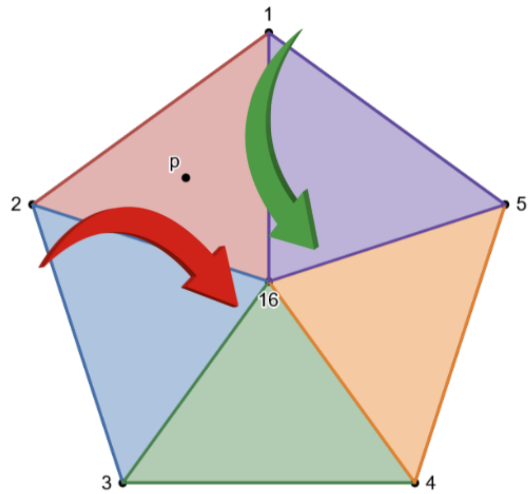

Figure 3.31: A depiction of the point p being restricted to lie in the first sector of the first layer in the 5-fold symmetry case.

$$\neg o_{1,16,p} \land o_{2,16,p} \tag{3.45}$$

The logic is similar for the remaining sectors. These constraints are applied in the same manner across all convex hull layers, ensuring that each point lies within its designated sector. This consistent enforcement preserves the intended rotational symmetry and radial organization throughout the layered structure of the hulls.

Symmetry

To enforce 5-fold rotational symmetry over the orientation triples, the encoding introduces clauses that equate each orientation triple with its image under a $72°$ rotation. Since the structure and placement of the points are fixed, this rotational mapping is deterministic and can be computed explicitly.

Each orientation triple $o_{a,b,c}$ is mapped to a corresponding triple $o_{a',b',c'}$, which results from applying a $72°$ counterclockwise rotation to each of the points involved. The mapping from a point p to its rotated image p' is defined based on the radial sector structure established earlier: Each point in a sector is mapped to the next point in the same convex hull layer under 5-fold rotation. This mapping is depicted in Figure 3.32 and follows a cyclic increment, so each fifth point will wrap around the first point in the same layer (e.g. $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 4$, $4 \mapsto 5$, $5 \mapsto 1$, etc.).

Just like in the 3- and 4-fold cases, for each orientation triple, clauses are added to enforce equivalence between the original triple and its rotated counterpart. Specifically, given variables $o_{a,b,c}$ and $o_{a',b',c'}$, the following biimplication is encoded:

$$o_{a,b,c} \iff o_{a',b',c'} \tag{3.46}$$

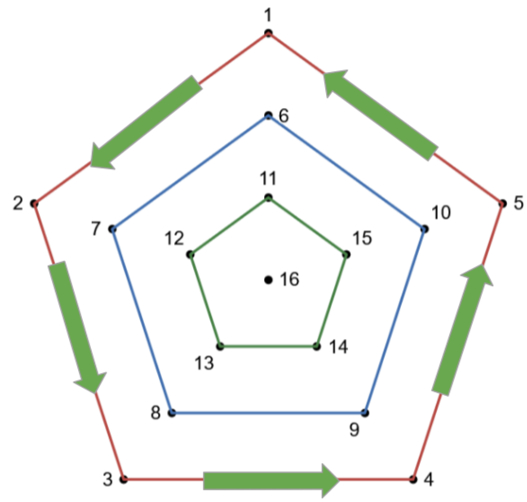

Figure 3.32: A depiction of the mapping of the points in the first layer of the 5-fold symmetry case, $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto 4$, $4 \mapsto 5$, $5 \mapsto 1$, etc.

This logic ensures that the relative orientation of any three points remains consistent across all $72°$ rotations. When applied to all orientation triples over the point set, these constraints collectively guarantee that the entire configuration exhibits global 5-fold rotational symmetry.

Results and Discussion

5-fold rotational symmetry was also found to be realizable, a particularly surprising result given the remarkably low satisfiability ratio observed in the layered 5-gon architecture, as discussed earlier. Despite the scarcity of satisfying configurations among the candidate placements, an explicit realization was discovered that satisfies the symmetry constraints, and is shown in Figure 3.33. As with the 4-fold symmetric case, the resulting structure exhibits a pinwheel-like form—an emergent consequence of the sector-enforcing constraints embedded in the SAT encoding. This formation not only reinforces the viability of the layered convex hull approach but also provides the first known instance of a 6-gon-free configuration on 16 points that admits full 5-fold rotational symmetry. The coordinates of the realization are provided below:

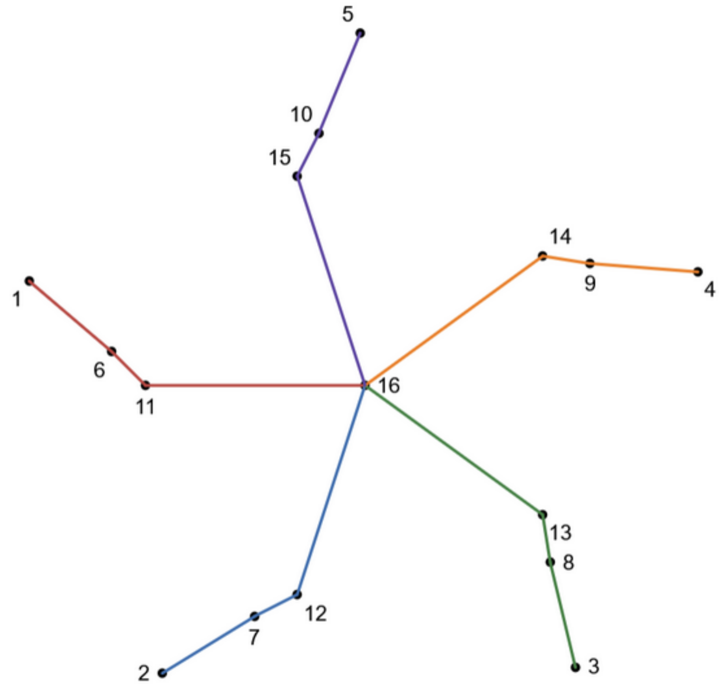

Figure 3.33: A 16 point realization exhibiting 5-fold symmetry.

Point 1:	$(-19.876200,\ 6.159400)$
Point 2:	$(-12.000000,\ -17.000000)$
Point 3:	$(12.459800,\ -16.666000)$
Point 4:	$(19.700600,\ 6.699900)$
Point 5:	$(-0.284100,\ 20.806700)$
Point 6:	$(-15.000000,\ 2.000000)$
Point 7:	$(-6.537400,\ -13.647800)$
Point 8:	$(10.959700,\ -10.434800)$
Point 9:	$(13.310800,\ 7.198700)$
Point 10:	$(-2.733100,\ 14.883900)$
Point 11:	$(-13.000000,\ 0.000000)$
Point 12:	$(-4.017200,\ -12.363700)$
Point 13:	$(10.517200,\ -7.641200)$
Point 14:	$(10.517200,\ 7.641200)$
Point 15:	$(-4.017200,\ 12.363700)$
Point 16:	$(0.000000,\ 0.000000)$

Chapter 4

Conclusion

In conclusion, both phases of this research achieved their intended objectives. The first phase focused on the structural analysis of candidate point placements by leveraging a generalized SAT-based encoding to exhaustively explore configurations consistent with various layered convex hull constructions. This approach enabled a comprehensive enumeration and evaluation of possible formations under each structural setting. The second phase centered on the realization of these configurations, with the goal of identifying explicit point placements that exhibit full rotational symmetry.

The structural analysis yielded meaningful insights into the relationship between convex hull layering and the presence of forbidden convex substructures, such as 6-gons. Meanwhile, the realization phase produced two novel and fully symmetric point sets, each corresponding to distinct layered architectures. These realizations not only demonstrate the feasibility of imposing rotational symmetry on such configurations but also contribute new examples to the landscape of extremal geometry. Together, the findings underscore the effectiveness of combining combinatorial encoding with geometric realization techniques in tackling long-standing open problems in discrete geometry.

4.1 Future Work

This research opens several promising avenues for future investigation, particularly concerning other instances of the Erdős–Szekeres Conjecture. The next outstanding case involves the conjectured extremal bound for convex 7-gon avoidance, which posits that no set of 33 points can be constructed without containing a convex 7-gon. While this remains unproven, constructions of 32-point sets avoiding all convex 7-gons have been established. Initial efforts within this project aimed to identify a fully 4-fold symmetric realization of such a placement. While satisfying assignments consistent with 4-fold symmetry were found in this initial attempt, realizations with no more than 24 points (corresponding to six layers of size four) have been produced.

Future research could pursue alternative strategies to achieve this result. One promising direction involves partitioning the point set into smaller substructures and attempting to stitch together partial realizations. Another possibility is the development or adaptation of realization tools specifically tailored to handle these large-scale geometric formulas and constraints with multiple satisfying assignments more efficiently. Advances in these areas may make it possible to fully resolve the 7-gon case with symmetry constraints and potentially discover a recursive general construction that can be used to resolve the Erdős-Szekeres Conjecture itself.

Bibliography

[1] Armin Biere, Tobias Faller, Katalin Fazekas, Mathias Fleury, Nils Froleyks, and Florian Pollitt. CaDiCaL 2.0. In Arie Gurfinkel and Vijay Ganesh, editors, *Computer Aided Verification - 36th International Conference, CAV 2024, Montreal, QC, Canada, July 24-27, 2024, Proceedings, Part I*, volume 14681 of *Lecture Notes in Computer Science*, pages 133–152. Springer, 2024. doi: 10.1007/978-3-031-65627-9_7. 1.2

[2] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935. URL `https://www.numdam.org/item/CM_1935__2__463_0/`. 1.1

[3] Marijn J. H. Heule and Manfred Scheucher. Happy ending: An empty hexagon in every set of 30 points, 2024. URL `https://arxiv.org/abs/2403.00737`. 2, 2.2

[4] Donald E. Knuth. *Axioms and Hulls*, volume 606 of *Lecture Notes in Computer Science*. Springer Berlin, Heidelberg, 1 edition, 1992. ISBN 978-3-540-55611-4. doi: 10.1007/3-540-55611-7. URL `https://doi.org/10.1007/3-540-55611-7`. Topics: Theory of Computation, Computer Applications, Discrete Mathematics, Computer Graphics, Algorithm Analysis and Problem Complexity, Combinatorics. 2, 2.1

[5] Bernardo Subercaseaux. Code for realizability of pseudo-points (through cc-systems / signotopes / triple orientations). `https://github.com/bsubercaseaux/point_realizer`, 2025. Accessed: 2025-05-01. 1.3, 3.2.3

[6] George Szekeres and Lindsay Peters. Computer solution to the 17-point erdős-szekeres problem. *The ANZIAM Journal*, 48(2):151–164, 2006. doi: 10.1017/S144618110000300X. 1.1

[7] Wikipedia contributors. Image: 16-point set with no convex hexagon (16nohexagon.svg). `https://en.wikipedia.org/wiki/File:16nohexagon.svg`, 2025. Accessed: 2025-05-01. Licensed under CC BY-SA 3.0. 3.1.1