

From Computational Certification to Exact Coordinates: Heilbronn’s Triangle Problem on the Unit Square Using Mixed-Integer Optimization

Nathan Sudermann-Merx*

March 9, 2026

Abstract

We develop an *optimize-then-refine* framework for the classical Heilbronn triangle problem that integrates global mixed-integer nonlinear programming with exact symbolic computation. A novel symmetry-breaking strategy, together with the exploitation of structural properties of determinants, yields a substantially stronger optimization model: for $n = 9$, the problem can be solved to certified global optimality in 15 minutes on a standard desktop computer, improving upon the previously reported effort of about one day by more than an order of magnitude. Combining the numerical certificate with exact symbolic computation, we provide the first proof that the configuration discovered by Comellas and Yebra in 2002 for $n = 9$ is globally optimal, and derive exact coordinates for all optimal configurations with $n = 5, \dots, 9$, confirming earlier best-known results and sometimes simplifying their presentation. An analysis of these configurations reveals structural patterns—notably the clustering of noncritical triangle areas around a small number of distinct values—which give rise to new research questions about the combinatorial geometry of extremal point sets. All configurations and code are publicly available to provide a reproducible foundation for further research.

1 Introduction

1.1 The problem

The Heilbronn triangle problem arises in discrete geometry and asks:

How should n points be placed in the unit square so that the smallest triangle determined by any three of them is as large as possible?

The problem is named after the mathematician Hans Heilbronn (and not after the city in southern Germany), who was born into a German-Jewish family and left Germany in 1933 during the rise of the Nazi regime. It first appeared in print in Roth’s 1951 paper [18], where the unit square was replaced by an arbitrary closed convex region in the plane of positive measure, and the emphasis was placed on the asymptotic behavior of the minimal triangle area rather than on explicit constructions for small values of n .

*Institute for Computer Science, DHBW Mannheim. Email: nathan.sudermann-merx@dhbw.de

1.2 Upper bounds

We denote the minimal triangle area in an optimal configuration of n points in the unit square by Δ_n . Heilbronn conjectured that Δ_n should satisfy an asymptotic upper bound of the form

$$\Delta_n = O\left(\frac{1}{n^2}\right).$$

As we will see in Section 1.3, this turned out to be incorrect. Roth observed in [18] that the problem admits a simple upper bound,

$$\Delta_n \leq \frac{1}{n-2}.$$

To see this, note that the convex hull of any configuration of n points can be partitioned into $n-2$ triangles with pairwise disjoint interiors and total area at most 1, implying that at least one of these triangles must have area no greater than $1/(n-2)$. Over the following decades, the upper bounds were steadily refined. In 1972, Schmidt [21] obtained a significant improvement, which was sharpened by Roth [19, 20] later that year and further strengthened by Komlós, Pintz, and Szemerédi [14] in 1981, establishing what remained the best-known estimate for over forty years. A major breakthrough came in 2023, when Cohen, Pohoata, and Zakharov [7] proved that for all sufficiently large n ,

$$\Delta_n \leq n^{-8/7-1/2000},$$

which is currently the strongest known asymptotic upper bound. Figure 1 illustrates that this bound provides a reasonably tight envelope for all best-known configurations with $n \geq 6$.

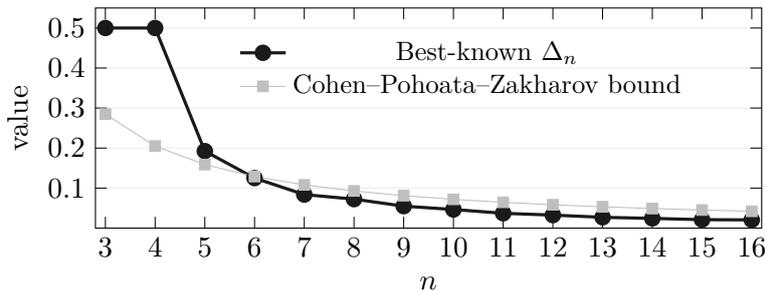


Figure 1: Best-known values of Δ_n for the Heilbronn problem in the unit square compared with the asymptotic upper bound $n^{-8/7-1/2000}$ of Cohen, Pohoata, and Zakharov [7].

1.3 Lower bounds

A simple lower bound of order c/n^3 for Δ_n can be obtained as follows. Without regard to the unit-square constraint, consider the n points (i, i^2) for $i = 0, \dots, n-1$. No three of these points are collinear, and hence every triangle determined by them has strictly positive area. Moreover, each triangle area can be expressed as half the absolute value of a 3×3 determinant formed from the vertices' coordinates (see Section 2.1). Since all coordinates are integers and the area is positive, the determinant has absolute value at least 1, implying that every triangle has area at least $1/2$. To place the configuration inside the unit square, scale the x -coordinates by $1/n$ and the y -coordinates by $1/n^2$. Areas then scale by $1/n^3$, giving

$$\Delta_n \geq \frac{1}{2n^3}.$$

As noted in the appendix of Roth's paper [18], Paul Erdős refined this idea by choosing a prime p with $n \leq p \leq 2n$ and considering the points $(i, i^2 \bmod p)$ for $i = 0, \dots, p-1$. This construction preserves the key properties used above—integral coordinates and the absence of collinear

triples—while having the advantage that the y -coordinates are bounded by p . Consequently, scaling to the unit square reduces areas only by a factor of $1/p^2$, leading to

$$\Delta_n \geq \frac{1}{2p^2} \geq \frac{1}{8n^2}$$

for all n .

Combined with Heilbronn’s conjecture, this would have yielded the correct asymptotic order of magnitude for the problem. However, in 1982, following a suggestion by Roth that the conjecture might be false, Komlós, Pintz, and Szemerédi proved in [15] that there exists a constant $c > 0$ such that

$$\Delta_n \geq \frac{c \log n}{n^2}$$

for infinitely many n , using a nonconstructive probabilistic argument. This gap between existence and construction was closed in 2000 by Bertram-Kretzberg, Hofmeister, and Lefmann [3], who presented a polynomial-time algorithm producing configurations with asymptotic lower bound $\Omega(\log n/n^2)$.

1.4 Results for small- n -configurations

All best-known configurations for the Heilbronn problem with $n \leq 16$ are documented on *Erich’s Packing Center*, a website maintained by Erich Friedman that tracks records for a variety of computational geometry problems (see [12]). The first case requiring a nontrivial proof, $n = 5$, was resolved in 1991 by Yang, Zhang, and Zeng [23], who proved $\Delta_5 = \sqrt{3}/9$. In the same work, they disproved a conjecture of Goldberg, who had suggested that optimal configurations for $n \leq 7$ arise from affine images of regular n -gons—a statement that turns out to hold only for $n = 6$. In 1995, Dress, Yang, and Zeng [11] established the optimal value $\Delta_6 = 1/8$.

In 2002, Comellas and Yebra [8] used simulated annealing followed by an analytical refinement procedure to discover best-known configurations and exact coordinates for many cases with $n \leq 12$. As we shall see, their configurations for $n = 7, 8, 9$ have since all been proved optimal, and their configurations for $n = 10$ and $n = 12$ remain the best known to date. The optimality of their configuration for $n = 7$, with $\Delta_7 \approx 0.083859$, was proved in 2011 by Zeng and Chen [24], who decomposed the problem into 226 nonlinear optimization subproblems. For $n = 9$, Chen, Xu, and Zeng [6] conducted in 2017 an extensive branch-and-bound computation using CPU and GPU clusters, requiring a total of 125 days of wall-clock time to obtain sharp upper bounds. In 2022, Dehbi and Zeng [10] combined numerical search with symbolic computation to prove that the configuration for $n = 8$ found by Comellas and Yebra, satisfying $\Delta_8 = (\sqrt{13} - 1)/36$, is indeed optimal.

For $n \geq 13$, the best-known configurations were found by Peter Karpov ($n = 13, 15$) and Mark Beyleveld ($n = 14, 16$) but have not been formally published; they are available only through personal communication via Erich’s Packing Center [12].

More recently, in 2025, Monji, Modir, and Kocuk [17] applied nonlinear global optimization techniques to compute an optimality certificate for Δ_9 , reducing the total wall-clock time to approximately one day using a state-of-the-art mixed-integer solver.

These results illustrate a broader trend: modern mixed-integer nonlinear programming (MINLP) solvers have matured to the point where they can tackle geometric optimization problems that were out of reach for general-purpose methods only a decade ago. A recent systematic study by Berthold et al. [2] confirms this observation across a range of combinatorial geometry problems, showing that straightforward MINLP formulations solved with off-the-shelf solvers can match or improve upon the best previously known solutions. The present paper continues this line of work for the Heilbronn problem.

1.5 Contributions

The contributions of this paper are fourfold.

A stronger mixed-integer formulation. We develop a mixed-integer optimization model for the Heilbronn problem that incorporates a novel symmetry-breaking strategy based on the boundary structure of optimal configurations. This strategy fixes several point coordinates a priori and, in addition, determines the orientation of certain triangles, which allows us to fix a number of binary sign variables before the solve. Together, these reductions lead to a formulation that is substantially stronger than existing models. In particular, the Heilbronn problem for $n = 9$ can be solved to certified global optimality in approximately 15 minutes on a standard desktop computer, improving upon the previously reported effort of about one day [17] by more than an order of magnitude.

Exact coordinates via an optimize-then-refine framework. We introduce a two-step methodology that combines global optimization (Step 1) with exact symbolic refinement (Step 2). The numerical solution from the solver identifies the combinatorial structure of an optimal configuration—specifically, which triangles are critical and which points lie on the boundary. This information yields a structured polynomial system that is then solved exactly using computer algebra. Applying this pipeline, we derive closed-form exact coordinates for all optimal configurations with $n = 5, \dots, 9$, confirming earlier results of Comellas and Yebra [8] and sometimes simplifying their presentation. In particular, we provide the first proof that their exact configuration for $n = 9$ is globally optimal: while Modir, Monji, and Kocuk [17] independently certified the optimal value of Δ_9 numerically, they did not connect it to the exact solution of Comellas and Yebra.

Structural observations and new research questions. An empirical analysis of the certified optimal configurations reveals two notable patterns: the number of critical triangles grows steadily with n , and the areas of the noncritical triangles cluster into a small number of distinct values. We document these observations and pose them as open research questions that may guide future investigations into the combinatorial geometry of extremal point sets.

Reproducibility and data. Many of the best-known configurations for the Heilbronn problem have circulated only through personal communication rather than formal publication, making independent verification difficult. We collect all best-known configurations for $n \leq 16$ in this paper and its appendix. The complete source code, input data, and resulting solution files are publicly available at <https://github.com/spiralulam/heilbronn>.

1.6 Paper outline

Section 2 establishes a formal setting for the Heilbronn problem and derives structural properties of optimal configurations that are exploited in the optimization model. Section 3 develops the mixed-integer formulation and its enhancements, including a symmetry-breaking strategy based on the boundary structure of optimal configurations. Section 4 presents the optimize-then-refine framework and illustrates it in detail for $n = 7$. Section 5 reports the exact optimal coordinates for $n = 5, \dots, 9$. Section 6 discusses the structural patterns and open questions. Section 7 concludes with an outlook. Appendix A collects the best-known configurations for $n = 10, \dots, 16$.

2 Original formulation and theoretical results

This section establishes a basic formulation of the Heilbronn problem in Section 2.1, which already suffices to prove its solvability. The theoretical results in Section 2.2 about the boundary points of optimal configurations and in Section 2.3 about a sign property of determinants are not only of theoretical interest but provide valuable structure that is exploited in the optimization model of Section 3, leading to a substantial runtime reduction.

2.1 The original formulation and its solvability

Let $n \geq 3$ denote the number of points $p_1 = (x_1, y_1), \dots, p_n = (x_n, y_n)$ we want to place in \mathbb{R}^2 , $I = \{1, \dots, n\}$, and

$$T = \{(i, j, k) \mid 1 \leq i < j < k \leq n\} \subset \mathbb{N}^3$$

the set of all triangles, where each triangle $t = (i, j, k) \in T$ is identified with the indices of the defining points p_i, p_j and p_k . By construction, $|T| = \binom{n}{3}$.

Remark 1. The enumeration of the points heavily influences the symmetry breaking and therefore becomes a crucial modeling choice, as discussed in Section 3.3.3.

Further, let A_t be the area of triangle t and $[0, 1]^2$ the unit square. Then the Heilbronn problem can be formulated as the following optimization problem:

$$P_\Delta : \quad \max_{x, y} \min_{t \in T} A_t \quad \text{s.t.} \quad (x_i, y_i) \in [0, 1]^2, \quad i \in I.$$

and Δ_n coincides with the optimal value of P_Δ . All further specifications of this problem amount to modeling decisions, of which there are many and which must be made carefully. A first fundamental modeling choice concerns the representation of the triangle areas. For a triangle $t = (i, j, k)$, the area can be expressed via the classical shoelace determinant formula,

$$A_t = \frac{1}{2} \left| \det \begin{pmatrix} x_i & y_i & 1 \\ x_j & y_j & 1 \\ x_k & y_k & 1 \end{pmatrix} \right| = \frac{1}{2} |x_i y_j + x_j y_k + x_k y_i - x_i y_k - x_j y_i - x_k y_j|.$$

With this representation, P_Δ becomes a nonsmooth, nonconvex optimization problem with box constraints. Before introducing additional modeling refinements in Section 3, we begin by establishing its solvability.

Proposition 1. *The optimization problem P_Δ admits a solution for every n , and its optimal value Δ_n is strictly positive.*

Proof. Existence of a solution follows from the Weierstraß theorem. The objective function is continuous, and the feasible set is nonempty, closed, and bounded; hence it is compact. Therefore, the maximum is attained. Strict positivity of Δ_n follows, for example, from Erdős' construction discussed in Section 1.3, which provides an explicit positive lower bound. \square

Remark 2. In [18], Roth argues that P_Δ is solvable because its feasible set “is a closed convex region”. However, closedness and convexity alone do not guarantee existence of a maximizer; compactness (or an equivalent coercivity argument) is required.

The following corollaries are immediate consequences of Proposition 1, since otherwise the configuration would contain a triangle of zero area.

Corollary 1. *An optimal configuration for the Heilbronn problem does not contain two coincident points.*

Corollary 2. *An optimal configuration for the Heilbronn problem does not contain three collinear points.*

2.2 Boundary structure of optimal configurations

Let $P = \{p_1, \dots, p_n\} \subseteq S = [0, 1]^2$ be an optimal configuration for the Heilbronn problem, and let

$$K := \text{conv}(P) = \left\{ \sum_{i=1}^n \lambda_i p_i \mid \lambda_i \geq 0, \sum_{i=1}^n \lambda_i = 1 \right\}$$

denote its convex hull. Let v_1, \dots, v_m be the vertices of K , listed in cyclic order, so that $K = v_1 v_2 \cdots v_m$ is a convex m -gon.

Lemma 1 (Minimal covering parallelogram). *The unit square $S = [0, 1]^2$ is a minimum-area covering parallelogram of K .*

Proof. Assume for contradiction that there exists a parallelogram $S' \supseteq K$ with $\text{area}(S') < 1$.

Since every parallelogram is an affine image of the unit square, there exists an affine bijection $T : S' \rightarrow S$ with positive determinant satisfying

$$|\det T| = \frac{\text{area}(S)}{\text{area}(S')} = \frac{1}{\text{area}(S')} > 1.$$

Because $K \subseteq S'$, we have $T(K) \subseteq S$, hence $T(P) \subseteq S$. For every triangle τ formed by points of P , its area scales by

$$\text{area}(T(\tau)) = |\det T| \text{area}(\tau).$$

Thus

$$\Delta(T(P)) = |\det T| \Delta(P) > \Delta(P),$$

contradicting the optimality of P . □

Proposition 2 (Boundary structure). *Let $n \geq 3$ and let $P \subseteq [0, 1]^2$ be an optimal configuration. Then:*

- (i) *Each edge of the unit square contains at least one point of P .*
- (ii) *At least three distinct vertices of $K = \text{conv}(P)$ lie on the boundary ∂S .*

Proof. Part (i). We prove the statement for the lower side $y = 0$; the other sides follow analogously. Assume that no point lies on $y = 0$ and let $y_{\min} := \min_i y_i > 0$. Now consider the affine map

$$T(x, y) = \left(x, \frac{y - y_{\min}}{1 - y_{\min}} \right).$$

This map sends the point set into $[0, 1]^2$ and scales all triangle areas by the constant factor

$$\det \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{1 - y_{\min}} \end{pmatrix} = \frac{1}{1 - y_{\min}} > 1.$$

Hence the minimal triangle area strictly increases, contradicting optimality. Therefore $y_{\min} = 0$, i.e., at least one point lies on the side $y = 0$. The same argument applies to the other three sides.

Part (ii). By part (i), every edge of S contains a point of P . With only two *distinct* boundary points, the pigeonhole principle forces each of them to lie on two edges, hence at two corners of S . Moreover, these two corners must be *opposite* (e.g. $(0, 0)$ and $(1, 1)$), since two adjacent corners share only one edge, leaving the opposite edge uncovered.

Let $\delta > 0$ be the minimum distance of the remaining $n - 2$ interior points to ∂S , and consider the linear map

$$T_\epsilon(x, y) = \left(\left(1 + \frac{\epsilon}{2}\right)x - \frac{\epsilon}{2}y, -\frac{\epsilon}{2}x + \left(1 + \frac{\epsilon}{2}\right)y \right).$$

One verifies:

1. $T_\epsilon(0, 0) = (0, 0)$ and $T_\epsilon(1, 1) = (1, 1)$, so both corner points are fixed.
2. $\det T_\epsilon = 1 + \epsilon > 1$.
3. For sufficiently small $\epsilon > 0$ (depending on δ), every interior point is mapped into S .

Hence $T_\epsilon(P) \subset S$ and $\Delta(T_\epsilon(P)) = (1 + \epsilon) \Delta(P) > \Delta(P)$, contradicting optimality. \square

Remark 3. Part (i) of Proposition 2 was proved in [17]; we reproduce the proof because the same affine scaling idea also underlies Lemma 1 and part (ii). The strengthening to three boundary vertices in part (ii) is new.

The following proposition sharpens the bound to five boundary vertices.

Proposition 3 (Boundary vertices). *Let $n \geq 5$ and let $P \subseteq [0, 1]^2$ be an optimal configuration. Then at least five distinct vertices of $K = \text{conv}(P)$ lie on the boundary ∂S of the unit square.*

Proof. By Lemma 1, S is a minimum-area covering parallelogram of K . Applying Lemma 2 of Zeng and Chen [24] yields three possible cases:

- (1) **No common vertices.** Then at least five vertices of K lie on the four edges of S .
- (2) **Exactly one common vertex.** Then each adjacent edge interior contains a vertex of K , and the remaining two edges each contain a vertex, yielding at least five distinct boundary vertices.
- (3) **Two common vertices forming a diagonal.** Then each of the four edges of S contains a vertex of K in its interior. Together with the two corner vertices, this yields at least six distinct boundary vertices.

In all cases, ∂S contains at least five distinct vertices of K , hence at least five distinct points of P . \square

Remark 4. We thank Amirali Modir for drawing our attention to Lemma 2 of [24] and for explaining the geometric mechanism that enforces multiple boundary vertices in optimal configurations.

2.3 On the signs of the triangle areas

The signed area of a triangle with vertices p_i, p_j, p_k is

$$A^\pm(i, j, k) := \frac{1}{2} \det \begin{pmatrix} 1 & x_i & y_i \\ 1 & x_j & y_j \\ 1 & x_k & y_k \end{pmatrix} = \frac{1}{2} [(x_j - x_i)(y_k - y_i) - (y_j - y_i)(x_k - x_i)].$$

It is well known that $A^\pm(i, j, k) > 0$ if and only if (p_i, p_j, p_k) are ordered counterclockwise, $A^\pm(i, j, k) < 0$ for a clockwise ordering, and $A^\pm(i, j, k) = 0$ for collinear points. In particular, the (unsigned) area satisfies $A(i, j, k) = |A^\pm(i, j, k)|$, and for counterclockwise-ordered triples $A(i, j, k) = A^\pm(i, j, k)$. This sign information plays a key role in the optimization model of Section 3.

3 Mixed-integer optimization models

After a brief introduction to mixed-integer optimization in Section 3.1, we present in Section 3.2 a baseline mixed-integer formulation that mirrors the original nonsmooth model as closely as possible; solving this model already certifies optimality for small instances but becomes intractable around $n = 8$. Section 3.3 then develops several enhancements—in particular a symmetry-breaking strategy based on boundary structure that, as a byproduct, also fixes the signs of certain triangle areas. Together, these improvements lead to a much stronger formulation. The resulting final model is collected in Section 3.4, and Section 3.5 compares it with the formulation of Modir, Monji, and Kocuk [17].

3.1 A brief primer on mixed-integer optimization

A *mixed-integer optimization problem* (MIP) is an optimization problem in which some of the decision variables are required to take integer values while others may vary continuously. This makes it possible to encode logical decisions, combinatorial constraints, and piecewise structure within a single mathematical model.

The practical power of MIP rests on the *branch-and-cut* algorithm, which interleaves two mechanisms:

- *Branching.* The algorithm recursively partitions the search space into smaller subproblems by fixing or restricting integer variables, and solves a continuous relaxation of each subproblem to obtain a *dual bound* (upper bound for a maximization problem). Subproblems whose relaxation bound cannot improve upon the best known feasible solution are discarded.
- *Cutting.* At each node of the search tree, the solver may add *cutting planes*—valid inequalities that are violated by the current relaxation solution but satisfied by all feasible integer points. These tighten the relaxation and often dramatically reduce the number of branches needed.

Throughout the search, the algorithm maintains a *primal bound* (lower bound for maximization) from the best feasible solution found so far. When the primal and dual bounds coincide, the solver has *proved* that no better solution exists—the result is a certificate of global optimality, not merely a best-found solution. This certifying capability makes mixed-integer optimization attractive for the Heilbronn problem, where the goal is to establish definitive values of Δ_n , not merely to find good configurations. For comprehensive introductions we refer to the textbooks by Wolsey [22] and Conforti, Cornuéjols, and Zambelli [9].

When the model involves nonlinear—in our case, quadratic—constraints or objective terms, it is called a *mixed-integer nonlinear problem* (MINLP). Solvers for such problems employ spatial branching and convex relaxations to handle nonconvexities, and have seen dramatic performance improvements over the past two decades; we refer to Belotti et al. [1] and Burer and Letchford [5] for surveys of the field, and to Bixby [4] for a historical perspective on the linear case. The formulations developed in this section fall into the MINLP class.

3.2 A first mixed-integer formulation

For each triangle $t = (i, j, k) \in T$ we introduce a continuous auxiliary variable $A_t^\pm \in [-\frac{1}{2}, \frac{1}{2}]$ representing the signed area

$$A_t^\pm = \frac{1}{2} (x_i y_j + x_j y_k + x_k y_i - x_i y_k - x_j y_i - x_k y_j).$$

To eliminate the minimum over triangles in the objective, we introduce an auxiliary variable z and impose

$$z \leq |A_t^\pm|, \quad t \in T.$$

This formulation remains nonsmooth because of the absolute value. To handle it, we introduce a binary variable $b_t \in \{0, 1\}$ for each triangle t that selects the sign of A_t^\pm , and replace the absolute-value constraint by

$$z \leq (2b_t - 1) A_t^\pm, \quad t \in T.$$

Maximizing z forces the model to choose the sign that yields the larger of $-A_t^\pm$ and A_t^\pm , thereby mimicking the absolute value $|A_t^\pm| = \max(A_t^\pm, -A_t^\pm)$.

Remark 5. This reformulation only mimics the absolute value for the so-called *critical triangles*, that is, the triangles of minimal area. This is sufficient since these determine the optimal value Δ_n .

Our first baseline mixed-integer optimization model is therefore

$$\begin{aligned}
P_{\Delta}^0 : \quad & \max_{z,x,y} \quad z \\
& \text{s.t.} \quad (x_i, y_i) \in [0, 1]^2, \quad i \in I, \\
& \quad A_t^{\pm} = \frac{1}{2} (x_i y_j + x_j y_k + x_k y_i - x_i y_k - x_j y_i - x_k y_j), \quad t \in T, \\
& \quad z \leq (2b_t - 1) A_t^{\pm}, \quad t \in T, \\
& \quad A_t^{\pm} \in \left[-\frac{1}{2}, \frac{1}{2} \right], \\
& \quad b_t \in \{0, 1\}, \quad t \in T, \\
& \quad 0 \leq z \leq \frac{1}{2}.
\end{aligned}$$

This formulation is a nonlinear mixed-integer model and can in principle be solved by any global MINLP solver.

Remark 6. This baseline model solves the case $n = 5$ within seconds using Gurobi and proves optimality for $n = 6$ in approximately three minutes. For $n = 7$, however, the solver did not terminate within 45 minutes, illustrating the rapid growth in computational difficulty.

3.3 Some improvements

3.3.1 Upper bounds on Δ_n

Besides the trivial bound $\Delta_n \leq 1/2$, we have $\Delta_n \leq \Delta_{n-1}$, which is straightforward (and formally proved in [17]): the Heilbronn problem with $n - 1$ points provides an upper bound for the problem with n points. This allows us to add the constraint $z \leq \Delta_{n-1}$ to the model.

3.3.2 Substitution of continuous products

It is usually advantageous for MINLP solvers to isolate the bilinear expressions, since they represent the main source of nonconvexity in the model and will be approximated during the solve (e.g., using McCormick inequalities [16]). We therefore introduce auxiliary variables $w_{ij} \in [0, 1]$ with $w_{ij} = x_i y_j$ for $i, j \in I$.

3.3.3 Symmetry breaking

The core obstacle in solving P_{Δ}^0 is its enormous symmetry. Two symmetry groups act simultaneously:

- The *dihedral group* D_4 of the unit square, consisting of four rotations (by 0° , 90° , 180° , 270°) and four reflections (about the two axes and the two diagonals), acts on the coordinates.
- The *symmetric group* S_n acts on the point labels: any permutation of $\{1, \dots, n\}$ yields an equivalent configuration.

Together, these give rise to a symmetry group of order $8 \cdot n!$ (e.g., $8 \cdot 9! = 2\,903\,040$ for $n = 9$), meaning that a single geometric configuration can be represented by millions of distinct variable assignments in the model. Without symmetry breaking, the solver must explore a vast number of equivalent branches, greatly inflating the search tree.

Our strategy exploits the boundary structure established in Proposition 3 to eliminate this symmetry group completely. We proceed in three steps.

Step 1: Boundary assignment and geometric symmetry elimination. By Proposition 3, every optimal configuration has at least five vertices of its convex hull on the boundary ∂S of the unit square. A counting argument shows that, among the four edges, at least one must carry two or more of these boundary vertices. We designate this edge to be the *left* edge ($x = 0$) without loss of generality, since any configuration can be mapped to one satisfying this convention by an appropriate element of D_4 . This eliminates all elements of D_4 except the horizontal reflection $(x, y) \mapsto (x, 1 - y)$, which preserves the left edge.

Step 2: Counterclockwise boundary labeling. We label the five designated boundary points p_1, \dots, p_5 in counterclockwise order starting from the left edge:

$$\begin{aligned}
 p_1 : x_1 = 0, & & (\text{left edge, lower}), \\
 p_2 : y_2 = 0, & & (\text{bottom edge}), \\
 p_3 : x_3 = 1, & & (\text{right edge}), \\
 p_4 : y_4 = 1, & & (\text{top edge}), \\
 p_5 : x_5 = 0, \quad y_5 \geq y_1, & & (\text{left edge, upper}).
 \end{aligned} \tag{1}$$

Because the assignment is counterclockwise from a fixed starting edge, this labeling is canonical: it maps every D_4 -orbit to a unique representative. The ordering constraint $y_5 \geq y_1$ resolves the remaining ambiguity between the two left-edge points. Furthermore, we impose $x_2 \leq x_4$, which eliminates the remaining horizontal reflection $(x, y) \mapsto (x, 1 - y)$ (this map swaps the bottom and top edges, hence swaps p_2 and p_4). Together with Step 1, this breaks the full dihedral symmetry D_4 .

Step 3: Interior point ordering. For $n \geq 6$, the remaining points p_6, \dots, p_n are ordered by increasing x -coordinate:

$$x_6 \leq x_7 \leq \dots \leq x_n.$$

This eliminates the $(n - 5)!$ -fold labeling symmetry among the non-boundary points.

The combined effect of these three steps is summarized in the following proposition.

Proposition 4 (Symmetry breaking). *For $n \geq 5$, the constraints (1) together with $x_2 \leq x_4$ and $x_6 \leq \dots \leq x_n$ select at least one representative from each equivalence class of configurations under the joint action of D_4 and S_n . In particular, they fix five of the $2n$ coordinate variables outright.*

Proof. Every equivalence class under the joint action of D_4 and S_n contains a representative in which (i) the edge carrying at least two convex-hull vertices is the left edge (eliminating D_4 up to the horizontal reflection), (ii) the boundary points are labeled counterclockwise starting from the left edge with $y_5 \geq y_1$ (fixing five boundary labels), (iii) $x_2 \leq x_4$ (eliminating the remaining horizontal reflection), and (iv) the interior points are sorted by x -coordinate (eliminating S_{n-5}). \square

Remark 7. The symmetry-breaking constraints do not always select a *unique* representative: if an edge other than the left edge carries more than one boundary vertex of K , any of them may serve as the designated boundary point (p_2 , p_3 , or p_4), while the remaining ones enter the interior ordering. This results in a small, bounded number of equivalent representations for such configurations, but does not affect the correctness of the model.

Table 1 summarizes the variables fixed by the symmetry-breaking constraints. The five fixed coordinates represent a reduction of the continuous search space by 5 out of $2n$ dimensions; simultaneously, the labeling symmetry is entirely eliminated, precluding the solver from revisiting equivalent branch-and-bound nodes. This is illustrated in Figure 2.

Point	Coord.	Value	Edge
p_1	x_1	0	left
p_2	y_2	0	bottom
p_3	x_3	1	right
p_4	y_4	1	top
p_5	x_5	0	left

Table 1: Coordinates fixed a priori by the symmetry-breaking strategy. The remaining free coordinates are y_1, x_2, y_3, x_4, y_5 , and all coordinates of p_6, \dots, p_n , subject to $y_5 \geq y_1, x_2 \leq x_4$, and $x_6 \leq \dots \leq x_n$.

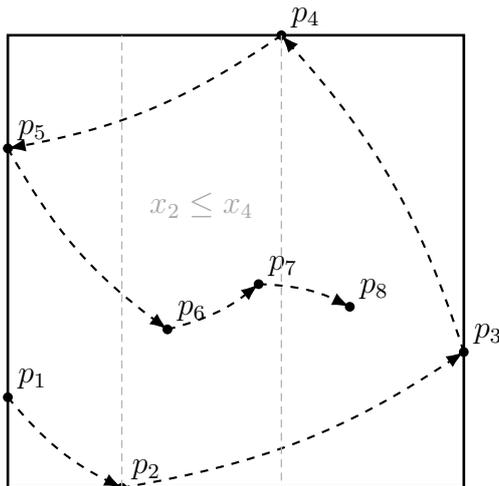


Figure 2: Symmetry breaking used in the optimization model. Five points are fixed on the boundary in counterclockwise order, while the remaining points are ordered by increasing x -coordinate; additionally $x_2 \leq x_4$.

Remark 8. This symmetry-breaking approach also yields correct results for $n < 5$, although the theoretical justification does not stem from Proposition 3 but from a separate investigation of the cases $n = 3$ and $n = 4$.

3.3.4 Sign fixing

The symmetry breaking approach from Section 3.3.3 allows us to fix the signs for two groups of triangles:

1. All triangles in the set $T^+ = \{(i, j, k) \in T \mid k \leq 5\}$ have their three vertices on the boundary, arranged counterclockwise (since $i < j < k$). Therefore $A_t^\pm \geq 0$ for all $t \in T^+$.
2. Since p_1 and p_5 are located on the left edge with $y_5 \geq y_1$, we have $A_t^\pm \leq 0$ for all $t \in T^- := \{(i, j, k) \in T \mid i = 1, j = 5\}$.

This allows us to add the constraints

$$b_t = 1, t \in T^+ \quad \text{and} \quad b_t = 0, t \in T^-$$

and therefore fix some of the binary variables in advance.

Remark 9. It is also valid to introduce only the binary variables that are not fixed a priori, but we did not experience a computational advantage from that approach.

3.4 The final model

Combining all ingredients from the previous subsection—product substitution, symmetry breaking, sign fixing, and the tighter bound on z —we arrive at the formulation P_Δ^* displayed in Figure 3. The colored boxes highlight the individual enhancement groups; the core model corresponds to the baseline formulation P_Δ^0 expressed in the substituted variables.

Remark 10. Using this final formulation, the Heilbronn problem for $n = 5$ is solved in 0.13 seconds, for $n = 6$ in 0.7 seconds, for $n = 7$ in 1.75 seconds, for $n = 8$ in 14.03 seconds, and for $n = 9$ in 908 seconds—improving upon the previously reported effort of about one day [17] by more than an order of magnitude.

Core Model

$$\begin{aligned}
 P_{\Delta}^* : \quad & \max_{z,x,y} \quad z \\
 & \text{s.t.} \quad (x_i, y_i) \in [0, 1]^2, \quad i \in I, \\
 & \quad A_{ijk}^{\pm} = \frac{1}{2} (w_{ij} + w_{jk} + w_{ki} - w_{ik} - w_{ji} - w_{kj}), \quad (i, j, k) \in T \\
 & \quad z \leq (2b_t - 1)A_t^{\pm}, \quad t \in T
 \end{aligned}$$

Product Substitution

$$w_{ij} = x_i y_j, \quad i, j \in I$$

Sign Fixing

$$\begin{aligned}
 b_t &= 1, \quad t \in T^+ \\
 b_t &= 0, \quad t \in T^-
 \end{aligned}$$

Symmetry Breaking

$$\begin{aligned}
 x_1 &= x_5 = y_2 = 0 \\
 x_3 &= y_4 = 1 \\
 x_3 &\leq x_4 \\
 y_1 &\leq y_5 \\
 x_i &\leq x_{i+1}, \quad i \geq 6
 \end{aligned}$$

Variable Domains

$$\begin{aligned}
 A_t^{\pm} &\in \left[-\frac{1}{2}, \frac{1}{2}\right], \quad t \in T \\
 w_{ij} &\in [0, 1], \quad i, j \in I \\
 b_t &\in \{0, 1\}, \quad t \in T, \\
 0 &\leq z \leq \Delta_{n-1}
 \end{aligned}$$

Figure 3: The final mixed-integer formulation P_{Δ}^* . Colored boxes indicate the core model P_{Δ}^0 (blue) and the enhancements: product substitution (teal), sign fixing (orange), symmetry breaking (red), and variable domains (gray).

3.5 Comparison with Monji, Modir and Kocuk

The concurrent work of Monji, Modir and Kocuk [17] was the first to formulate the Heilbronn problem as a mixed-integer nonlinear problem and to solve it with a global optimization solver. Their approach certifies the known optimal values for $n \leq 8$ in seconds and certifies $n = 9$ within one day—a dramatic improvement over the months-long CPU and GPU campaign of Chen et al. [6]. While both their work and ours ultimately rely on solving an MINLP with Gurobi, the two formulations differ in several respects that we discuss below.

Handling the absolute value. The core modeling challenge is the constraint $z \leq |A_t^{\pm}|$. Monji et al. linearize it with a classical big- M disjunction: for every triangle t they impose the pair of inequalities $(1 - b_t)(M + \frac{1}{2}) + A_t^{\pm} \geq z$ and $b_t(M + \frac{1}{2}) - A_t^{\pm} \geq z$ (together with their tightening counterparts), where M is an upper bound on the optimal value. In contrast, our formulation replaces the absolute value by the single bilinear constraint $z \leq (2b_t - 1)A_t^{\pm}$, which avoids a big- M parameter altogether and lets the solver handle the nonconvexity through its native spatial branching.

Symmetry breaking. Monji et al. break the labeling symmetry by sorting the y -coordinates, i.e., $y_1 \leq y_2 \leq \dots \leq y_n$, and by fixing $x_1 \leq 1/2$ and $x_2 \geq x_1$. Our strategy, by contrast, exploits the boundary structure established in Proposition 3: five points are placed on the boundary of the unit square and arranged in counterclockwise order, which fixes five coordinate values a priori ($x_1 = x_5 = y_2 = 0$, $x_3 = y_4 = 1$); by comparison, the y -sorting convention imposes only ordering inequalities without pinning any coordinate to a specific value. Moreover, the induced sign fixing (Section 2.3) eliminates $\binom{5}{3} + (n - 5) = n + 5$ binary variables from the model, a reduction that has no counterpart in the sorting approach.

No additional auxiliary variables. Monji et al. introduce three groups of enhancements that progressively add auxiliary variables to the model: binary variables c_{1i}, c_{2i} for edge proximity (Corollary 1), binary strip-assignment variables r_{pi} for a horizontal partition into thin rectangles each hosting at most two points (Corollary 3), and binary grid-cell variables u_{pqi} for a square subgrid in which each cell hosts at most one point (Corollary 4). All three families strengthen the relaxation, but they also increase the model size. Our formulation requires none of these extra variables: the boundary fixing already constrains the first five points to specific edges, and the tighter upper bound $z \leq \Delta_{n-1}$ narrows the objective range directly in the variable bounds without additional binary decisions.

No optimization-based bound tightening. Monji et al. compute a heuristic lower bound H_n by sampling 10^6 random placements and feed it into the strip-width and grid-size parameters that govern the auxiliary binary variables described above. In addition, they solve a series of restricted Heilbronn sub-problems in smaller rectangles to compute explicit bounds on individual y_i -coordinates, and their proof of the near-boundary property for $n = 9$ requires a further auxiliary solve of about 23 hours. These pre-processing steps account for the bulk of the reported one-day computational effort. Our formulation avoids both heuristic seeding and optimization-based pre-processing: the only problem-specific bounds are the proven upper bound $\Delta_n \leq \Delta_{n-1}$ and the coordinate constraints implied by the boundary structure (Proposition 3), both of which follow from rigorous arguments.

No unproven structural assumptions in certified results. For $n = 9$, Monji et al. prove a near-boundary property (their Proposition 5) through a pair of auxiliary solves that compare the case of eight near-boundary points against the case of at most seven, using an $\varepsilon = 10^{-2}$ tolerance. For $n = 10$, their best result uses two unproven structural conjectures. By contrast, every enhancement in our certified results for $n \leq 9$ is backed by a mathematical proof; we flag conjectural assumptions explicitly whenever they are used (see Section 5).

Remark 11. Monji et al. conjecture that every optimal configuration for $n \geq 9$ places exactly eight points on the boundary (their Conjecture 1). The best-known configurations for $n = 13$ and $n = 15$ [12] have fewer than eight boundary points. These configurations are not proven to be optimal, so the conjecture could still be true; however, the counterexamples among the best known solutions suggest that it is at least not obviously so.

In summary, our final model P_Δ^* is more compact than the formulation of Monji et al.—it contains fewer auxiliary binary variables and no big- M constants—and still achieves faster certified solves without requiring any optimization-based pre-processing. Both approaches demonstrate that mixed-integer nonlinear optimization is a viable tool for the Heilbronn problem.

4 The optimize-then-refine framework

4.1 The framework

Our approach proceeds in two stages. In the first step, we solve the MINLP formulation P_{Δ}^* from Section 3.4 to global optimality using Gurobi. The solver returns a numerical point placement together with matching lower and upper bounds, thereby certifying the optimal value Δ_n . From these numerical coordinates we identify the *critical triangles*—those whose area equals the certified minimum—and record which points lie on which edges of the unit square. In the second step, we treat these structural observations as an ansatz for deriving exact symbolic coordinates: setting the areas of the critical triangles equal yields a system of polynomial equations that can be solved in closed form with a computer algebra system. The result is a fully symbolic optimal configuration whose optimality is backed by the numerical certificate from the first step.

A key practical concern is the identification of critical triangles and boundary points in the first step, since these decisions rest on numerical data subject to solver tolerances. In all instances considered here, the gaps are unambiguous: the area of the smallest non-critical triangle exceeds that of the critical triangles by a margin that is several orders of magnitude larger than the solver tolerance, and every point classified as lying on an edge has a coordinate that deviates from the boundary value by less than 10^{-6} . This clear separation makes the structural input to Step 2 unambiguous in all cases considered.

The pipeline is illustrated in Figure 4.

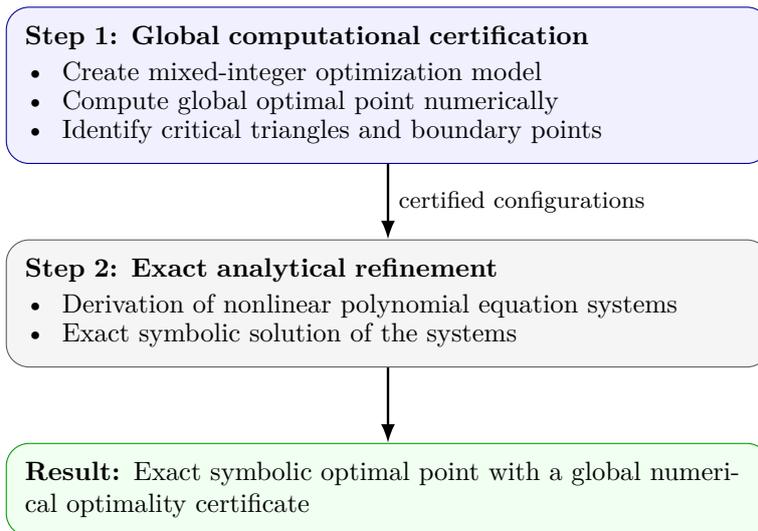


Figure 4: Methodological pipeline of the paper: global computational certification (Step 1), exact analytical refinement (Step 2), yielding exact symbolic optimal points with global numerical optimality certificates.

While Step 1 is a standard (though computationally challenging) optimization task, Step 2 is more involved. The equal-area conditions on the critical triangles form a system of polynomial equations whose unknowns are the free coordinate parameters remaining after boundary constraints and coordinate coincidences have been exploited. The number k of these parameters ranges from 4 (for $n = 5$) to 8 (for $n = 9$), and we employ three solving strategies, tried in succession.

1. *Direct symbolic solve.* For small systems ($k \leq 6$, covering $n = 5, 6, 8$), SymPy’s symbolic solver produces exact solutions directly, including the detection of the one-parameter family for $n = 6$.

2. *Gröbner basis.* For $n = 7$, the solution lies in a cubic extension field; a lexicographic Gröbner basis computation yields a triangular system from which all coordinates follow by back-substitution through the minimal polynomial $19f^3 - 27f^2 + 11f - 1 = 0$ (see Section 4.3).
3. *Numeric-to-symbolic refinement.* For $n = 9$, where $k = 8$, neither a direct solve nor a Gröbner basis computation terminates in reasonable time. Instead, we apply a numeric-to-symbolic approach: from the certified numerical coordinates we first identify the underlying number field $\mathbb{Q}(\sqrt{65})$ by recognizing $\sqrt{65}$ in the simplest coordinate, then express all coordinates as elements of this field using SymPy’s `nsimplify` function with the known algebraic extension.

The numeric-to-symbolic step involves a risk: numerical proximity does not guarantee algebraic correctness, since Gurobi’s certified values carry only about six significant digits. For instance, for $n = 9$ the spurious expression $\frac{19}{166} + \frac{3\sqrt{65}}{166} \approx 0.26017$ is numerically indistinguishable from the true value $\frac{9}{16} - \frac{3\sqrt{65}}{80} \approx 0.26017$ at solver precision, yet only the latter satisfies the equal-area equations. We therefore verify every candidate solution symbolically by substituting into the full polynomial system. Any coordinate that fails this check is corrected by fixing the validated coordinates and solving the resulting reduced (and much smaller) subsystem exactly.

The entire pipeline is implemented as a reproducible Python script, available together with all input data in the accompanying repository.¹

4.2 Software and hardware

The models are implemented in Python 3.11 using the Gurobi 13.0 solver interface (`gurobipy`); the exact symbolic refinements use SymPy 1.14. Gurobi is a commercial solver; free academic licenses are available. Gurobi runs use default settings with `MIPFocus = 2` to prioritise the dual bound. Computations were carried out on a standard desktop PC equipped with an AMD Ryzen 9 9950X3D processor (16 cores, 4.30 GHz), 96 GB DDR5 RAM, running Windows 11.

4.3 Illustration of the methodology for $n = 7$

We now illustrate the two-step pipeline on the case $n = 7$, for which Comellas and Yebra [8] obtained conjectured optimal coordinates via simulated annealing followed by a local optimality analysis. Global optimality was subsequently proved by Zeng and Chen [24], who decomposed the problem into 226 nonlinear optimization subproblems. Our approach provides an independent proof and yields simpler coordinate expressions.

Step 1: Global computational certification

Running our final MIP model (Section 3.4) with $n = 7$, Gurobi terminates after 1.75 seconds with a certified optimal value of $H_7^* \approx 0.0838590$ and the numerical configuration shown in Table 2.

Figure 5 shows the configuration together with its eight critical triangles, which all have areas within 5×10^{-7} of the optimal value. The ninth-smallest triangle area is roughly 0.0948, leaving a comfortable gap of about 13% above the optimum.

Two structural observations guide the symbolic step that follows:

1. Five of the seven points lie on the boundary (p_4, p_5 at corners, p_1 on the left edge, p_2 on the bottom edge, p_3 on the right edge).
2. Points p_2 and p_7 share the same x -coordinate, suggesting a hidden algebraic relation.

¹<https://github.com/spiralulam/heilbronn>

Pt	x	y	Edge
p_1	0	0.1352	left edge
p_2	0.7127	0	bottom edge
p_3	1	0.1808	right edge
p_4	1	1	corner
p_5	0	1	corner
p_6	0.1939	0.4926	interior
p_7	0.7127	0.5839	shared x with p_2

Table 2: Numerical coordinates returned by Gurobi for $n = 7$ (rounded to four decimal places). Five of the seven points lie on the boundary of the unit square.

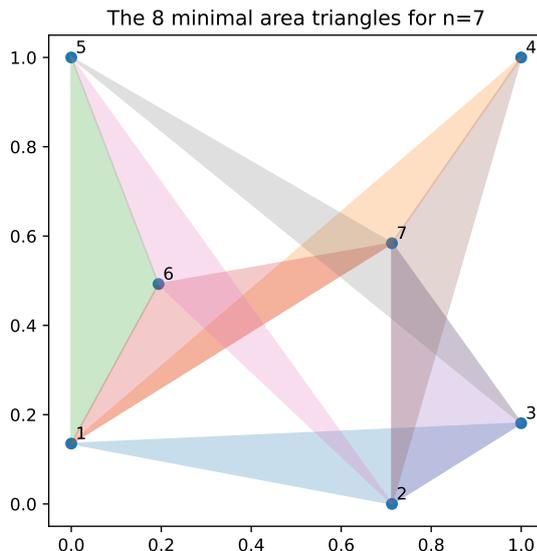


Figure 5: Optimal configuration for $n = 7$ with its eight critical triangles highlighted.

Step 2: Exact analytical refinement

Guided by these observations, we set up a parametric ansatz. The boundary structure fixes seven coordinates ($x_1 = 0, y_2 = 0, x_3 = 1, x_4 = y_4 = 1, x_5 = 0, y_5 = 1$), and the shared x -coordinate of p_2 and p_7 introduces an additional coupling, leaving six unknowns a, b, c, d, e, f (Table 3).

Pt	x	y	Rationale
p_1	0	a	left edge
p_2	b	0	bottom edge
p_3	1	c	right edge
p_4	1	1	corner
p_5	0	1	corner
p_6	d	e	interior
p_7	b	f	shared x with p_2

Table 3: Parametric ansatz for $n = 7$. Boundary constraints fix seven coordinates; the shared x -coordinate of p_2 and p_7 introduces the coupling $x_7 = b$.

Setting the areas of the eight critical triangles equal yields seven polynomial equations in six unknowns. A lexicographic Gröbner basis computation eliminates all unknowns except f , which must satisfy the irreducible cubic

$$19f^3 - 27f^2 + 11f - 1 = 0.$$

This cubic has three real roots: $f_1 \approx 0.1269$, $f_2 \approx 0.5839$, and $f_3 \approx 0.7103$. The first root gives $b = 19f_1^2 - 27f_1 + 10 \approx 6.88 > 1$, and the third gives $a = 19f_3^2 - 16f_3 + 3 \approx 1.22 > 1$; both violate the unit-square constraint. The unique feasible solution corresponds to $f = f_2$. Back-substitution through the Gröbner basis yields the exact symbolic coordinates in Table 4.

The optimal value simplifies to

$$H_7^* = f - \frac{1}{2} \approx 0.0838590090 \dots$$

Pt	x	y
p_1	0	$19f^2 - 16f + 3$
p_2	$19f^2 - 27f + 10$	0
p_3	1	$\frac{-19f^2 + 10f + 1}{2}$
p_4	1	1
p_5	0	1
p_6	$-19f^2 + 8f + 2$	$57f^2 - 41f + 5$
p_7	$19f^2 - 27f + 10$	f

Table 4: Exact symbolic coordinates of the optimal seven-point configuration, where $f \approx 0.5839$ is the second real root of $19f^3 - 27f^2 + 11f - 1$.

Comparison with Comellas and Yebra [8]. The exact coordinates we obtain are notably simpler than those reported by Comellas and Yebra, who express each coordinate as a degree-two polynomial in a parameter z with rational coefficients over denominators 19 and 38 and integer numerators as large as 223, for instance $x_1 = -\frac{50}{19}z - \frac{17}{38}z^2 + \frac{37}{38}$. In our formulation every coordinate is a degree-two polynomial in f with integer coefficients, and the optimal value has the particularly clean form $f - \frac{1}{2}$. The simplification arises because our boundary-aware labeling (assigning points to edges before solving) fixes more structure upfront, whereas the simulated annealing approach of Comellas and Yebra did not exploit this structure.

5 Optimal solutions for $5 \leq n \leq 9$

For $n = 3$ and $n = 4$ the optimal configurations are trivially given by placing the points at the corners of the unit square, yielding smallest triangle areas of 0.5, which must be optimal.

5.1 $n = 5$: Four critical triangles, $H_5^* = \frac{\sqrt{3}}{9}$

The configuration has four critical triangles. Setting these four areas equal yields a system whose unique feasible solution gives the closed-form optimal value $H_5^* = \frac{\sqrt{3}}{9} \approx 0.19245$.

Pt	x	y
1	0	$\frac{1}{3}$
2	$\frac{\sqrt{3}}{3}$	0
3	1	$1 - \frac{\sqrt{3}}{3}$
4	$\frac{2}{3}$	1
5	0	1

Table 5: Exact coordinates for $n = 5$.

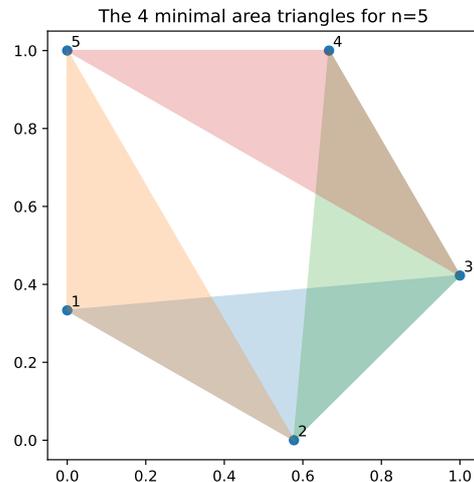


Figure 6: Optimal configuration for $n = 5$.

5.2 $n = 6$: Six critical triangles, $H_6^* = \frac{1}{8}$

The configuration has six critical triangles. Remarkably, there is a one-parameter family of optimal solutions parameterized by $c \in [0, \frac{1}{4}]$, all achieving $H_6^* = 1/8 = 0.125$. The full distribution of triangle areas is independent of the choice of c .

Pt	x	y
1	0	c
2	$\frac{1}{2}$	0
3	1	$\frac{1}{2} - c$
4	$\frac{1}{2}$	1
5	0	$c + \frac{1}{2}$
6	1	$1 - c$

Table 6: Exact coordinates for $n = 6$, $c \in [0, \frac{1}{4}]$.

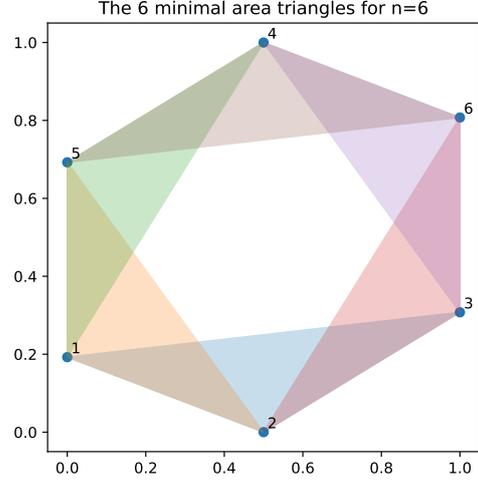


Figure 7: Optimal configuration for $n = 6$.

5.3 $n = 7$: Eight critical triangles, $H_7^* = f - \frac{1}{2}$

The detailed derivation for $n = 7$ is given in Section 4.3. The configuration has eight critical triangles, and the exact coordinates are parameterized by the root of $19f^3 - 27f^2 + 11f - 1$ near 0.5839, yielding $H_7^* \approx 0.08386$.

Pt	x	y
1	0	$19f^2 - 16f + 3$
2	$19f^2 - 27f + 10$	0
3	1	$\frac{-19f^2 + 10f + 1}{2}$
4	1	1
5	0	1
6	$-19f^2 + 8f + 2$	$57f^2 - 41f + 5$
7	$19f^2 - 27f + 10$	f

Table 7: Exact coordinates for $n = 7$.

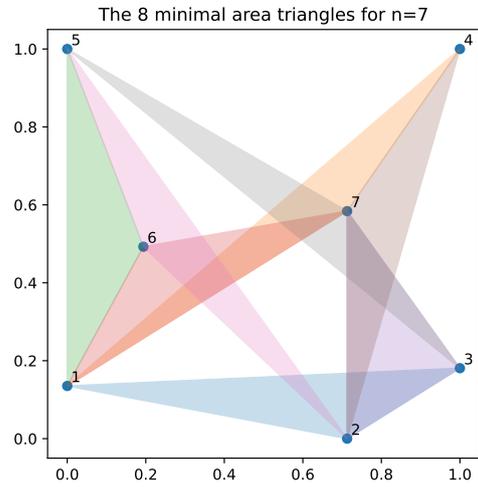


Figure 8: Optimal configuration for $n = 7$.

5.4 $n = 8$: Twelve critical triangles, $H_8^* = -\frac{1}{36} + \frac{\sqrt{13}}{36}$

The configuration exhibits a point-symmetric structure with twelve critical triangles and a unique optimal solution involving $\sqrt{13}$.

Pt	x	y
1	0	0
2	$\frac{1}{6} + \frac{\sqrt{13}}{6}$	0
3	1	$\frac{7}{18} - \frac{\sqrt{13}}{18}$
4	1	1
5	0	$\frac{11}{18} + \frac{\sqrt{13}}{18}$
6	$\frac{5}{6} - \frac{\sqrt{13}}{6}$	1
7	$\frac{5}{6} - \frac{\sqrt{13}}{6}$	$\frac{7}{9} - \frac{\sqrt{13}}{9}$
8	$\frac{1}{6} + \frac{\sqrt{13}}{6}$	$\frac{2}{9} + \frac{\sqrt{13}}{9}$

Table 8: Exact coordinates for $n = 8$.

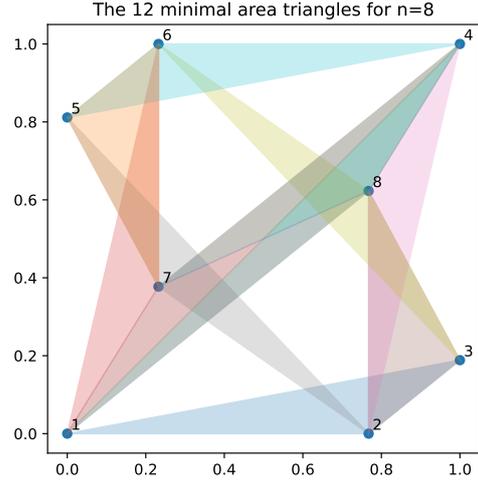


Figure 9: Optimal configuration for $n = 8$.

5.5 $n = 9$: Eleven critical triangles, $H_9^* = -\frac{11}{64} + \frac{9\sqrt{65}}{320}$

The configuration has eleven critical triangles and is invariant under reflection in the anti-diagonal $(x, y) \mapsto (1 - y, 1 - x)$. The unique optimal solution involves $\sqrt{65}$.

Pt	x	y
1	0	$1 - \frac{\sqrt{65}}{10}$
2	$\frac{3}{8} - \frac{\sqrt{65}}{40}$	0
3	1	$\frac{9}{16} - \frac{3\sqrt{65}}{80}$
4	$\frac{3}{8} - \frac{\sqrt{65}}{40}$	1
5	0	$\frac{5}{8} + \frac{\sqrt{65}}{40}$
6	$\frac{1}{4} + \frac{\sqrt{65}}{20}$	$\frac{3}{4} - \frac{\sqrt{65}}{20}$
7	$\frac{7}{16} + \frac{3\sqrt{65}}{80}$	0
8	$\frac{\sqrt{65}}{10}$	1
9	1	$\frac{5}{8} + \frac{\sqrt{65}}{40}$

Table 9: Exact coordinates for $n = 9$.

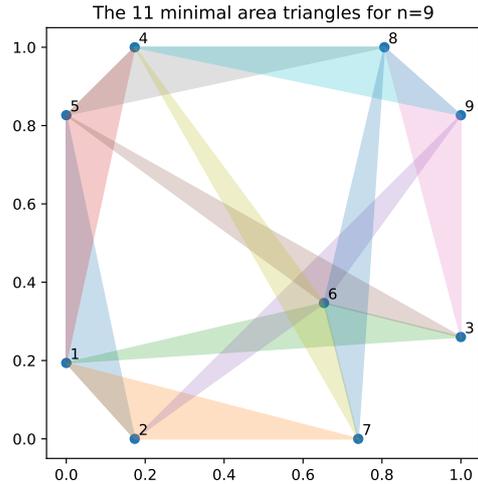


Figure 10: Optimal configuration for $n = 9$.

6 Further observations and research questions

6.1 Growth of the number of critical triangles

A structural feature of the optimal configurations worth noting is the number of critical triangles, i.e., triangles whose area equals the minimum. Table 10 and Figure 11 summarize the counts for $n = 5, \dots, 9$.

The data suggest a broadly increasing trend, although the decrease from $n = 8$ to $n = 9$ shows that the growth is not strictly monotone. A natural question arises:

Can one establish a nontrivial lower bound on the number of critical triangles in an optimal Heilbronn configuration as a function of n ?

n	Critical triangles
5	4
6	6
7	8
8	12
9	11

Table 10: Number of critical triangles for $n = 5, \dots, 9$.

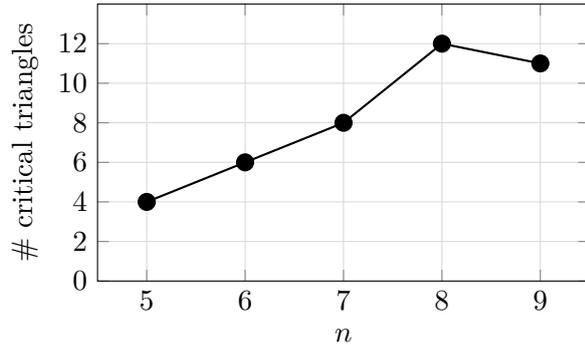


Figure 11: Critical triangles as a function of n .

Any such bound would have implications for the structure of the polynomial systems that arise in Step 2 of the framework, since each critical triangle contributes an equality constraint. Even a linear lower bound of the form $\Omega(n)$ would be significant, as it would guarantee that optimal configurations are constrained by at least as many active area equalities as there are free coordinate variables.

6.2 Clustering of noncritical triangle areas

A second observation concerns the *noncritical* triangles—those whose area strictly exceeds the minimum. Inspecting the degeneracy plots (Figures 12 and 13), one notices that the triangle areas do not spread out continuously but instead cluster around a small number of distinct values.

For $n = 5$, there are only three distinct noncritical area levels beyond the four critical triangles (see Figure 12). For $n = 8$, the pattern is particularly clear: after the twelve critical triangles, the remaining $\binom{8}{3} - 12 = 44$ triangle areas visibly group into a handful of narrow clusters (see Figure 13). Similar clustering is observed for $n = 6$, $n = 7$, and $n = 9$.

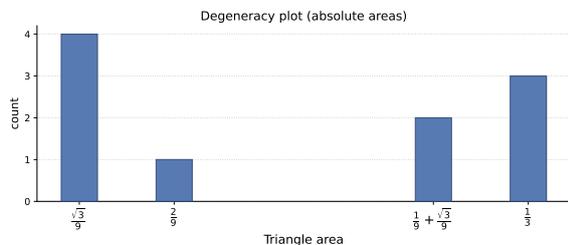


Figure 12: Degeneracy plot for $n = 5$: all $\binom{5}{3} = 10$ triangle areas sorted in ascending order. Four critical triangles are followed by three clearly separated noncritical area levels.

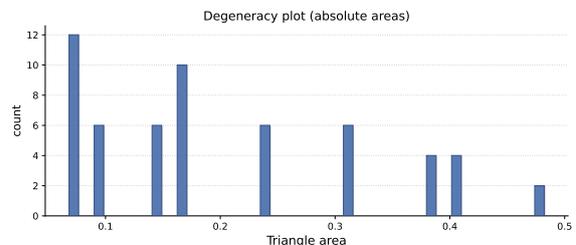


Figure 13: Degeneracy plot for $n = 8$: all $\binom{8}{3} = 56$ triangle areas sorted in ascending order. Twelve critical triangles are followed by distinctly clustered noncritical area levels.

To the best of our knowledge, this clustering phenomenon has not been noted in the literature. This suggests a structural rigidity in the extremal configurations that goes beyond the optimality of the minimum area.

Can one explain the clustering of noncritical triangle areas in optimal Heilbronn configurations? In particular, is it possible to show that, for each n , the $\binom{n}{3}$ triangle areas of an optimal configuration take only $O(n)$ (or even $O(1)$) distinct values?

A positive answer would reveal strong hidden symmetries in the optimal configurations and could considerably simplify the algebraic systems used to derive exact coordinates.

7 Outlook and final remarks

We have presented an *optimize-then-refine* framework for the Heilbronn triangle problem on the unit square that combines a mixed-integer optimization model with exact symbolic computation. A symmetry-breaking strategy and a product-form determinant reformulation yield a model that Gurobi solves to certified global optimality for all $n \leq 9$, and the subsequent symbolic step recovers exact algebraic coordinates from the numerical certificate.

Several directions remain open. Extending the global certification to $n = 10$ appears within reach given continued improvements in solver technology and hardware, but will require substantial computational effort. On the theoretical side, establishing nontrivial bounds on the number of critical triangles (Section 6.1) or explaining the clustering of noncritical triangle areas (Section 6.2) would deepen our understanding of optimal Heilbronn configurations.

More broadly, the two-step methodology illustrated here—global certification via mixed-integer programming followed by exact coordinate recovery through computer algebra—is applicable to other packing and covering problems in discrete geometry where the conjectured optimal configurations involve algebraic numbers of moderate degree.

References

- [1] P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan. “Mixed-integer nonlinear optimization”. In: *Acta Numerica* 22 (2013), pp. 1–131.
- [2] T. Berthold, D. Kamp, G. Mexi, S. Pokutta, and I. Pólik. *Global Optimization for Combinatorial Geometry Problems Revisited in the Era of LLMs*. 2026.
- [3] C. Bertram–Kretzberg, T. Hofmeister, and H. Lefmann. “An Algorithm for Heilbronn’s Problem”. In: *SIAM Journal on Computing* 30.2 (2000), pp. 383–390.
- [4] R. E. Bixby. “A brief history of linear and mixed-integer programming computation”. In: *Optimization Stories*. European Mathematical Society-EMS-Publishing House GmbH, 2012, pp. 107–121.
- [5] S. Burer and A. N. Letchford. “Non-convex mixed-integer nonlinear programming: a survey”. In: *Surveys in Operations Research and Management Science* 17.2 (2012), pp. 97–106.
- [6] L. Chen, Y. Xu, and Z. Zeng. “Searching approximate global optimal Heilbronn configurations of nine points in the unit square via GPGPU computing”. In: *Journal of Global Optimization* 68.1 (2017), pp. 147–167.
- [7] A. Cohen, C. Pohoata, and D. Zakharov. “Lower bounds for incidences”. In: *Inventiones Mathematicae* 240 (2025), pp. 713–752.
- [8] F. Comellas and J. L. A. Yebra. “New lower bounds for Heilbronn numbers”. In: *The Electronic Journal of Combinatorics* 9.1 (2002), R6.
- [9] M. Conforti, G. Cornuéjols, and G. Zambelli. *Integer Programming*. Vol. 271. Graduate Texts in Mathematics. Springer, 2014.
- [10] L. Dehbi and Z. Zeng. “Heilbronn’s problem of eight points in the square”. In: *Journal of Systems Science and Complexity* 35.6 (2022), pp. 2452–2480.
- [11] A. W. M. Dress, L. Yang, and Z. Zeng. “Heilbronn Problem for Six Points in a Planar Convex Body”. In: *Minimax and Applications*. Ed. by D.-Z. Du and P. M. Pardalos. Boston, MA: Springer US, 1995, pp. 173–190.
- [12] E. Friedman. *Packing Center*. <https://erich-friedman.github.io/packing/>. Accessed: 2026-02-20.

- [13] M. Goldberg. “Maximizing the Smallest Triangle Made by N Points in a Square”. In: *Mathematics Magazine* 45.3 (1972), pp. 135–144.
- [14] J. Komlós, J. Pintz, and E. Szemerédi. “On Heilbronn’s triangle problem”. In: *Journal of the London Mathematical Society*. 2nd ser. 24.3 (1981), pp. 385–396.
- [15] J. Komlós, J. Pintz, and E. Szemerédi. “A lower bound for Heilbronn’s problem”. In: *Journal of the London Mathematical Society*. 2nd ser. 25.1 (1982), pp. 13–24.
- [16] G. P. McCormick. “Computability of Global Solutions to Factorable Nonconvex Programs: Part I — Convex Underestimating Problems”. In: *Mathematical Programming* 10.1 (1976), pp. 147–175.
- [17] A. Monji, A. Modir, and B. Kocuk. *Solving the Heilbronn Triangle Problem using Global Optimization Methods*. 2025.
- [18] K. F. Roth. “On a Problem of Heilbronn”. In: *Journal of the London Mathematical Society* s1-26.3 (1951), pp. 198–204.
- [19] K. F. Roth. “On a problem of Heilbronn, II”. In: *Proceedings of the London Mathematical Society*. 3rd ser. 25.2 (1972), pp. 193–212.
- [20] K. F. Roth. “On a problem of Heilbronn, III”. In: *Proceedings of the London Mathematical Society*. 3rd ser. 25.3 (1972), pp. 543–549.
- [21] W. M. Schmidt. “On a problem of Heilbronn”. In: *Journal of the London Mathematical Society*. 2nd ser. 4.3 (1972), pp. 545–550.
- [22] L. A. Wolsey. *Integer Programming*. 2nd. John Wiley & Sons, 2021.
- [23] L. Yang, J. Zhang, and Z. Zeng. *Heilbronn problem for five points*. Tech. rep. 1991.
- [24] Z. Zeng and L. Chen. “On the Heilbronn Optimal Configuration of Seven Points in the Square”. In: *Automated Deduction in Geometry*. Ed. by T. Sturm and C. Zengler. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 196–224.

A Best-known configurations for $10 \leq n \leq 16$

The best-known configurations for $10 \leq n \leq 12$ are due to Comellas and Yebra [8] (for $n = 10$ and $n = 12$) and Goldberg [13] (for $n = 11$). We reproduce their exact coordinates below.

The configurations for $13 \leq n \leq 16$ listed below were communicated to us by Erich Friedman (personal communication), to whom we are grateful. The configurations for $n = 13$ and $n = 15$ were found by Peter Karpov, and those for $n = 14$ and $n = 16$ by Mark Beyleveld. None of these have been published. All configurations in this section are obtained by heuristic methods and are not certified to be globally optimal.

For $n = 10$, let $z = \frac{3}{4} - \frac{(63+8\sqrt{62})^{1/3}}{12} - \frac{1}{12(63+8\sqrt{62})^{1/3}} \approx 0.3156$, $x = z/2$, and $y = 1 - 3z + 2z^2$.

For $n = 12$, let $x = 1 - \frac{(27+3\sqrt{57})^{2/3}+6}{6(27+3\sqrt{57})^{1/3}} \approx 0.1154$ and $y = 2x^2 - 3x + \frac{1}{2}$.

Table 11: Best-known configurations for $n = 10, \dots, 16$. The configurations for $n = 10, 12$ are due to Comellas and Yebra [8], $n = 11$ to Goldberg [13], $n = 13, 15$ to Peter Karpov (unpublished), and $n = 14, 16$ to Mark Beyleveld (unpublished).

$n = 10$			$n = 11$			$n = 12$		
#	x	y	#	x	y	#	x	y
0	x	0	0	1/3	0	0	x	0
1	$1-y$	0	1	2/3	0	1	$1-x$	0
2	0	x	2	0	2/9	2	0	x
3	1	y	3	1	2/9	3	1	x
4	$1-z$	z	4	1/3	4/9	4	1/2	y
5	z	$1-z$	5	2/3	4/9	5	y	1/2
6	0	$1-y$	6	0	2/3	6	$1-y$	1/2
7	1	$1-x$	7	1	2/3	7	1/2	$1-y$
8	y	1	8	1/2	7/9	8	0	$1-x$
9	$1-x$	1	9	1/6	1	9	1	$1-x$
			10	5/6	1	10	x	1
						11	$1-x$	1

$\Delta_{10} \geq \frac{5}{8}z^2 - \frac{1}{2}z^3 \approx 0.04654$
 $\Delta_{11} \geq 1/27 \approx 0.03704$
 $\Delta_{12} \geq \frac{1}{4}x + \frac{1}{2}xy - \frac{1}{2}x^2 \approx 0.03260$

Table 12: Best-known configurations for $n = 13, 14$ (continued).

$n = 13$			$n = 14$		
#	x	y	#	x	y
0	0.964815	0.087630	0	0.077620	0
1	0	1	1	0.922380	1
2	0.896939	0.902546	2	0.922380	0
3	0.761346	0.441996	3	0.077620	1
4	0.655161	1	4	0	0.186886
5	0.748551	0	5	1	0.813114
6	0	0.099250	6	1	0.186886
7	1	0.461332	7	0	0.813114
8	0.328490	0.633357	8	0.292333	0.321345
9	0.087939	0.614507	9	0.707667	0.678655
10	0.345014	0.901507	10	0.707667	0.321345
11	0.087938	0	11	0.292333	0.678655
12	0.500181	0.149235	12	0.5	0.138278
			13	0.5	0.861722

$\Delta_{13} \geq 0.02702$
 $\Delta_{14} \geq 0.02430$

Table 13: Best-known configurations for $n = 15, 16$ (continued). All coordinates for $n = 16$ are rational.

$n = 15$			$n = 16$		
#	x	y	#	x	y
0	0.934094	1	0	2/31	0
1	0.287119	0.302829	1	29/31	1
2	0.342286	0.701349	2	23/31	0
3	0.963064	0.095730	3	8/31	1
4	0.066630	0.633568	4	0	10/33
5	0.648909	0	5	1	23/33
6	0.277707	1	6	1	2/33
7	0.066641	0	7	0	31/33
8	0.589972	0.272487	8	8/31	4/11
9	0.603055	0.928222	9	23/31	7/11
10	0.895664	0.684290	10	10/31	2/33
11	0	0.192215	11	21/31	31/33
12	0.670814	0.614942	12	21/31	10/33
13	0	0.924975	13	10/31	23/33
14	1	0.399875	14	29/31	4/11
			15	2/31	7/11
$\Delta_{15} \geq 0.02111$			$\Delta_{16} \geq 7/341 \approx 0.02053$		