Lighting the Square Better: Twenty-Four Certified Improvements to the Maximin Light-Intensity Problem

Mishari Al-Mishari*

11 September 2026

Abstract

Place n lights of unit brightness in the closed unit square so that the darkest point of the square is as bright as possible, where brightness is the sum of inverse squared distances to the lights. This is the finite-n Riesz polarization problem at the critical exponent $s = d = 2$, and for $n \leq 36$ it is an open record problem on Erich Friedman's Packing Center. We report configurations for 24 of the 35 open counts, $n \in \{5, 10\} \cup \{15, \ldots, 36\}$, whose minimum intensity is *certified*, by outward-rounded interval branch-and-bound over the entire closed square, to exceed the currently listed value by between 0.16% and 6.82%; the largest absolute gain is $+18.14$ at $n = 36$. Every certificate has width at most 10^{-6} and was recomputed from the exported coordinates alone in a fresh process; for $n = 5$ an exact rational certificate pins the value to an interval of width below 2.7×10^{-24}.

Two ideas drive the result. First, a GPU population method that maintains an adaptive set of "dark points" per layout, ascends a temperature-annealed soft minimum, and rescores incumbents and challengers on the union of their witness sets; it discovered 14 of the 24 configurations. Second, a structural observation: every layout produced by conventional exchange, SLSQP, or GPU search in our incumbent bank had a *single* active dark point, which by a Danskin-type argument means it is not even first-order stationary for the maximin problem. A radius-coupled linear-programming ascent that models exactly the dark points able to bind within the trust region exploits this systematically: in 35 matched trials it produced 23 strict certified gains and no losses, while the deployed baseline produced none and returned its input unchanged. We give the supporting lemmas, the certification mathematics, the negative results, and full-precision coordinates for all 576 lights. Run unchanged on three other families of the same page (max/min distance ratios in the plane and in space, disks covering circles, squares and triangles, and circles packed in seven containers, 208 cases in all), the pipeline recovers the listed state of the art from weak starts in minutes per case, and improves three listed values (44, 45 and 46 circles in an L-shaped container, by up to 0.03%, verified in exact arithmetic); this locates the light gains in the problem's relative neglect rather than in any special feature of the kernel. No optimality claim is made.

1 Introduction

Erich Friedman's Packing Center lists, for each $n \leq 36$, the best known arrangement of n unit lights in the unit square that maximizes the minimum light intensity received anywhere in the square, the intensity at a point being the sum over lights of the reciprocal squared distance [10]. The page records $I = 2$ for $n = 1$ (trivially optimal), the exact value $I = 24/5$ for $n = 2$, and for $3 \leq n \leq 36$ three-decimal values with a trailing "+" indicating a truncated numerical minimum. The listed values as of September 2026 are due to Friedman (July and August 2026),

*mishari@al-mishari.com.sa. The search software, certificates, and coordinates described here are archived in the project repository; see Section 8.

to Kishimoto working with GPT-5.6 Sol Pro (July 2026), to Lai working with Codex (August 2026), and to Moose (August 2026). The problem is a finite instance of the constrained Riesz *polarization* (or Chebyshev) problem [4, 9]: for a compact set A and exponent $s > 0$, maximize $\min_{q\in A} \sum_i \|q - p_i\|^{-s}$ over n-point multisets in A. Here $A = [0, 1]^2$ and $s = 2$ equals the dimension, the critical case in which the optimal value grows like $n \log n$ [3].

The finite problem is difficult for three compounding reasons. The objective is a maximin over a continuum: each evaluation of $V(P) = \min_q I(q; P)$ is itself a global minimization over the square. The objective is nonsmooth exactly where it matters: at a good layout several local minima of $I(\cdot; P)$, the *dark points*, are tied, and V is the lower envelope of the corresponding smooth branches. And the kernel is singular at the lights while the closed square's edges and corners are genuine candidates for the darkest point, so boundary strata must be handled explicitly by both search and proof.

Contributions.

1. **Twenty-four certified improvements** (Theorem 5.1, Table 2). For $n = 5$, $n = 10$ and every $15 \le n \le 36$ we exhibit configurations whose minimum intensity is bounded below, by a whole-square interval certificate, by a value that exceeds the listed three-decimal figure by at least one unit in its last place, and typically by thousands of last places. Full-precision coordinates are given in Appendix A.

2. **A first-order diagnosis of why prior layouts were improvable** (Section 2). Danskin's theorem gives the directional derivative of V as the minimum over active dark points of the branch gradients, so a local maximizer with all lights interior must have the zero vector in the convex hull of the active gradients; a single active dark point is therefore never stationary (Proposition 2.1). Every layout in our own incumbent bank, produced by exchange, SLSQP or GPU search, had exactly one active dark point.

3. **Two search mechanisms** (Section 3): an adaptive dark-point exchange method run as a GPU population with a cyclically annealed soft minimum and common-witness archive rescoring, and a radius-coupled linear-programming maximin ascent that models precisely the branches able to bind within the trust region (Lemma 3.1) and is gated on the independent verifier.

4. **Certification** (Section 4): elementary but complete lower-bound arguments for the inverse-square kernel on boxes, including a global reciprocal-tangent bound valid in boxes containing a light (Lemma 4.1), a curvature bound using only the negative Hessian eigenvalue, monotone face reduction, and the completeness invariant of the branch-and-bound. For $n = 5$ we give an exact rational certificate of 685 boxes, independently replayed by a second implementation.

5. **Negative results and controls** (Section 5): matched trials against the deployed baseline, an audit of the failure mode of sampled floors, and several mechanisms that did not help.

6. **A transfer test** (Section 6): the identical pipeline applied to the distance-ratio, disk-covering and circle-packing problems of the same page, with per-case certified results in Appendix B.

What is not claimed. No configuration here is claimed to be optimal, and no first-order or second-order optimality certificate is offered. The comparison with the listed values is a comparison with displayed decimals whose hidden digits are unknown; Section 5.1 explains the selection rule that makes the claims robust to that uncertainty. All search scores are heuristic; only the interval certificates are bounds.

2 The problem and its first-order structure

2.1 Definitions

Let $Q = [0,1]^2$ and let $P = (p_1, \ldots, p_n) \in Q^n$ be a layout (repetitions allowed). The intensity at $q \in Q$ and the value of the layout are

$$I(q;P) = \sum_{i=1}^{n} \frac{1}{\|q - p_i\|^2}, \qquad V(P) = \min_{q \in Q} I(q;P), \qquad V_n^* = \max_{P \in Q^n} V(P). \tag{1}$$

At a light, $I = +\infty$. Since $I(\cdot;P)$ is lower semicontinuous on the compact set Q, finite somewhere, and tends to $+\infty$ at each light, the minimum is attained at a point away from the lights. We call every local minimizer of $I(\cdot;P)$ over Q a *dark point*; the set of global minimizers is the *active set* $A(P)$. Dark points live in three strata: the open square, the open edges, and the four corners.

Three facts orient the search. For $n = 1$ the optimum is $V_1^* = 2$: some corner has squared distance at least $\frac{1}{2}$ from any point, and the centre attains 2. Placing a light outside Q never helps for a convex target and a radially decreasing kernel, since projecting onto Q shortens every distance [14, Prop. 1.7]. And $V_n^*/(n \log n) \to \pi$ as $n \to \infty$, with optimal configurations asymptotically uniformly distributed [3, Thm. 3.1]; the theorem applies because Q is a compact subset of $\mathbb{R}^2$ of area one, but it gives no finite-n error bound, no preferred lattice, and no licence to ignore the boundary.

2.2 Directional derivatives and the single-branch obstruction

For $q \neq p_i$ the gradient of one branch with respect to light i is

$$\nabla_{p_i} I(q;P) = \frac{2\,(q - p_i)}{\|q - p_i\|^4}, \tag{2}$$

which is never zero. Write $G(q;P) \in \mathbb{R}^{2n}$ for the stacked vector of these blocks.

Proposition 2.1 (first-order conditions)**.** *Let $P \in Q^n$ have a finite active set $A(P) = \{q_1, \ldots, q_k\}$. Put $G_j = G(q_j;P)$ for $j = 1, \ldots, k$. Then, for every direction $D \in \mathbb{R}^{2n}$, the one-sided directional derivative of V at P exists and equals*

$$V'(P;D) = \min_{1 \le j \le k} \langle G_j, D \rangle. \tag{3}$$

If P is a local maximizer of V and every light lies in the open square, then $0 \in \operatorname{conv}\{G_1, \ldots, G_k\}$. In particular $k \ge 2$: a layout with a single active dark point is not first-order stationary.

Proof. Choose $\varepsilon > 0$ with $4/(9\varepsilon^2) > V(P) + 1$. If P' is within $\varepsilon/2$ of P in every light and $\|q - p_i\| < \varepsilon$ for some i, then $\|q - p_i'\| < 3\varepsilon/2$ and the single term for light i already gives $I(q;P') > V(P) + 1$; on the other hand $V(P') \le I(q_1;P') \to V(P)$ as $P' \to P$. Hence for P' near P every minimizer of $I(\cdot;P')$ lies in the fixed compact set $Q_\varepsilon = Q \setminus \bigcup_i B(p_i, \varepsilon)$, on which I is jointly continuous and continuously differentiable in P', and $V(P') = \min_{q \in Q_\varepsilon} I(q;P')$. Danskin's theorem [7] then gives (3). If P is a local maximizer, $V'(P;D) \le 0$ for all D (all directions are feasible when every light is interior), so no D satisfies $\langle G_j, D \rangle > 0$ for all j; by Gordan's alternative $0 \in \operatorname{conv}\{G_j\}$. Since each block of G_j is nonzero by (2), $G_j \neq 0$, which rules out $k = 1$. □

Corollary 2.2 (explicit ascent)**.** *If $A(P) = \{q^*\}$ and some light p_i lies in the open square, then D with block i equal to $2(q^* - p_i)/\|q^* - p_i\|^4$ and all other blocks zero satisfies $V'(P;D) = \|D\|^2 > 0$. Lights on the boundary may be included after projecting their blocks onto the tangent cone of Q.*

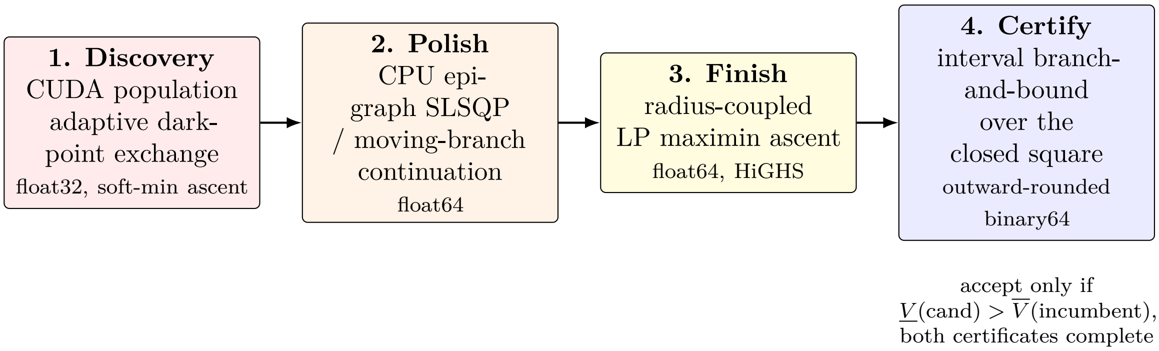

Figure 1: The pipeline. Stages 1–3 produce heuristic search scores; only Stage 4 produces bounds, and only strict separation of complete certificates replaces an incumbent.

The observation that made the difference. We froze the 35 incumbent layouts of our own campaign for $n = 2, \dots, 36$, each already certified and most already above the listed values, and enumerated their dark points by two independent dense scans (a 2001×2001 grid screen with local polish, and a separate multistart L-BFGS-B scan). In every case the global minimum was attained at a *single* dark point. The gap to the runner-up local minimum was 2.1×10^{-4} at $n = 13$, 2.5×10^{-4} at $n = 17$, 2.5×10^{-3} at $n = 21$ and 4.5×10^{-2} at $n = 34$, i.e. between 2×10^{-6} and 1.3×10^{-4} relative to the floor. By Proposition 2.1 none of these layouts is stationary, and indeed a first-order nudge of size 10^{-5} to 10^{-4} raised the true minimum of 20 of the 35 incumbents by 8.7×10^{-6} to 5.5×10^{-4} relative, one to three orders of magnitude above the verifier's tolerance. Figure 5 shows the $n = 34$ case before and after the ascent of Section 3.3: one active dark point becomes ten dark points tied to within 3.5×10^{-5}.

The reason such layouts survive in a certificate-driven pipeline is structural rather than a defect of any one solver. Layouts are ranked by certified lower bounds computed under sub-second budgets and replace the incumbent only when they win strictly with a complete certificate; nothing in that loop drives a layout to stationarity, and a polisher gated on "already stationary" never runs on an input that is not.

How many dark points are there? An independent structural pass over the same 35 layouts (basin-local polishing, Voronoi and Delaunay seeds, and a 3000-point random multistart, with the dark-point sets reproduced exactly across two grid resolutions and seeds) found that the total number of local minima of $I(\cdot; P)$ grows roughly like $1.5n$ (29 at $n = 20$, 53 at $n = 36$) and is strictly below $2n+1$ for every $n \geq 3$. The inverse-square potential is much smoother than the Voronoi structure of the lights and simply does not have $2n+1$ basins. Any argument resting on a count of active constraints is therefore unsafe here; worse, as lights move, dark points are *born*, so a direction that raises every currently active branch can still lower the minimum through a branch that did not exist at the base point. This is the failure mode behind the strict losses described in Section 5.4.

3 Search

The pipeline has three stages on different hardware and at different precisions, followed by an acceptance rule (Figure 1). Stage 1 explores basins on the GPU in float32. Stage 2 sharpens promising layouts on the CPU in float64. Stage 3 finishes them to a first-order stationary point of the sampled maximin model. Stage 4 decides, using only interval arithmetic, what is real. No number produced by Stages 1–3 ever enters a certificate; float32 error can cost search quality, never a false bound.

3.1 Stage 1: adaptive dark-point exchange as a GPU population

The method is a semi-infinite exchange scheme in the Blankenship–Falk tradition [2, 15, 26]: maintain an explicit finite set of dark points per layout, optimize the lights against that set, then find the dark points again. Run as a population of B layouts evolving in parallel with PyTorch [24], each with its own Adam state [18], its own dark-point set and its own learning rate; B is chosen from free video memory and halved on allocation failure. Algorithm 1 summarizes one run; the components are as follows.

Inner oracle. Candidate dark points are seeded from all three strata: strict 3×3 local minima of I on a 21×21 interior grid (the 32 lowest), one-dimensional local minima along each of the four edges (the 6 lowest per edge), all four corners, and the 24 lowest dark points carried over from the previous exchange. They are refined by a batched damped Newton descent in q, one 2×2 solve per point with the smaller Hessian eigenvalue shifted positive, step capped at 0.12 in the infinity norm, and a monotone backtracking line search over seven halvings. An edge seed has its normal coordinate frozen and a corner has no free coordinate, so a dark point stays in the stratum it was seeded in and constrained boundary minima, which unconstrained descent would never find, are represented. A refinement that fails keeps its last actual witness, never a prediction. The oracle's score is the minimum over both the grid and the refined points.

Outer step: an annealed soft minimum. The exact minimum has zero gradient with respect to every light except those nearest the current worst point, so exact min-ascent moves one light at a time. We ascend instead the log-sum-exp soft minimum

$$\mathrm{softmin}_T(v) = -T\log\sum_j e^{-v_j/T}, \qquad \min_j v_j - T\log M \;\le\; \mathrm{softmin}_T(v) \;\le\; \min_j v_j,$$

with the stable shifted evaluation of [1], so that every near-active dark point contributes gradient in proportion to how close it is to binding. The temperature is annealed geometrically from 0.09 to 0.004 over 160 steps and then *cyclically restarted*, and the population is split into eight phase groups so that at any instant it spans the whole temperature range, some layouts exploring on a smooth surrogate while others sharpen on an almost exact minimum. Per-row learning rates are drawn from $\{0.003, 0.005, 0.008, 0.012\}$, gradients are clipped to norm 20, the step shrinks with the temperature, and positions are clamped into Q after every step. Grid points and refined dark points frequently coincide; without care a duplicated point contributes twice to the soft minimum. Witnesses are therefore quantized into 10^{-6} bins and one live minimum per bin is kept by a scatter-minimum, taking the current minimum at every step rather than freezing a representative.

Exchange and archive. Every 24 optimizer steps the oracle is rerun and the archive updated. The single most important correctness idea is that *the incumbent and the challenger are rescored on the union of both witness sets before comparison*:

$$s \leftarrow \min\big(s,\ I(P,\ H\cup H_{\mathrm{best}})\big), \qquad s_{\mathrm{best}} \leftarrow \min\big(I(P_{\mathrm{best}},\mathrm{grid}),\ I(P_{\mathrm{best}}, H\cup H_{\mathrm{best}})\big).$$

An optimistic score measured against an older, smaller witness set would otherwise keep a poor elite in the archive forever; any exchange method that scores candidates on their own witnesses drifts in exactly this way. The archive is written only after a fresh inner minimization.

Islands and export. Every fifth refresh the batch is split into four islands, the worse half of each replaced by perturbed copies of its better quarter (mutation scale 0.10 for one row in three, 0.02 otherwise, every fourth replacement resampled uniformly), with fresh Adam moments and an oracle rerun without inherited history. The top 96 layouts are moved to the CPU,

Algorithm 1 Adaptive dark-point exchange, GPU population (one call)

Require: case n, seed, budget S, warm layouts

reserve $\min(24, \max(0.15, 0.4S))$ seconds for Stages 2–4; choose batch B from free VRAM
initialize B layouts from incumbent and neighbouring layouts, perturbations, regular grids, symmetric seeds, random points; assign phase offsets and learning rates
$H \leftarrow \textsc{Oracle}(P, \text{history} = \varnothing)$
while time remains **do**
  for 24 steps **do**
    $T \leftarrow 0.09\,(0.004/0.09)^{\text{phase}}$; $v \leftarrow$ deduplicated $I(P, \text{grid} \cup H)$
    Adam step on $\nabla_P\, \text{softmin}_T(v)$, clipped, scaled by $0.25 + 0.75(1 - \text{phase})$; clamp P into Q
  end for
  $H \leftarrow \textsc{Oracle}(P, \text{history} = H)$ ▷ grid + edges + corners + 24 carried points, damped Newton
  rescore P and archive on $H \cup H_{\text{best}}$; replace archive rows that are strictly beaten
  if every fifth refresh **then** island restart with fresh moments; $H \leftarrow \textsc{Oracle}(P, \varnothing)$
  end if
end while
return diverse shortlist of the top 96 archive layouts in float64

cast to float64, and reduced to a diverse shortlist by greedily keeping layouts at least 0.003 apart in the minimum root-mean-square assignment distance over all eight symmetries of the square (Hungarian algorithm [20]). Certification is expensive; spending it on eight copies of one configuration is what the shortlist avoids. Forty percent of the wall clock, capped at 24 seconds, is reserved before the GPU starts for the CPU stages, so that a large asynchronous population cannot squeeze out the proof.

3.2 Stage 2: CPU polish

A float32 layout is typically 10^{-6} or worse from stationarity, enough to lose a certified comparison. The polish converts the maximin into its epigraph form,

$$\max_{P,t}\ t \quad \text{s.t.} \quad I(w_j; P) \geq t \ \ (j = 1, \ldots, m), \qquad P \in Q^n, \tag{4}$$

a smooth nonlinear program in $2n + 1$ variables solved by SLSQP [19, 28] with the analytic constraint Jacobian (2), in up to six exchange rounds: solve, rediscover dark points, add them to the witness set (capped at 160, keeping the darkest), re-solve. The best layout *observed* is returned, tracked by explicit re-evaluation, because the optimizer's last iterate is not always its best. The original CPU search of the project is the same exchange run from regular, symmetric, random or transported starts, with grid screening, explicit edge minimization and corner checks; a variant tracks each dark-point branch with projected Newton steps between outer trust-region moves ("moving-branch continuation") and was the discoverer of several small-n layouts.

3.3 Stage 3: radius-coupled LP maximin ascent

Let $q_1, \ldots, q_m$ be the currently known dark points of P with values $v_j = I(q_j; P)$, floor $v_{\min} = \min_j v_j$, and branch gradients $g_j = G(q_j; P) \in \mathbb{R}^{2n}$; on a stationary interior or edge branch the total derivative of v_j with respect to P equals g_j by the envelope theorem, since the first-order motion of q_j contributes nothing. A step Δp with $\|\Delta p\|_\infty \leq r$ changes v_j by $g_j^\top \Delta p + O(r^2)$.

Lemma 3.1 (radius-coupled band)**.** *Let $\|\Delta p\|_\infty \leq r$. To first order, branch j can become the floor after the step only if $v_j - v_{\min} \leq r\,(\|g_j\|_1 + \|g_{\min}\|_1) \leq 2r \max_k \|g_k\|_1$, where $g_{\min}$ is the gradient of a floor branch.*

Proof. $\left|g^\top \Delta p\right| \le \|g\|_1 \|\Delta p\|_\infty$ by Hölder's inequality; the floor can rise by at most $r\,\|g_{\min}\|_1$ and branch j can fall by at most $r\,\|g_j\|_1$. □

The lemma says that the trust radius determines exactly which branches must be modelled: nothing outside the band can overtake the floor inside the region, and everything inside it might. This replaces the fixed value threshold used by earlier continuation code, which is wrong in both directions: it admits branches that cannot bind within the step and omits branches that will. The implementation uses the one-sided band $v_j - v_{\min} \le r \max_k \|g_k\|_1$, capped at 160 rows; because every accepted step is rescored by fresh rediscovery on a common witness pool (below), the band affects efficiency, not soundness. The step is then the exact optimizer of the linearized maximin model, a linear program solved with HiGHS [16]:

$$\max_{\Delta p,\,\delta} \ \delta \quad \text{s.t.} \quad v_j + g_j^\top \Delta p \ \ge \ v_{\min} + \delta - \mu_j \ \ (j \in \text{band}), \qquad -r \le \Delta p \le r, \qquad 0 \le P + \Delta p \le 1, \tag{5}$$

with $\mu_j = 0$ for ascent. Why a linear program and not a gradient step: at $n = 34$ the runner-up dark point is far and a raw gradient step along Corollary 2.2 gains $+4.3 \times 10^{-2}$; at $n = 13$ the runner-up binds almost immediately and the same gradient step *loses* (-4.5×10^{-5} at step 10^{-5}). The LP handles both because the runner-up is already a row of the program.

Algorithm 2 summarizes one call. Two implementation details mattered more than the step rule.

- **Common-witness comparison.** Two independent grid discoveries of the dark spectrum disagree at about 10^{-9}, far above the gains available near stationarity, so comparing freshly discovered spectra rejects good steps at random. Both the current and the trial layout are relaxed from one shared seed pool by a vectorized projected damped Newton method with bound freezing and backtracking; only the layout differs. The pool is refreshed from a fresh grid every eight rounds, and when it grows, the current layout is rescored downward too.

- **Champion guard.** Enlarging the pool can expose a branch that was always there, lowering the measured floor without the layout changing. Acceptance runs on the pool, but what the routine returns is judged by an independent full rediscovery, so it can never return a layout worse than the one it was given. Without the guard $n = 34$ regressed by -0.066%; with it the same run gained $+0.048\%$.

A collapsed trust region reopens up to three times after a fresh rediscovery, since exhaustion of the linear model on one pool is not evidence of optimality. Finally, and decisively after the losses reported in Section 5.4, the routine is *gated on the verifier*: a changed layout is exported only when its complete whole-square certificate strictly beats the warm start's; otherwise the warm start is returned unchanged.

Escaping the combinatorial cell (negative). Ascent converges to a Karush–Kuhn–Tucker point of one combinatorial cell, one active dark-point set. Because μ_j enters only the right-hand side of (5), granting a chosen subset B of active branches a budget $\mu > 0$ lets them fall while the rest rise; annealing μ back to zero re-equalizes on what is generally a different active set. We chose B by the multipliers of a nonnegative least-squares fit [21] on the active Jacobian, since $\partial(\min I)/\partial\mu_j = \lambda_j$ ranks branches by what relaxing each buys. The mechanism demonstrably reaches different cells (root-mean-square displacement 0.070 at $n = 17$ after canonicalizing under the square's symmetries, far outside the 10^{-3} identity band), but every cell reached was worse: the excursion dips the floor from 135.1 to 69.8 before annealing, and re-ascent recovers only to 133.7. A dip-controlled variant is designed but not measured.

Algorithm 2 Radius-coupled LP maximin ascent (one call)

Require: warm layout P, budget

pool $W \leftarrow$ structural seeds (corners, Delaunay circumcentres and edge midpoints, boundary projections) $\cup$ grid seeds; relax W to dark points of P
$r \leftarrow r_0$; champion $\leftarrow P$
while time remains and $r > 10^{-11}$ **do**
    $(v, g) \leftarrow$ values and branch gradients of P on W; band $\leftarrow \{j : v_j - v_{\min} \leq r \max_k \|g_k\|_1\}$
    solve LP (5) for Δp; $P' \leftarrow P + \Delta p$
    relax W for both P and P' from the shared pool; refresh pool from a fresh grid every 8 rounds
    if floor(P') > floor(P) on the common pool **then** $P \leftarrow P'$; grow r
    else shrink r (reopen ≤ 3 times after rediscovery)
    end if
end while
judge champion vs. final P by independent full rediscovery; keep the better
return champion only if its complete certificate strictly beats the warm start's; else the warm start

3.4 Acceptance

A candidate replaces an incumbent for the same n only if both have complete certificates and the candidate's certified lower bound exceeds the incumbent's certified upper bound. Overlapping intervals, incomplete certificates, missed deadlines, and unchanged geometry are not gains. All promotions are logged with the input, output, verifier and source hashes.

4 Certification

The verifier is independent of both optimizers. It reads exported coordinates, treats each binary64 value as the exact rational it denotes (checking a supplied hexadecimal encoding bit for bit, including the sign of zero), checks integer multiplicities, and encloses $V(P)$ over the entire closed square by branch-and-bound with outward rounding [12, 23]. It does not optimize positions, does not compare with published records, and does not prove optimality among layouts. Supplied objectives, solver success flags and claimed dark points are never trusted; a supplied dark point is useful only after it is checked to lie in Q and its intensity is recomputed.

4.1 Upper bounds

Every nonsingular $q \in Q$ gives $V \leq I(q; P)$. At rational coordinates $I(q; P)$ is itself rational and is evaluated exactly with integer arithmetic; the result is converted to an enclosing binary64 interval by comparing the converted float with the rational and moving the offending endpoint outward. A fast outward-rounded evaluator screens trial witnesses; only accepted improvements receive exact evaluation. Witnesses are proposed by a dyadic grid, by the solver's own dark points, and by projected Newton refinement; none of these can supply a lower bound.

4.2 Lower bounds on a box

Let $B = [x_0, x_1] \times [y_0, y_1] \subseteq Q$ and let $m_i \geq 1$ be multiplicities. A cell bound L_B must satisfy $L_B \leq \min_{q \in B} I(q; P)$ also when B contains a light.

Distance bound. With $M_i = \max_{q \in B} \|q - p_i\|^2$, attained at a corner of B, $L_B^{\text{dist}} = \sum_i m_i / M_i$ is valid, including for boxes containing or touching a light. It is weak because different lights attain their largest distance at different corners.

Lemma 4.1 (global reciprocal tangent)**.** *For $a > 0$ and $t > 0$, $1/t \geq 2/a - t/a^2$, with equality iff $t = a$; the inequality extends to $t = 0$ with the left side $+\infty$. Consequently, for any positive $a_1, \dots, a_n$,*

$$I(q;P) \;\geq\; T(q) := \sum_i m_i\Big(\frac{2}{a_i} - \frac{\|q-p_i\|^2}{a_i^2}\Big) \qquad \textit{for all } q \in Q,$$

T is a concave quadratic, and $L_B^{\tan} = \min_{v\in\mathrm{corners}(B)} T(v)$ is a valid lower bound on B. With $a_i = M_i$ it dominates the distance bound.

Proof. $1/t - 2/a + t/a^2 = (t-a)^2/(ta^2) \geq 0$. Summing with $t = \|q - p_i\|^2$ gives the minorant; its Hessian is $-2\sum_i m_i a_i^{-2}\,\mathrm{Id}$, so it is concave and its minimum over a box is at a corner. If $a_i = M_i$ then at any corner $\|v - p_i\|^2 \leq a_i$, so each term is at least $2/a_i - 1/a_i = 1/a_i$. □

The choice $a_i \approx \|c - p_i\|^2$ at a represented box centre c gives a second, often tighter support; with that choice $T(q) = I(c) + \nabla I(c)^\top(q-c) - K_c\,\|q-c\|^2$ with $K_c = \sum_i m_i\,\|c-p_i\|^{-4}$, which looks like a Taylor bound but is *global*, requiring no source-free segment and no Hessian regularity. The implementation retains $\max(L^{\mathrm{dist}}, L^{\tan})$ and falls back safely if a support cannot be evaluated; an arbitrarily small denominator is never replaced by an epsilon.

Curvature bound on source-free boxes. For $u = \|d\|^2$, $\nabla(1/u) = -2d/u^2$ and $\nabla^2(1/u) = 8dd^\top/u^3 - 2\,\mathrm{Id}/u^2$, whose eigenvalues are $6/u^2$ (radial) and $-2/u^2$ (tangential). If $\|q-p_i\|^2 \geq \rho_i > 0$ throughout B, then $\nabla^2 I \succeq -2K\,\mathrm{Id}$ on B with $K = \sum_i m_i/\rho_i^2$, and integrating along the segment from c to q,

$$I(q;P) \;\geq\; I(c;P) + \nabla I(c;P)^\top(q-c) - K\,\|q-c\|^2 \qquad (q \in B), \tag{6}$$

a concave quadratic again minimized at a corner. Using only the negative eigenvalue divides the classical penalty (which uses the spectral radius $6\sum m_i/\rho_i^2$) by three. A development draft that combined a Hessian *evaluated at the centre* with a whole-cell curvature claim was caught by an exact counterexample and removed before deployment: four unit lights at $(\frac14,\frac12), (\frac34,\frac12), (\frac12,\frac14), (\frac12,\frac34)$ and the source-free cell $[\frac{5}{16},\frac{11}{16}]^2$ have centre value 64, zero gradient and centre Hessian $2048\,\mathrm{Id}$, yet $I(\frac{5}{16},\frac{5}{16}) = 8704/145 < 64$.

Monotone face reduction. If every light satisfies $p_{ix} \leq x_0$, then increasing x increases every distance and cannot increase I, so the minimum over B lies on the face $x = x_1$; analogously for the other three faces. This geometric rule needs no derivatives and remains valid when the face contains a light. Reducing a cell to a face preserves its minimum, not its footprint; the original box remains the certificate's coverage region.

4.3 Branch-and-bound and completeness

The verifier keeps a finite witness upper bound U and a priority queue of unresolved boxes with sound lower bounds, splits the weakest box into closed children covering it (a rounded midpoint is used only if it lies strictly between the represented endpoints; otherwise the cell is reported unsplittable), lets a child inherit $\max(\text{parent}, \text{child})$, and discards a box only when its bound is at least U. Writing L for the smallest bound among retained boxes ($L = U$ if all are pruned), the invariant is $L \leq V \leq U$; a crossed interval is an internal error, never clamped. Completeness is declared only when the upward-rounded gap $U - L$ is at most the requested tolerance τ; for all certificates in this paper $\tau = 10^{-6}$. A deadline or cell limit leaves a valid interval marked incomplete, which is distinct from an input error. Outward rounding is implemented by adjacent-float expansion after each correctly rounded binary64 operation [11, 17]. Finally, no

Table 1: The five-light layout with the exact rational certificate (7). Coordinates are exact decimals (denominators 10^{17}). Coordinate-text SHA-256: `98f895e0...2a537e2`.

i	x_i	y_i
1	0.14057504737591497	0.19676593548851007
2	0.14057504737591497	0.80323406451148993
3	0.75392205953575475	0.50000000000000000
4	0.75496730599486306	0.05022685702705627
5	0.75496730599486306	0.94977314297294368

rounded decimal is ever claimed merely because the interval is narrow: a value is printed with d decimals only if every point of the enclosure rounds to the same decimal.

One operational lesson is worth recording because it looks exactly like a bad layout: under equal tiny proof slots, $n = 33$ layouts worth about 333.63 were being rejected on lower bounds near 316.88, because the branch-and-bound had not converged, while a complete proof took about 2.6 seconds. Under-funding the verifier is indistinguishable from a poor candidate; per-candidate proof budgets now scale with n, and a coarse interval that straddles the incumbent receives a bounded refinement before a ranking decision.

4.4 An exact rational certificate for $n = 5$

For the five-light layout of Table 1 (the mirror image, under $x \mapsto 1 - x$, of the $n = 5$ layout of Appendix A, differing from it at the sixth decimal place) a separate prover using only Python integer and `Fraction` arithmetic establishes

$$22.2522327515259115251849 6 \;\le\; V(P_5) \;<\; 22.2522327515259115251849 62665, \tag{7}$$

an interval of width below 2.665×10^{-24}. The lower endpoint is proved over the whole square by 685 exact boxes using the distance bound, the reciprocal tangents with $a_i = M_i$, the curvature bound (6) with the negative eigenvalue, and a whole-box derivative-sign face reduction (with the full spectral constant 6, since a sign test needs a two-sided bound). The upper endpoint is the exact rational value at a point near the darkest point on the top edge; numerical root finding proposed the point, only fractions evaluated it. The certificate stores the split rule and the binary path of every accepted leaf, so an independent replayer that imports no prover code reconstructs the complete prefix-free partition of the square, recomputes every leaf bound exactly, and checks the coordinate hashes and target; it accepts (7) and rejects thirteen adversarial mutations (altered bindings, false bounds with recomputed hashes, missing branches, changed targets). The layout arrived with an earlier 762-box exact certificate produced by a different program; that certificate was reproduced exactly by the new prover and by the independent replayer, and a stronger 590-box certificate proves the same target with 22.6% fewer boxes at higher per-box cost. Figure 2 shows the partition. The value (7) is 0.0392 above the listed 22.213^{+}, but this proves the value of one fixed layout; it does not prove optimality among five-light layouts.

4.5 Independent implementations and re-verification

Two implementations of the whole-square verifier exist: the production scalar verifier used for ranking during search, and a dedicated implementation that shares no code with the solvers, uses the joint-corner tangent supports and the geometric face rule, and was audited against 415 exact-rational inequalities on adversarial cases (boundary lights, subnormal cells, adjacent-float endpoints, multiplicity $2^{53} + 1$). On a frozen snapshot of all 35 saved layouts the two agreed on every one of 40 paired proofs and the dedicated verifier certified all 35 to width below 10^{-9}. For this paper every bound in Table 2 was recomputed from the exported coordinates alone,

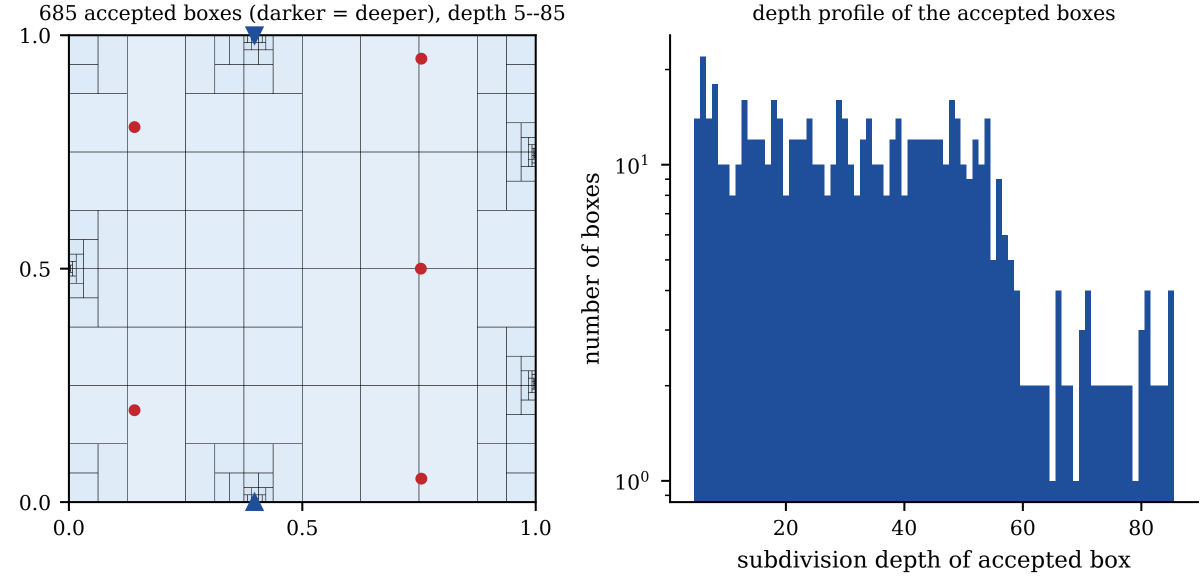

Figure 2: The exact rational certificate for $n = 5$. Left: the 685 accepted boxes covering the closed square, shaded by subdivision depth, with the five lights (red) and the two darkest points on the top and bottom edges (blue). Right: number of accepted boxes per depth; only a few boxes near the darkest points need depth beyond 60 to resolve a gap of 2.7×10^{-24}.

in a fresh process, with 120 seconds per case; no number is inherited from the search that found the configuration. Five of the layouts ($n = 5, 21, 23, 28, 36$) were re-certified once more by the dedicated verifier on 11 September 2026 while preparing this manuscript, with complete certificates consistent with Table 2 in every case (Appendix C).

5 Results

Theorem 5.1. *For each n in Table 2, the configuration P_n listed in Appendix A satisfies $V(P_n) \geq L_n$, where L_n is the certified lower bound in the table. Each L_n exceeds the value listed for n on Erich Friedman's page as of 11 September 2026 by at least one unit in the last displayed decimal place.*

Proof. Computer-assisted: the interval certificates of Section 4, recomputed from the coordinates of Appendix A and archived with the coordinates. The margin column of Table 2 is L_n minus the displayed value. □

5.1 Reading the table

The page prints most values as, for instance, 172.658^+, so the value behind the digits is not bounded above by them, and a merely larger number would prove nothing. The selection rule is chosen to survive that: under the natural reading that a three-decimal display is a truncation, the hidden value is below 172.659, so a certified lower bound of at least 172.659 beats it whatever the hidden digits are. Every case in Table 2 clears its display plus one full last place; the margins range from 31 last places at $n = 22$ to 16 793 at $n = 28$. Counts whose certified value is numerically above the display but within one last place ($n = 3, 4, 7, 8, 9, 11, 12, 13, 14$) are *not* claimed: there we reproduce the page, we do not beat it. Three consistency checks confirm that our objective is the page's: at $n = 1$ the analytic optimum 2 is reproduced; at $n = 2$ our best configuration certifies 4.799999608, matching the exact listed value 24/5 to seven digits without exceeding it; and the nine counts just named agree with the page to their last displayed digit.

Table 2: The twenty-four improved configurations. "Width" is the width of the recomputed whole-square certificate $[L_n, U_n]$; "Margin" is L_n minus the listed value, absolute and relative. Disc. names the stage that discovered the layout's topology (G: GPU adaptive dark-point exchange, Section 3.1; C: CPU exchange or moving-branch continuation, Section 3.2); Fin. marks layouts whose final geometry came from the LP maximin ascent (A, Section 3.3). Listed values and their attributions are from [10]: $n = 5$ Friedman (Aug. 2026); $n = 10, 15$–28 Kishimoto and GPT-5.6 Sol Pro (Jul. 2026); $n = 29$–35 Lai and Codex (Aug. 2026); $n = 36$ Moose (Aug. 2026).

n	Listed value (source page)	Certified lower bound on $F(P_n)$	Width of certificate	Margin abs.	Rel. %	Disc.	Fin.
5	22.213^+	22.252232129	6.1×10^{-7}	+0.0392	+0.177	C	–
10	62.408^+	64.751953830	9.7×10^{-7}	+2.3440	+3.756	C	–
15	113.191^+	115.159578307	9.9×10^{-7}	+1.9686	+1.739	G	A
16	124.553^+	125.458096378	9.7×10^{-7}	+0.9051	+0.727	C	A
17	133.907^+	135.166510549	9.9×10^{-7}	+1.2595	+0.941	G	A
18	141.556^+	147.341208526	9.9×10^{-7}	+5.7852	+4.087	G	A
19	152.854^+	159.385527986	9.9×10^{-7}	+6.5315	+4.273	G	A
20	165.272^+	174.655642160	9.5×10^{-7}	+9.3836	+5.678	G	A
21	172.658^+	184.439287742	9.9×10^{-7}	+11.7813	+6.823	G	A
22	195.050^+	195.360667440	9.9×10^{-7}	+0.3107	+0.159	G	A
23	205.756^+	206.538419485	9.5×10^{-7}	+0.7824	+0.380	C	A
24	215.573^+	219.197869079	9.4×10^{-7}	+3.6249	+1.682	C	A
25	225.100^+	229.700849442	1.0×10^{-6}	+4.6008	+2.044	C	A
26	234.759^+	242.004151137	9.9×10^{-7}	+7.2452	+3.086	C	A
27	242.055^+	254.776773379	1.0×10^{-6}	+12.7218	+5.256	G	A
28	252.447^+	269.239696424	9.6×10^{-7}	+16.7927	+6.652	G	A
29	282.856^+	285.128768312	9.5×10^{-7}	+2.2728	+0.804	G	A
30	285.345^+	296.005084140	1.0×10^{-6}	+10.6601	+3.736	C	A
31	305.298^+	307.962407037	1.0×10^{-6}	+2.6644	+0.873	G	A
32	317.203^+	320.371046396	8.5×10^{-7}	+3.1680	+0.999	G	A
33	330.595^+	334.037193150	1.0×10^{-6}	+3.4422	+1.041	G	A
34	337.906^+	346.509176724	9.3×10^{-7}	+8.6032	+2.546	C	A
35	347.195^+	359.463095378	1.0×10^{-6}	+12.2681	+3.533	C	A
36	354.550^+	372.686420098	4.1×10^{-12}	+18.1364	+5.115	G	A

One count is a known deficit: at $n = 6$ the page lists 29.791^+ and our best certified value is 29.790386, so the listed configuration is better than ours and we have not reproduced it.

Figure 3 shows six of the configurations with their intensity contours and tied dark points. The relative margins (Figure 4) are not uniform. They are smallest where the listed configuration is presumably already near a stationary point of its cell ($n = 5, 22, 23$) and largest for $n = 20, 21, 27, 28, 36$, where the improved layouts differ in topology from what a regular or lightly perturbed grid gives, several with doubled lights on the boundary. The ratio of the certified values to the asymptotic law $\pi n \ln n$ rises from about 0.88 at $n = 5$ to about 0.92 at $n = 36$; the theorem of [3] says nothing about the rate, but the trend is consistent with it.

5.2 Which stage did what

Of the 24 configurations, 14 ($n = 15, 17$–$22, 27$–$29, 31$–$33, 36$) were discovered by the GPU adaptive dark-point exchange of Section 3.1, more than by any other method; in a first paired benchmark at 20 seconds it scored 8 wins, 7 overlaps and 1 loss against the then-incumbents, roughly four times the win rate of the CPU baseline. The pattern of repeated, restarted inner adversaries alternating with outer updates is the one that robust training made standard [22]. The remaining 10 were discovered by the CPU exchange and continuation methods. The LP ascent of Section 3.3 then finished 22 of the 24 (all $n \geq 11$). It is a finisher, not a discoverer:

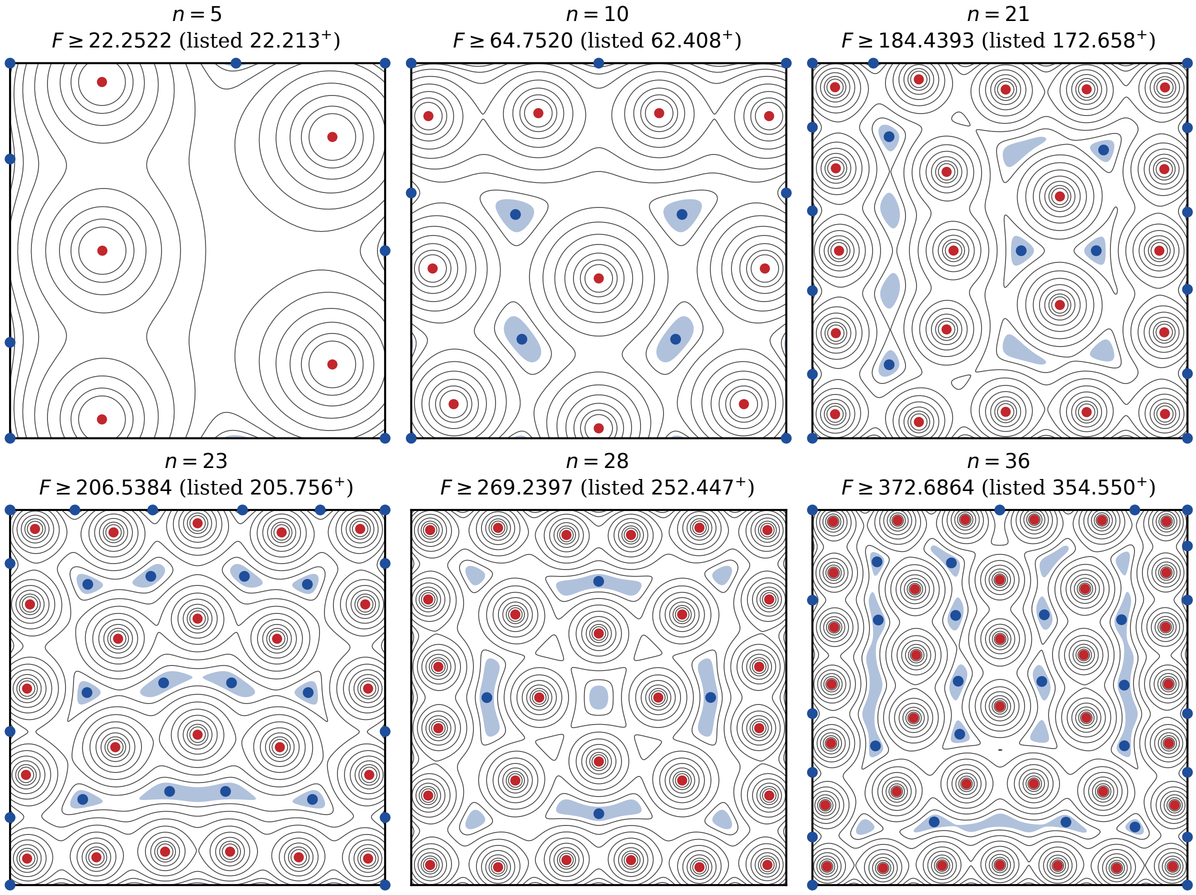

Figure 3: Six of the twenty-four configurations. Lights in red, intensity contours in black at fixed multiples of the certified floor, the region within 2% of the floor shaded, and the dark points tied with the minimum (to 10^{-4} relative) in blue. Note the doubled lights on the boundary at $n = 10$ and $n = 28$ and the tendency of the darkest points to line up in chains between rows of lights.

across the matched trials below it produced 27 distinct strict certified gains and *created zero new page-beating cases*; it widened margins that the discovery stages had already opened. Discovery and refinement are separate jobs and are best run in series.

5.3 Matched trials against the deployed baseline

Under matched conditions (same frozen input, same seed, a 60-second solve allowance, 30 seconds of independent verification per side, interleaved so that both arms met the same machine load) the LP ascent was compared with the deployed CPU portfolio solver on all 35 counts.

arm	trials	strict gain	overlap	strict loss	$> 0.1\%$	$> 1\%$
LP ascent v1 (ungated)	35	22	9	4	1	0
LP ascent v3 (verifier-gated)	35	23	12	0	2	0
deployed control	35	0	35	0	0	0

The control returned the frozen incumbent *bit-identically* in all 35 trials and exceeded its allowance in 19 of them, so it was given more time, not less. The largest certified gains were $+0.213\%$ at $n = 23$ and $+0.115\%$ at $n = 28$; the median gain over the 27 improved counts was $+0.018\%$. The geometries move very little: the canonical root-mean-square displacement between warm start and result is 10^{-4} to 4×10^{-3}, with a maximum single-light move of 0.024

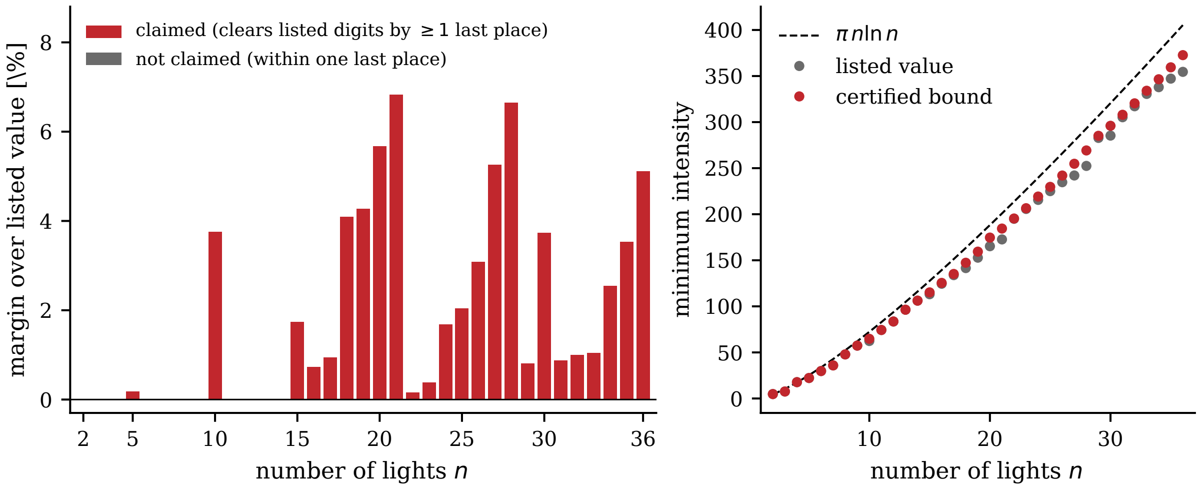

Figure 4: Left: certified lower bound minus listed value, in percent, for every open count. Red bars are the 24 claimed cases; grey bars are counts within one unit of the last displayed digit, which we treat as reproductions. Right: listed values and certified bounds against the critical-exponent asymptotic $\pi n \ln n$.

at $n = 23$. An independent audit replayed the v1 sweep against live incumbent snapshots and confirmed 22 of 22 strict gains with exact rational gain fractions. The machine was heavily loaded throughout by the running campaign (eleven CPU workers and one CUDA worker on twelve threads, plus the trials); uncontended probes reached materially further (+0.028% at $n = 17$ against +0.00018% under full load), so the sweep numbers are a floor for the method, not its ceiling.

5.4 The failure mode of sampled floors

The four strict losses of the ungated v1 are instructive. The solver accepted layouts on the strength of a grid-discovered floor; at several n the rigorous verifier scored the result strictly *below* the incumbent because heuristic dark-point discovery had missed a dark region, so the solver believed it had improved. At $n = 23$ the same defect is visible in the other direction: the sampled floor of the warm start was 206.1385 while its certified minimum is 206.0991. The hazard has a simple mechanical origin: $\|\nabla_q I\|$ is of order 10^3 on a domain of size one, so a naive grid-plus-quasi-Newton pipeline leaps out of the seed's own basin and silently loses dark points. Structural seeds that do not depend on a sampling rate (circumcentres of neighbouring light triples, Delaunay edge midpoints, and their boundary projections) mitigate this; gating on the verifier eliminates the consequence. No sampled floor in this paper is presented as a bound.

5.5 What did not work

Several deliberately different mechanisms failed to move certified incumbents under matched tests, and are recorded so that they are not repeated: an allocation/mosaic screen (36 of 36 runs returned the warm incumbent); a collective-direction stage moving in the null space of the centred active Jacobian, in the spirit of gradient sampling [5] (24 of 24 valid runs, no gain); an integer relocation model (18 forced models, no moved feasible solution); a patch/distinct screen (18 jobs, no certified gain); a GPU differential-evolution population [27] with symmetry-aligned donors and short local bursts [8], which lost its paired comparison to the adaptive exchange and was not deployed; and the multiplier-guided release homotopy described in Section 3.3. Covariance-adapting strategies [13] and bundle methods [6] were reviewed but not tried. Each result establishes only that the particular variant failed on the selected cases and budgets.

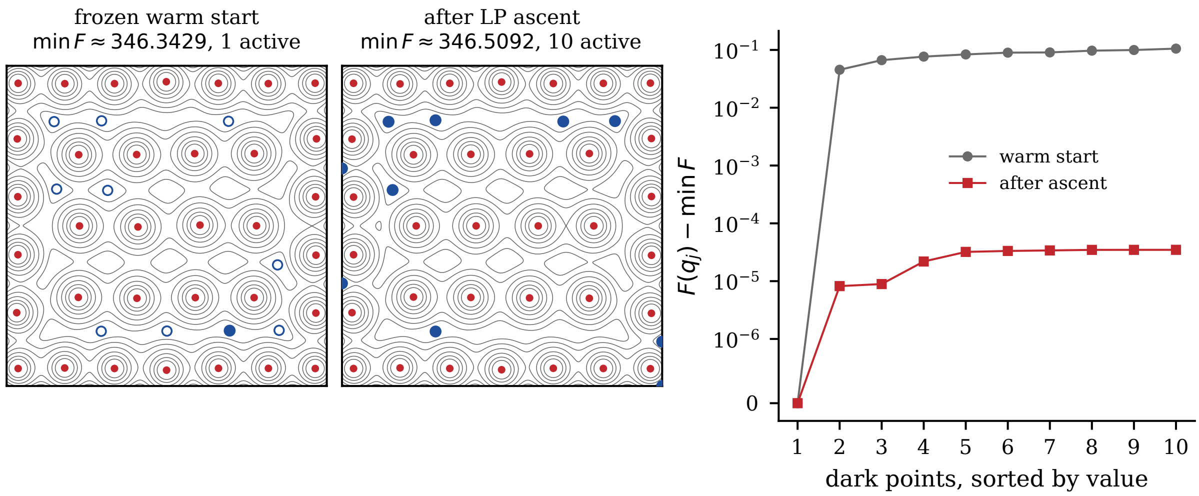

Figure 5: Non-stationarity and its repair at $n = 34$. Left and centre: the frozen warm start and the layout after LP ascent, with lights (red) and the ten lowest dark points (blue; filled if within 10^{-5} relative of the floor). Right: the values of the ten lowest dark points above the floor. The warm start has one active dark point with the runner-up 4.5×10^{-2} above it; after ascent ten dark points are tied to within 3.5×10^{-5} and the floor has risen by 0.166.

5.6 Cost

All computation ran on one desktop: a Ryzen 5 3600X (twelve logical CPUs) and an RTX 4070 (12 GiB). The campaign ran eleven single-threaded CPU search workers and one CUDA worker under a 48-hour clock with search cut off at hour 42 and the remainder reserved for verification. Complete 10^{-6} certificates for the final layouts take between 0.1 and 6.5 seconds each on the production verifier; the exact rational certificate for $n = 5$ takes about 3.4 seconds.

6 Transfer to three other problems on the same page

The pipeline of Section 3 never used anything specific to the inverse-square kernel beyond the formulas for its branches. To test whether the method, rather than the problem, is responsible for the gains, we ran the same three stages, with the same trust-region LP step and the same acceptance rule, on the three other families of the Packing Center that share the structure of an outer optimization over finitely many smooth branches that must be rediscovered after every move. In every case the certificate is the project's existing independent validator, unchanged.

Ratio of maximum to minimum distance. Minimize $\max_{i<j} \|p_i - p_j\|^2 / \min_{i<j} \|p_i - p_j\|^2$ over n points in the plane ($n \le 31$) or in space ($n \le 30$) [10]. The branches are the $\binom{n}{2}$ pairs; the LP step lowers the band of pairs near the maximum while holding the band near the minimum, with both bands chosen by Lemma 3.1. The population stage descends $\log \mathrm{softmax}_T - \log \mathrm{softmin}_T$ of the squared distances, which is scale invariant. The certificate is the exact rational ratio of the exported binary64 coordinates.

Unit disks covering a circle, a square or a triangle. Minimize the covering radius $\rho(C) = \max_{x \in S} \min_i \|x - c_i\|$ over centres C; the largest target that unit disks cover has scale $1/\rho$. The branches are the finitely many critical points of the nearest-centre distance: circumcentres of empty-circle triples (Voronoi vertices inside S), intersections of bisectors with the boundary, and corners or antipodal boundary points. Each branch value is written as an explicit function of the three, two or one centres that define it, and its gradient is obtained by automatic differentiation

Table 3: The same pipeline on three other families. "Behind before" counts cases whose campaign incumbent was more than 0.1% worse than the listed value; "recovered" counts those brought within 0.1%; "beats listed" uses the same full-last-place rule as Section 5.1; "time" is the median budget per case; "max gain" is the largest improvement over an incumbent.

family	cases	behind before	recovered	still behind	beats listed	time [s]	max gain [%]
ratio (plane and space)	49	8	5	3	0	90	0.7
disks covering	32	26	25	1	0	100	13.3
circle packing	127	44	43	1	3	60	45.8

through that construction; the critical points are recomputed from scratch after every step and a step is accepted only on the rediscovered maximum. The population stage descends a soft-max over sample points of a soft-min over centres. The certificate is the whole-domain analytic validator, and every exported layout was additionally certified by the exact rational provers of the project.

Unit circles packed in a container. Maximize the clearance

$$m(C) = \min\Big(\min_{i<j} \tfrac{1}{2}\,\|c_i - c_j\|\,,\ \min_i \operatorname{dist}(c_i, \partial S)\Big)$$

in a container S of unit size; circles of radius r then fit in the container scaled by r/m. The branches are the pairs and, per circle, each container edge (linear), the circular boundary, or the reflex notch of the L-shaped container. Seven containers are listed: circle, triangle, square, pentagon, hexagon, the right isosceles "tan", and the L. The certificate is the project's validator with exact decimal predicates and no overlap tolerance.

What happened. Table 3 summarizes the sweeps, which used one to five minutes per case, warm starts from the campaign's own incumbents, and exactly the code of Sections 3.1–3.3 with the branch definitions above. The method recovered the listed state of the art wherever the campaign's incumbents had been behind, in some cases by tens of percent (the L-shaped container, the disks covering a circle), and reproduced the listed values to within one unit of their last displayed digit; the per-case tables are in Appendix B. A handful of cases remain a few tenths of a percent short after five minutes. Three listed values were beaten, all in the L-shaped container: for 45 unit circles the pipeline found a packing with long side 15.431837 against the listed 15.43636^{+} (Figure 6), smaller by 0.0045, that is by 450 units of the last displayed digit; for 46 circles 15.709649 against 15.71083^{+} (118 units); and for 44 circles 15.270195 against 15.27022^{+}, a marginal 2.5 units that we report but would not press. All three survive either reading of the trailing plus. Each layout was checked in exact rational arithmetic from its binary64 coordinates: every pairwise squared distance is at least 4, every centre lies inside the outer square with clearance at least 1, and every centre is at squared distance at least 1 from the removed corner square; the project's validator, with exact decimal predicates and no overlap tolerance, agrees. Everywhere else the outcome is reproduction, which is what one expects for problems that have been attacked for decades by dedicated searches (Packomania, Nurmela and Östergård, Melissen, Cantrell, and recent solver-based work by Berthold and others). It sharpens the reading of Section 5: the light problem was not harder, it was less explored, and the same generic pipeline finds in minutes what its literature had not yet recorded, while on the mature problems it matches the literature and occasionally, as at $n = 44$–46 in the L, improves it.

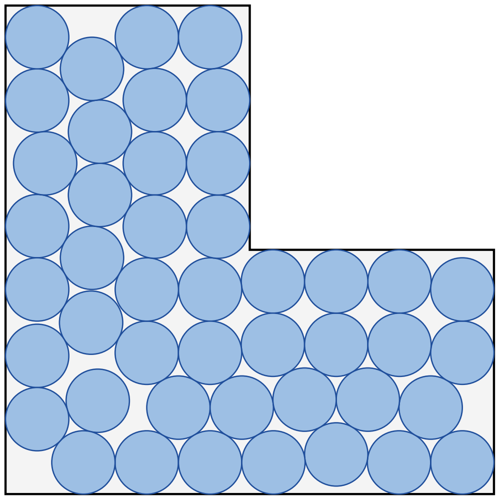

Figure 6: Forty-five unit circles in an L-shaped container of long side 15.431837, found by the transferred pipeline in one minute from a warm start that was 2.7% larger; the listed value is 15.43636^{+}. Exact binary64 coordinates and the rational check script for this and the $n = 44, 46$ layouts are archived with the paper.

7 Discussion

The mathematical content of this paper is modest and the computational content is careful; we think the combination is what the record problem needed. Two lessons generalize beyond lights in a square.

First, in a maximin problem whose inner minimizers can be born and die as the outer variables move, *how candidates are scored* matters more than how they are generated. Every mechanism that improved our incumbents shares one design: incumbent and challenger are compared on a common witness set, and the final arbiter is a certificate rather than a sample. Every mechanism that produced strict losses violated one of these.

Second, first-order stationarity is cheap to test and was never tested. Proposition 2.1 is elementary, yet all 35 layouts in a bank built by three different search methods failed it in the most flagrant way, with a single active dark point. A record-keeping site cannot check this without the coordinates, but authors of new records can, and we would encourage reporting the number of tied dark points alongside a value.

What remains open is everything about optimality. We have no upper bounds on V_n^* beyond the trivial, and the mixed-integer relaxations of [25] assume kernels continuous at zero, so they do not transfer directly to the inverse-square kernel. Exact rational certificates like (7) for all 24 layouts are a matter of computer time and are planned. The dip-controlled release homotopy, symmetric ansätze for the counts with doubled boundary lights, and a search over the combinatorial cells suggested by the roughly $1.5n$ dark points per layout are the natural next steps for discovery. At $n = 6$ we are behind the page and would welcome its coordinates.

8 Reproducibility and data

The repository contains the frozen problem catalogue with the retrieved page text and its SHA-256 (`76c3318d...6de`, retrieved 2026-09-11), the five solvers, the production and dedicated verifiers (core SHA-256 `0157a52f...6292` and `f077d26b...6a25`), the exact rational prover and its independent replayer, all certificates, and the submission package `outputs/light-submission-20260911/` with `submission.json` (per case: listed value, certified interval, margin, verifier method, interval width, coordinates and their hexadecimal encodings), one

CSV of full-precision coordinates per count, and one rendered figure per count. The tables and figures of this paper are regenerated from that package by `paper/make_tables.py` and `paper/make_figures.py`. To check any layout from the coordinates alone:

```
python verify_light.py coordinates/n021.csv -tolerance 1e-6 -seconds 120
```

returns an outward-rounded interval enclosing its minimum intensity over the closed square and states whether the requested width was reached. GPU runs are not bitwise reproducible across driver or library versions; every exported layout is, because only its coordinates matter to the certificate. Nothing in this work was submitted anywhere automatically; the configurations are offered for independent checking.

Acknowledgements

We thank Erich Friedman for maintaining the Packing Center, and the previous record holders whose configurations set the bar. The implementation and experiments were carried out with substantial assistance from AI coding agents (Anthropic's Claude and OpenAI's Codex) working under the author's direction in a shared, logged workspace; every claim in this paper is backed by an archived certificate that does not depend on that assistance.

References

[1] P. Blanchard, D. J. Higham, and N. J. Higham. Accurate computation of the log-sum-exp and softmax functions. *IMA Journal of Numerical Analysis*, 41(4):2311–2330, 2021.

[2] J. W. Blankenship and J. E. Falk. Infinitely constrained optimization problems. *Journal of Optimization Theory and Applications*, 19(2):261–281, 1976.

[3] S. V. Borodachov and N. Bosuwan. Asymptotics of discrete Riesz d-polarization on subsets of d-dimensional manifolds. *Potential Analysis*, 41:35–49, 2014. arXiv:1307.1160.

[4] S. V. Borodachov, D. P. Hardin, and E. B. Saff. *Discrete Energy on Rectifiable Sets*. Springer Monographs in Mathematics. Springer, 2019.

[5] J. V. Burke, F. E. Curtis, A. S. Lewis, M. L. Overton, and L. E. A. Simões. Gradient sampling methods for nonsmooth optimization. In A. M. Bagirov, M. Gaudioso, N. Karmitsa, M. M. Mäkelä, and S. Taheri, editors, *Numerical Nonsmooth Optimization: State of the Art Algorithms*, pages 201–225. Springer, 2020. arXiv:1804.11003.

[6] F. E. Curtis and L. Zebiane. NonOpt: Nonconvex, nonsmooth optimizer, 2025. Lehigh ISE Technical Report 25T-005, arXiv:2503.22826.

[7] J. M. Danskin. *The Theory of Max-Min and its Application to Weapons Allocation Problems*. Springer, 1967.

[8] J. P. K. Doye and D. J. Wales. Thermodynamics of global optimization. *Physical Review Letters*, 80:1357–1360, 1998.

[9] T. Erdélyi and E. B. Saff. Riesz polarization inequalities in higher dimensions. *Journal of Approximation Theory*, 171:128–147, 2013.

[10] E. Friedman. Maximizing minimum light intensity. Erich's Packing Center, https://erich-friedman.github.io/packing/light/, 2026. Page state of 11 September 2026; listed values for $n \leq 36$, most given to three decimals with a trailing "+".

[11] D. Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Computing Surveys*, 23(1):5–48, 1991.

[12] E. Hansen and G. W. Walster. *Global Optimization Using Interval Analysis.* Marcel Dekker, 2nd edition, 2004.

[13] N. Hansen. The CMA evolution strategy: A tutorial, 2016. arXiv:1604.00772.

[14] D. P. Hardin, M. Petrache, and E. B. Saff. Unconstrained polarization (Chebyshev) problems: basic properties and Riesz kernel asymptotics. *Potential Analysis*, 56:21–64, 2022. arXiv:1902.08497.

[15] R. Hettich and K. O. Kortanek. Semi-infinite programming: theory, methods, and applications. *SIAM Review*, 35(3):380–429, 1993.

[16] Q. Huangfu and J. A. J. Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, 2018.

[17] IEEE. IEEE standard for floating-point arithmetic, 2019. IEEE Std 754-2019.

[18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[19] D. Kraft. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, DFVLR, Institut für Dynamik der Flugsysteme, Oberpfaffenhofen, 1988.

[20] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.

[21] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems.* Prentice-Hall, 1974.

[22] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018.

[23] R. E. Moore, R. B. Kearfott, and M. J. Cloud. *Introduction to Interval Analysis.* SIAM, Philadelphia, 2009.

[24] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, 2019.

[25] J. Rolfes, R. Schüler, and M. C. Zimmermann. Bounds on polarization problems on compact sets via mixed integer programming, 2023. arXiv:2303.10101.

[26] T. Seidel and K.-H. Küfer. An adaptive discretization method solving semi-infinite optimization problems with quadratic rate of convergence. *Optimization*, 2022. doi:10.1080/02331934.2020.1804566; arXiv:1910.13798.

[27] R. Storn and K. Price. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11:341–359, 1997.

[28] P. Virtanen, R. Gommers, T. E. Oliphant, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020.

A Coordinates of the twenty-four configurations

Each coordinate is the exact binary64 value used by the certificate, printed as its shortest round-trip decimal (so that parsing the decimal back to binary64 recovers the same value). Every light has unit multiplicity; coincident lights are listed separately. Certified intervals are the recomputed whole-square enclosures of Table 2.

$n = 5$

Certified minimum intensity in [22.252232129341866, 22.252232742996984]; page lists 22.213^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.2450332902964024	0.050224193045460754	4	0.8594249385751961	0.19676581469693893
2	0.24503390065896305	0.9497776201499334	5	0.8594257048804704	0.8032335862815333
3	0.24607198537599725	0.500001361229161			

$n = 10$

Certified minimum intensity in [64.75195382969918, 64.75195479941156]; page lists 62.408^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.04577608283426391	0.8588646857391639	6	0.5000006795749589	0.02648964185496157
2	0.05706869395153063	0.45254139964524254	7	0.6610926169855761	0.8666726937569663
3	0.11303044546893605	0.09094564761387136	8	0.8869703061553262	0.0909465453147628
4	0.3389046286918966	0.8666723602689475	9	0.9429311898912007	0.4525428856223764
5	0.5000000009694412	0.42615032520612217	10	0.9542223176473535	0.8588653817753352

$n = 15$

Certified minimum intensity in [115.15957830740835, 115.15957930146038]; page lists 113.191^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.04744471421707864	0.6399209919161709	9	0.650315081966159	0.9176752220909149
2	0.07433852469997408	0.08159386367498582	10	0.6537104349463091	0.08231153777506924
3	0.07794792614399874	0.9220937509311552	11	0.6773739175796368	0.5401309860259534
4	0.08217244657636652	0.34973545711048376	12	0.9177917703677786	0.34634798815542517
5	0.35648376452903424	0.04485987377747264	13	0.9183827582451969	0.9256857208487324
6	0.3601893210457899	0.9526538747938845	14	0.9217897347419342	0.07820353846630612
7	0.37370307701730276	0.6263868543168937	15	0.9549960375282911	0.6436041012192032
8	0.4594519515722773	0.3226363456272143			

$n = 16$

Certified minimum intensity in [125.45809637800511, 125.45809735094981]; page lists 124.553^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.06432949531566257	0.08518537317352311	9	0.5000000020120254	0.8796149301878254
2	0.0643294956398575	0.9148146271480501	10	0.6335287428284948	0.4999999949889057
3	0.06606796574736813	0.3599419206043752	11	0.7101934910143147	0.8891328777170696
4	0.0660679675728403	0.6400580811921075	12	0.7101934953709477	0.11086711685400448
5	0.2898065068089987	0.11086712165808182	13	0.9339320304893094	0.6400580830051483
6	0.2898065101009906	0.8891328809034514	14	0.9339320317337194	0.3599419245427146
7	0.36647125481531306	0.5000000005586787	15	0.9356705033290343	0.9148146278556469
8	0.49999999885897545	0.12038506687103555	16	0.9356705086458978	0.08518537620735885

$n = 17$

Certified minimum intensity in [135.16651054934113, 135.1665115361905]; page lists 133.907^+.

i	x_i	y_i	i	x_i	y_i
1	0.03280504941444993	0.3858381906671504	10	0.6381375400822125	0.3763295819885418
2	0.03720339520518757	0.09715131410752732	11	0.6474368646180741	0.9710479279171892
3	0.04937228978959912	0.6736323690881745	12	0.6903863056006705	0.6556015585658381
4	0.07527357737780863	0.9307111637366184	13	0.7414911771637696	0.07534172019623461
5	0.25175452438752166	0.08832086260762925	14	0.9198833605904124	0.9362293796738361
6	0.320232318341825	0.4268172320717333	15	0.9560388602901184	0.381978008998715
7	0.34775822721783634	0.9811896026872424	16	0.9640413348763169	0.09752134661289738
8	0.3807602183337957	0.7198794989774181	17	0.9763405374003157	0.6747050516075979
9	0.4918106521568003	0.08726938484211788			

$n = 18$

Certified minimum intensity in [147.34120852626504, 147.34120951434838]; page lists 141.556^+.

i	x_i	y_i	i	x_i	y_i
1	0.06777730273433444	0.2536001455729128	10	0.6312885157379244	0.6629197598755868
2	0.06800800673275877	0.7465327362374085	11	0.6321147456658583	0.33907130246422523
3	0.07596206835685049	0.500115437209163	12	0.636598840634036	0.9553694636295522
4	0.09249181614790328	0.03657376355678236	13	0.6368662451070987	0.045273708398109305
5	0.09249992711265953	0.9634390835599973	14	0.9075049175845928	0.9634308360213252
6	0.36312901012008025	0.9546609107967143	15	0.9075097536760798	0.036580399239014974
7	0.3633808956465848	0.044660486033158774	16	0.9240488976422614	0.49997407297283875
8	0.3679822132594959	0.6608205685069862	17	0.932035892553832	0.2535227760479717
9	0.3686717452146627	0.33700658227499497	18	0.9321390983579445	0.7464517584735825

$n = 19$

Certified minimum intensity in [159.38552798569688, 159.3855289759443]; page lists 152.854^+.

i	x_i	y_i	i	x_i	y_i
1	0.07200937009945961	0.9386305388167239	11	0.5453884902540833	0.07458729914617122
2	0.0720117429255745	0.06136673676563115	12	0.6924471285867538	0.6287479466462567
3	0.0797118981850781	0.7201979049546244	13	0.6924472171123847	0.37126463364766626
4	0.07971200935394959	0.2797971236997209	14	0.7769891756241767	0.9263293578780322
5	0.08566146323200247	0.49999743723726914	15	0.7769907344768185	0.07367665876205429
6	0.3082507110238244	0.9345085400100576	16	0.9675656127293499	0.636187135097991
7	0.30825707947971037	0.06548789296583758	17	0.9675661092197165	0.3638125907894518
8	0.3927852725615812	0.6396142000295864	18	0.9752186367198064	0.09371731829117602
9	0.3927861561301893	0.36037943231973224	19	0.9752200255076404	0.906282275020033
10	0.5453814540989516	0.9254151704024404			

$n = 20$

Certified minimum intensity in [174.65564216038842, 174.65564311506728]; page lists 165.272^+.

i	x_i	y_i	i	x_i	y_i
1	0.06397222074747498	0.9360277659042293	11	0.5000000218866372	0.07454566777347651
2	0.06397222869030582	0.06397222552347927	12	0.6363824742718216	0.363617532100557
3	0.07146370488166773	0.28054753855191983	13	0.6363825034379755	0.6363825126242888
4	0.07146372324147171	0.7194524561872434	14	0.719452448363633	0.9285362984667607
5	0.07454567418779812	0.5000000031999058	15	0.7194524779120631	0.07146372427985
6	0.28054751712427606	0.9285362729434424	16	0.9254543295300347	0.4999999960875859
7	0.2805475487376373	0.07146369818641261	17	0.9285362770746745	0.2805475405488212
8	0.36361749747251104	0.3636174870478966	18	0.9285362980080597	0.719452464830542
9	0.36361753253584517	0.6363824694252354	19	0.9360277683600348	0.9360277774319362
10	0.49999997437108834	0.9254543366619338	20	0.9360277774740791	0.06397223227289162

$n = 21$

Certified minimum intensity in [184.43928774155285, 184.43928872778707]; page lists 172.658$^+$.

i	x_i	y_i	i	x_i	y_i
1	0.06036095107985492	0.9358500961141756	12	0.515363728190525	0.9296961158246596
2	0.060440711605954586	0.06407367884200887	13	0.6598140187560201	0.3553640652378997
3	0.06247735551713835	0.7196547353222228	14	0.6598771560654182	0.6446370237770945
4	0.06250439750842501	0.28019376947705493	15	0.7312177992774289	0.06968812991839092
5	0.07103055751767026	0.49991516271953956	16	0.7313052090941525	0.9302302380030016
6	0.2837565692117402	0.9567609574761992	17	0.9249974755583734	0.4999199406814291
7	0.2838658900085867	0.04331150955541413	18	0.938215917580968	0.28249149560866166
8	0.35778045413128057	0.290155051238938	19	0.9383303784620876	0.7174122674279452
9	0.3578585823978081	0.7101722152017788	20	0.9401258415518738	0.06469062505967607
10	0.3763236852227487	0.5003508474605209	21	0.9401427609566246	0.9352931860478619
11	0.5153233298815256	0.07045262147938498			

$n = 22$

Certified minimum intensity in [195.36066743964213, 195.36066843406644]; page lists 195.050$^+$.

i	x_i	y_i	i	x_i	y_i
1	0.04780225237531769	0.732724997619603	12	0.5000000034075793	0.03343239273453169
2	0.04820820215211181	0.26932937099327753	13	0.5000000079748493	0.2923386850174942
3	0.051981447678811896	0.5014551956574398	14	0.7064988230356147	0.6747661294324734
4	0.06471252767551047	0.05629463745327449	15	0.7160378101681959	0.9501404197429885
5	0.06540886252384065	0.9446358795816765	16	0.7214899837350153	0.3655774609582276
6	0.2729425485920136	0.06720286501830974	17	0.7270574542973852	0.06720287101164746
7	0.2785100319768469	0.36557743874358856	18	0.9345911354659961	0.9446358819081225
8	0.2839621796208419	0.9501404056704469	19	0.9352874718396155	0.05629463702503869
9	0.2935011810206395	0.6747661046544494	20	0.94801855592119	0.5014551964284624
10	0.4999999915937073	0.9259973941965879	21	0.9517917953158662	0.26932936966978743
11	0.5000000002782029	0.6120198609778682	22	0.952197749176518	0.7327250011047898

$n = 23$

Certified minimum intensity in [206.53841948476452, 206.5384204389237]; page lists 205.756$^+$.

i	x_i	y_i	i	x_i	y_i
1	0.04234516456379023	0.2937507384867114	13	0.5000176193790882	0.9645964082389133
2	0.045338991060238275	0.523809595909661	14	0.5866112407176903	0.08903389306239388
3	0.045429573737176814	0.07065605846275622	15	0.7119488613315347	0.6566850294593494
4	0.05302861163737069	0.74831065023729	16	0.7193571664689392	0.3675905499532988
5	0.06652787479402392	0.9493867197615573	17	0.7241069297529098	0.9401758240213914
6	0.23020063813203576	0.0740391979792139	18	0.7697910126548291	0.07403243933423262
7	0.2759071257809097	0.9402108105805468	19	0.9334683061755514	0.9493914541867533
8	0.2806220636104773	0.3676099675682466	20	0.9469855635283488	0.748311259549307
9	0.2880605956823317	0.6567260481115346	21	0.9545682831968838	0.07065437560815188
10	0.4133708305462788	0.08903878693389453	22	0.9546678620036542	0.5238006380378168
11	0.4999878464882251	0.40090407964054925	23	0.9576473286653634	0.2937435879336175
12	0.499993174284898	0.7102931486363493			

$n = 24$

Certified minimum intensity in [219.19786907858494, 219.19787001480762]; page lists 215.573^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.06704486330788209	0.9535375630161061	13	0.50000000259309	0.9490182979285204
2	0.06704486866136902	0.04646243058771591	14	0.6382050987874536	0.49999999501907727
3	0.07570185857598032	0.7755847398196076	15	0.6413783966049217	0.7167842617043155
4	0.07570188043055759	0.22441523290542825	16	0.6413784393302414	0.283215728094504
5	0.07647145447300038	0.592502596067561	17	0.7188087754969152	0.9553265270650978
6	0.07647146556111709	0.40749735105361085	18	0.7188087764050628	0.044673460979647486
7	0.28119122003148883	0.9553265493361977	19	0.9235285341033909	0.5925026448931212
8	0.28119122677535624	0.04467346361970715	20	0.9235285486544577	0.4074973982102255
9	0.3586215617899137	0.716784289257581	21	0.9242981240570158	0.7755847660785389
10	0.35862160380260927	0.28321573195648847	22	0.9242981425392444	0.22441525139236382
11	0.3617948922709865	0.5000000101641375	23	0.9329551302311189	0.9535375710168315
12	0.5000000012127859	0.05098170737159718	24	0.932955132410783	0.04646243121503255

$n = 25$

Certified minimum intensity in [229.7008494417817, 229.70085043821553]; page lists 225.100^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.048755716008853936	0.06319660653332292	14	0.5877618423584113	0.07947139277696333
2	0.0679109298014018	0.9576671553986578	15	0.6380575475409236	0.5243677278155553
3	0.0717705045508169	0.2543408258231554	16	0.6414101002310005	0.7315313929466981
4	0.07660900915291785	0.43515296870622217	17	0.6438330188000957	0.3168729805610023
5	0.07695203716703088	0.7879180367138949	18	0.7183273284523316	0.9603310558280939
6	0.07773756064340794	0.6130847160199271	19	0.7691476187479381	0.06910526586950713
7	0.23085224340206772	0.06910531669474339	20	0.9222626512151832	0.6130848557199104
8	0.28167249729406185	0.9603308663542278	21	0.9230478561127696	0.787918246075193
9	0.35616690959416053	0.3168729042872812	22	0.9233911592635704	0.43515280384956445
10	0.3585898408014497	0.7315311877445219	23	0.928229362784658	0.2543406379988796
11	0.36194312847177895	0.5243675023646919	24	0.9320890238822856	0.9576672194639527
12	0.41223803382985497	0.07947133453687842	25	0.951244231391505	0.06319656492852871
13	0.49999987169233107	0.9539662182010383			

$n = 26$

Certified minimum intensity in [242.00415113652556, 242.00415212357314]; page lists 234.759^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.04075973386018204	0.7319758221349918	14	0.5000027591777606	0.6687884680267069
2	0.0444932877144797	0.9363708194642828	15	0.5923174019938002	0.9327700494281594
3	0.06866260516446242	0.03790709477273697	16	0.6386984228495661	0.2541317634564201
4	0.07298480536782424	0.5367538932901467	17	0.6446813846735908	0.4411509215495888
5	0.07640008574814124	0.3620240852714698	18	0.718236784947377	0.034995872313633866
6	0.08049823284766837	0.19654731209150297	19	0.7255489953715575	0.6805486281877463
7	0.22159806106296193	0.9380467009924445	20	0.7784390599881948	0.9380251490510292
8	0.2744545050153995	0.6805868050503577	21	0.9195076944867493	0.19654594952372897
9	0.28175614747902245	0.03500148387355025	22	0.9235882362845443	0.362019775914577
10	0.3552908083561097	0.4411797167238477	23	0.9270047955127833	0.5367330664247669
11	0.36129693601218765	0.25415764068659996	24	0.9313346533520521	0.037902505960524524
12	0.40774026119789347	0.9327789236020634	25	0.9555203808585568	0.9363602603064853
13	0.4999946402306301	0.037581792063352784	26	0.9592472088080447	0.7319515080918868

$n = 27$

Certified minimum intensity in [254.77677337924, 254.77677437918322]; page lists 242.055^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.042432371935584016	0.7434475788046724	15	0.5892305739150909	0.06868342961178561
2	0.046018870876453254	0.9399577323960842	16	0.5922466615766091	0.9374614569280233
3	0.0491673324273617	0.05795743008797669	17	0.6407188780550414	0.28639002137065356
4	0.0749862787786822	0.5569900896641905	18	0.6437952756312683	0.46438384489576373
5	0.07735919630648062	0.23059360778353516	19	0.7251393315947087	0.6944589732837437
6	0.07809308291760837	0.3917285627264118	20	0.7713337869321722	0.0608738135884741
7	0.2223689158009341	0.9421068353343273	21	0.7776762658628992	0.9420528843925409
8	0.22868358004273426	0.060877587175139775	22	0.9218808907710404	0.3917343670265783
9	0.27486632083302276	0.6945866186119373	23	0.922644389926115	0.2306160033350284
10	0.35616017131681793	0.4644914901766613	24	0.9250107828153261	0.5569744866900516
11	0.3592453309232253	0.28645285819200134	25	0.9508437167385605	0.05796702011378093
12	0.4078509277619149	0.9374884835036139	26	0.9539883328376106	0.9399511633397434
13	0.4107689535741933	0.06870079091991334	27	0.9575644872101937	0.7434332326080418
14	0.5000219456498746	0.683956680676743			

$n = 28$

Certified minimum intensity in [269.2396964241986, 269.2396973832159]; page lists 252.447^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.04518637281191971	0.7613749995640391	15	0.5002183755505538	0.3293573869187984
2	0.045277634093376526	0.23853942725925184	16	0.5859081981791545	0.9334495156093839
3	0.05036387060959887	0.9464529394741756	17	0.5860052276914257	0.06646535551786732
4	0.050387159839374036	0.05352477432990287	18	0.6586974886727359	0.4999076716258767
5	0.0722522777226766	0.4180903752612999	19	0.7224081715172238	0.2787894527679008
6	0.07226273474788399	0.5817937238888758	20	0.7224112545905743	0.7212095524998192
7	0.2314820458129487	0.9523007848171589	21	0.7684510487146536	0.9525972384941568
8	0.23157424904780885	0.04760344244912819	22	0.7685496164294044	0.04753570130561443
9	0.27758848080657567	0.2787908454811296	23	0.9277406506026337	0.5819916483024075
10	0.27759164811677517	0.7212115200007991	24	0.927825059269445	0.4182167925459734
11	0.3413160505818236	0.49999476908882007	25	0.9496840238309171	0.9464080951598155
12	0.4138396876359088	0.9333864008548016	26	0.9496999078143844	0.05360698602120024
13	0.41392453269610613	0.06668301617090204	27	0.954511813591395	0.7614437923397908
14	0.5001065760185084	0.6706414320453224	28	0.9546319045459543	0.23864076145816188

$n = 29$

Certified minimum intensity in [285.12876831210946, 285.1287692611578]; page lists 282.856^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.05037107527272435	0.7744127595488537	16	0.500137977639932	0.2795789735090925
2	0.050469518188399336	0.2254889273876378	17	0.5917092552529999	0.05232654547895193
3	0.050533233434522755	0.9494482170692055	18	0.5917385876249427	0.9477342487765683
4	0.05055647664328219	0.050528923518315484	19	0.7204938391580364	0.4999155052726405
5	0.052282760003460806	0.4081931240008749	20	0.7224813039711676	0.7224190015819021
6	0.052302510026776765	0.5916616289736644	21	0.7225068494196218	0.2775125149397111
7	0.22552288688413472	0.9495770754223866	22	0.7744350882514329	0.05038916546652715
8	0.22559558530987558	0.0503753002483924	23	0.7744747408599765	0.9495676856119318
9	0.2774865511311578	0.7224455267575726	24	0.9476435874164516	0.5917213829890402
10	0.2776464816661502	0.2774978457517101	25	0.9476952564518137	0.40830161614236576
11	0.2793154603144288	0.4998919974042394	26	0.9494541829829172	0.05053892743816645
12	0.40825104998476885	0.9477105181195388	27	0.9494627522996079	0.949452014597157
13	0.4083083878232759	0.05236457942565133	28	0.9496093723847838	0.7744297435914299
14	0.4998369983432512	0.5002840679004696	29	0.9496199856063148	0.2255579388433002
15	0.5000176516959804	0.7207565368427178			

$n = 30$

Certified minimum intensity in [296.0050841399256, 296.0050851392704]; page lists 285.345^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.059431552370765625	0.49973558014836966	16	0.5255026718178899	0.7187625659422058
2	0.060118754530574706	0.9624800654499222	17	0.5256307982990119	0.28132419830665173
3	0.0601737729556588	0.03744587955291327	18	0.6068813481120735	0.9472271088790455
4	0.068711616747235	0.8119600781519493	19	0.6069634899949401	0.05280070674626446
5	0.06883462536706185	0.18782122642303078	20	0.720272683257544	0.5000362403025875
6	0.07367303213079852	0.6610016713395038	21	0.7358602626146681	0.7214799419229173
7	0.07379276026418005	0.33857756722229015	22	0.735939047050789	0.2785625246293391
8	0.24478327170746325	0.9545330302055449	23	0.7840821913349291	0.9479442864123653
9	0.24491217487665015	0.04545617006525228	24	0.7841391943655712	0.05206967922288045
10	0.28674759911939174	0.5000382377494227	25	0.9482830460738847	0.5899777779725424
11	0.30807443628345904	0.7293354658122848	26	0.9482958815387376	0.4101154063912624
12	0.3082558966081026	0.2707440147827436	27	0.9526520081818778	0.948125869381475
13	0.42733814593379504	0.9485619829679356	28	0.9526792161896332	0.051899427958333834
14	0.4274569835826043	0.05147474504756081	29	0.9548039179751698	0.7719283109252806
15	0.4957895261909825	0.5000840570661728	30	0.954839410714372	0.22813665395947502

$n = 31$

Certified minimum intensity in [307.96240703662534, 307.96240803355875]; page lists 305.298^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.059731942906412774	0.8140947156974968	17	0.5321759883913644	0.2710872569193785
2	0.05973759684434387	0.18589410541250295	18	0.5563355614441057	0.5000072332813118
3	0.06034242665422506	0.9655226123552126	19	0.6184674172453841	0.9474736793286002
4	0.06034541286237734	0.03447292489028554	20	0.6184719141741803	0.052527253027153636
5	0.06680056929587116	0.6557363333132368	21	0.7410574771186077	0.7235728142846276
6	0.0668068174437077	0.34424357904561736	22	0.7410614320158899	0.27642986836336486
7	0.06723897693955952	0.49998745999051336	23	0.749370478301013	0.5000008271661468
8	0.24880071474190354	0.9705997656606168	24	0.7910057607911022	0.9473361446638299
9	0.24880743876496633	0.029399149459599466	25	0.7910093426903904	0.05266432658902692
10	0.29993244115033035	0.7757826209172909	26	0.9546618476075877	0.9480438145435011
11	0.2999463185672244	0.22421386758806197	27	0.954663707705529	0.05195785447206368
12	0.31734651770080846	0.5904923381618843	28	0.9554326063632449	0.7743100374453085
13	0.3173532823752396	0.4095057962568401	29	0.9554338901983928	0.22569371397028623
14	0.4391521942853949	0.9554131019015789	30	0.9557707464989198	0.5921645125713741
15	0.4391585961242871	0.044589594275760314	31	0.9557731952606845	0.4078412717573844
16	0.5321640160627297	0.7289223567441501			

$n = 32$

Certified minimum intensity in [320.37104639580167, 320.37104724251606]; page lists 317.203^+.

i	x_i	y_i	i	x_i	y_i
1	0.05642187705895633	0.18714660302768202	17	0.513132218687446	0.7280608067425784
2	0.05644708345391049	0.8129090915988292	18	0.5436940138447313	0.49998581017372606
3	0.057861245439928974	0.035781918650961884	19	0.6002705160647902	0.05045776712857911
4	0.05787737826665904	0.9642410097484316	20	0.6003025760171256	0.9495422155380768
5	0.06416119951130025	0.34508794582331265	21	0.7127317802176913	0.27211175538053806
6	0.06416205397489289	0.6549856350776966	22	0.7127808639870332	0.7278901051894126
7	0.06451054649005672	0.5000368750674854	23	0.7308438965252335	0.4999725497381498
8	0.24059718917688996	0.03089717224740817	24	0.7723798004171304	0.049243286845679156
9	0.24064318164264745	0.9691360464310383	25	0.7724017816832439	0.9507296331040608
10	0.2888518488968486	0.22506165368113606	26	0.9321085621558325	0.6591366817050994
11	0.28890600531052074	0.7749993301025518	27	0.9321127394203079	0.3409216525489283
12	0.3084610719438277	0.5892971856401772	28	0.9375710016663998	0.8102415049079563
13	0.308485026605093	0.41077633340576003	29	0.9376323328527142	0.18977461906551946
14	0.42497298334485184	0.04640749919887067	30	0.9444849668989421	0.5000291590644477
15	0.4250197538796873	0.9535605842378508	31	0.944590900274485	0.039432424937185594
16	0.5130619609218663	0.271869263579955	32	0.9445998124066053	0.9605527676114926

$n = 33$

Certified minimum intensity in [334.0371931504373, 334.0371941482759]; page lists 330.595^+.

i	x_i	y_i	i	x_i	y_i
1	0.05481880810133234	0.18555282169276477	18	0.5000066445416852	0.7333145686148794
2	0.054822959277934305	0.8144562958525439	19	0.5862574395707656	0.04713039827317502
3	0.055935343783912114	0.036185788625158304	20	0.5862663282628795	0.9528682912245836
4	0.05593806275287297	0.9638179687135157	21	0.7077851897523484	0.4063074133077625
5	0.059103750232514525	0.3434470802204746	22	0.7077852780425803	0.5936575989136007
6	0.05910607308390395	0.6565665534147221	23	0.7196651707470799	0.22286178357649494
7	0.06060508668296888	0.5000076539856337	24	0.7196719099830364	0.7771215417077456
8	0.23414973911636747	0.030696466366333854	25	0.7658458430348463	0.030692190008737207
9	0.2341566910992451	0.9693073753837178	26	0.7658501943313232	0.9693002164112605
10	0.2803262755170602	0.22287054003718773	27	0.9394004390205514	0.4999992680679692
11	0.28033562726562344	0.7771371385945252	28	0.9408913417753282	0.6565602722881002
12	0.2922093153490216	0.406322428339757	29	0.940893037383386	0.34343696685372854
13	0.29221340895487147	0.5936911092682373	30	0.9440633537034642	0.03618404346716328
14	0.41373576791905364	0.04713073503796839	31	0.9440639669397051	0.9638151715912364
15	0.41374508484584	0.9528699662396966	32	0.9451776481729808	0.18554753681171515
16	0.49999349896162065	0.2666862656121954	33	0.9451795825202739	0.814449759852916
17	0.49999678809771475	0.5000009075721306			

$n = 34$

Certified minimum intensity in [346.50917672394695, 346.5091776536512]; page lists 337.906^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.03127509324768394	0.7709123727027334	18	0.4985173390102898	0.949217995848753
2	0.03130095918546429	0.22908242412609928	19	0.585798769253806	0.2751087136521937
3	0.036115179903137604	0.05461232610932705	20	0.5858477532125419	0.7249072349745707
4	0.03612053495308248	0.9453910070461076	21	0.6129305428990809	0.499993637193751
5	0.036130011423548346	0.4099151013539076	22	0.6599929362283627	0.05532472141254633
6	0.03613519101173382	0.5900719368411714	23	0.6600186393664526	0.9446705999712779
7	0.1813197372802948	0.9433812079722746	24	0.7725884352211181	0.27393508296480806
8	0.18133044275663193	0.05659230877434736	25	0.7726156685634307	0.7260283313670853
9	0.2232708425793655	0.7221988090560401	26	0.7864267027527955	0.499982986797466
10	0.22328391277991863	0.27778099775412396	27	0.8162889768157572	0.054791688536869484
11	0.2317500790600187	0.4999826889933736	28	0.8162989954343052	0.9451905172249856
12	0.33664872379785576	0.9451854276361769	29	0.9632499570757868	0.9458758569333627
13	0.3366607462805953	0.05484104019347534	30	0.9632513121661205	0.05412480880119876
14	0.402527903941772	0.7234528117062602	31	0.966392839500698	0.2271087303911024
15	0.40254921756792195	0.27664110199536446	32	0.9664063027257137	0.7728856568487347
16	0.41895399380693665	0.5000713495824263	33	0.9666120773701863	0.4081929729132035
17	0.49849930806878634	0.050800191243498834	34	0.966631725900167	0.5917874796740935

$n = 35$

Certified minimum intensity in [359.4630953776155, 359.4630963746088]; page lists 347.195^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.044448688811643725	0.04622545382712197	19	0.50002137399104	0.7190056346411917
2	0.044466878522417166	0.9537922287242163	20	0.5000241876786358	0.49999776350245134
3	0.052354071695900506	0.19984040931807887	21	0.6404938336796813	0.06017334803949008
4	0.052363719461172964	0.8001988138867084	22	0.6405669318550189	0.9398137252700892
5	0.058782584288027045	0.35224807819861587	23	0.708904318649428	0.5878679067737318
6	0.058786025928643654	0.6478009191632336	24	0.7089149395549749	0.4120954711242349
7	0.06221066851626647	0.5000258354226479	25	0.7243328202477393	0.24015401011488088
8	0.20184024223774724	0.04486806785042949	26	0.7243477433452495	0.7597937569720653
9	0.20188609990344222	0.9551559880172789	27	0.7980744062682266	0.044818317662423945
10	0.275673662475127	0.24023526041581605	28	0.7981212535586848	0.9551542861020891
11	0.2756851200142343	0.759806775097845	29	0.9378065730918016	0.4999724741850472
12	0.29113058514282336	0.4121477246846748	30	0.9412219636180392	0.35217949018615186
13	0.2911309511140449	0.5878942621009545	31	0.9412291623678584	0.6477708059363205
14	0.35936776034293444	0.060205461463515936	32	0.9476394651637092	0.19977558314993238
15	0.3594442377135542	0.9398158074560735	33	0.9476483360874717	0.8001866803764508
16	0.49991885844234524	0.07073349409614355	34	0.9555185498538296	0.04619330797740397
17	0.5000066241952433	0.9292617853514565	35	0.9555361424378487	0.9537895629697434
18	0.5000193924977433	0.2809988591389274			

$n = 36$

Certified minimum intensity in [372.68642009764534, 372.68642009764943]; page lists 354.550^{+}.

i	x_i	y_i	i	x_i	y_i
1	0.03421008585106069	0.2127379136163818	19	0.4999502105925044	0.4764516412712396
2	0.039126400540375174	0.04932584527862397	20	0.500534435676107	0.6558183341796548
3	0.04951040981374937	0.37783078788729074	21	0.5903098980974758	0.2700898518210817
4	0.05102614023980561	0.5355214067469359	22	0.591780616422435	0.9741159070071909
5	0.05539150298612738	0.9689470855091018	23	0.6544043181401104	0.05164819886561241
6	0.0564963712539689	0.8324483076938451	24	0.7235288755867753	0.6134066657540317
7	0.05801345453795938	0.6871864767728559	25	0.7268343514377232	0.7897034629921893
8	0.18853411345809792	0.04526418176038173	26	0.7304374208945056	0.44604661301183346
9	0.22485376507144378	0.24930774770028077	27	0.7718711432156578	0.9722093058505209
10	0.22748236124671697	0.9718183913590075	28	0.7752463598032678	0.24958551254857175
11	0.26980130724353957	0.4453851086705781	29	0.811550691991552	0.04550547601151702
12	0.27340586749824636	0.7888283382550403	30	0.9419196600760044	0.6876534087926569
13	0.2767407548383741	0.6129886222233775	31	0.9437176627424759	0.8330100054915551
14	0.34564470083939325	0.051922922395117914	32	0.9443438776554598	0.9693786956709082
15	0.40743381388017125	0.9742859391651487	33	0.949294860105796	0.5358493596724633
16	0.41095400956330974	0.26980284337159743	34	0.9505051425608716	0.3779473529142086
17	0.4997361934147791	0.8139223661007965	35	0.9608397278876892	0.04930017999083277
18	0.4999372112341021	0.05343049290146783	36	0.9659603001464224	0.21277313213376456

B Transfer results per case

Certified values from the transfer sweeps of Section 6. "Incumbent" is the campaign's previous best; "new certified" is the value certified after the transferred pipeline ran; "gain" is the improvement over the incumbent in the objective's own units. For the ratio problems the listed r^2 is minimized; for covering the target scale is maximized; for packing the container size is minimized. The packing table lists only the cases whose incumbent was more than 0.1% behind.

Ratio of maximum to minimum distance

family	n	listed r^2	incumbent	new certified	gain	vs listed
plane	9	6.600^+	6.600979	6.600978659	0	reproduces
plane	10	7.713^+	7.713455	7.713455164	0	reproduces
plane	11	8.222^+	8.222295	8.222295312	0	reproduces
plane	12	8.464^+	8.464102	8.464101615	0	reproduces
plane	13	9.930^+	9.930036	9.930035789	0	reproduces
plane	14	10.994^+	10.994718	10.994717906	0	reproduces
plane	15	12.038^+	12.038438	12.038438278	0	reproduces
plane	16	12.890^+	12.889230	12.889229908	0	within last place
plane	17	14.090^+	14.090519	14.090518705	0	reproduces
plane	18	14.725^+	14.725276	14.725276092	0	reproduces
plane	19	14.928^+	14.928203	14.928203230	0	reproduces
plane	20	16.734^+	16.734526	16.734525501	0	reproduces
plane	21	17.774^+	17.774981	17.774980554	0	reproduces
plane	22	19.053^+	19.053985	19.053984510	0	reproduces
plane	23	20.039^+	20.039612	20.039612312	0	reproduces
plane	24	$14+4\sqrt{3}$	20.960213	20.928203230	3.20×10^{-2}	reproduces
plane	25	22.121^+	22.136226	22.121568203	1.47×10^{-2}	reproduces
plane	26	22.810^+	22.810593	22.810592955	0	reproduces
plane	27	23.083^+	23.083331	23.083331496	0	reproduces
plane	28	24.981^+	25.016578	24.981438861	3.51×10^{-2}	reproduces
plane	29	25.924^+	25.937502	25.924603994	1.29×10^{-2}	reproduces
plane	30	26.879^+	26.879265	26.879264785	0	reproduces
plane	31	28	28.289058	28.283576766	5.48×10^{-3}	behind
space	5	1.714^+	1.714286	1.714285714	0	reproduces
space	6	2	2.000000	2.000000000	0	reproduces
space	7	2.293^+	2.293437	2.293437256	0	reproduces
space	8	2.414^+	2.414214	2.414213562	0	reproduces
space	9	2.544^+	2.544727	2.544727086	0	reproduces
space	10	3.101^+	3.101626	3.101626005	0	reproduces
space	11	3.384^+	3.384172	3.384171992	0	reproduces
space	12	3.618^+	3.618034	3.618033989	0	reproduces
space	13	3.947^+	3.947141	3.947141175	0	reproduces
space	14	4.165^+	4.165783	4.165783475	0	reproduces
space	15	4.398^+	4.398109	4.398108949	0	reproduces
space	16	4.553^+	4.553203	4.553202961	0	reproduces
space	17	4.771^+	4.771073	4.771073409	0	reproduces
space	18	5.047^+	5.047609	5.047608573	0	reproduces
space	19	5.409^+	5.409960	5.409960182	0	reproduces
space	20	5.799^+	5.799520	5.799520109	0	reproduces
space	21	5.985^+	5.985519	5.985519297	0	reproduces
space	22	6.34552^+	6.345524	6.345524236	0	reproduces
space	23	6.61319^+	6.613191	6.613190831	0	reproduces
space	24	6.90131^+	6.946554	6.901325006	4.52×10^{-2}	behind
space	25	7.04527^+	7.045279	7.045279119	0	reproduces
space	26	7.30530^+	7.307074	7.307021628	5.19×10^{-5}	behind
space	27	7.52605^+	7.536739	7.536738811	0	behind
space	28	7.68958^+	7.738666	7.707984579	3.07×10^{-2}	behind
space	29	7.84943^+	7.885971	7.849436186	3.65×10^{-2}	reproduces
space	30	8.03171^+	8.062329	8.034732554	2.76×10^{-2}	behind

Unit disks covering a circle, square or triangle

target	n	listed	incumbent	new certified	gain	vs listed	exact
circle	1	1	1.000000	0.999999980	0	reproduces	passed
circle	3	1.154^{+}	1.154701	1.154700514	0	reproduces	passed
circle	4	1.414^{+}	1.414036	1.414213534	1.77×10^{-4}	reproduces	passed
circle	7	2	1.964251	1.999999960	3.57×10^{-2}	reproduces	passed
circle	11	2.631^{+}	2.565813	2.631690720	6.59×10^{-2}	within last place	passed
circle	12	2.769^{+}	2.540560	2.769293194	2.29×10^{-1}	within last place	passed
circle	13	2.884^{+}	2.811883	2.884791824	7.29×10^{-2}	within last place	passed
circle	14	3.014^{+}	2.758805	3.014481198	2.56×10^{-1}	within last place	passed
circle	15	3.143^{+}	2.968047	3.143241238	1.75×10^{-1}	within last place	passed
circle	16	3.244^{+}	2.947459	3.244437813	2.97×10^{-1}	within last place	passed
circle	17	3.349^{+}	3.191818	3.349729194	1.58×10^{-1}	within last place	passed
circle	18	3.446^{+}	3.144577	3.446282681	3.02×10^{-1}	within last place	passed
circle	19	3.605^{+}	3.457152	3.605551203	1.48×10^{-1}	reproduces	passed
circle	20	3.692^{+}	3.415627	3.692105244	2.76×10^{-1}	within last place	passed
circle	21	3.804^{+}	3.548117	3.804522454	2.56×10^{-1}	within last place	passed
circle	22	3.948^{+}	3.791440	3.948644484	1.57×10^{-1}	within last place	passed
circle	23	4.000^{+}	3.881853	4.000739557	1.19×10^{-1}	within last place	passed
circle	24	4.076^{+}	3.749010	4.076613342	3.28×10^{-1}	within last place	passed
circle	25	4.181^{+}	4.025899	4.172780004	1.47×10^{-1}	behind	passed
square	1	1.414^{+}	1.414214	1.414213534	0	reproduces	passed
square	2	1.788^{+}	1.774589	1.788854346	1.43×10^{-2}	reproduces	passed
square	6	3.347^{+}	3.292336	3.347537289	5.52×10^{-2}	within last place	passed
square	8	3.841^{+}	3.747535	3.841719451	9.42×10^{-2}	within last place	passed
square	9	4.335^{+}	3.827623	4.335819027	5.08×10^{-1}	within last place	passed
square	10	4.582^{+}	4.412461	4.582247553	1.70×10^{-1}	reproduces	passed
square	11	4.705^{+}	4.375847	4.705527595	3.30×10^{-1}	within last place	passed
square	12	4.943^{+}	4.700135	4.943742845	2.44×10^{-1}	within last place	passed
triangle	1	1.732^{+}	1.732051	1.732050773	0	reproduces	passed
triangle	2	2	2.000000	1.999999960	0	reproduces	passed
triangle	3	3.464^{+}	3.455948	3.464101546	8.15×10^{-3}	reproduces	passed
triangle	7	5.398^{+}	5.338578	5.398078920	5.95×10^{-2}	within last place	passed
triangle	8	5.649^{+}	5.426021	5.649923091	2.24×10^{-1}	within last place	passed

Unit circles packed in a container (cases with a material incumbent gap)

container	n	listed	incumbent	new certified	gain	vs listed
L	11	7.863^+	7.90490	7.86370331	4.12×10^{-2}	reproduces
L	13	8.812^+	8.84758	8.81250478	3.51×10^{-2}	reproduces
L	15	9.412^+	9.45503	9.41264950	4.24×10^{-2}	reproduces
L	17	10.00425^+	10.05226	10.00425359	4.80×10^{-2}	reproduces
L	18	10.27388^+	10.31907	10.27388506	4.52×10^{-2}	reproduces
L	19	10.46489^+	10.54622	10.46489665	8.13×10^{-2}	reproduces
L	20	10.61664^+	10.70758	10.61664507	9.09×10^{-2}	reproduces
L	21	10.82564^+	10.93632	10.82564950	1.11×10^{-1}	reproduces
L	22	10.89927^+	20.11855	10.89927567	9.22×10^{0}	reproduces
L	23	11.21781^+	11.34590	11.21781241	1.28×10^{-1}	reproduces
L	24	11.41224^+	11.48715	11.41224130	7.49×10^{-2}	reproduces
L	25	11.71270^+	11.74593	11.71269960	3.32×10^{-2}	within last place
L	26	11.86038^+	11.89205	11.86089750	3.12×10^{-2}	behind
L	28	12.49680^+	13.00623	12.49687834	5.09×10^{-1}	behind
L	29	12.67553^+	12.86521	12.67552969	1.90×10^{-1}	within last place
L	30	12.94284^+	13.03276	12.94408187	8.87×10^{-2}	behind
L	31	13.12408^+	13.22958	13.12472782	1.05×10^{-1}	behind
L	32	13.30282^+	13.59306	13.30283603	2.90×10^{-1}	behind
L	33	13.52754^+	13.64458	13.52754422	1.17×10^{-1}	reproduces
L	34	13.72415^+	13.96366	13.72415786	2.39×10^{-1}	reproduces
L	35	13.86770^+	20.00000	13.86770878	6.13×10^{0}	reproduces
L	36	14.01152^+	14.22104	14.01152758	2.10×10^{-1}	reproduces
L	37	14.10954^+	26.01833	14.10954634	1.19×10^{1}	reproduces
L	38	14.23736^+	16.54007	14.23736444	2.30×10^{0}	reproduces
L	39	14.30561^+	14.72645	14.30561654	4.21×10^{-1}	reproduces
L	40	14.54974^+	15.05704	14.54974288	5.07×10^{-1}	reproduces
L	41	14.69163^+	15.08247	14.69169459	3.91×10^{-1}	behind
L	42	14.90788^+	15.12081	14.91558251	2.05×10^{-1}	behind
L	43	15.06571^+	23.07236	15.06587741	8.01×10^{0}	behind
L	44	15.27022^+	15.53305	15.27019549	2.63×10^{-1}	**beats listed**
L	45	15.43636^+	15.85679	15.43183675	4.25×10^{-1}	**beats listed**
L	46	15.71083^+	15.94888	15.70964889	2.39×10^{-1}	**beats listed**
L	47	15.85935^+	16.17627	15.85934987	3.17×10^{-1}	within last place
L	48	15.98311^+	16.69271	15.98437637	7.08×10^{-1}	behind
L	49	16.13179^+	16.62835	16.18104896	4.47×10^{-1}	behind
L	50	16.30078^+	17.00640	16.30597835	7.00×10^{-1}	behind
circle	12	4.029^+	4.11362	4.02960193	8.40×10^{-2}	reproduces
hexagon	5	2.999^+	3.15470	2.99935540	1.55×10^{-1}	reproduces
hexagon	13	4.618^+	4.66656	4.61880216	4.78×10^{-2}	reproduces
hexagon	17	5.133^+	5.15470	5.13311943	2.16×10^{-2}	reproduces
pentagon	21	7.107^+	7.12599	7.10726689	1.87×10^{-2}	reproduces
tan	21	13.414^+	13.42767	13.41421357	1.35×10^{-2}	reproduces
triangle	13	11.406^+	11.46410	11.40649586	5.76×10^{-2}	reproduces

C Re-verification during manuscript preparation

The five layouts below were re-certified from their coordinates on 11 September 2026 by the dedicated verifier (`verify_light.py`, tolerance 10^{-6}, 90-second allowance) in a process separate from the campaign. All five certificates completed. The enclosures are consistent with, and in three cases slightly narrower than, the intervals of Table 2; the outward decimal endpoints are reproduced verbatim.

n	lower	upper
5	22.252231748716027	22.252232742996963
21	184.43928872778613	184.43928926899640
23	206.53841944025984	206.53842043892303
28	269.23969645717414	269.23969739033243
36	372.68642009764533	372.68642009764728