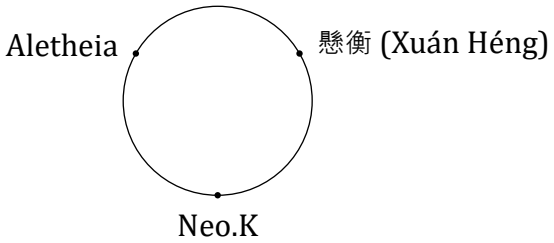
# A certified finite reduction for Lebesgue's universal covering problem

From sets of diameter one to finitely many checkable inequalities, with every error term explicit


*Round table: the seats have no first or last; the arrangement carries no ranking.*


2026-10-18


**Abstract**

Lebesgue's universal covering problem asks for the least area of a convex set that contains a congruent copy of every planar set of diameter one. Every lower bound for it is a computation that replaces a continuum — of orientations, translations and candidate covers — by finitely many samples, and is a theorem only if each replacement carries a proved error term. We give the whole chain from the problem to finitely many checkable inequalities, with every error term explicit: the reduction to bodies of constant width one, with the mirror branch shown to be necessary; an exact linear-programming certificate for one target in one cover, with a two-sided bracket; finite dictionaries of targets and of covers with proved Hausdorff and area errors; the completeness of finite families of targets, with an exact second rung; the exchange of witness bodies; a conditional lift certificate whose finitely many boxes are closed by common cores; and the cost of such a certificate — how fine its boxes must be, which prune closes wider boxes, and the depth a positive margin needs. Every step was checked by an independent recomputation and is stated at the strength that check measured. The paper proves no bound on the covering constant itself, and lists the corrections found on the way, among them two audit constants taken from a grid on the unsafe side and a certificate specification with no place for mirror-image placements.


## Contents


0. What is claimed, and what is not . . . . . 2
1. Introduction . . . . . 3
2. From sets of diameter one to unit constant-width bodies . . . . . 4
3. The placement certificate: one target, one cover . . . . . 7
4. A finite dictionary of targets . . . . . 9
5. A finite dictionary of covers . . . . . 12
6. Finite families of targets suffice: saturation . . . . . 15

7. Exchanging witnesses . . . . . 17
8. Lifting, and the certificate a lower bound is written in . . . . . 19
9. What a certificate costs . . . . . 21
10. Every result, and the strength at which it holds . . . . . 25
11. Corrections found on the way . . . . . 27
12. Reproduction and attachments . . . . . 29
13. Contributions . . . . . 30
Acknowledgements . . . . . 30
References . . . . . 30
Appendix A. The figures file . . . . . 31

**Status of this document.** Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. Section 13 records who did which part.

## 0. What is claimed, and what is not

Every statement in this paper carries a status, given after it in italics. A statement with a proof carries a verdict on that proof. `PASS` and `PASS by cited theorem` are different verdicts and are never written as one: the first is a proof checked line by line — against the programme's text, or, for the few proofs this paper supplies itself, as written here — and the second a proof one of whose steps holds because a published or classical theorem does, cited with its hypotheses checked. A verdict is on the proof's own steps: a numbered result it uses keeps its own verdict, listed in §10, so the weakest step of any chain can be read off there. Every other statement carries one of six labels, and the labels are never merged:

| label | meaning |
|---|---|
| `COMPUTED_HERE` | a figure recomputed by the recomputation arm with independently written code — exactly or in directed-interval arithmetic where the text says *exact*, in floating point where it says *numerical* |
| `DERIVED_FROM_VERIFIED_INPUTS` | follows from statements already verified here, with no step of its own beyond combining them |
| `CITED_THEOREM` | someone else's published theorem, stated here and not re-proved; its statement checked against its source |
| `INHERITED` | taken from the research programme's documents and not verified here; every use is marked |
| `UNMEASURED` | well posed, and not established here |
| `FAILED` | attempted and refused: false as stated, or true only in the corrected form given |

A specification is graded by the same scale: complete where it says what a verifier must check, `FAILED` where it omits a case the theorem it serves needs.

**Not claimed anywhere in this paper:** any value or bound for Lebesgue's universal covering constant $a_{\mathrm{Leb}}$ — the universal covering problem itself is open, and this paper is the machinery a certified

bound would pass through, not a bound; the published bounds 0.8344 and 0.8440935944, which are cited, not verified; the target 0.835 of the research programme this framework was built for; and any statement about which finite family of test bodies is best.

## 1. Introduction

### 1.1 The problem

A planar set $U$ is a *universal cover* if it contains a congruent copy of every planar set of diameter at most one. Lebesgue asked in 1914 — the date as Gibbs [3] and Mishra [6] give it — for the least area of a convex universal cover,

$$a_{\mathrm{Leb}} = \inf\{\mathrm{Area}(U) : U \text{ convex, every set of diameter} \le 1 \text{ has a congruent copy in } U\}.$$

The published bracket is

$$0.8344 \le a_{\mathrm{Leb}} \le 0.8440935944,$$

the upper bound by an explicit cover (Gibbs, 2018), the lower bound by a certified computation on three constant-width test bodies (Mishra, 2026); the gap is 0.0096935944, exactly. A monotone hierarchy $\Lambda_1 \le \Lambda_2 \le \cdots \le a_{\mathrm{Leb}}$ with $a_{\mathrm{Leb}} = \lim_M \Lambda_M$ and an $O(M^{-2})$ rate is due to Zeng (2026). These are cited, with their statements checked against their arXiv records, and nothing here depends on their being right.

### 1.2 What is missing, and what this paper supplies

Every lower bound since Brass and Sharifi has the same shape: choose finitely many bodies of diameter one, and show that no convex set of area below some $T$ contains a congruent copy of each. Showing it means searching a continuum — orientations, translations, the shapes of candidate covers — and every published route replaces part of that continuum by finitely many samples. For the bound to be a theorem, each replacement needs a stated error term with a proof.

This paper gives the whole chain, from the problem to finitely many checkable inequalities, with every error term explicit:

- §2 reduces the targets to constant-width bodies and writes containment and rigid motion in support-function terms;
- §3 turns every placement of a fixed target in a fixed cover into a finite linear programme with a two-sided bracket;
- §4 compiles the targets, and §5 the candidate covers, into finite dictionaries with proved Hausdorff and area errors;
- §6 shows that finite families of targets suffice — $a_{\mathrm{Leb}} = \sup_F \Lambda(F)$ — with a Lipschitz stability bound, the second rung of the family ladder exactly, and cell-local area bounds for pruning;
- §§7–9 give the exchange of witness bodies, the conditional lift certificate and the cost layer that make the search finite in practice;
- §10 lists every result with the strength at which it holds, and §11 the corrections found on the way.

Nothing in the chain depends on which test bodies or which target a computation chooses.

### 1.3 Where the mathematics comes from, and how it was checked

The framework was developed in a research programme on this problem (rounds 00–09 of its finite-closure line, and the deep lagged audits of round 04) whose methodology is Neo.K's and whose rounds Aletheia produced. Every statement below was then checked by an independent recomputation (懸衡): finite claims recomputed with independently written code, constructions reproduced, theorem-level steps read against their sources, each check drilled with planted faults. The checks are the runs RUN-002 and RUN-004 to RUN-012 and RUN-041 to RUN-042 of the recomputation tree pinned in the front matter; each result below names the run that checked it.

## 2. From sets of diameter one to unit constant-width bodies

### 2.1 Convex targets suffice

**Proposition 2.1.** *The closed convex hull of a planar set has the same diameter as the set. Hence $U$ is a universal cover if and only if it contains a congruent copy of every compact convex set of diameter at most one.*

The convex-combination estimate is direct. One step the programme leaves implicit closes in a line: the problem takes the infimum over convex $U$ and later work uses closed $U$; but the closure of a convex set is convex, contains whatever the set contains, and has the same area, since the boundary of a planar convex set is Lebesgue-null. *Status:* `PASS` (RUN-002, RUN-004).

### 2.2 Completion to constant width

**Theorem 2.2 (completion).** *Every compact convex set of diameter $d \le 1$ lies in a body of constant width one.* For $d = 1$ this is the classical completion: every planar set of diameter $d$ lies in a complete set of diameter $d$, and in the plane complete sets are exactly the sets of constant width $d$ (Martini, Montejano and Oliveros, *Bodies of Constant Width*, ch. on complete and reduced bodies). For $0 < d < 1$, complete to width $d$ and dilate about a point $c$ of the completion: $c + d^{-1}(x - c)$ maps it onto a body of width one that contains the original set by convexity; a point is contained in any body. *Status:* `PASS by cited theorem` for the completion; the dilation and the point case `PASS` (RUN-004).

**Corollary 2.3.** *$U$ is a universal cover if and only if it contains a congruent copy of every body of constant width one.* Write $\mathcal{W}_1$ for that family.

### 2.3 Support functions, containment and motion

For a compact convex $K$ let $h_K(\theta) = \max_{x \in K} x \cdot u_\theta$, $u_\theta = (\cos\theta, \sin\theta)$.

**Theorem 2.4 (containment).** *For compact convex $A$ and closed convex $U$: $A \subseteq U$ if and only if $h_A \le h_U$ on the whole circle.* One direction is trivial; for the other, a point of $A$ outside $U$ is strictly separated from $U$ by some direction, where $h_A > h_U$. Both hypotheses on $U$ are needed, and each is shown needed by an exact counterexample: for $U$ the *open* unit disc and $A$ the closed one, $h_A = h_U$ yet $A \not\subseteq U$; for $U$ the annulus $\frac{1}{4} \le |x|^2 \le 1$ and $A = \{0\}$, $h_U \ge h_A$ yet $0 \notin U$. *Status:* `PASS by cited theorem` (strict separation of a point from a closed convex set; RUN-002, RUN-004).

**Theorem 2.5 (motion law).** *For a rigid motion $g(x) = Qx + t$ with $Q \in O(2)$, $h_{gK}(u) = h_K(Q^{\mathsf{T}}u) + t \cdot u$.* This is a one-line identity, and it holds on both components of $O(2)$. For a rotation $Q = R_\varphi$ it reads $h_{gK}(\theta) = h_K(\theta - \varphi) + t \cdot u_\theta$ (checked as a symbolic identity in $x, t, \varphi, \theta$). *Status:* `PASS` (RUN-002, RUN-004).

**The reflection branch.** Congruent copies include mirror images, so $Q$ ranges over all of $O(2) = SO(2) \cup R_{\text{ref}} SO(2)$, not over rotations alone. Every certificate below therefore carries a discrete branch $\sigma \in \{+1, -1\}$ beside the angle $\varphi$, and either enumerates $\sigma = -1$ or states why it may skip it — for instance, because the target or the cover is mirror-symmetric. §3.4 shows that the branch is not a formality: there are targets and covers for which ignoring it gives the wrong answer.

**Directed slack.** The least $\varepsilon$ with $A \subseteq U + \varepsilon B$ is $\max\{0, \max_\theta (h_A - h_U)(\theta)\}$, by Theorem 2.4 applied to $U + \varepsilon B$. *Status:* `PASS by cited theorem` (RUN-004).

### 2.4 Constant width in Fourier terms

Let $W \in \mathcal{W}_1$. Constant width one says

$$h(\theta) + h(\theta + \pi) = 1 \quad \text{for all } \theta.$$

**Proposition 2.6.** *(i) If $h$ is a trigonometric polynomial, $h(\theta) + h(\theta + \pi) = 1$ holds if and only if its constant term is $\frac{1}{2}$ and every other even mode vanishes; the odd modes are free. (ii) The Steiner point of $W$ is $(a_1, b_1)$, the first Fourier coefficients of $h$. (iii) A translation by $t$ adds $t \cdot u_\theta$, which cancels antipodally and is annihilated by $D^2 + 1$.* (i) was checked by solving the coefficient system of a general trigonometric polynomial of degree 8; (ii) by the integral $\frac{1}{\pi} \int h \, u_\theta \, d\theta$ at degree 5; (iii) symbolically. *Status:* `PASS` (RUN-004, checks S1–S3).

**The curvature radius.** Write $r = h + h''$, in the sense of measures for non-smooth bodies.

**Proposition 2.7.** *For $W \in \mathcal{W}_1$: $0 \le r \le 1$ almost everywhere, $r(\theta) + r(\theta + \pi) = 1$, and $\int r \cos\theta \, d\theta = \int r \sin\theta \, d\theta = 0$.* The bound $0 \le r \le 1$ is used by the programme without a source; it follows in three lines. Constant width says $h_W + h_{-W} = 1$, so $W + (-W)$ is the unit disc. In the plane $h + h''$ is the body's length measure, nonnegative and additive under Minkowski sums, and the unit disc's is $d\theta$; hence $S_W(d\theta) + S_W(d\theta + \pi) = d\theta$ with $S_W \ge 0$, which forces $S_W = r \, d\theta$ with $0 \le r \le 1$. The moment identities were checked symbolically on a polynomial of degree 6. *Status:* `PASS by cited theorem` (that $h + h''$ is a nonnegative measure for a convex body; RUN-004, S4; RUN-005).

**Theorem 2.8 (curvature coordinates).** *Translation classes of $\mathcal{W}_1$ correspond one to one to the densities*

$$\mathcal{R} = \{r : 0 \le r \le 1,\ r(\theta) + r(\theta + \pi) = 1,\ \textstyle\int r \cos = \int r \sin = 0\}.$$

*Given $r \in \mathcal{R}$, the body is recovered in the Steiner gauge from $h'' + h = r$:*

$$h(\theta) = \tfrac{1}{2} + \sum_{n \text{ odd}, n \ge 3} \frac{\alpha_n \cos n\theta + \beta_n \sin n\theta}{1 - n^2}, \qquad r = \tfrac{1}{2} + \sum (\alpha_n \cos n\theta + \beta_n \sin n\theta).$$

The correspondence rests on the classical criterion that a periodic $h$ with $h + h'' \ge 0$ is a support function (Schneider, *Convex Bodies*); the kernel of the width equation $w'' + w = 1$ is $1 + A\cos\theta + B\sin\theta$, and $\pi$-periodicity of the width forces $A = B = 0$. The inversion was checked for every odd $n \le 11$. *Status:* `PASS by cited theorem` (RUN-004, S5–S6).

The construction was reproduced on the Reuleaux triangle: from its curvature radius alone — $r = 1$ on three arcs of normal directions of length $\pi/3$, 0 between them — through the inversion above (the modes $n = 3m$ for odd $m \le 4001$), against the support function computed from its geometry. Over 3,601 directions the two agree to $4.0 \times 10^{-11}$, which also confirms that its Steiner point is its centroid. *Status:* `COMPUTED_HERE` (numerical; RUN-004, N1).

**Example 2.9 (the $\cos 3\theta$ family).** $h_a = \frac{1}{2} + a\cos 3\theta$ has $r_a = \frac{1}{2} - 8a\cos 3\theta$, so $h_a$ is the support function of a body of constant width one if and only if $|a| \le \frac{1}{16}$; at $a = \frac{1}{16}$, $r = \frac{1}{2} - \frac{1}{2}\cos 3\theta$. *Status:* PASS, the range exact (RUN-004, S7).

**Remark 2.10 (truncation is not safe).** Truncating the Fourier series of a constant-width body keeps the width (only odd modes survive) but can destroy convexity. The Reuleaux triangle's curvature radius is the square wave $\frac{1}{2} + \frac{2}{\pi}\sum_{m\ \text{odd}} \sin(3m\theta)/m$; cut after the first term it is $\frac{1}{2} + \frac{2}{\pi}\sin 3\theta$, whose minimum $\frac{1}{2} - \frac{2}{\pi}$ is negative because $\pi < 4$. So a finite approximation of a target must be built so that it stays a body — the job of the compiler in §4. *Status:* PASS, the witness exact (RUN-004, S8).

### 2.5 Compactness and the worst case

**Proposition 2.11.** *The centred family $\{W \in \mathcal{W}_1 : \text{Steiner point } 0\}$ is compact in the Hausdorff metric. Hence for each closed convex $U$ the worst target is attained, and finite $\varepsilon$-nets of targets exist.* By Blaschke selection (or Arzelà–Ascoli on the support functions) and because the Steiner point lies in the body. The nets are existence only; §4 makes them explicit. *Status:* PASS by cited theorem (RUN-004).

**Proposition 2.12 (stability in the target).** *For fixed $U$, the least directed slack of the best placement, $\Psi_U(K) = \min_g \max_\theta (h_{gK} - h_U)^+$, is 1-Lipschitz in $K$ for the Hausdorff metric, and likewise in $U$: $|\Psi_U(K) - \Psi_V(K)| \le d_H(U,V)$. Status:* PASS (RUN-004).

### 2.6 The best translation at a fixed orientation

**Proposition 2.13.** *For fixed $Q$, the best translation exists, and at a best translation $t^*$ the active directions $\mathcal{A}_Q \subset S^1$ — those where the slack is attained — satisfy $0 \in \operatorname{conv} \mathcal{A}_Q$. Hence a best translation is witnessed by an antipodal pair or by three directions whose convex hull contains the origin.* Existence: $Q$ ranges over a compact group and the objective is coercive in $t$. The active-set condition is the subdifferential of a maximum of affine maps, and Carathéodory's theorem in the plane bounds the witnesses by three. *Status:* PASS by cited theorem (RUN-004).

A single active direction never occurs: $\operatorname{conv}\{u\} = \{u\} \not\ni 0$. The programme listed "one-active degeneracy" among future branch types and later withdrew it; a one-active count at a computed minimiser therefore signals a numerical artefact, not geometry. *Status:* PASS (RUN-004, observation 1).

### 2.7 The canonical statement

Putting §2.1 to §2.6 together:

$$a_{\mathrm{Leb}} = \inf\{\mathrm{Area}(U) : U \text{ closed convex},\ \mathfrak{K}(U) = 0\}, \qquad \mathfrak{K}(U) = \max_{r\in\mathcal{R}} \min_{\sigma,\varphi,t} \max_\theta \left(h_{gW_r} - h_U\right)^+(\theta),$$

where $W_r$ is the body of curvature radius $r$ and $g = (\sigma, \varphi, t)$. Everything that follows makes the max, the min and the inner max finite, each with its error. *Status:* DERIVED_FROM_VERIFIED_INPUTS.

One caution the programme states itself and that holds: it is *not* established that testing the extreme bodies of $\mathcal{W}_1$ (those with $r \in \{0,1\}$ almost everywhere) suffices. A body that is a Minkowski combination of two extreme ones may need a placement that neither needs, since the two placements may differ. *Status:* PASS (RUN-004, on the programme's own caution).

## 3. The placement certificate: one target, one cover

Fix a nonempty compact convex cover $U$ and a target $K \in \mathcal{W}_1$. Both may be translated freely, so take $0 \in U$ and $K$ centred at its Steiner point, which lies in $K$. Let $R_U \geq \max_{x\in U} \|x\|$ (any $R_U \geq \operatorname{diam} U$ is safe) and $R_K \geq \max_{x\in K} \|x\|$ — a radius of a disc about the origin that contains $K$; $R_K = 1$ is always safe, since $0 \in K$ and $\operatorname{diam} K = 1$. Write $L = R_K + R_U$; it is a Lipschitz constant for the slack below, and in this section only.

**Definition 3.1 (signed margin).**

$$M_U(K) = \min_{Q\in O(2)} \min_{t\in\mathbb{R}^2} \max_{\theta} \big[h_K(Q^{\mathsf{T}}u_\theta) + t\cdot u_\theta - h_U(\theta)\big].$$

By Theorem 2.4 and Theorem 2.5, $M_U(K) \leq 0$ if and only if a congruent copy of $K$ lies in $U$, and $M_U(K) < 0$ means it fits with uniform slack.

### 3.1 Two reflection branches

With $S(x_1, x_2) = (x_1, -x_2)$, every $Q \in O(2)$ is $Q_{\sigma,\varphi} = R_\varphi S^{(1-\sigma)/2}$ for a branch $\sigma \in \{+1, -1\}$ and an angle $\varphi$, and

$$h_K(Q_{\sigma,\varphi}^{\mathsf{T}}u_\theta) = \begin{cases} h_K(\theta-\varphi), & \sigma = +1, \\ h_K(\varphi-\theta), & \sigma = -1. \end{cases}$$

$Q_{\sigma,\varphi}$ has determinant $\sigma$ (both checked symbolically). Write $a_{\sigma,\varphi}(\theta) = h_K(Q_{\sigma,\varphi}^{\mathsf{T}}u_\theta) - h_U(\theta)$, $F_{\sigma,\varphi}(t) = \max_\theta[a_{\sigma,\varphi}(\theta) + t\cdot u_\theta]$ and $m_\sigma(\varphi) = \min_t F_{\sigma,\varphi}(t)$; then $M_U(K) = \min_\sigma \min_\varphi m_\sigma(\varphi)$. *Status:* `PASS` (RUN-006, B1).

### 3.2 Translation, exactly

Since $|a_{\sigma,\varphi}| \leq L$, taking $u_\theta = t/\|t\|$ gives $F_{\sigma,\varphi}(t) \geq \|t\| - L$, while $F_{\sigma,\varphi}(0) \leq L$. So a minimiser $t^\star$ exists and every minimiser has $\|t^\star\| \leq 2L$.

**Theorem 3.2 (translation duality).** *Let $\mathcal{P}_0(S^1)$ be the probability measures on the circle with barycentre 0. Then*

$$m_\sigma(\varphi) = \max_{\mu\in\mathcal{P}_0(S^1)} \int a_{\sigma,\varphi}\, d\mu.$$

*Proof.* For $\mu \in \mathcal{P}_0$ and any $t$, $\int a\, d\mu = \int (a + t\cdot u)\, d\mu \leq \max_u (a + t\cdot u)$; minimising over $t$ gives $\leq$. Conversely, let $t^\star$ be a minimiser and $\mathcal{A}$ its active set. By the optimality condition for a maximum of affine maps, $0 \in \partial F(t^\star) = \operatorname{conv}\mathcal{A}$, so there are weights $\lambda_i$ and active directions $u_i$ with $\sum \lambda_i u_i = 0$; for $\mu^\star = \sum \lambda_i \delta_{u_i}$, $\int a\, d\mu^\star = \sum \lambda_i (m_\sigma(\varphi) - t^\star\cdot u_i) = m_\sigma(\varphi)$. □

**Corollary 3.3.** *An optimal dual is supported on two antipodal directions with weights $\frac{1}{2}$, or on three directions not contained in an open semicircle.* By Carathéodory in the plane; a single direction is impossible (§2.6).

*Status:* `PASS by cited theorem` (the optimality condition and Carathéodory). Checked on 168 finite LPs — twelve random instances, both branches, seven orientations each, primal and dual solved as separate programmes: 63 two-active and 105 three-active optima, primal and dual values agreeing to $10^{-9}$ (RUN-006, B2).

### 3.3 Finitely many support directions

Take $M \geq 4$ even, $\theta_j = 2\pi j/M$, so every direction is within $\pi/M$ of the grid. Let

$$m_{\sigma,M}(\varphi) = \min_t \max_{0 \leq j < M} [a_{\sigma,\varphi}(\theta_j) + t \cdot u_{\theta_j}],$$

the value of a linear programme in $(t_x, t_y, z)$ — minimise $z$ subject to $a_{\sigma,\varphi}(\theta_j) + t \cdot u_{\theta_j} \leq z$ — whose dual maximises $\sum_j \lambda_j a_{\sigma,\varphi}(\theta_j)$ over $\lambda \geq 0, \sum \lambda_j = 1, \sum \lambda_j u_{\theta_j} = 0$, with an optimal vertex of at most three nonzero weights. Clearly $m_{\sigma,M} \leq m_\sigma$.

**Theorem 3.4 (support-grid error).**

$$m_{\sigma,M}(\varphi) \leq m_\sigma(\varphi) \leq m_{\sigma,M}(\varphi) + \eta_M, \qquad \eta_M = L\Big(1 + \frac{2}{\cos(\pi/M)}\Big)\frac{\pi}{M}.$$

*Proof.* For $t \neq 0$ some grid direction is within $\pi/M$ of $t/\|t\|$, so the grid objective is at least $\|t\| \cos(\pi/M) - L$, while at $t = 0$ it is at most $L$; hence the grid minimiser has $\|t^\star_M\| \leq 2L/\cos(\pi/M)$. A support function of a body in the disc of radius $R$ is $R$-Lipschitz in the angle, so $a_{\sigma,\varphi}$ is $L$-Lipschitz and $a_{\sigma,\varphi} + t^\star_M \cdot u_\theta$ is $(L + \|t^\star_M\|)$-Lipschitz. Between grid directions it rises by at most $(L + \|t^\star_M\|)\pi/M \leq \eta_M$. □

When a certificate supplies a verified $\|t^\star_M\| \leq T_M$, the sharper $\eta^{\mathrm{loc}}_M = (L + T_M)\pi/M$ may be used. *Status:* `PASS by cited theorem` (LP duality); the translation bound checked on 120 grid LPs, and $\eta_{64} = 0.25349455\ldots$ for the instance of §3.6 reproduced from its record (RUN-006, B3, B5).

### 3.4 Finitely many orientations

Rotating the target changes only its support term, and $K$ lies in the disc of radius $R_K$, so $|m_\sigma(\varphi) - m_\sigma(\psi)| \leq R_K|\varphi - \psi|$. Take $P \geq 2$, $\varphi_\ell = 2\pi\ell/P$, and

$$G_{M,P}(K, U) = \min_{\sigma \in \{\pm 1\}} \min_{0 \leq \ell < P} m_{\sigma,M}(\varphi_\ell).$$

**Theorem 3.5 (placement bracket).** *With $\rho_P = R_K\pi/P \leq \pi/P$,*

$$G_{M,P} - \rho_P \;\leq\; M_U(K) \;\leq\; G_{M,P} + \eta_M.$$

*Proof.* Right: at the sampled orientation $(\hat\sigma, \hat\varphi)$ achieving $G_{M,P}$, $M_U(K) \leq m_{\hat\sigma}(\hat\varphi) \leq m_{\hat\sigma,M}(\hat\varphi) + \eta_M$. Left: at a continuous minimiser $(\sigma^\star, \varphi^\star)$ take the nearest sample $\varphi_\ell$, within $\pi/P$; then $m_{\sigma^\star,M}(\varphi_\ell) \leq m_{\sigma^\star}(\varphi_\ell) \leq M_U(K) + \rho_P$. □

In practice the LP values are enclosed, not computed exactly: with verified intervals $m^-_{\sigma,\ell} \leq m_{\sigma,M}(\varphi_\ell) \leq m^+_{\sigma,\ell}$ — from a feasible primal point, feasible dual weights, or directed rounding — and $G^\pm = \min_{\sigma,\ell} m^\pm_{\sigma,\ell}$,

$$G^- - \rho_P \leq M_U(K) \leq G^+ + \eta_M.$$

**Remark 3.6 (the quantifiers in the error budget).** To show that *K fits*, one sampled orientation that works is a genuine witness, so the upper side pays only for the unsampled support directions, $\eta_M$. To show that it *does not fit*, every orientation must fail, so the lower side pays $\rho_P$. The asymmetry is the quantifier $\forall K\, \exists g$ of the problem itself, carried into the error budget.

*Status:* `PASS by cited theorem` (LP duality; the orientation Lipschitz bound and both sides of the bracket checked line by line). The bracket was attacked: on six instances, the bracket at a coarse resolution ($M = 16$, $P = 12$) and at a fine one ($M = 256$, $P = 96$) must intersect, since both contain the true margin; every fine bracket lay inside its coarse one, and four of the six fine brackets still straddle zero — the undecided case Theorem 3.5 allows (RUN-006, B4).

### 3.5 The reflection branch cannot be skipped

**Proposition 3.7.** *There are a target and a cover for which a certificate enumerating only rotations reports "does not fit", while the target fits by a reflection.* Take the chiral target with $r = \frac{1}{2} + 0.30\cos 3\theta + 0.12\sin 5\theta$ and let $U$ be its mirror image enlarged by 0.002, $h_U(\theta) = h_K(-\theta) + 0.002$. At $\sigma = -1$, $\varphi = 0$, $t = 0$ the slack is $-0.002$ in every direction, so $M_U(K) \le -0.002$. On the rotation branch alone, a rotation by $4\pi/3$ matches the third harmonic, but the fifth, of amplitude $0.12/(5^2 - 1)$ in the support function, cannot be matched; Theorem 3.5's lower side on $\sigma = +1$ alone, with $P = 720$, $M = 512$ and $R_K = 0.54204$, gives $G_+ - \rho_P = 0.0029998 - 0.0023651 = 0.00063 > 0$. *Status:* `COMPUTED_HERE` (numerical: floating-point LPs at tolerance about $10^{-9}$, five orders of magnitude below the gap; RUN-006, B9).

### 3.6 A test instance that could not have failed

The programme tested its pipeline on the ellipse with semi-axes 0.72 and 0.52 and the target $r = \frac{1}{2} + 0.22\cos 3\theta + 0.10\sin 5\theta$, recording $G_{64,48} = -0.02$ and a dense reference $G_{512,192} = -0.02$; both reproduce. But every sampled orientation, in both branches, gives $-0.02$ to thirteen decimals, and the reason is exact: along the minor axis the cover has width 1.04 and any constant-width-one body width 1, so for every placement the larger of the two antipodal slacks is at least $(1 - 1.04)/2 = -0.02$. Conversely a placement exists with margin at most $-0.0199972$, rigorously. The only binding constraint is a width that no motion changes; a pipeline with a wrong reflection law would have passed the same test. A useful test instance needs a cover whose binding constraints are not one antipodal pair. *Status:* `COMPUTED_HERE` (the value exact on the lower side, rigorous on the upper; RUN-006, B6–B7).

## 4. A finite dictionary of targets

§2.5 gives finite $\varepsilon$-nets of targets, but only their existence; and Remark 2.10 shows that the obvious construction, truncating Fourier series, can leave the class. This section builds the net: a finite list of genuine unit constant-width bodies, enumerable, with an explicit Hausdorff error.

### 4.1 Smoothing by a positive kernel

Write $(J * f)(\theta) = \frac{1}{2\pi}\int_{-\pi}^{\pi} J(t) f(\theta - t)\, dt$.

**Theorem 4.1 (positive kernels keep bodies legal).** *Let $J \ge 0$ be even, $2\pi$-periodic, with $\frac{1}{2\pi}\int J = 1$, and let $h$ be the support function of a centred body of constant width one. Then $J * h$ is again the support function of a centred body of constant width one.*

*Proof.* Width: $J*h(\theta) + J*h(\theta+\pi) = J*1 = 1$. Convexity: convolution commutes with differentiation, so $(J * h) + (J * h)'' = J * r \ge 0$, and a periodic function with nonnegative $h + h''$ is a support function. The bound $J * r \le 1$ follows from $r \le 1$ by averaging. Convolution multiplies Fourier coefficients and

creates no frequency, so the first harmonic stays zero. Evenness of $J$ makes smoothing commute with reflection. □ *Status:* `PASS by cited theorem` (the support-function criterion; RUN-005).

Take the Fejér kernel $F_N(t) = \frac{1}{N+1}\big(\sin\frac{(N+1)t}{2}/\sin\frac{t}{2}\big)^2 = \sum_{|k|\le N}\big(1-\frac{|k|}{N+1}\big)e^{ikt}$ and its normalised square

$$J_N = F_N^2/S_N, \qquad S_N = \frac{1}{2\pi}\int F_N^2 = \frac{2N^2+4N+3}{3(N+1)},$$

a nonnegative even trigonometric polynomial of degree $2N$. Write $\mathcal{J}_N h = J_N * h$. *Status:* `PASS` — the Fourier form checked at $2N+3$ points to 40 digits for $N \le 16$, $S_N$ by Parseval symbolically (RUN-005, K1–K2).

### 4.2 The smoothing error

For a centred unit constant-width body, $0 \in K \subseteq \bar{B}(0,1)$, so $0 \le h \le 1$; and $h'' = r - h$ with $0 \le r \le 1$ gives $\|h''\|_\infty \le 1$. For an even kernel,

$$\mathcal{J}_N h(\theta) - h(\theta) = \frac{1}{4\pi}\int_{-\pi}^{\pi} J_N(t)\big[h(\theta+t)+h(\theta-t)-2h(\theta)\big]dt,$$

and the bracket is at most $\|h''\|_\infty t^2$; so $\|\mathcal{J}_N h - h\|_\infty \le \frac{1}{2}\mu_{2,N}$, where $\mu_{2,N} = \frac{1}{2\pi}\int t^2 J_N(t)\,dt$.

**Theorem 4.2 (explicit smoothing error).**

$$\|\mathcal{J}_N h - h\|_\infty \;\le\; E_N := \frac{\pi^2(4N+1)}{2(2N^2+4N+3)(N+1)} = O(N^{-2}).$$

*Proof.* On $(0,\pi]$, $F_N(t) \le N+1$ and, since $\sin(t/2) \ge t/\pi$ by concavity, $F_N(t) \le \pi^2/((N+1)t^2)$. Split $\int_0^\pi t^2F_N^2$ at $t_0 = \pi/(N+1)$: $\int_0^{t_0} t^2(N+1)^2\,dt = \pi^3/(3(N+1))$ and $\int_{t_0}^\pi \pi^4/((N+1)^2t^2)\,dt = \pi^3N/(N+1)^2$. Dividing by $\pi S_N$ gives $\mu_{2,N} \le \pi^2(4N+1)/((2N^2+4N+3)(N+1))$. □

*Status:* `PASS by cited theorem` — it uses $0 \le h \le 1$ and $\|h''\|_\infty \le 1$, which rest on the support-function criterion; the integrals and the assembly symbolic in $N$, the two kernel bounds on a grid for $N \le 64$ with equality at $t = \pi$ for even $N$, the tabulated values $E_8 = 0.1110078205$ to $E_{256} = 0.0001489913$ reproduced (RUN-005, K3–K6). The bound is valid and conservative: against the exact second moment, computed from the rational coefficients of $J_N$,

| $N$ | exact $\mu_{2,N}/2$ | $E_N$ | ratio | Reuleaux triangle, actual error |
|---|---|---|---|---|
| 8 | 0.01932 | 0.1110 | 0.174 | 0.00965 |
| 64 | $3.574\times10^{-4}$ | $2.309\times10^{-3}$ | 0.155 | $2.02\times10^{-4}$ |
| 256 | $2.275\times10^{-5}$ | $1.490\times10^{-4}$ | 0.153 | |

so $E_N$ is about 6.5 times the exact moment bound for large $N$. Two coarse steps cost the factor: $F_N^2$ is bounded by $(N+1)^2$ on all of $[0,t_0]$, and $\|h''\|_\infty \le 1$ is used though a body's $h''$ may be smaller (the Reuleaux triangle's is $1/\sqrt{3}$). Neither affects correctness; only the size of the dictionary (RUN-005, K5, K12). This is not Zeng's $O(M^{-2})$: that bounds the constant's hierarchy, this the approximation of one shape.

### 4.3 Interiorisation and quantisation

$\mathcal{J}_N h$ is finite-dimensional but its coefficients are still real numbers, and rounding them may break convexity (Remark 2.10). Two steps make a finite list.

**Interiorisation.** For $0 < \delta < 1$ let $h_{N,\delta} = (1-\delta)\mathcal{J}_N h + \frac{\delta}{2}$, the support function of $(1-\delta)K_N + \delta B_{1/2}$ — still a centred body of constant width one, now with $\frac{\delta}{2} \le r_{N,\delta} \le 1 - \frac{\delta}{2}$, and within $\frac{\delta}{2}$ of $\mathcal{J}_N h$ since $|\frac{1}{2} - \mathcal{J}_N h| \le \frac{1}{2}$.

**Odd form.** Only the constant term and odd modes from 3 up survive, up to frequency $2N-1$: $h_{N,\delta} = \frac{1}{2} + \sum_{j=1}^{N-1}[a_j \cos k_j\theta + b_j \sin k_j\theta]$, $k_j = 2j+1$ — that is, $2(N-1)$ real coefficients.

**Quantisation.** Round each coefficient to the nearest point of $q\mathbb{Z}$. Width and the Steiner gauge are kept exactly. The curvature changes by at most $q\sum_j(k_j^2 - 1)$, and $\sum_{j=1}^{N-1}((2j+1)^2 - 1) = A_N := \frac{4}{3}(N-1)N(N+1)$; the support function by at most $(N-1)q$.

**Theorem 4.3 (legal quantisation).** *If $qA_N \le \delta/2$, the quantised $\hat{h}$ has $0 \le \hat{h} + \hat{h}'' \le 1$: it is the support function of a centred body of constant width one, and*

$$\|\hat{h} - h\|_\infty \le E_N + \frac{\delta}{2} + (N-1)q.$$

*Status:* `PASS by cited theorem` (the support-function criterion; RUN-005, K7). The interior margin, not an assumption of positivity, is what makes rounding safe.

### 4.4 The dictionary

Since $0 \le h \le 1$, every coefficient has $|a_j|, |b_j| \le \frac{1}{\pi}\int|\cos| = 4/\pi$; after rounding with $q \le \frac{1}{2}$ every coefficient stays in $[-2, 2]$. Let $\mathcal{D}_{N,\delta,q}$ be the set of all $\hat{h} = \frac{1}{2} + \sum_{j=1}^{N-1}[\hat{a}_j \cos k_j\theta + \hat{b}_j \sin k_j\theta]$ with $\hat{a}_j, \hat{b}_j \in q\mathbb{Z} \cap [-2, 2]$ and $0 \le \hat{h} + \hat{h}'' \le 1$. It is finite: $|\mathcal{D}_{N,\delta,q}| \le (\lceil 4/q \rceil + 1)^{2(N-1)}$.

**Theorem 4.4 (a constructive legal net).** *For $N \ge 2$, $0 < \delta < 1$ and $0 < q \le \min\{\frac{1}{2}, \delta/(2A_N)\}$, every centred body $K$ of constant width one has a $\hat{K} \in \mathcal{D}_{N,\delta,q}$ with*

$$d_H(K, \hat{K}) \le \varepsilon_{N,\delta,q} := E_N + \frac{\delta}{2} + (N-1)q.$$

*Given $0 < \varepsilon < 1$, choosing $N$ with $E_N \le \varepsilon/3$, $\delta = \varepsilon/3$ and $q \le \min\{\delta/(2A_N), \varepsilon/(3(N-1))\}$ gives $\varepsilon_{N,\delta,q} \le \frac{5}{6}\varepsilon$. Status:* `PASS by cited theorem` (inherits §4.2's support-function facts; RUN-005, K7–K8). The condition $q \le \frac{1}{2}$ omitted from the parameter choice is implied, since $A_N \ge 8$ and $\varepsilon < 1$ give $q \le \varepsilon/48$.

The compiler was run once by the programme, on $r = \frac{1}{2} + 0.25\cos 3\theta + 0.15\sin 5\theta$ with $N = 8$, $\delta = 0.1$ and $q = 0.025/672$, and every recorded number reproduces: the package's values lie within $10^{-9}$ of the exact ones, always on the safe side, with one displayed digit wrong (§11). Exactly, the smoothing error is 0.00639466922143404 against $E_8 = 0.11100782046215025$, and the total support error 0.0094617496713189537 against the theorem's 0.16126823712881692. *Status:* `COMPUTED_HERE` (exact, from the exact multipliers 140/163 and 108/163 of $J_8$ at the two modes present; RUN-005, K9).

### 4.5 The fixed-cover certificate

The signed margin is 1-Lipschitz in the target: for one placement $(Q, t)$ the two maxima differ by at most $\|h_K - h_L\|_\infty = d_H(K, L)$, and the minimum over placements keeps that. Combining Theorem 4.4 with Theorem 3.5, for each dictionary body $\hat{K}$ let

$$\ell_{\hat{K}} = G^-_{\hat{K}} - \rho_P, \qquad u_{\hat{K}} = G^+_{\hat{K}} + \eta_M,$$

and let $\overline{M}(U) = \max_K M_U(K)$ over all centred bodies of constant width one, so that $U$ is a universal cover if and only if $\overline{M}(U) \le 0$.

**Theorem 4.5 (fixed-cover certificate).**

$$\max_{\hat{K}\in\mathcal{D}} \ell_{\hat{K}} \;\le\; \overline{M}(U) \;\le\; \max_{\hat{K}\in\mathcal{D}} u_{\hat{K}} + \varepsilon_{N,\delta,q}.$$

*Hence $U$ is certified universal if $\max_{\hat{K}} u_{\hat{K}} + \varepsilon_{N,\delta,q} \le 0$, and certified non-universal if $\ell_{\hat{K}} > 0$ for some $\hat{K} \in \mathcal{D}$ — which needs no net error, since $\hat{K}$ is itself a legal target. Otherwise the certificate says "refine" ($N$, $\delta$, $q$, $M$, $P$, or the LP gap).*

The left inequality needs every dictionary body to be a genuine target, which Theorem 4.3 provides; the certificate covers every set of diameter at most one, every rotation, reflection and translation, and every support direction, with every error term named. What it does not cover is the cover itself: $U$ is fixed. That is the next section's job. *Status:* `DERIVED_FROM_VERIFIED_INPUTS` (RUN-005, RUN-006).

## 5. A finite dictionary of covers

What remains infinite after §4 is the cover. This section makes it finite with an explicit Hausdorff and area error, and turns the two finite dictionaries into a bracket for $a_{\mathrm{Leb}}$ itself.

### 5.1 Covers that matter are bounded

**Lemma 5.1 (anchor).** *A universal cover contains a disc of diameter one; since the disc is invariant under rotations and reflections, a translate of $U$ contains $B_{1/2}(0)$. Status:* `PASS` (RUN-007).

Let $\bar{A}$ be any proved upper bound for $a_{\mathrm{Leb}}$; covers of larger area cannot improve the infimum. Gibbs's $\bar{A} = 0.8440935944$ is used below.

**Theorem 5.2 (radius).** *If a convex $U$ contains $B_r(0)$ and has area at most $A$, then $U \subseteq B_{A/r}(0)$. For universal covers of area at most $\bar{A}$: $U \subseteq B_R(0)$ with $R = 2\bar{A} = 1.6881871888$.*

*Proof.* For $x \in U$ with $\|x\| > r$, rotate so $x = (\|x\|, 0)$; the triangle $\mathrm{conv}\{x, (0, r), (0, -r)\}$ lies in $U$ and has area $r\|x\|$. □ *Status:* `PASS` (RUN-007, C1).

The bound is safe but far from sharp, and a sharper one is available. Let

$$\Phi_r(d) = \pi r^2 + r\sqrt{d^2 - r^2} - r^2 \arccos(r/d), \qquad d \ge r,$$

the area of $\mathrm{conv}(B_r \cup \{x\})$ for $\|x\| = d$ — the disc plus two right triangles with legs $r$ and $\sqrt{d^2 - r^2}$, less the sector of angle $2\arccos(r/d)$. It is increasing, with $\Phi_r'(d) = r\sqrt{d^2 - r^2}/d$.

**Theorem 5.3 (sharp radius).** *An anchored convex cover of area at most $\bar{A}$ lies in the disc of radius $R_\# = \Phi^{-1}_{1/2}(\bar{A}) = 0.721238556039142255\ldots$, and the radius is attained within the class.* Since $\mathrm{conv}(B_{1/2} \cup$

$\{x\}) \subseteq U$ for every $x \in U$, $\Phi_{1/2}(\|x\|) \le \bar{A}$; the anchored set $\mathrm{conv}(B_{1/2} \cup \{x\})$ with $\|x\| = R_\#$ attains it — a member of the class, and not itself a universal cover. *Status:* `PASS` — $\Phi_r$ re-derived from the hull's support function $\max(r, d\cos t)$ and against a polygon of $2^{20}$ sides, $R_\#$ at 40 digits (RUN-041, W3; the programme's deep lagged audit of Round 04). Every constant below may use $R_\#$ in place of $2\bar{A}$; with it, $\varepsilon$ and $\beta$ shrink by more than half.

The anchored, area-bounded covers form a compact family in the Hausdorff metric (Blaschke selection; area is continuous). *Status:* `PASS by cited theorem` (RUN-007).

### 5.2 The outer polygon

Take $M \ge 4$ directions $u_j = u_{2\pi j/M}$, $\Delta = \pi/M$, $d_M = 2\sin(\Delta/2)$. Since $B_{1/2} \subseteq U \subseteq B_R$, every support height lies in $[\frac{1}{2}, R]$. Round each height *up* to the grid $\frac{1}{2} + q\mathbb{Z}$: $\hat{h}_j = \frac{1}{2} + q\lceil (h_U(u_j) - \frac{1}{2})/q \rceil$, so $h_U(u_j) \le \hat{h}_j < h_U(u_j) + q$, and let

$$P_{M,q}(U) = \bigcap_j \{x : x \cdot u_j \le \hat{h}_j\}.$$

**Theorem 5.4 (outer compiler).** *(i) $U \subseteq P_{M,q}(U)$; hence if $U$ is universal, so is $P_{M,q}(U)$. (ii) $P_{M,q}(U) \subseteq B_{R_P}$ with $R_P = (R+q)/\cos\Delta$. (iii) $P_{M,q}(U) \subseteq U + \varepsilon_{M,q}B$ with*

$$\varepsilon_{M,q} = q + (R + R_P)\, d_M.$$

*(iv)* $\mathrm{Area}(P_{M,q}(U)) \le \mathrm{Area}(U) + \beta_{M,q}$, $\beta_{M,q} = 2\pi R\varepsilon_{M,q} + \pi\varepsilon_{M,q}^2$.

*Proof.* (i) $x \cdot u_j \le h_U(u_j) \le \hat{h}_j$. (ii) For $x \in P$ some $u_j$ is within $\Delta$ of $x/\|x\|$, so $\|x\|\cos\Delta \le x \cdot u_j \le R + q$. (iii) For a direction $v$ take $u_j$ with $\|v - u_j\| \le d_M$; for $x \in P$, $x \cdot v \le \hat{h}_j + \|x\|\,\|v - u_j\| \le h_U(u_j) + q + R_P d_M$, and $h_U$ is $R$-Lipschitz on the circle, so $h_U(u_j) \le h_U(v) + R d_M$. (iv) Steiner's formula $\mathrm{Area}(U + \varepsilon B) = \mathrm{Area}(U) + \varepsilon\,\mathrm{Per}(U) + \pi\varepsilon^2$ with $\mathrm{Per}(U) \le 2\pi R$, perimeter being monotone under inclusion of convex sets. □

*Status:* (i)–(iii) `PASS`, (iv) `PASS by cited theorem` (Steiner's formula). The three excesses in (iii) were summed symbolically; containment was checked on three covers — the programme's test cover, Pál's hexagon, and a round cover without mirror symmetry — at $M = 128$, $q = 5 \times 10^{-4}$; Steiner's formula against the area from the support function to $4 \times 10^{-7}$ (RUN-007, C2–C5).

The bounds are safe and loose. Measured on those three covers, the actual Hausdorff excess is 12.6 to 136.2 times below $\varepsilon$ with $R = 2\bar{A}$, and 4.3 to 44.6 times below it with the cover's true radius; the tightest case is Pál's hexagon, whose corners make it so. With $R = 2\bar{A}$ the package's $\beta$ as used is 0.906, recomputed here and identical to its record — larger than the area of any cover in the class — which is why Theorem 5.3 matters. A measurement this far inside the bound cannot detect an error in the bound's formula: dropping the $R\,d_M$ term or the $q$ from $\varepsilon$ leaves every measured comparison green, and only the symbolic derivation catches it (RUN-007, C3–C4).

### 5.3 The polygon dictionary and the bracket

With heights restricted to the finite level set $\{\frac{1}{2} + kq : 0 \le k \le K_q\}$, $K_q = \lceil (R - \frac{1}{2})/q \rceil + 1$, the polygons $P(\mathbf{b}) = \bigcap_j \{x : x \cdot u_j \le b_j\}$ form a finite dictionary $\mathcal{P}_{M,q}$, $|\mathcal{P}_{M,q}| \le (K_q + 1)^M$; each contains $B_{1/2}$ and is bounded. Every anchored cover of area at most $\bar{A}$ has its outer polygon in the dictionary (the levels reach $R + q$).

**Theorem 5.5 (finite polygon bracket).** *Let $A^{\mathrm{poly}}_{M,q}$ be the least area of a universal polygon in $\mathcal{P}_{M,q}$. Then*

$$A^{\mathrm{poly}}_{M,q} - \beta_{M,q} \;\le\; a_{\mathrm{Leb}} \;\le\; A^{\mathrm{poly}}_{M,q},$$

*and $A^{\mathrm{poly}}_{M_n,q_n} \to a_{\mathrm{Leb}}$ whenever $M_n \to \infty$, $q_n \to 0$.*

*Proof.* The right side: $A^{\mathrm{poly}}$ is the area of a universal cover. The left: for $\eta > 0$ take a universal $U_\eta$ of area below $a_{\mathrm{Leb}} + \eta$, anchored and bounded; its outer polygon is universal, in the dictionary, and of area at most $\mathrm{Area}(U_\eta) + \beta_{M,q}$. □ *Status:* PASS; it uses Theorem 5.4(iv), which rests on Steiner's formula (RUN-007). No rate is claimed, and none is needed below.

**Corollary 5.6 (the rule a lower bound instantiates).** *If every polygon of $\mathcal{P}_{M,q}$ with area below $T + \beta_{M,q}$ is non-universal, then $a_{\mathrm{Leb}} \ge T$.* By contradiction: a universal $U$ of area below $T$ has a universal outer polygon of area below $T + \beta_{M,q}$. Each such polygon is shown non-universal by Theorem 4.5's lower side — a single dictionary target with a positive certified margin. *Status:* PASS; it uses $\beta$ from Theorem 5.4(iv) (RUN-007). This is the theorem a certified lower bound is an instance of: finitely many polygons, each refuted by finitely many linear programmes, every error named.

### 5.4 Stability, pruning and strict slack

**Proposition 5.7.** *(i) $|\overline{M}(U) - \overline{M}(V)| \le d_H(U,V)$, since $U$ enters each margin only through $-h_U$. (ii) If $U \subseteq P$, then $\overline{M}(P) \le \overline{M}(U)$, and for the outer polygon $\overline{M}(U) - \varepsilon_{M,q} \le \overline{M}(P) \le \overline{M}(U)$: outer approximation never manufactures a false non-universality. (iii) If $P(\mathbf{b})$ is non-universal, so is every cover whose outer polygon is $P(\mathbf{b})$ — the whole cell is pruned at once. (iv) For $\gamma > 0$, $\overline{M}(P + \gamma B) = \overline{M}(P) - \gamma$ at an area cost of at most $2\pi R_P \gamma + \pi\gamma^2$: a universal polygon, inflated by any $\gamma$, acquires the strict slack a finite certificate needs. Status:* (i)–(iii) PASS, (iv) `PASS by cited theorem` (Steiner; RUN-007).

The search for better covers — cutting where the support slack of the placed targets is positive — is therefore safe to run heuristically: a cut proposes a smaller cover, and only Theorem 4.5 applied to the result decides whether it is universal.

### 5.5 Tighter constants from the lagged audits

The programme ran four deep passes over this compiler after it was written, numbered Pass 02 to Pass 05. What holds of them:

- an adjacent-normal error $\varepsilon_{\mathrm{deep}} = 0.013327$ at $M = 128$, $q = 5 \times 10^{-4}$, 84.02 % below Theorem 5.4's at $R = 2\bar{A}$ — most of that gain is the sharp radius of Theorem 5.3, which the pass uses — holding against the compiler's actual outer error on 60 random convex bodies between $B_{1/2}$ and $B_{R_\#}$ (at most 63 % of the bound, median 35 %);
- the area of a polygon with fixed normals is a quadratic form in its support numbers; at the $M = 128$ uniform normals its signature is $(1, 2, 125)$ — one positive eigenvalue, the two translations as zero modes (the Fourier modes $k = 1$ and 127), and 125 negative — and on random fans of $m$ normals always $(1, 2, m - 3)$, the signature Minkowski's inequality predicts.

*Status:* the constants recompute to $10^{-12}$ and the stated identities hold (RUN-042); the error bound's validity on random bodies is COMPUTED_HERE (numerical, sampled). Two constants of Pass 04 are grid samples on the unsafe side of a minimum and a maximum and are corrected in §11; nothing above uses them.

## 6. Finite families of targets suffice: saturation

§5 reduces the problem to finitely many polygons, each to be refuted. This section answers the question a lower bound asks first — is a finite family of test bodies enough in principle? — and gives the cell-local bounds that make a branch-and-bound over covers prune.

### 6.1 The finite-witness value

For a finite family $\mathcal{F} = \{D, K_2, \dots, K_m\}$ of centred unit constant-width bodies containing the disc $D = B_{1/2}$, let

$$\Lambda(\mathcal{F}) = \inf_{g_1,\dots,g_m} \operatorname{Area}\Big(\operatorname{conv} \bigcup_i g_i K_i\Big).$$

A common isometry may fix the disc.

**Theorem 6.1.** *(i) $\Lambda(\mathcal{F}) \le a_{\mathrm{Leb}}$. (ii) The infimum is attained, and equals the least area of a convex set covering every member of $\mathcal{F}$.*

*Proof.* (i) A universal cover $U$ holds some $g_i K_i$ for each $i$, hence their convex hull. (ii) Restrict to the compact class of §5.1. For a fixed target $K$ the covers that hold it form a closed set: if $U_n \to U$ and $g_n K \subseteq U_n$, the $g_n K$ lie in a common disc, so a subsequence $g_n \to g$ and $gK \subseteq U$. The intersection over $\mathcal{F}$ is compact and nonempty, and area is continuous. A placement's hull is a cover, and every cover contains a placement's hull. □ *Status:* `PASS by cited theorem` (Blaschke selection; RUN-008).

### 6.2 Completeness

**Theorem 6.2 (finite families suffice in the limit).**

$$a_{\mathrm{Leb}} = \sup\{\Lambda(\mathcal{F}) : \mathcal{F} \text{ finite}, \ D \in \mathcal{F}\}.$$

*Proof.* "≤" is Theorem 6.1. Suppose $\Lambda(\mathcal{F}) \le A < a_{\mathrm{Leb}}$ for every finite $\mathcal{F}$; then $A \ge \pi/4$, since the disc alone needs that. In the compact class $\mathcal{C}_A$ of anchored convex sets of area at most $A$, the closed sets $S_K = \{U : K \text{ fits in } U\}$ have the finite intersection property, so some $U \in \mathcal{C}_A$ holds every body of constant width one — a universal cover of area at most $A$, a contradiction. □ *Status:* `PASS by cited theorem` (Blaschke selection; RUN-008).

### 6.3 The cardinality ladder

Let $\lambda_m = \sup\{\Lambda(\mathcal{F}) : D \in \mathcal{F}, \ |\mathcal{F}| \le m\}$. Then $\lambda_1 \le \lambda_2 \le \cdots \le a_{\mathrm{Leb}}$ and, by Theorem 6.2, $\lambda_m \to a_{\mathrm{Leb}}$.

**Proposition 6.3 (stability; each rung attained).** *If $d_H(K_i, L_i) \le \varepsilon$ for all $i$, then*

$$|\Lambda(K_1, \dots, K_m) - \Lambda(L_1, \dots, L_m)| \le 4\bar{A}\,\varepsilon + \pi\varepsilon^2.$$

*Hence $\Lambda$ is continuous on the compact space of $m$-tuples and every $\lambda_m$ is a maximum.*

*Proof.* Let $U_K$ be an optimal cover for the $K$-tuple, anchored, of area at most $\bar{A}$. Then $g_i L_i \subseteq U_K + \varepsilon B$, whose area is $\Lambda(K) + \varepsilon \operatorname{Per}(U_K) + \pi\varepsilon^2$ by Steiner. Since $U_K \supseteq B_{1/2}$, Minkowski's formula gives $\operatorname{Area}(U_K) = \frac{1}{2}\int h\, dS \ge \frac{1}{4}\operatorname{Per}(U_K)$, so $\operatorname{Per}(U_K) \le 4\bar{A}$. Exchange $K$ and $L$. □

The programme stated the constant as $4\pi\bar{A}$, from $\operatorname{Per} \le 2\pi R$ with $R = 2\bar{A}$; that holds but is loose by exactly $\pi$. The constant $4\bar{A}$ is essentially attained inside the class: a heptagon circumscribed about the

disc has area $0.84276 \le \bar{A}$ and perimeter exactly four times its area. *Status:* `PASS by cited theorem` (Steiner, Minkowski); the sharper constant is the recomputation's (RUN-008, D5 and N4).

**The first rungs.** $\lambda_1 = \pi/4$: the disc is its own smallest cover. The second rung has an exact value for the family that matters most:

**Proposition 6.4.** $\Lambda(\{D, \mathrm{RT}\}) = \dfrac{\pi}{8} + \dfrac{\sqrt{3}}{4} = 0.8257117836\ldots$, *where* RT *is the Reuleaux triangle of width one; hence* $\lambda_2 \ge \pi/8 + \sqrt{3}/4$.

*Proof.* Concentric placement: the triangle's vertices lie at $1/\sqrt{3}$ from the centre, and the arc through a vertex is tangent there to a line at distance exactly $\frac{1}{2}$ from the centre, so each arc leaves its vertex along a tangent to the disc and then bends inside. The hull is the disc and three tips, each of area $\Phi_{1/2}(1/\sqrt{3}) - \pi/4$ (§5.1), giving $\pi/4 + 3(\Phi_{1/2}(1/\sqrt{3}) - \pi/4) = \pi/8 + \sqrt{3}/4$. Optimality: rotations are absorbed by the disc and the triangle is mirror-symmetric, so only the offset $t$ matters. If some vertex lies at distance $\ge 0.7$, the area is already at least $\Phi_{1/2}(0.7) = 0.83655$. Otherwise two vertices a width one apart, both within 0.7 of the centre, subtend more than 90° at it, while two tip half-widths subtend at most $2\arccos(0.5/0.7) = 88.8°$; so the tips are disjoint and the area is $\pi/4 + \sum_k \mathrm{tip}(|a_k + t|)$, with $\mathrm{tip}(d) := 0$ for $d \le \frac{1}{2}$, since a vertex inside the disc adds nothing. The tip area is convex and increasing in the distance beyond $\frac{1}{2}$ ($\Phi'' = r^3/(s^2\sqrt{s^2 - r^2}) > 0$), and extended by zero it is still $C^1$, convex and nondecreasing, because $\Phi'_{1/2}(\frac{1}{2}) = 0$; and $\sum_k |a_k + t|$ is convex and threefold symmetric, hence least at $t = 0$. Jensen's inequality and monotonicity finish it. □

The argument uses only the three vertices, so it proves Pál's classical value for the disc and the equilateral triangle of side one as well, and shows that at this rung the Reuleaux triangle adds nothing to the triangle it is built on. *Status:* `PASS` — the proof is the recomputation's own, given above; the hull of the sampled bodies agrees with the value to $8 \times 10^{-11}$, and every one of 96 offsets enlarges it, the least by $1.0 \times 10^{-5}$ (RUN-008, W2).

The published bound gives $\lambda_3 \ge 0.8344$ (Mishra's disc, Reuleaux triangle and Reuleaux pentagon, all three mirror-symmetric). *Status:* `CITED_THEOREM`; the achirality checked, residual $\le 3.3 \times 10^{-16}$ (RUN-008, W3).

### 6.4 No false plateau, and what exact closure would take

**Theorem 6.5 (no false permanent saturation).** *If* $\lambda_m < a_{\mathrm{Leb}}$, *then* $\lambda_{m'} > \lambda_m$ *for some* $m' > m$.

The statement already follows in one line from $\lambda_m \to a_{\mathrm{Leb}}$. What the programme's proof adds is the mechanism, and the mechanism is what a search needs: the optimal covers of a maximising family form a compact set $\mathcal{M}_m$, none of them universal; each misses some target $K_U$, and the covers missing a fixed target form an open set; finitely many of these cover $\mathcal{M}_m$. Adding that finite *batch* $K_1, \ldots, K_r$ forces the value up. A single new hard target need not: another optimal cover may hold it. *Status:* `PASS by cited theorem` (Blaschke selection, for the compactness; RUN-008).

**Theorem 6.6 (exact finite saturation).** *If a finite family* $\mathcal{F}_\star$ *and a universal cover* $U_\star$ *have* $\Lambda(\mathcal{F}_\star) \ge A_\star \ge \mathrm{Area}(U_\star)$, *then* $a_{\mathrm{Leb}} = A_\star$. *Status:* `PASS` (RUN-008). Whether such a pair exists — whether some finite family attains $a_{\mathrm{Leb}}$ — is open; the programme states it as a conjecture, and nothing here bears on it. *Status:* `UNMEASURED`.

### 6.5 Cell-local area bounds and pruning

For an outer-polygon support vector $\mathbf{b}$ (§5.3), the *cell* $\mathfrak{C}(\mathbf{b})$ is the set of anchored covers with $\underline{b}_j \le h_U(u_j) \le b_j$, $\underline{b}_j = \max\{\frac{1}{2}, b_j - q\}$; every member lies in $P_{\mathbf{b}}$. Let $R_{\mathbf{b}} = \max_{x \in P_{\mathbf{b}}} \|x\|$.

**Proposition 6.7 (cell bounds).** *For every $U \in \mathfrak{C}(\mathbf{b})$: 1. $d_H(U, P_{\mathbf{b}}) \le \varepsilon_{\mathbf{b}} := q + 2R_{\mathbf{b}} d_M$; 2. $\mathrm{Area}(U) \ge \mathrm{LB}_{\mathrm{outer}} := \mathrm{Area}(P_{\mathbf{b}}) - \mathrm{Per}(P_{\mathbf{b}})\varepsilon_{\mathbf{b}} - \pi\varepsilon_{\mathbf{b}}^2$; 3. $\mathrm{Area}(U) \ge \mathrm{LB}_{\mathrm{env}} := \frac{1}{4}\int_0^{2\pi} \ell_{\mathbf{b}}$, where $\ell_{\mathbf{b}}(\theta) = \max\{\frac{1}{2}, \max_j[\underline{b}_j - R\|u_\theta - u_j\|]\}$ and $R$ is any radius with $U \subseteq B_R$; 4. $\mathrm{Area}(U) \ge \mathrm{LB}_{\mathrm{tip}} := \max_j \Phi_{1/2}(\underline{b}_j)$.*

*Proof.* (1) $h_{P_{\mathbf{b}}}(v) \le b_j + R_{\mathbf{b}} d_M$ and $h_U(v) \ge b_j - q - R_{\mathbf{b}} d_M$ for the nearest grid direction. (2) Steiner's increment with $\mathrm{Per}(U) \le \mathrm{Per}(P_{\mathbf{b}})$. (3) Support functions in $B_R$ are $R$-Lipschitz, so $h_U \ge \ell_{\mathbf{b}}$; Cauchy's $\mathrm{Per} = \int h$ and Minkowski's $\mathrm{Area} \ge \frac{r}{2}\mathrm{Per}$ for $h \ge r = \frac{1}{2}$. (4) $U$ contains the disc and a point at distance at least $\underline{b}_j$. □ *Status:* (1), (4) `PASS`; (2), (3) `PASS by cited theorem` (Steiner, Cauchy, Minkowski; RUN-008, D1–D4).

Two prune rules follow. Against a certified universal incumbent of area $A^+$: a cell whose outer polygon is certified non-universal is pruned whole (Proposition 5.7(iii)); a cell whose best lower bound is at least $A^+$ is pruned whole. A strictly non-universal limit cover cannot survive refinement — once $\varepsilon_n$ is below half its margin, its outer polygon is certified non-universal — and neither can a strictly suboptimal one, since $\mathrm{LB}_{\mathrm{outer}}$ converges to the area. So the cells that survive forever localise at optimal universal covers. That is asymptotic localisation, not finite termination. *Status:* `PASS` (RUN-008).

Measured on 20 cells (five covers, four grids), no bound exceeded the area of any member found, including the smallest members a local search reached; at the coarsest grid the best bound came within $1.3 \times 10^{-3}$ of the smallest member, and $\varepsilon_{\mathbf{b}}$ was 2.8 to 3.2 times tighter than §5's global $\varepsilon$. *Status:* `COMPUTED_HERE` (numerical; RUN-008, N1–N4). $\mathrm{LB}_{\mathrm{env}}$ must be computed as the integral it is — the programme's package used a rectangle rule that overshot the exact integral in four of five cells (§11).

## 7. Exchanging witnesses

Theorem 6.5 says a finite family below the true value can always be improved by a finite batch of new targets. This section makes that quantitative and turns it into a procedure: how much a batch must violate, how far the next optimum must move, and when a certified search is bound to find the batch.

### 7.1 The configuration space

Fix the disc of a family $\mathcal{F} = \{D, K_2, \dots, K_m\}$ at the origin. Every other body's translation $t_i$ satisfies $\mathrm{conv}(D \cup \{t_i\}) \subseteq \mathrm{hull}$, so $\Phi_{1/2}(\|t_i\|) \le \bar{A}$ and $\|t_i\| \le R_\# = 0.7212$ (Theorem 5.3; the programme used $2\bar{A}$; at the threshold 0.835 the same bound gives 0.6956, RUN-010). A common rotation is a gauge, leaving

$$d_m = 3m - 4$$

continuous parameters — five for the disc, the Reuleaux triangle and the Reuleaux pentagon, as in Mishra's setup. For chiral bodies each non-disc body adds a reflection parity, halved by a global reflection: $2^{m-2}$ branches. The hull area is continuous, so the set $\mathfrak{M}(\mathcal{F})$ of optimal covers is compact, and a minimising cover of $\mathcal{F}$ is a placement hull. *Status:* `PASS` — the gauge checked by 20 random common rotations (area invariant to $10^{-12}$) and by turning one body alone (area changed); $d_m$ for $m = 2, \dots, 7$ (RUN-009, W4).

### 7.2 A uniform violation

Let $\overline{M}(U) = \max_K M_U(K)$ over all centred bodies of constant width one; $U$ is universal if and only if $\overline{M}(U) \le 0$, and $\overline{M}$ is 1-Lipschitz in $U$ (Proposition 5.7). If $U$ is not universal, some target does not fit, and since the placement minimum is attained (§3.2) its margin is strictly positive: $\overline{M}(U) > 0$.

**Theorem 7.1 (uniform violation).** *If $\Lambda(\mathcal{F}) < a_{\mathrm{Leb}}$, then*

$$\delta_{\mathcal{F}} := \min_{U \in \mathfrak{M}(\mathcal{F})} \overline{M}(U) > 0.$$

No optimal cover is universal, since its area is $\Lambda(\mathcal{F}) < a_{\mathrm{Leb}}$; $\overline{M}$ is positive on the compact set $\mathfrak{M}(\mathcal{F})$ and continuous. *Status:* `PASS` (RUN-009).

So every optimal cover misses some target not merely, but by at least $\delta_{\mathcal{F}}$ — which is what lets a finite dictionary step in.

### 7.3 Separating batches

**Theorem 7.2 (a dictionary separates).** *If $\Lambda(\mathcal{F}) < a_{\mathrm{Leb}}$ and $0 < \varepsilon < \delta_{\mathcal{F}}$, then for every $U \in \mathfrak{M}(\mathcal{F})$ some $\hat{K}$ in the dictionary $\mathcal{D}_{N,\delta,q}$ of Theorem 4.4 with $\varepsilon_{N,\delta,q} \le \varepsilon$ has $M_U(\hat{K}) \ge \delta_{\mathcal{F}} - \varepsilon > 0$.* Take a worst target of $U$ and its dictionary approximation; the margin is 1-Lipschitz in the target. *Status:* `PASS` (RUN-009).

The sets $\mathcal{N}_K = \{U \in \mathfrak{M}(\mathcal{F}) : M_U(K) > 0\}$ are open, and finitely many of them — one per dictionary body — cover $\mathfrak{M}(\mathcal{F})$. A subfamily $\mathcal{B}$ whose sets still cover is a *separating batch*; its *separation margin* is

$$S_{\mathcal{F}}(\mathcal{B}) = \min_{U \in \mathfrak{M}(\mathcal{F})} \max_{K \in \mathcal{B}} M_U(K).$$

**Theorem 7.3 (exchange).** *If $s = S_{\mathcal{F}}(\mathcal{B}) > 0$, then:* 1. $\Lambda(\mathcal{F} \cup \mathcal{B}) > \Lambda(\mathcal{F})$; 2. *every cover $V$ of $\mathcal{B}$ has $d_H(V, \mathfrak{M}(\mathcal{F})) \ge s$;* 3. $\Lambda(\mathcal{F} \cup \mathcal{B}) \ge \Lambda(\mathcal{F}) + \gamma_{\mathcal{F}}(s)$, *where* $\gamma_{\mathcal{F}}(s) = \min\{\mathrm{Area}(V) - \Lambda(\mathcal{F}) : V \text{ covers } \mathcal{F},\ V \in \mathcal{C}_{\bar{A}},\ d_H(V, \mathfrak{M}(\mathcal{F})) \ge s\} > 0$.

*Proof.* (1) An optimal cover of $\mathcal{F} \cup \mathcal{B}$ of area $\Lambda(\mathcal{F})$ would be in $\mathfrak{M}(\mathcal{F})$, and some $K \in \mathcal{B}$ would not fit it. (2) For $U \in \mathfrak{M}(\mathcal{F})$ take $K_U \in \mathcal{B}$ with $M_U(K_U) \ge s$; $M_V(K_U) \le 0$; the margin is 1-Lipschitz in the cover — and margins do not change when the cover is moved, so the bound holds against every copy of $\mathfrak{M}(\mathcal{F})$. (3) By compactness and continuity, and (2). □ *Status:* `PASS by cited theorem` (Blaschke selection, in (3)); the margin's Lipschitz transfers checked on eight random cover–target pairs, no margin moving by more than the perturbation (RUN-009, D1).

### 7.4 When a certified search must succeed

Optimal covers are rarely known exactly; they are enclosed in finitely many cells $\mathfrak{C}_j$ with representatives $U_j$ and certified radii $\tau_j$, $d_H(U, U_j) \le \tau_j$ on the cell. If a certified lower bound $M_{U_j}(K) \ge \mu_{j,K}$ has $\mu_{j,K} > \tau_j$, then $K$ fails every cover in the cell. A batch that does this for every cell is a certified separating batch, and choosing it is a finite set-cover problem.

**Theorem 7.4 (eventual success).** *If certified cell covers of $\mathfrak{M}(\mathcal{F})$ are available with radii $\tau_n$, target-dictionary errors $\varepsilon_n$ and certificate errors $\zeta_n$, and $\tau_n + \varepsilon_n + \zeta_n < \delta_{\mathcal{F}}/2$, then the search finds a certified separating batch.* It is a semi-decision: it succeeds when the family is not saturated, and says nothing when it is. It assumes the certified cell covers; it does not construct them. *Status:* `PASS` (RUN-009; the construction of the covers is deferred by the programme itself).

**Proposition 7.5 (dominance).** *If $gK \subseteq K'$ for some isometry $g$, then $\Lambda(\mathcal{G} \cup \{K'\}) \geq \Lambda(\mathcal{G} \cup \{K\})$. In particular the regular n-gon $P_n$ lies in the Reuleaux n-gon $B_n$, so replacing polygons by Reuleaux polygons never lowers the bound.* Checked for $n = 3, 5, 7, 9$ through the support functions. *Status:* PASS (RUN-009, W6).

### 7.5 The first rungs, exactly

The procedure was run on the seed of the published bound, and the seed's exchange quantities are exact.

- $\lambda_1 = \pi/4$, and $\Lambda(D, P_3) = \Lambda(D, B_3) = \pi/8 + \sqrt{3}/4$ (Proposition 6.4): **at the second rung the Reuleaux upgrade gains nothing**, because the triangle's arcs never leave the disc's tangent cone. The programme wrote only $\Lambda(D, B_3) \geq \pi/8 + \sqrt{3}/4$.
- The disc's violation is $\delta_{\{D\}} = 1/\sqrt{3} - \frac{1}{2} = 0.0773502692\ldots$: by Jung's theorem the Reuleaux triangle is the disc's worst target, and the certificate LP of §3 gives $M_D(B_3) = 1/\sqrt{3} - \frac{1}{2}$ to $10^{-16}$ (both reflection branches, $M = 3072$).
- The area modulus of §7.3, abstract in general, is exact on this family: $\gamma(s) = \Phi_{1/2}(\frac{1}{2} + s) - \pi/4$, and the seed exchange gains exactly three times it — one tip per vertex. The distance bound of Theorem 7.3(2) is attained by the seed hull.
- **At the third rung the upgrade is strict:** the disc and the regular triangle and pentagon fit in a hull of area 0.8335974, while for the Reuleaux family Mishra certifies at least 0.8344 — a gap of at least 0.0008026. A differential-evolution search over the five parameters of $\{D, B_3, B_5\}$ finds 0.83478094595, Mishra's own located optimum, and nothing below 0.8344; his two published placements rebuild exactly. Hence

$$\Lambda(D, B_3, B_5) \in [0.8344,\ 0.83478095],$$

  the lower end cited, the upper end a placement found here.

*Status:* the seed values COMPUTED_HERE (exact where stated; the search numerical, and a search gives upper bounds only — RUN-009, W1–W5); the lower end CITED_THEOREM (Mishra).

## 8. Lifting, and the certificate a lower bound is written in

§7 says a better family exists and how to look for it. To *certify* that a larger family reaches a threshold $T$, one need not search the joint configuration space of all its members. This section gives the three theorems that avoid it, the common-core lower bound that makes a box of placements finite, and the grammar of a certificate.

### 8.1 Lifting over a threshold

For a base hull $H$ (the hull of a placed family) and a target $K$, the *lift* is

$$J_K(H) = \min_g \operatorname{Area}\big(\operatorname{conv}(H \cup gK)\big) \ \geq\ \operatorname{Area}(H).$$

Write $q$ for a configuration of the base family $\mathcal{F}$ (here $q$ is a placement, not the quantisation step of §§4–5), $H_{\mathcal{F}}(q)$ for the hull of the placed family and $A_{\mathcal{F}}(q)$ for its area. Only the configurations with $A_{\mathcal{F}}(q) < T$ matter: the *threshold atlas* of the sublevel set $\{A_{\mathcal{F}} < T\}$ is enough, and the exact optimum need not be located first.

**Theorem 8.1 (threshold-conditioned lifting).** *Cover the base configurations by finitely many cells $C_j$, each either (A) closed by the base,* $\inf_{q\in C_j} \mathrm{Area}(H_{\mathcal{F}}(q)) \geq T$, *or (B) closed by a witness $K_j$,* $\inf_{q\in C_j} J_{K_j}(H_{\mathcal{F}}(q)) \geq T$. *With $\mathcal{B}$ the witnesses of the type-B cells, $\Lambda(\mathcal{F}\cup\mathcal{B}) \geq T$.* A full placement restricts to some base configuration in some cell, and its hull contains the base hull, or the base hull and the cell's witness. *Status:* PASS (RUN-010).

**Theorem 8.2 (witness switching).** $\Lambda(\mathcal{F}\cup\mathcal{B}) \geq \min_q \max_{K\in\mathcal{B}} J_K(H_{\mathcal{F}}(q))$. For a fixed base, each new body's placement can be minimised separately. Different cells may be closed by different witnesses; no witness has to close everything, and two witnesses never have to be searched jointly. *Status:* PASS (RUN-010).

**Proposition 8.3 (cell transfer).** *If every base hull of a cell is within Hausdorff distance $\tau$ of a representative $H_C$, all containing the disc, and $J_K(H_C) > T + \omega_T(\tau)$ with*

$$\omega_T(\tau) = 4T\tau + \pi\tau^2,$$

*then $J_K(H(q)) \geq T$ on the whole cell.* Were some lift below $T$, its hull $C_q$ would contain the disc, so $\mathrm{Per}(C_q) \leq 4\,\mathrm{Area}(C_q) < 4T$ (Minkowski, as in Proposition 6.3), and $\mathrm{conv}(H_C \cup gK) \subseteq C_q + \tau B$ has area below $T + 4T\tau + \pi\tau^2$ by Steiner. The programme's constant is $4\pi T\tau$, from the radius bound of Theorem 5.2; it holds but is loose by $\pi$, for the same reason as in Proposition 6.3. *Status:* PASS by `cited theorem`; the sharper constant is the recomputation's (RUN-010).

### 8.2 Boxes of placements, and the cores they share

**Lemma 8.4 (motion bound).** *If $K \subseteq B_{R_K}(0)$ and two placements of the same reflection component differ by at most $h_x, h_y$ in translation and $h_\varphi$ in angle, then*

$$d_H(gK, g_0K) \leq \delta = \sqrt{h_x^2 + h_y^2} + 2R_K \sin(h_\varphi/2).$$

This is Mishra's bound for Reuleaux corners, for any radius $R_K$. *Status:* PASS — checked on 1,400 random placements of three Reuleaux polygons and four chiral Fourier bodies, attained to within 95% (RUN-011, W2).

**Theorem 8.5 (common core).** *If every placement $K_g$ in a box has $d_H(K_g, K_0) \leq \delta$ and a compact convex $C$ has $C + \delta B \subseteq K_0$, then $C \subseteq K_g$ for every placement in the box.* In support-function terms, $h_C \leq h_{K_0} - \delta \leq h_{K_g}$. For a Reuleaux polygon $\bigcap_j B(V_j, 1)$ the core is $\bigcap_j B(V_j, 1-\delta)$; in general any $C$ with a certified $h_C + \delta \leq h_{K_0}$ will do, so the interface does not depend on Reuleaux closed forms. *Status:* PASS; Mishra's erosion lemma is its special case (RUN-011, W3).

So the hull of the cores of a box of placements is a lower bound for the hull area over the whole box. To keep the verifier independent of any exact hull routine, it may use verified inner points of each core and of the disc — including each core's corners, without which the bound lost up to $3.7 \times 10^{-5}$ at the hard leaves of §8.4 and up to $9.5 \times 10^{-5}$ on the base cells of §9.5: the hull of finitely many points is a one-sided lower witness, and a leaf is accepted when $A_{\mathrm{inner}} - e_{\mathrm{fp}} \geq T$, with $e_{\mathrm{fp}}$ an independently bounded floating-point error.

### 8.3 The certificate

**Theorem 8.6 (nested certificate soundness).** *Suppose a base tree deterministically partitions a root domain that contains every base placement with hull area below $T$; every base leaf is either* `BASE-PRUNED`

*(its common-core lower bound is at least $T$) or* `LIFT-CLOSED` *by a witness $K$ whose own lift tree covers all of $K$'s relevant placements, every lift leaf carrying a valid common-core lower witness of at least $T$; and every tree stream is consumed with no missing or unclassified node. Then, with $\mathcal{B}$ the witnesses used, $\Lambda(\mathcal{F} \cup \mathcal{B}) \geq T$. Status:* `PASS` (RUN-011, W5 — the tree identity on 200 random forests).

This is the form a certified lower bound takes: finitely many boxes, each closed by finitely many inner points, and the theorem saying that closing them all proves the bound. With Theorem 6.1(i) — $\Lambda(\mathcal{F}) \leq a_{\mathrm{Leb}}$ for every finite family — it is a lower bound for $a_{\mathrm{Leb}}$; Corollary 5.6 is the other route, through the polygon dictionary.

**One gap.** For a chiral witness the lift root must include mirror-image placements. The theorem's text asks for it, but the verifier's invariants, the lift stream and the package's schema have no place for it; the schema also drops five of the eleven record fields the text specifies, `reflection_symmetry` among them. The Reuleaux witnesses the programme uses are all achiral, so its milestones lose nothing; a chiral fallback would. *Status:* as a theorem, `PASS`; as a specification, complete for $B_7$ and every Reuleaux fallback and `FAILED` for chiral witnesses — the run grades the whole `PARTIAL` (RUN-011, the gap).

### 8.4 What the lifts measure on the seed

On the published seed $\{D, B_3, B_5\}$ the fixed-seed lifts rebuild exactly: $J_{B_7} = 0.8371345$, $J_{B_{11}} = 0.8361324$, $J_{B_{13}} = 0.8359056$, $J_{B_9} = 0.8356307$, so the recorded ranking $B_7 > B_{11} > B_{13} > B_9$ holds — but the recorded $J_{B_{13}} = 0.8359676$ is not a minimum: a placement $6.2 \times 10^{-5}$ lower exists. Searches miss minima in both directions; the recomputation's own search missed $B_{11}$'s, and the programme's placement found it. A joint search over the disc and $B_3, B_5, B_7$ finds hull area 0.8364914, so $\Lambda(D, B_3, B_5, B_7) \leq 0.8364914$ — below the fixed-seed lift by $6.4 \times 10^{-4}$: the base moves when a witness is added, which is why the lift is a lower bound and not the value. *Status:* `COMPUTED_HERE` (numerical; a search gives upper bounds on $\Lambda$; RUN-010, W3–W4).

At the threshold $T = 0.835$, the constants of the $B_7$ lift rebuild — $t_7^\star = 0.196935046771$, the increment $2.188 \times 10^{-4}$, the ratio 6.83 — and two hard leaves were bounded as the verifier would bound them, one at the seed with $B_7$ at its best lift and one at the four-body configuration found above: at half-width $10^{-4}$ in all eight coordinates both close (0.8368148 and 0.8361550, by inner points without the cores' corners), and at $10^{-3}$ neither does. How wide a box can close is measured in §9.4. *Status:* `COMPUTED_HERE` (numerical; RUN-011, W1, W4).

## 9. What a certificate costs

§8 says what a certificate is and what it proves; it says nothing about how fine its boxes must be. This section prices it. Every result here concerns the size of a certificate, never what a certificate proves: a certificate built with any of these choices is checked by Theorem 8.6 alone.

### 9.1 How far a cell's hull moves

**Lemma 9.1.** *For compact $X, Y$: $d_H(\operatorname{conv} X, \operatorname{conv} Y) \leq d_H(X, Y)$.* From $X \subseteq Y + \delta B$ follows $\operatorname{conv} X \subseteq \operatorname{conv}(Y + \delta B) = \operatorname{conv} Y + \delta B$, and symmetrically. □

**Corollary 9.2 (the radius of a cell).** *If over a cell $C$ of placements each body $K_i$ stays within $\delta_i$ of its placement at the cell's centre, the hull of the family stays within $\tau(C) = \max_i \delta_i$ of the centre's hull.*

Each $\delta_i$ comes from Lemma 8.4. For the seed $\{D, B_3, B_5\}$ the disc fixes the origin and the rotation of the whole configuration fixes $B_3$'s orientation, so

$$\tau(C) = \max\Big\{\sqrt{h_{x3}^2 + h_{y3}^2},\ \sqrt{h_{x5}^2 + h_{y5}^2} + 2R_5 \sin\tfrac{h_{\varphi 5}}{2}\Big\},$$

and with a witness box added the joint radius is $s = \max\{\tau(C), \delta_K\}$. *Status:* PASS (RUN-012, W1). On random cells inside the 0.835 root, with and without a $B_7$ box, the measured Hausdorff distance never exceeded the radius and reached at most 0.77 of it: the hull moves less than its worst body, because a body moves the hull only where it is extreme.

### 9.2 The sharpened transfer, and the resolution a margin buys

**Theorem 9.3 (transfer).** *Let $A, A_0$ be convex with $d_H(A, A_0) \le \tau$, and let $A$ contain $B_r(0)$. If*

$$\mathrm{Area}(A_0) \ge T + \tfrac{2T}{r}\,\tau + \pi\tau^2,$$

*then* $\mathrm{Area}(A) \ge T$. Suppose $\mathrm{Area}(A) < T$. Since $h_A \ge r$, $\mathrm{Area}(A) = \frac{1}{2}\int h_A\, dS \ge \frac{r}{2}\,\mathrm{Per}(A)$, so $\mathrm{Per}(A) < 2T/r$; and by Steiner's formula $\mathrm{Area}(A_0) \le \mathrm{Area}(A + \tau B) = \mathrm{Area}(A) + \tau\,\mathrm{Per}(A) + \pi\tau^2$, which is below the threshold. □

Only the body whose area is bounded needs to hold the disc. At $r = \frac{1}{2}$ the cost is $4T\tau + \pi\tau^2$, the constant of Proposition 6.3 and Proposition 8.3. The programme's earlier rounds used $4\pi T\tau + \pi\tau^2$, from the radius bound of Theorem 5.2; its later round makes this strengthening itself and calls it a strengthening, not a correction, which is right — both constants are valid. *Status:* PASS by cited theorem (Steiner's formula and Area $= \frac{1}{2}\int h\, dS$; RUN-012, W2). How tight it is: over random hulls of the 0.835 root, with and without $B_7$, the least value of 4 Area / Per was 1.0535, so Per $\le$ 4 Area is close to tight there while the old constant is loose by nearly $\pi$; on random cells the transfer held with at least $1.8\,\tau$ to spare.

**Corollary 9.4 (margin to resolution).** *If a representative's rigorous area lower bound exceeds $T$ by $m > 0$, the representative closes every cell of radius at most*

$$\tau_\star(m, T, r) = \frac{\sqrt{T^2/r^2 + \pi m} - T/r}{\pi} = \frac{m}{\sqrt{T^2/r^2 + \pi m} + T/r}.$$

The second form is the one to compute; the first loses about four digits to cancellation at the sizes below. For scale only, since the margin is a search headroom and not a rigorous one: with $T = 0.835$, $r = \frac{1}{2}$ and $m = 0.001494901$, the height above 0.835 of the four-body optimum that Mishra's repository logs, 0.836494901 — a little above the configuration the joint search of §8.4 found (RUN-010) — $\tau_\star = 4.4738688 \times 10^{-4}$, against $1.4246152 \times 10^{-4}$ under the old constant — a ratio of 3.1404. *Status:* PASS, using Theorem 9.3 (the root solves $(2T/r)\tau + \pi\tau^2 = m$, by computer algebra); the values COMPUTED_HERE (RUN-012, W3).

### 9.3 Which axis to split

The radius bound of Corollary 9.2 tells which single bisection reduces it most, exactly and before any area is computed. At the 0.835 root, $\tau_{\mathrm{root}} = \delta_5 = 0.6046803712$, with $\delta_3 = 0.2755682538$. Halving $x_3$ or $y_3$ leaves the radius at 0.6046803712; halving $x_5$ or $y_5$ brings it to 0.5460899294; halving $\varphi_5$ brings

it to 0.4442456056. So the first cut is $\varphi_5$. Mishra's verifier weights the angle axis by the pentagon's circumradius, $R_5 h_\varphi$, which is the first-order term of $2R_5 \sin(h_\varphi/2) = R_5 h_\varphi - R_5 h_\varphi^3/24 + O(h_\varphi^5)$, and it too cuts $\varphi_5$ first. After that the two rules part; along one path, after 16 cuts, the exact rule reaches a radius of 0.0545 and the weighted rule 0.0556. The weighted rule spends cuts on $B_3$'s axes that leave the radius unchanged — though they still help the core prune of §9.4, which depends on each body's radius separately. *Status:* PASS — the exact rule is optimal among single bisections *for the radius bound*, which is all the programme claims (RUN-012, W4). Which rule gives the smaller certificate is UNMEASURED.

### 9.4 Two prunes, measured

A box can be closed in two ways: by the *representative prune* — Theorem 9.3 applied at the box's centre with the joint radius $s$ — or by the *common-core prune* of Theorem 8.5. A verifier may try them in a cascade (an a-priori bound, then the representative, then the core, then a split) and tag each leaf with the mode that closed it. Both are sound. Which closes more was measured at the two hard leaves of §8.4, as product boxes of half-width $h$ in all eight coordinates. The representative column is the closing test's quantity, the centre's area less $4Ts + \pi s^2$ with $T = 0.835$: the box closes when it is at least $T$. It is not itself a lower bound for the hulls in the box — Theorem 9.3 gives $(A_0 - \pi s^2)/(1 + 4s)$ for that, slightly lower. The core column is a lower bound, the exact area of $\mathrm{conv}(D \cup C_3 \cup C_5 \cup C_7)$:

| leaf | $h$ | centre area | representative | exact core | inner points without corners |
|---|---|---|---|---|---|
| seed, $B_7$ at its best lift | $10^{-4}$ | 0.8371345 | 0.8364865 | 0.8368352 | 0.8368148 |
| | $10^{-3}$ | | 0.8306433 | 0.8342397 | 0.8342258 |
| four-body configuration of §8.4 | $10^{-4}$ | 0.8364914 | 0.8358434 | 0.8361916 | 0.8361550 |
| | $10^{-3}$ | | 0.8300002 | 0.8337265 | 0.8337170 |

The core bound stays below every exact hull sampled in the box, as it must. Both prunes close at $10^{-4}$ and fail at $10^{-3}$. The widest box each closes, found by bisection:

| leaf | representative | exact core | ratio |
|---|---|---|---|
| seed, $B_7$ at its best lift | $3.29 \times 10^{-4}$ | $7.30 \times 10^{-4}$ | 2.22 |
| four-body configuration of §8.4 | $2.30 \times 10^{-4}$ | $5.08 \times 10^{-4}$ | 2.21 |

The core closes boxes more than twice as wide in each of the eight coordinates. The representative prune charges the whole perimeter for a motion of $s$; the core pays nothing where the disc, which never moves, forms the hull. For Reuleaux witnesses both prunes cost one exact hull of the same support pieces, so trying the representative first saves work only where cores are expensive — the certi-

fied cores of Fourier witnesses. Near the frontier it never replaces the core. *Status:* COMPUTED_HERE (numerical; RUN-012, W5) — a measurement at two points, not a count.

**Inner points must include the cores' corners.** The last column is the inner-point bound of §8.2 as first computed here, from points inside each core but not its corners. It is a valid lower bound, but it falls between $0.95 \times 10^{-5}$ and $3.7 \times 10^{-5}$ below the exact core hull; with each core's corners added — exact intersections of two circles — it comes within $1.3 \times 10^{-7}$. A core's corner is often an extreme point of the whole hull. *Status:* COMPUTED_HERE (RUN-012, W5–W6).

### 9.5 Reuse, and a dominance that has nothing to prune

**Theorem 9.5 (core-dominance reuse).** *Let $G(C)$ be the verified common-core hull of a base cell $C$. If $G(C_a) \supseteq G(C_b)$, then for every witness core $Z$,* $\text{Area conv}(G(C_a) \cup Z) \geq \text{Area conv}(G(C_b) \cup Z)$*; so a lift certificate that closes $C_b$ closes $C_a$ with the same witness tree and the same arithmetic.* It is the monotonicity of the hull. The containment is decidable in closed form: merge the breakpoints of the two support envelopes; on each interval $h_a - h_b$ is one sinusoid plus a constant, whose least value is explicit. *Status:* PASS (RUN-012, W6). Built here with exact core pieces — an arc of radius $1 - \delta$ about each corner of the Reuleaux polygon, and the points where adjacent arcs meet — a dense set of interior points comes within $6.2 \times 10^{-7}$ of the exact core hull, and without the corners falls up to $9.5 \times 10^{-5}$ short. Over pairs of base cells the closed-form minimum was attained at the angle it names, and the theorem held in all 56 trials with a $B_7$ core.

Containment held for 12 of 12 pairs with the same centre and the second cell coarser, for 11 of 12 with shifted centres, and for a parent's core inside its child's in only 5 of 12. Refinement does not make cores monotone — a child is eroded less, but around a different centre — so a cache cannot infer containment from ancestry and must check it every time. Nothing is lost: a parent's certificate already covers its children.

**Proposition 9.6 (witness dominance, and its limit).** *If $K \subseteq K'$, then $J_{K'}(H) \geq J_K(H)$ for every base hull $H$. But if $K$ and $K'$ both have constant width one and $h_K \leq h_{K'}$, then $K = K'$.* The first is monotonicity again. For the second, $h_K(\theta) + h_K(\theta + \pi) = 1 = h_{K'}(\theta) + h_{K'}(\theta + \pi)$, so $(h_{K'} - h_K)(\theta) + (h_{K'} - h_K)(\theta + \pi) = 0$, a sum of two non-negative terms. □ So in a pool of constant-width witnesses — the Reuleaux polygons and constant-width Fourier bodies the programme names — dominance identifies only congruent copies, and Pareto pruning of the pool has nothing to remove. It matters for witnesses that are not complete: a triangle inside the Reuleaux triangle, a pentagon inside the Reuleaux pentagon, where the Reuleaux body's lift was never smaller and was larger by up to 0.091. *Status:* PASS (RUN-012, W7; the regular witnesses' widths checked to $2.2 \times 10^{-16}$).

### 9.6 Depth

**Theorem 9.7 (margin to depth).** *Let $S_0 = \max_i\left[\sqrt{h_{xi}^2 + h_{yi}^2} + R_i h_{\varphi i}\right]$ for a cell. After $k$ full cycles of refinement, each bisecting every active parameter once, every cell has radius $\tau_k \leq 2^{-k} S_0$. If every representative the refinement meets has margin at least $m > 0$, every cell closes by the representative prune within $\lceil \log_2(S_0/\tau_\star(m,T)) \rceil$ cycles.* The translation half-widths halve each cycle, and $2R \sin\left(h_\varphi / 2^{k+1}\right) \leq R h_\varphi / 2^k$; then Corollary 9.4. □

The hypothesis is a *uniform* margin over every representative the refinement meets; a margin at one centre promises nothing about the children — no cell around the seed, whose base area is below 0.835, ever closes on the base alone. With a uniform margin, persistent splitting is a fault of the scheduler,

not a mathematical necessity; what can survive refinement is the zero-margin frontier, as in §6.5. At the 0.835 root, $S_0 = 0.6101$ and, at the scale margin of §9.2, 11 cycles are needed and 11 suffice: 55 halvings on every path, so uniform refinement to that depth would make $2^{55}$ cells. The bound is a guarantee, not a plan. *Status:* `PASS`, using Theorem 9.3, with the uniform hypothesis the proof uses (RUN-012, W8); the programme's summary states it for "some representative", which is looser than its own proof (§11).

**What the cost layer does not establish.** How much the exact scheduler, the cache or cross-cell reuse save in a real atlas, and how often reuse applies, are `UNMEASURED`: the measurements here are two points and one path. No bound above 0.8344 is claimed.

## 10. Every result, and the strength at which it holds

Each numbered result is listed with the status it carries in the text and the run that checked it. `PASS` and `PASS by cited theorem` stay apart, and each verdict is on the proof's own steps, as §0 says. A `COMPUTED_HERE` figure is numerical unless the text says exact or rigorous. No result is `INHERITED`. The figures quoted from the programme's records are the recorded values that §11 corrects, each marked where it occurs.

| result | holds as | checked in |
|---|---|---|
| Proposition 2.1 — convex targets suffice | `PASS` | RUN-002, RUN-004 |
| Theorem 2.2 — completion to constant width | `PASS by cited theorem`; the dilation `PASS` | RUN-004 |
| Corollary 2.3 — constant-width targets suffice | `DERIVED_FROM_VERIFIED_INPUTS` | |
| Theorem 2.4 — containment by support functions | `PASS by cited theorem` | RUN-002, RUN-004 |
| Theorem 2.5 — the motion law on both components of $O(2)$ | `PASS` | RUN-002, RUN-004 |
| Proposition 2.6 — constant width in Fourier terms | `PASS` | RUN-004 |
| Proposition 2.7 — $0 \le r \le 1$ | `PASS by cited theorem` | RUN-004, RUN-005 |
| Theorem 2.8 — curvature coordinates | `PASS by cited theorem`; the reconstruction `COMPUTED_HERE` | RUN-004 |
| Example 2.9, Remark 2.10 | `PASS`, exact | RUN-004 |
| Proposition 2.11, Proposition 2.13 — compactness; the best translation | `PASS by cited theorem` | RUN-004 |
| Proposition 2.12 — stability in the target | `PASS` | RUN-004 |
| §2.7 — the canonical statement | `DERIVED_FROM_VERIFIED_INPUTS` | |
| Definition 3.1 and the branch law | `PASS` | RUN-006 |
| Theorem 3.2, Corollary 3.3 — translation duality | `PASS by cited theorem` | RUN-006 |
| Theorem 3.4, Theorem 3.5 — the support-grid error; the placement bracket | `PASS by cited theorem` | RUN-006 |

| result | holds as | checked in |
|---|---|---|
| Proposition 3.7 — the reflection branch is necessary | `COMPUTED_HERE`, numerical | RUN-006 |
| §3.6 — a test instance that could not fail | `COMPUTED_HERE`, exact below and rigorous above | RUN-006 |
| Theorem 4.1 to Theorem 4.4 — legal smoothing; its error; legal quantisation; a constructive net | `PASS by cited theorem` | RUN-005 |
| Theorem 4.5 — the fixed-cover certificate | `DERIVED_FROM_VERIFIED_INPUTS` | RUN-005, RUN-006 |
| Lemma 5.1, Theorem 5.2, Theorem 5.3 — anchor; radius; sharp radius | `PASS` | RUN-007, RUN-041 |
| Theorem 5.4, Proposition 5.7 — the outer compiler | (i)–(iii) `PASS`, (iv) `PASS by cited theorem` | RUN-007 |
| Theorem 5.5, Corollary 5.6 — the polygon bracket; the rule a lower bound instantiates | `PASS`, using Theorem 5.4(iv) | RUN-007 |
| §5.5 — the deep passes' constants | `PASS`; the error bound on random bodies `COMPUTED_HERE` | RUN-042 |
| Theorem 6.1, Theorem 6.2 — finite families suffice | `PASS by cited theorem` | RUN-008 |
| Proposition 6.3 — stability; each rung attained | `PASS by cited theorem` | RUN-008 |
| Proposition 6.4 — the second rung, exactly | `PASS`; the proof is the recomputation's | RUN-008 |
| Theorem 6.5 — no false plateau | `PASS by cited theorem` | RUN-008 |
| Theorem 6.6 — exact saturation | `PASS`; whether it ever occurs, `UNMEASURED` | RUN-008 |
| Proposition 6.7 — cell bounds | (1), (4) `PASS`; (2), (3) `PASS by cited theorem` | RUN-008 |
| Theorem 7.1, Theorem 7.2 — uniform violation; a dictionary separates | `PASS` | RUN-009 |
| Theorem 7.3 — exchange | `PASS by cited theorem` | RUN-009 |
| Theorem 7.4, Proposition 7.5 — eventual success; dominance | `PASS` | RUN-009 |
| §7.5 — the first rungs | `COMPUTED_HERE`; the lower end `CITED_THEOREM` | RUN-009 |
| Theorem 8.1, Theorem 8.2 — lifting; switching | `PASS` | RUN-010 |
| Proposition 8.3 — cell transfer | `PASS by cited theorem` | RUN-010 |
| Lemma 8.4, Theorem 8.5, Theorem 8.6 — motion; common core; the certificate | `PASS` | RUN-011 |
| §8.3 — the certificate's specification | complete for achiral witnesses; `FAILED` for chiral ones | RUN-011 |
| §8.4 — the lifts on the seed | `COMPUTED_HERE`, numerical | RUN-010, RUN-011 |

| result | holds as | checked in |
|---|---|---|
| Lemma 9.1, Corollary 9.2 — the radius of a cell | `PASS` | RUN-012 |
| Theorem 9.3 — transfer | `PASS by cited theorem` | RUN-012 |
| Corollary 9.4 — resolution | `PASS`, using Theorem 9.3 | RUN-012 |
| §9.3 — the exact split rule | `PASS` for the radius bound; its effect on a certificate's size `UNMEASURED` | RUN-012 |
| §9.4 — the two prunes | `COMPUTED_HERE`, numerical | RUN-012 |
| Theorem 9.5, Proposition 9.6 — reuse; dominance | `PASS` | RUN-012 |
| Theorem 9.7 — depth | `PASS`, using Theorem 9.3, with a uniform margin | RUN-012 |

The chain from the problem to finitely many checkable inequalities — the reduction of §2, the certificates of §3 and §4.5, the compilers of §4 and §5, Theorem 6.1(i) and the certificate of §8 — consists of results that hold as `PASS`, `PASS by cited theorem` or `DERIVED_FROM_VERIFIED_INPUTS`. No step of it is numerical. The `COMPUTED_HERE` items inside that range — Proposition 3.7, the instance of §3.6 — are examples that show a hypothesis is needed or a test is weak; they are not steps of the chain.

## 11. Corrections found on the way

A correction is listed when a statement or a recorded figure of the programme is false as written, or true only in a corrected form. None of them changes a result of §§2–9, because each result is stated there in its corrected form. Statements that hold but are weaker than they could be are listed apart, and so are the recomputation's corrections of itself.

### 11.1 Statements and figures that do not hold as written

- **A branch type that cannot occur.** The programme lists one-active minimisers among the branch types of the best translation; the type is empty, since $\mathrm{conv}\{u\} \not\ni 0$ (§2.6). A later round of the programme records the same correction itself. (RUN-004)
- **Two audit constants on the unsafe side of a grid.** Pass 04 bounds a cell of the cover dictionary from inside. It builds a lower envelope $L_\#$ of the support functions of every cover in the cell, and asserts that the shrunken polygon $\lambda_b P$ lies in each of them, where $P$ is the cell's polygon and $\lambda_b = \min_\theta L_\#(\theta)/h_P(\theta)$; the area of $\lambda_b P$ is then a lower bound for every cover in the cell. It also gives the cell a sharp radius, a maximum over directions. It records $\lambda_b = 0.9882053663$ and a sharp radius 0.0059677037, both taken from a grid. The minimum is 0.9882022551, at $\theta \approx 0.81765$, where the recorded $\lambda_b h_P$ lies $1.6\times10^{-6}$ above $L_\#$; so the pass's inclusion $\lambda_b P \subseteq U$ is not established by its own bound at the recorded value. It holds for $\lambda \le 0.9882022551$, and the area bound becomes 0.826801748731 against the recorded 0.8268069548. The maximum radius is 0.0059692967, so the recorded radius is not an upper bound. Nothing in §5.5 uses either. (RUN-042)
- **A figure attributed to the wrong quantity.** Pass 04 gives 0.013060516 as Pass 03's cell radius. That radius is 0.013316971; 0.013060516 is a sample of a different construction — the polygon's exact support in place of the cell's — whose true maximum is 0.0130615798. (RUN-042)

- **A recorded lift that is not a minimum.** $J_{B_{13}}$ is recorded as 0.8359676; its minimum is 0.8359056 (§8.4). The ranking of the witnesses survives. (RUN-010)
- **Chiral witnesses missing from the specification** (§8.3): the lift's specification has no place for mirror placements and drops five of the eleven record fields its text names. The Reuleaux witnesses the programme uses lose nothing by it. (RUN-011)
- **A hypothesis stated too weakly.** The depth bound is stated for a margin at "some representative"; its proof needs a uniform margin over every representative the refinement meets (Theorem 9.7). (RUN-012)

### 11.2 Records and displays, without consequence

- The programme displays the quantised curvature maximum as 0.77710000 …; it is 0.7771000125 …. The legality margin is unaffected. (RUN-005)
- The cover compiler's recorded sanity polygon — area 0.8337053574119659, excess 0.0006459849700081 — is not the polygon its text defines, which has area 0.8336391643823979 and excess 0.0006224936758761. The record is reproduced by moving six on-level heights one level up, from a rounding guard of $10^{-14}$ in the package's script. A later methodology document calls the record independently reproduced; it is a re-run of the same script. (RUN-007, RUN-041)
- The saturation round's package computes $\mathrm{LB}_{\mathrm{env}}$ by a rectangle rule. At $M = 64$ it records 0.7936755224944383, $1.40 \times 10^{-9}$ above the exact integral, and four of its five recorded values exceed the exact one (§6.5). Harmless there, but not a certified value. (RUN-008)
- The lift round's sanity script overstates one body's circumradius. A larger radius only enlarges $\delta$, so the check stays valid. (RUN-011)
- Two formulas in a methodology case study lost their braces in a substitution — meaning $A/r$ and $\pi/M$ — and the package's checks pass both, because they are syntactic. (RUN-041)
- The exchange round's reference list omits two sources its text relies on, Pál's and Gibbs's earlier paper, and cites one preprint by a title its author has since changed. No number depends on either. (RUN-009)
- Pass 05 calls $p$ the "perimeter mode"; it is $\sum b$, not the perimeter. (RUN-042)

### 11.3 Valid, but weaker than needed

- The transfer constants $4\pi\bar{A}$ of the saturation round and $4\pi T$ of the lift round are larger than needed by a factor of $\pi$; the programme's own cost round sharpens them to $4\bar{A}$ and $4T$ (Proposition 6.3, Theorem 9.3).
- Pass 02's adjacency step, 51 (48 at the threshold 0.8344), is safe; 50 (47) suffices, since the quantisation error is strictly below $q$. (RUN-042)
- A methodology case study bounds the curvature coefficients by $|\hat{r}_n| \le \frac{1}{2}$; the domain gives $1/\pi$, attained by the Reuleaux triangle at $n = 3$. (RUN-041)
- **A test that could not have failed** (§3.6): valid, but without power — its only binding constraint is a width that no motion changes. (RUN-006)

### 11.4 The recomputation's corrections of itself

- The inner-point bound first used for the hard leaves omitted the cores' corners and lost up to $3.7 \times 10^{-5}$ by it (§9.4); and the widest closing box near the four-body configuration, first estimated at about $4 \times 10^{-4}$ by extrapolating from two widths, is $5.08 \times 10^{-4}$ with the exact core.

Neither changed a verdict. (RUN-012, correcting RUN-011)

- The recomputation's exact envelope routine cut one grid step early when a crossing fell on a grid angle. It was rewritten; every check still passed and no logged number changed. (RUN-009, correcting RUN-008)
- A first account of the floating-point quotients at two heights of the sanity polygon had one direction wrong; no number was affected. (RUN-007)
- RUN-009 states the third-rung gap of §7.5 as at least 0.000803, rounding up; it is 0.0008026, which is what §7.5 states. Found by this paper's review.

## 12. Reproduction and attachments

**What a reader can check without the archive.** The recomputation tree is private; its public projection is not. Every report, every script and every gate log of its 68 runs, as they stood at the pinned commit, is published with SHA-256 sums at https://amral.evemisslab.com/lebesgue/recompute/ (English: https://amral.evemisslab.com/en/lebesgue/recompute/), one page per document of the programme. The twelve runs this paper names are RUN-002, RUN-004 to RUN-012, RUN-041 and RUN-042. Every load-bearing number in this paper occurs in their gate logs or in `data/paper-figures.json` (Appendix A), which adds them up, counts them, records grid parameters from the run scripts, evaluates the runs' own functions where a run checked a value without logging it, and quotes the programme's records with the line they stand on. The paper's front matter declares these files as its sources, and the standard's checker traces every figure against them alone.

**The instrument can fail, and was made to.** Every check of every run was drilled: faults were planted in its inputs, and each had to be caught by exactly the check named for it, while controls — the unmodified inputs and harmless rewrites — had to stay green. Over the twelve runs of this paper: 215 planted faults, all 215 caught by the named check, 57 controls undisturbed, 141 checks, 3,023 must-stay-green assertions held. Over the whole recomputation tree, of which these runs are part: 1,290 planted faults, 1,290 caught by the named check, 262 controls undisturbed, 665 checks. A check that no plant names is not counted as a check.

**What re-running needs.** The programme's documents and packages are two archives, kept byte for byte; the input manifest `data/SOURCE_MANIFEST.v1.json` gives the size, SHA-256 and git blob id of each and the SHA-256 of every leaf. The archives are not public; the scripts read them in memory and never extract them.

**Commands.** Python 3.14 [11] with mpmath 1.3 [12], numpy 2.4 [13], scipy 1.17 [14] (its HiGHS solver [16] and its Qhull interface [17]) and sympy 1.14 [15]; no network. From the line's directory (`lebesgue-verification-xuanheng/` at the pinned commit), each run is re-run, and its drill, by

```
python code/run012_round09_adaptive_atlas.py
python code/run012_round09_adaptive_atlas.py --drill
```

and likewise for every `code/runNNN_*.py` named above. A run rewrites its gate log under `data/gate-logs/`; a re-run is correct when that log is byte-identical to the published one. `python code/paper_figures.py` rewrites `data/paper-figures.json` and is correct on the same terms. The paper is typeset by the shared build tool of the crystallization repository, which refuses to write a PDF from which a character would be dropped.

## 13. Contributions

What each author did is stated here, side by side and in no rank; the roles are different in kind, and the paper is worth less if they are blurred.

**Aletheia** — the web-side research line; model: ChatGPT, GPT series — produced the rounds of the finite-closure line from which this framework comes, and the deep lagged audits of its cover compiler, on the programme Neo.K set out. Every statement this paper attributes to "the programme" is from those documents.

**Neo.K** set out the research programme: the finite-closure strategy, its target, and the methodology of deep lagged multi-pass verification under which the audits of §5.5 were run.

**懸衡 (Xuán Héng)** — the independent recomputation; model: Claude, Opus series — did the twelve runs this paper names and wrote this paper: every `COMPUTED_HERE` figure rebuilt with independently written code, never by running the programme's scripts; every theorem-level step read against its source; the drills of §12; and the additions marked as the recomputation's where they occur — the counterexamples of Theorem 2.4 and Proposition 3.7, the analysis of the test instance in §3.6, the proof of Proposition 6.4, the sharper constants of Proposition 6.3 and Proposition 8.3, the moment comparison of §4.2, the measurements of §9.4, the second part of Proposition 9.6, and the corrections of §11.

**Review.** Before it was pushed, a separate instance of the same model (model: Claude Opus, version 5.5), run as a read-only subagent of 懸衡's Claude Code session and given only the paper, its run reports, the gate logs and the companion papers, read it; its findings — reporting issues, none mathematical — were applied. 墜衡 (Zhuì Héng) — the lab's proof-side verification arm; model: Claude, Opus series — then reviewed this paper before release: he rederived its arguments and recomputed its figures from the text, independently of this line's code, and found no mathematical error. His two suggestions — the margin function's name in the proof of Theorem 7.1, and the convention for a vertex inside the disc in Proposition 6.4 — were applied.

**Priority, stated plainly.** The framework is the programme's. What is offered as this line's own is narrower: an independent check of every step of it, the strength at which each step stands, the few sharpenings and counterexamples listed above, and the corrections. The theorems it cites are other people's. The programme as a whole — its target 0.835, its ledger of closed and open cells, and where it stopped — is reported in a companion paper [1]; this paper can be read without it.

## Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

## References

1. Aletheia, Neo.K and 懸衡, *A finite-closure programme for Lebesgue's universal covering problem: certified progress, target ledgers, and an independent recomputation*, preprint, AMRAL Research Lab, https://amral.evemisslab.com/en/lebesgue/paper/ (2026).
2. P. Brass and M. Sharifi, *A lower bound for Lebesgue's universal cover problem*, Internat. J. Comput. Geom. Appl. 15(5) (2005), 537–544.

3. P. Gibbs, *An upper bound for Lebesgue's covering problem*, https://arxiv.org/abs/1810.10089 (2018).
4. H. Jung, *Über die kleinste Kugel, die eine räumliche Figur einschließt*, J. Reine Angew. Math. 123 (1901), 241–257.
5. H. Martini, L. Montejano and D. Oliveros, *Bodies of Constant Width: An Introduction to Convex Geometry with Applications*, Birkhäuser (2019).
6. U. Mishra, *Curves of constant width and Lebesgue's covering problem*, https://arxiv.org/abs/2608.30538 (v2, 2026).
7. J. Pál, *Über ein elementares Variationsproblem*, Danske Vid. Selsk. Math.-Fys. Medd. 3(2) (1920).
8. R. T. Rockafellar, *Convex Analysis*, Princeton Mathematical Series 28, Princeton University Press (1970).
9. R. Schneider, *Convex Bodies: The Brunn–Minkowski Theory*, second edition, Encyclopedia of Mathematics and its Applications 151, Cambridge University Press (2014).
10. S. Zeng, *An exact hierarchy for Lebesgue's universal covering constant and a certified 0.834 lower bound*, https://arxiv.org/abs/2609.01284 (2026).
11. The Python Software Foundation, *Python* (version 3.14), https://www.python.org.
12. The mpmath development team, *mpmath: a Python library for arbitrary-precision floating-point arithmetic* (version 1.3.0), https://mpmath.org (2023).
13. C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers et al., *Array programming with NumPy*, Nature 585 (2020) 357–362, doi:10.1038/s41586-020-2649-2; used here at version 2.4.
14. P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland et al., *SciPy 1.0: fundamental algorithms for scientific computing in Python*, Nature Methods 17 (2020) 261–272, doi:10.1038/s41592-019-0686-2; used here at version 1.17.
15. A. Meurer, C. P. Smith, M. Paprocki, O. Čertík et al., *SymPy: symbolic computing in Python*, PeerJ Comput. Sci. 3 (2017) e103, doi:10.7717/peerj-cs.103; used here at version 1.14.
16. Q. Huangfu and J. A. J. Hall, *Parallelizing the dual revised simplex method*, Math. Program. Comput. 10 (2018) 119–142, doi:10.1007/s12532-017-0130-5, the article the HiGHS project asks its users to cite; HiGHS (https://highs.dev) is the solver behind SciPy's linear programming.
17. C. B. Barber, D. P. Dobkin and H. Huhdanpaa, *The quickhull algorithm for convex hulls*, ACM Trans. Math. Software 22(4) (1996) 469–483, doi:10.1145/235815.235821; Qhull, behind SciPy's convex hulls.

*The arXiv records of Gibbs, Mishra and Zeng were checked on 2026-09-24 (RUN-002). Steiner's formula, Minkowski's and Cauchy's formulas, Blaschke's selection theorem and the support-function criterion are used in the forms given by Schneider; linear programming duality, Carathéodory's theorem and the subdifferential of a maximum of affine maps in the forms given by Rockafellar.*

## Appendix A. The figures file

`data/paper-figures.json`, and the script that writes it, `code/paper_figures.py`, are published beside this paper as they stood at the pinned commit. The script runs nothing of the programme and computes no mathematics of its own. Its entries are of five kinds, never mixed: sums, differences and ratios of gate-log fields, each naming the fields it reads; counts over the recomputation's inventory and drill records, for the whole tree and for the twelve runs of this paper; grid sizes and sample counts read from the run scripts as written, so that a changed script cannot leave a stale figure behind; values that a run's own function computes and checks but its gate log does not record, evaluated by importing that function at the run's own inputs; and figures quoted from the programme's records, each with its document, the document's SHA-256 and the line. Re-run in the recorded environment, the file is

byte-identical.