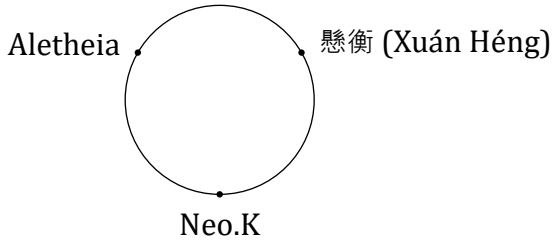
A finite-closure programme for Lebesgue's universal covering problem

Certified progress, target ledgers, and an independent recomputation



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 a_{Leb} of a convex set containing a congruent copy of every planar set of diameter 1; the published bracket is $0.8344 \le a_{\mathrm{Leb}} \le 0.8440935944$. We report on a finite-closure programme that set out to prove $a_{\mathrm{Leb}} \ge 167/200 = 0.835$: reduce the problem to finite families of constant-width test bodies, partition the space of their placements into cells, close each cell by an exact certificate that the hull of the placed bodies is too large, and keep a ledger of what remains. Every finite claim of the programme's 42 documents was recomputed independently, every construction reproduced, and every theorem-level step read against its source, in 68 runs. We state each result at the strength those runs measured. The reduction chain and the certificate machinery hold; exact certificates verified in independent integer arithmetic close several cells of the main atlas, a far-set gate, 644 roots of one branch and part of the hardest root; but the target 0.835 is not certified, eleven of 77 necessity cells are not even geometrically complete, and near root 6217 a configuration of the witness family has hull area $0.8349861105592 < 0.835$ (high-precision floating point). The recomputation also corrects or refutes a number of the programme's statements — among them a claimed stationarity argument, a germ stiffness off by a factor of almost two, and a claim that every other region had been closed by certificate. We list the open problems in the order the programme would have to meet them.

Contents

0. What is claimed, and what is not 2
1. Introduction 3
2. Setting 4
3. Finite witnesses and atlases 5
4. Certificates at scale 5
5. The closure loop 8

6. Local analysis at root 6217 9
7. Overlap decomposition 11
8. Independent recomputation 11
9. Status ledger 13
10. Open problems 15
11. Reproduction and attachments 17
12. Contributions 18
References 18
Appendices 18

Status of this document. Preprint, version 0.1 — the first crystallized paper from this line: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. The contributions section records who did which part, including the parts the recomputation did not do.

0. What is claimed, and what is not

This paper separates six kinds of statement and never merges them.

label	meaning
`COMPUTED_HERE`	recomputed by the recomputation arm with independently written code, never by running the programme's scripts — in exact or directed-interval arithmetic where the text says *certified*, in floating point where it says *numerical*
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities the recomputation already computed
`CITED_THEOREM`	a published theorem, used with its hypotheses checked; the step holds because the theorem does
`INHERITED`	taken from the programme's documents or packages and never verified here — the programme's MPFR replays and audits, and regions closed only in its search logs; every use is marked
`UNMEASURED`	the statement is well posed and the recomputation has not established it
`FAILED`	attempted and refused: the statement does not hold, or holds only in the corrected form given

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one: a step that holds because a classical theorem (Blaschke selection, Fáry–Rédei, linear-programming duality) does is not a step the recomputation proved.

Not claimed anywhere in this paper: any value or new bound for Lebesgue's universal covering constant a_{Leb} — the universal covering problem itself is open, and nothing here closes it; the programme's target $a_{\mathrm{Leb}} \geq 167/200$; the published bounds 0.8344 and 0.8440935944, which are cited,

not verified; any conjecture about the shape of a minimal cover; the stationarity of the root-6217 barrier; and the closure of any region that only a search log, not a certificate, covers.

Inherited and marked throughout: the programme's MPFR replays and dominance audits (recorded, not recomputed), and the parts of the F3 branch and of root 6217 closed only by searches whose trees are not shipped.

1. Introduction

The problem. Lebesgue's universal covering constant a_{Leb} is the least area of a convex set in the plane that contains a congruent copy of every set of diameter 1. Its value is unknown. The best published bounds are $a_{\mathrm{Leb}} \leq 0.8440935944$ [1] and $a_{\mathrm{Leb}} \geq 0.8344$ [2]. Finite certificates in fixed test-set frameworks have given 0.833 (a disc, an equilateral triangle and a regular pentagon) [4] and 0.834 (with an exact hierarchy) [3]; a dual formulation of the problem is in [5]. These results are cited here, not re-proved; a search of arXiv on 2026-10-18 found no improvement of either published bound.

The programme. In September 2026 a research programme on this problem produced 42 documents: a methodology freeze and 37 rounds of finite-closure work, and four methodology tracks with their applications to the problem — a deep-lag verification method, an unbounded-expansion closure loop with seventeen cycles, a variational shape-closure line of eleven papers and a handoff, and an overlap decomposition method with six application documents. Its strategy is the classical one for lower bounds on covering problems, made finite and exact. Every planar set of diameter 1 lies in a body of constant width 1, so a cover must contain congruent copies of finitely many chosen test bodies; the least area of a convex set that does is a lower bound for a_{Leb}, and it increases to a_{Leb} as the family grows. The programme fixes a family — the disc and Reuleaux polygons — partitions the space of their relative placements into cells, and closes each cell by an exact certificate that the convex hull of the placed bodies has area above the target for every placement in the cell. Where the base family cannot close a cell, a further witness is lifted in. Its target was $167/200 = 0.835$, which would improve the published lower bound.

This paper. We do not claim a new bound. We report what of the programme stands, at the strength at which an independent recomputation measured it: every finite claim recomputed from definitions, every construction reproduced, every theorem-level step read against the source it cites, and every certificate re-verified in arithmetic of our own rather than the programme's (§8). §2–§7 follow the programme in its own order; §9 is the ledger — what is proved, what is certified, what is only numerical, what is corrected, what is open — and §10 lists the open problems.

In brief.

- The reduction to constant-width test bodies, the finite shape and placement compilers, the saturation theory and the witness-exchange and lifting theorems hold, in part by cited classical theorems (§2–§3). Two exact values come out on the way: $\Lambda(D, B_3) = \Lambda(D, P_3) = \pi/8 + \sqrt{3}/4$, and $\Lambda(D, B_3, B_5) \in [0.8344, 0.83478095]$ (§3).
- The certificates for cells 0, 47, 48, 68, 75 and 76 of the main atlas — cells the programme counts as closed — check in our own integer arithmetic, independently of the programme's floating-point and MPFR backends, and three more cells carry depth-32 lower bounds checked the same way; a far-set gate holds by a dual certificate; and 644 of the 22,516 roots of the F3 branch, and part of the hardest root, 6217, are closed by certificates whose 56,149 and 182,438 area leaves hold in exact arithmetic (§4–§5, §9.2).

- The headline target is not certified. Eleven of the 77 necessity cells are not geometrically complete; the F3 branch and root 6217 keep open parts; and near root 6217 a barrier configuration has hull area 0.8349861105592, below 0.835 — a high-precision floating-point value, 1.4×10^{-5} below the target — so, numerically, that target cannot be closed there by the witness family in use (§6, §9.5–§9.6).
- Among the statements the recomputation corrects: the barrier's first-order criticality holds numerically but not by the argument given; a germ stiffness is 0.2366, not 0.4543; and the statement that all other regions had been closed by certificate leaves 44 % of root 6217's box closed only by a search whose trees are not kept (§6–§7, §9.4).

2. Setting

The problem. Lebesgue's universal covering constant a_{Leb} is the least area of a convex set that contains a congruent copy of every planar set of diameter 1. The published bracket the programme adopts, $0.8344 \le a_{\mathrm{Leb}} \le 0.8440935944$, is reported correctly — titles, authors, dates, numbers and quoted statements agree with their arXiv records, and the gap 0.0096935944 is exact — but the bounds themselves stand here only as cited results (RUN-002). Round 00 writes the placement group as $E(2)$ but parametrises only rotations and translations; that is immaterial for test bodies with a mirror symmetry and material for chiral ones, and every later certificate has to enumerate the reflection branch where its witnesses lack one (RUN-002, RUN-006).

A note on reading the programme: its own claim-status rule does not describe its documents — only 89 of 758 judgment statements use the ten statuses it prescribes — so the strength of each round's claims has to be read from the round's text (RUN-003). This paper reports the strength the recomputation measured.

The reduction (Round 01). Planar sets of diameter at most 1 reduce to bodies of constant width 1; these are parametrised by curvature-radius densities; containment of a body in a convex set is expressed through support functions; and a placement compiler turns a cover candidate into finitely many containment checks. All of this holds — in part by proofs checked line by line, in part by cited classical theorems (completion to constant width, halfspace representation, Blaschke selection) (RUN-004). The support-function law and the containment interface are correct for proper motions, the interface by the cited strict separation theorem (RUN-002).

Finite shapes (Round 02). The squared-Fejér shape compiler is correct: the smoothing error bound $E_N = \pi^2(4N+1)/\big(2(2N^2+4N+3)(N+1)\big)$ is valid — about 6.5 times the exact moment bound for large N, so not tight — and interiorization and quantization preserve legality (RUN-005).

Placements and covers (Rounds 03–04). For a fixed cover the placement certificate — two reflection branches, P orientations, M-direction linear programmes, the bracket $G-\rho_P \le M_U(K) \le G+\eta_M$ — holds, its key steps by cited linear-programming duality; a chiral pair constructed in the recomputation shows that the reflection branch is necessary (RUN-006). The cover compiler holds: anchored covers lie in radius $R = 2\bar{A} = 1.6881871888$, upward-quantized outer polygons satisfy $U \subseteq P \subseteq U + \varepsilon B$ and $\mathrm{Area}(P) \le \mathrm{Area}(U) + \beta$, giving $A^{\mathrm{poly}} - \beta \le a_{\mathrm{Leb}} \le A^{\mathrm{poly}}$ — a sound but loose rule ($\beta = 0.906$ as used) (RUN-007).

Saturation (Round 05). For a finite family F of test bodies let $\Lambda(F)$ be the least area of a convex set containing a congruent copy of each member. Then $a_{\mathrm{Leb}} = \sup_F \Lambda(F) = \lim_m \lambda_m$, where $\lambda_m = \sup\{\Lambda(F) : |F| \le m\}$ is attained and increases to a_{Leb}, partly by cited compactness and mixed-area results; $\lambda_1 = \pi/4$; $\lambda_3 \ge 0.8344$ only by Mishra's cited theorem; exact finite saturation is an open

conjecture (RUN-008). The recomputation proves in passing that $\Lambda(\{D, \mathrm{RT}\}) = \pi/8 + \sqrt{3}/4 \approx 0.8257$ for the disc D and the Reuleaux triangle RT, and that the continuity constant of the saturation gate can be improved by a factor π (RUN-008).

3. Finite witnesses and atlases

Write D for the unit-width disc, B_n for the Reuleaux n-gon of width 1 and P_3 for the equilateral triangle of side 1.

Witness exchange (Round 06). The witness-exchange compiler holds as an argument given Rounds 02–04; its Theorem 24.1 is conditional on certified covers of the minimizer cells, which it does not construct. On the seed ladder the recomputation proves

$$\Lambda(D, B_3) = \Lambda(D, P_3) = \frac{\pi}{8} + \frac{\sqrt{3}}{4}$$

exactly — Round 06 writes "≥" — and finds $\Lambda(D, B_3, B_5) \in [0.8344, 0.83478095]$, the lower end only by Mishra's cited certificate (RUN-009).

The minimizer atlas (Round 07). The lifting theorems hold — cell-local lifting, witness switching, and robust transfer with cost $4\pi T\tau + \pi\tau^2$, loose by a factor π — and every external number agrees with Mishra's paper and repository. The four-body search gives only $\Lambda(D, B_3, B_5, B_7) \leq 0.8364914$; it finds no counterexample to the gate "≥ 0.835", which remains uncertified (RUN-010).

Conditional lifts and cost (Rounds 08–09). The proof grammar for the 0.835 milestone holds as proved — the motion bound, the common core, Reuleaux erosion, nested soundness — and two hard lift leaves close at half-width 10^{-4}, in floating point and without an error term; the verifier invariants lack a mirror-image root for chiral witnesses, which does not affect the achiral B_7 path (RUN-011). The cost layer is sound: hull non-expansiveness, the sharpened transfer $T + 4T\tau + \pi\tau^2$ (from Per $\leq 4\cdot$Area), the resolution $\tau^\star \approx 4.47 \times 10^{-4}$, the scheduler, core dominance and the depth bound; measured at two leaves, the exact common core closes cells about 2.2 times wider than the representative prune (RUN-012).

The programme's own lagged verification (DLMVC). DLMVC's case studies recompute — its compiler legality, E_M exactly for $M = 3, \ldots, 129$, the coercivity check, $R_{\mathrm{sharp}} = 0.721238556039142255$ and Round 04's constants — and the coefficient bound improves from $1/2$ to $1/\pi$. But its Round 04 case study reports "independently reproduced" for a rerun of the author's own script, which reproduced the script's area 0.8337053574119659, not the defined polygon's 0.833639164382 (RUN-041). Of its four deep passes, all constants recompute except three Pass 04 figures: λ_b and the sharpened cell constant are grid samples on the unsafe side (true values 0.9882022551 and 0.0059692967), and its "Pass 03 cell radius" is a different quantity (RUN-042).

4. Certificates at scale

Rounds 10–37 turn the reduction of §2–§3 into a production line: a box of placements is split until each leaf is certified — the hull of the witness placements over the whole leaf has area above the target — or left pending; certificates are sharded, checkpointed, merged and replayed. This section reports what of that line stands, at the strength each run measured. Throughout, *floating point* means the programme's or the recomputation's double or high-precision floating-point arithmetic without enclosure; *certified* means re-derived in exact or directed integer-interval arithmetic by the recomputation.

4.1 From a local dry run to sharded certificates (Rounds 10–15)

The reference certificate of Round 10 is a local dry run on a box holding 1.1×10^{-18} of the base root's volume: its four lift trees (11,372 nodes, 5,688 leaves) replay and re-emit byte for byte, and with the exact common core every leaf clears 0.835 by at least 3.79×10^{-4} — in floating point, with no derived error term (RUN-013). The frontier and sharding theorems of Round 11 — coverage, replacement, partition merge and $N = 2L - S$ — hold on random frontiers (tested, not proof-checked); the reference shards are Round 10's local certificate regrouped, and nothing global is established (RUN-014). The checkpoint, merge and capacity machinery of Rounds 12–15 works where it was checked: the merge theorems (certified ancestor contraction, atomic merge) hold on 60 random sequences of 80 primitive steps (RUN-016), and the capacity recursion $C(p) = \max(D(p), \min(C(p_0), C(p_1)))$ and its repair agree with exhaustive enumeration on random trees (RUN-018). Two of its outputs do not: stored aggregates in three checkpoints contradict their own ledgers undetected, and the correction rule re-admits a rejected certificate's unchanged bytes (RUN-015); and the local proof's capacity is 0.835, not the recorded 0.83500035, because the B_7 lift root's a-priori exclusion caps it (RUN-017, RUN-018).

4.2 The global pilot and the D_3 quotient (Rounds 16–18)

The global pilot's counts, its scheduler comparison and its decisions rebuild (RUN-019), but only as floating-point measurements, and two of its exact counts depend on choices made in code and not in the text. Its near-frontier profile measures the 40-direction sampling (the exact least centre is 0.8385552, not 0.8356828), and its cycle proxy is not conservative: 19 HARD cells hold configurations below 0.835 (RUN-019). The residual D_3 quotient of Round 17 is exact and its canonical wedge meets every orbit — under the a-priori exclusion $|t_5| \leq t_5(0.8365)$, so the quotient root is sound up to 0.8365; at equal cell size it removes about 73 % of nodes and HARD leaves, not the reported 47.32 % and 32.73 % (RUN-020). Round 18's production checkpoints decode from the text's grammar and re-emit byte for byte; they are pilot forests with open leaves, not bounds (RUN-021).

4.3 Necessity cells and witness escalation (Rounds 19–23)

Round 19's lazy proof cascade is right as stated, its necessity theorem holds (its proof is one line), and its continuations rebuild — in the pilot's floating-point arithmetic; the "witness-only floor" and the switch parameter belong to the 40-direction sampled lift bound (RUN-022). This is where the *necessity cells* enter. Theorem 13.1: a cell holding a placement whose area is below the target has no base-only certificate, and neither does any descendant holding that placement; a *marker* stores the placement and its bound (RUN-022). Of 85 exact centres below 0.835, 77 also have their outer bracket below it (RUN-022), and the 77 cells holding them are the necessity cells of everything that follows: each needs a further witness body lifted in. Round 20's counterexample theorem holds, and its witness B_{23} is disqualified on cell 0 with outer area 0.8349443918, below 0.835 by 5.6×10^{-5} (floating point; RUN-023). The harmonic route does not survive: the legality rules for odd Fourier witnesses hold — odd modes keep width 1; the budget $\sum_k (k^2 - 1)\sqrt{a_k^2 + b_k^2} \leq \lambda/2$ is sufficient (checked on random draws at the full budget) and single-mode sharp; a pure harmonic is legal exactly for $|a| \leq 1/(2(k^2 - 1))$ (RUN-025) — but in floating-point exact-core runs the Reuleaux heptagon B_7 closes every cell tried while no harmonic witness closes any (RUN-023 to RUN-026).

4.4 The B_7 wave over the necessity cells (Rounds 24–27)

The first complete edge is cell 0 with B_7 at lift depth 30: 18,555 nodes, 9,278 certified leaves, none pending (RUN-027). Twelve reference cells then close with a single B_7 witness, rebuilt node for node in floating-point exact-core runs, the least common budget being depth 35, not the recorded 36 (RUN-028). Over all 77 necessity cells, under the package's pad (hull -10^{-10}), 51 close by lift depth 36 — the closure curve at depths 22, 24, …, 36 is 0, 0, 0, 0, 7, 19, 38, 51 — over 4,481,119 nodes, with every leaf still pending its arithmetic certificate; the least slack over all complete trees is 8.47×10^{-11}, not the audited 2.8982×10^{-9} (RUN-029). At a common depth the residual curve continues exactly, 57 at depth 38 and 63 at 40, with 14 cells open at 40 (RUN-030). All of this is floating-point geometry; the recomputation's own interval replay of Round 27 certifies individual leaves only (margins 3.1949×10^{-8} to 4.0147×10^{-7}), not a cell (RUN-030).

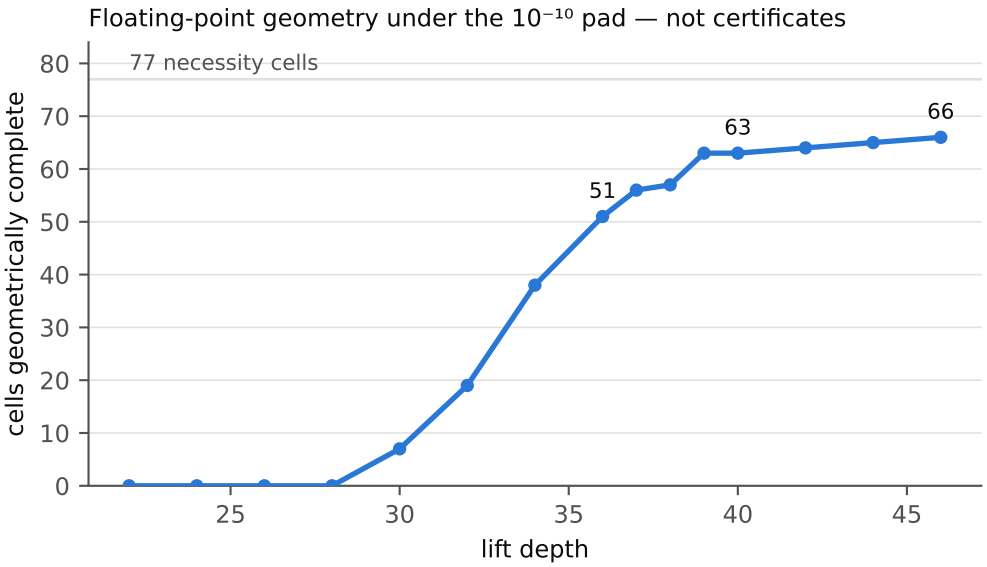


Figure 1: Necessity cells geometrically complete against lift depth: depths 22–36 at a common budget (RUN-029), 37–40 continued at a common depth (RUN-030), 42–46 continued from the frontier at 40 (RUN-035). Floating-point exact-core geometry under the package's 10^{-10} pad; not certificates.

4.5 Exact arithmetic and cross-backend replay (Rounds 28–33)

Round 28 brackets the exact master root — $d^\star$, t_3, t_5, t_7 — and the recomputation's own directed integer intervals, 10^{-90} wide, lie inside the programme's; all 9,278 leaves of cell 0 then check in that arithmetic, least margin $3.1948922738 \times 10^{-8}$ (RUN-031). Each of the 77 necessity markers has an outer polygon below 167/200, certified with 128 to 256 normals once the hull's bridging- edge normals are included; each marker certifies one configuration in its cell (RUN-032). Cell 0's first whole shard — marker 70's outer bound 0.8349075014501105 together with cell 0's lower leaves — checks; its upper half shares the emitter's support code, so implementation-level independence holds for the lower half only (RUN-033). The MPFR backend nests the root in a 2^{-221} dyadic interval and replays the shard, which becomes the programme's first publication candidate (RUN-034). The cross-backend dominance theorem and its polygon corollary are correct as read, and all 9,285 leaves of cell 48 are certified directly, without mpmath or MPFR; the MPFR audits themselves are recorded, not recomputed (RUN-036).

4.6 Certificate production (Rounds 32–37)

Continued under the package's pad from the recomputation's own frontier at depth 40, cells 31, 27 and 7 close at depths 42, 43 and 45 (floating point; RUN-035). The RHCERT format decodes from its specification alone, and the certificates of cells 47 and 48 — 4,272,769 stored vertices — check in integer-interval arithmetic independent of mpmath and MPFR (RUN-037); so do cell 68's (9,255 leaves, 2,126,336 vertices, equal leaf for leaf to the recomputation's own; RUN-038), cells 75 and 76's (2,627,357 and 2,617,252 vertices; RUN-039), and depth-32 lower bounds for cells 49, 64 and 66 (RUN-040). Cell 33 keeps 78,507 open boxes at depth 42, and ten residual frontiers are not recovered (RUN-040). At the end of Round 37 the atlas stands at 66 of 77 cells geometrically complete, 51 of 77 at the only complete common budget, and seven cells closed or better; $a_{\mathrm{Leb}} \geq 0.835$ is not certified (RUN-035, RUN-038, RUN-040).

5. The closure loop

UESFCM is a protocol, not a computation: fix a canonical target, expand without a preset bound, link every artifact back to the target, close each round into a finite state, and recentre on what remains. Its twelve defined parts are defined; which states end the loop is stated differently in eight places; and the canonical target its one application anchors to, with its first state, is not in the archive (RUN-043). The application — seventeen cycles on this problem — chains cleanly with a consistent ledger and never marks evidence as proof (RUN-043).

Cycles 01–06. The cycles' theorems and the numbers they compute from definitions recompute — the B_7 grid linear programme to seventeen digits, the domain constants to the last unit in the last place (RUN-044). Two readings do not: the "sharp root domain" is a remark already in Mishra's paper and saves no boxes, and Cycle 03's separator constant $s_7 \approx 0.00938484$ sits 1.52×10^{-7} above the numerical minimum 0.009384687909 (numerical, not certified; RUN-044). Gate A — the far-set bound $M_{U_0}(B_7) > 0.0088366789091384$ — holds as a finite dual certificate, re-enclosed in 400-bit integer intervals with least lower bound 0.00927495913424859; the shipped verifier does not check two premises the bound rests on (positive sine-gap weights; all 1,024 orientations), which hold when checked (RUN-045).

Cycles 07–11: the F3 frontier. The saved F3 frontier has 22,516 roots; its membership, its least centre area 0.835188375167486, its skeletons and signatures reproduce from the recomputation's own hull, but its Taylor margins are finite-difference diagnostics, not bounds, and the frontier itself comes from a prototype that is not shipped (RUN-046). Two erosion certificates (27 roots and 2,699 nodes; 617 roots and 109,037 nodes) decode exactly, and all 56,149 of their area leaves hold in exact arithmetic with exact cores, least margin 2.28×10^{-5}: 644 of the 22,516 roots are closed by certificate. The split rule is recoverable only from the verifier, the 10^{-18} tiling gaps are closed by a continuity argument the documents state only in words, and Cycle 11's verifier cannot run from its own package (RUN-047).

Cycles 12–17: root 6217. Root 6217's three certificates (124 subroots; a two-dimensional rectangle; eight tail subroots — 66,927 area leaves) hold in exact arithmetic, least margin 2.60×10^{-5}, and the symmetric triple contact reproduces to 80 digits at area 0.8349861105592 (RUN-048). But the $\rho_5 = 0$ reduction drops the F3 constraint, so "Gate 6217-A: closed" is not established: the arc side of that slice is covered by no certificate (RUN-048). Cycles 15–17's lemmas hold — radial convexity by the Fáry–Rédei theorem, the secant lemmas, the two-colour criterion, prefix-tree closure — and all 115,511 area leaves of their certificates hold in exact arithmetic (least 2.66×10^{-5}), taking the root's open

cells from 68 to 61 (RUN-049). The statement that "all other regions have been eliminated by finite certificates" is false: 44 % of root 6217's box was closed only by a search whose trees are not shipped (RUN-049).

6. Local analysis at root 6217

The LESR papers attack the region the closure loop could not close by describing its least-area shape.

The calculus (Papers 00–05). Every formula of Paper 00 holds — the area, its first and second variations, the translation kernel, the Fourier multiplier $1-m^2$ — but its No-Slack Boundary Principle needs the hypothesis $\mu_h(I) > 0$: at a corner the first variation of an admissible inward move vanishes and area drops only at second order, $A(\varepsilon) = 1 - (3\pi/8)\varepsilon^2$ for the unit square (RUN-050). Paper 01's two theorems hold on the problem's own kernel: each pair of support pieces crosses at most twice per turn independently of placement, so signatures are finite; and the hull area is analytic within a contact stratum, with a kink of 3.7×10^{-6} across Cycle 12's contact switch (RUN-051). Paper 02's group-layer facts hold, and on the kernel's hull no pure mode $m \geq 2$ is admissible; the mirror-slice polish of the configuration $q^{\dagger}$, area 0.8347809458213 (floating point), is a strict local minimum (RUN-052). The variational calculus of Paper 03 holds symbolically and on bodies, but on the kernel — where the container is the hull of its witnesses — the affine first-order model is wrong: its linear programme reports optimum 0 at a generic placement whose area gradient has length 0.390, so the finite-mode closure certificate built on it is not sound there (RUN-053). The same placement stops Paper 04's projected flow, which never starts though the area falls at 0.390 per unit along steepest descent (RUN-054), and Paper 05's finite dual certificate, which exists at every placement solved, including that one — so it does not mean stationarity; its smooth No-Slack theorem holds with the curvature hypothesis Paper 00 lacked (RUN-055).

The barrier (Papers 06–07). The barrier shape at root 6217 reproduces: $\theta^{\star} \approx 0.454045$, hull area 0.834986110559, stabilizer exactly D_1; four-contact dual stationarity at the barrier is not merely open but impossible once $\cos 2\theta$ is in the library (RUN-056). Every printed constant of the barrier — $\theta^{\star}$, $x_5^{\star}$, $\alpha^{\star}$, the bridge lengths, $b^{\star}$, g_3, the area to 65 digits — and its 13-interval signature reproduce from the contact equations; the translation KKT condition is not established as argued — the stated subgradient $(g_3, 0)$ fails — though numerically the barrier minimizes on $x_3 \leq x_3^{\star}$ with least feasible rate 0.0354 and multiplier about 0.047 (measured), not 0.1228 (RUN-057).

The rotation (Papers 08–09B). Paper 08's two-gate reduction and its arithmetic reproduce ($A_{\mathrm{ref}} - m(0) = 3.8944077590 \times 10^{-7}$; the split 10 / 141 / 44 of its 195 roots), but the gates are open, holding at 26 profile samples; first-order criticality of the rotation holds numerically, not by the argument given; and the 195-root ledger omits 14.6 % of the box (RUN-058). Paper 09A's blow-up chart and its branch velocity limit $(5.157, -0.1892, 0)$ reproduce in floating point; numerically $\Phi_1 \geq 0$ on a 280-point grid with a single zero; none of its four certificates is established (RUN-059). Paper 09B's six wall constants, $u^{\star} = 5.15712624526841370759\ldots$ and $\Phi_1^{\mathrm{primary}} = \gamma_w w$ with $\gamma_w = 0.18389963923727719444\ldots$ reproduce in high-precision floating point; the full Φ_1 is non-negative along $-w$ by a margin of only 0.0354 per unit, and the second-order slivers open at $\pi/2$, $\theta^{\star}$ and $2\pi - \theta^{\star}$ (RUN-060). The midpoint handoff's recorded state — the barrier to 34 decimals, its area to 37 — is right; its listing of B_3's subgradient as done is wrong; and "evenness implies $m'(0) = 0$" is an invalid inference, though $m'(0^{+}) = 0$ holds numerically (RUN-061).

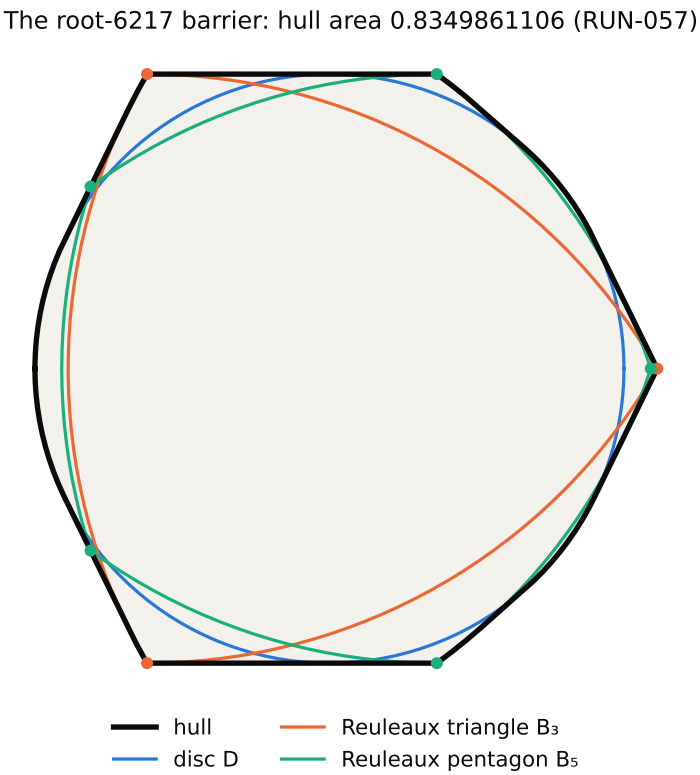


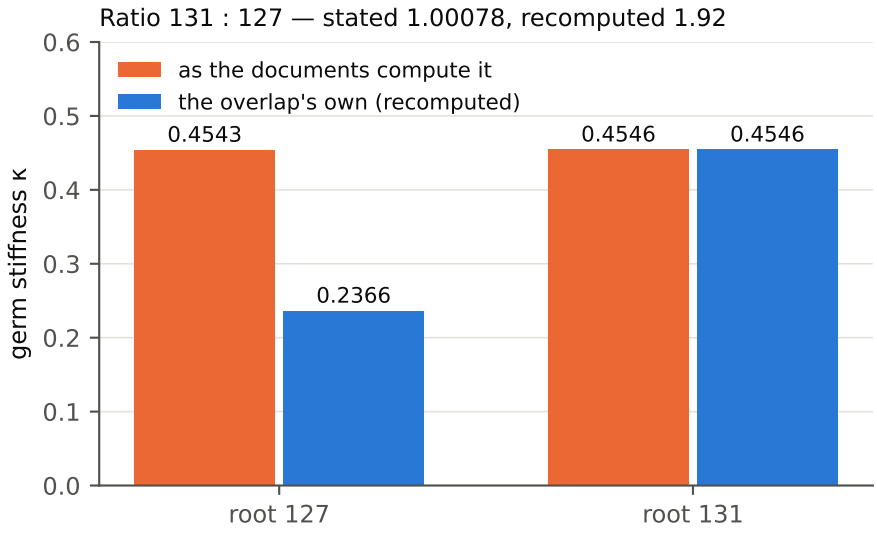
Figure 2: The root-6217 barrier: the convex hull of the disc D of width 1, the Reuleaux triangle B_3 translated to $x_3^\star$ and the Reuleaux pentagon B_5 translated to $x_5^\star$ along the first axis (RUN-057). Its area, 0.8349861105592 ..., is a high-precision floating-point value, not an enclosure; the drawing samples the three boundaries and checks the sampled hull's area against it.

7. Overlap decomposition

BCODR computes nothing of its own; where its mathematics can be instantiated it holds — D_μ and D_J are metrics on discs, and the arrangement has $O(N^2)$ cells, tight at $N(N-1)+2$ (RUN-062). Its application to root 6217 starts from the 195-root ledger, which omits 14.6 % of the box (RUN-062).

LUC 01's arrival-time field over the 195 cells rebuilds cell for cell, with cell 131 the unique zero-radius nexus holding the whole critical ray; on the exact dyadic grid the active and pruned boxes touch in 64 pairs, not 35, and the frontier is mostly set by the ray's cut at $\rho = 0.003$ (RUN-063). LUC 02's 288 children, their signatures and the 17 recomposed classes rebuild exactly; root 127's provenance "flip" from 144 to 155 is a tie (exactly 1/8 from both), and its "thick corridor" is a box face meeting a point (RUN-064). LUC 03's true overlap times rebuild to 40 digits; the seven outward-limited gaps are exactly zero, and 101, 103 and 104 are exactly provenance-degenerate, a four-way tie, not "close to" one (RUN-065).

At root 127, LUC 04's meeting point, rank and critical tangent space rebuild at 50 digits, and its straight-line stiffness is 0.45429481026314; but the overlap's actual second-order law is the Lagrangian $\alpha Q_+ + (1-\alpha)Q_-$, whose softest stiffness is 0.2366, so the overlap lens grows like $1.036\sqrt{\varepsilon}$, not $0.7476\sqrt{\varepsilon}$, and a stiffness target above 0.4 would fail (RUN-066). At root 131 the stiffness is 0.4546, so the two germs' ratio is 1.92, not 1.00078; and because the inward seed is a box, the stiffness cannot be negative, so the "unstable" class of the isolation rule cannot occur (RUN-067).



LUC 07's midpoint separation theorem and mirror algebra hold; the zero-rotation barrier is exactly up/down symmetric, body by body, hence achiral; and the candidate vector the document calls unrecorded is in Cycle 17's diagnostic, with the barrier at 131's corner 1.379×10^{-7} below the candidate (RUN-068). Throughout these documents no area is certified.

8. Independent recomputation

8.1 What was recomputed

The programme is documented in 42 canonical documents: 38 rounds (Round 00, a methodology freeze, and Rounds 01–37) and four methodology tracks with their applications — DLMVC (deep-lag multi-pass verification), UESFCM (unbounded-expansion self-referential finite closure, with seven-

teen cycles on this problem), LESR (from skew fields to variational shape closure, Papers 00–09B and a midpoint handoff) and BCODR (bidirectional circular overlap decomposition and recomposition, with its LUC documents 01–05 and 07). The documents and their packages — 1,459 leaf files, 401.6 MB — are archived byte for byte; every leaf's SHA-256 is recorded, and the recomputation reads the archive in memory without extracting it, so the material it checks cannot drift (Appendix D).

Each document was taken in its own order and given one run (RUN-001 to RUN-068). A run does three things:

1. **Recompute every finite claim independently.** Numbers are rebuilt from the definitions the document states — in exact rational or integer arithmetic where the claim is exact, at stated precision (up to 90 digits) where it is not. The programme's own scripts are never run. A package's code is read only after the run's own reconstruction exists, and then only to find where the two differ.
2. **Reproduce every construction**: trees, frontiers, certificates and their byte formats are rebuilt from the specification the document gives, and a shipped certificate is accepted only when the rebuilt one equals it, or when each of its leaves is re-verified on its own box.
3. **Read every theorem-level step against the source it cites**, instantiating each hypothesis.

8.2 Verdicts

Verdicts are kept apart and never merged into one grade: `PASS`, `PASS by cited theorem` (the step holds because a named published result does, which is cited and not re-proved), `PARTIAL`, `UNMEASURED` (the run could not measure it; a real verdict, often the honest one) and `FAIL`. Where a document's statement holds but not as stated — a different constant, a narrower range, a different reason — the run records both versions. §9 sorts every item by the strength at which it was measured.

8.3 Checks on the checks

A run's checks are themselves tested. Each run carries a mutation drill: deliberate errors ("plants") are introduced into copies of the inputs and each must be caught by exactly the checks that name it; controls (inputs that are right) must not trip anything; and a set of must-stay-green assertions must hold throughout. Each check is also blinded in turn and must go blind in exactly as many cases as the plants that name it. Before a run is recorded, it is rerun cold and its gate log must be byte-identical to the first.

Over the 68 runs the drills plant 1,290 defects; all 1,290 are caught, each by the check named for it; none is caught by a wrong check or missed. The 262 controls disturb nothing, and 13,068 must-stay-green assertions hold, across 665 checks. These are the sums of the `drill` records in the 68 gate logs.

8.4 What the recomputation is not

It is not a proof of anything beyond what each run measured, and it does not replace a formal proof. A `PASS` says that the run's independent reconstruction agrees with the document on what the run checked; where the reconstruction and the document agree only because both make the same choice the text leaves open, the run says which choice.

8.5 Corrections

The corrections table of the inventory (Appendix B) has 236 rows, each giving what a source says against what the recomputation finds; a few rows, marked as such, correct an earlier version of a report rather than a document. By strength, 124 items stand only in a corrected form and 27 do not hold at all. Those that bear on what the programme can claim are in §9.4.

9. Status ledger

Every item below is cited to the run that measured it (RUN-*nnn*, the reports in the recomputation tree). "Certified" means re-derived in exact rational or directed integer-interval arithmetic by the recomputation; "numerical" means computed in floating point or in high-precision floating point without enclosure; neither is a proof of the programme's target.

9.1 Counts

The inventory of the 68 runs (Appendix A) holds 1,218 items, classified there as: 136 theorem-level steps whose proofs were checked; 52 that hold by a cited published theorem; 115 claims re-derived in exact arithmetic (finite certificates, and symbolic identities checked exactly); 389 numerical facts; 124 statements that hold only as corrected; 27 refuted; 42 open; 40 not measured; and 293 labels (status words, names, definitions and schedules that assert nothing to measure).

The classification was made while extracting the items from the reports and then audited against the reports item by item for the two strongest classes. Of the 111 items first classified as certified, 18 were moved: 17, checked in floating point only, to "numerical", and one to proof-checked. Of the 199 first classified as proof-checked, 64 were moved: 23 to labels, 22 to exact re-derivation (symbolic identities), 16 to numerical (tested on samples, not proved), and three to corrected or open. The counts above are after the audit. This paper cites only items it rechecked in their reports.

9.2 Certified

- **The seed identity.** $\Lambda(D, B_3) = \Lambda(D, P_3) = \pi/8 + \sqrt{3}/4$ exactly; the next rung, $\Lambda(D, B_3, B_5)$, lies in $[0.8344, 0.83478095]$, its lower end only by Mishra's cited certificate (RUN-009).
- **The exact master root.** $d^\star, t_3, t_5, t_7$ bracketed to width 10^{-90}, inside the programme's own intervals (RUN-031).
- **Cell-level certificates of the main atlas.** The programme counts seven of its 77 necessity cells as closed or better — 47, 48, 68, 69, 70, 75, 76 (RUN-038). The recomputation verified, in its own integer arithmetic and independently of the programme's mpmath and MPFR backends:
  - the RHCERT certificates of cells 47 and 48 (4,272,769 stored vertices; least margins 6.75×10^{-9} and 1.03×10^{-7}; RUN-037), with all 9,285 leaves of cell 48 also certified directly (RUN-036);
  - cell 68's certificate, 9,255 leaves and 2,126,336 vertices, equal leaf for leaf to the recomputation's own (RUN-038), and cells 75 and 76's, 2,627,357 and 2,617,252 vertices (RUN-039);
  - cell 0's 9,278 leaves — cell 0 of the reference rounds is the cell of marker 70 — least margin $3.1948922738 \times 10^{-8}$ (RUN-031);
  - depth-32 lower bounds for cells 49, 64 and 66 (9,986, 10,539 and 10,636 leaves; RUN-040).

Every certified polygon is strictly convex with area above 167/200. Cell 69's depth-30 lower ledger (7,856 of 7,856 leaves, least margin 1.6525×10^{-9}) was reproduced in floating point, to 10^{-16} of the ledger's margins, not re-certified (RUN-035).

- **Upper markers.** Each of the 77 necessity cells contains a configuration whose outer polygon has area below 167/200, certified with 128–256 normals (RUN-032); each marker certifies one configuration in its cell.
- **The far-set gate.** Gate A, $M_{U_0}(B_7) > 0.0088366789091384$, holds as a finite dual certificate re-enclosed in 400-bit integer intervals, least lower bound 0.00927495913424859 (RUN-045).
- **The F3 branch.** Two erosion certificates close 644 of the 22,516 roots of the F3 frontier; all 56,149 area leaves hold in exact arithmetic, least margin 2.28×10^{-5} (RUN-047). At root 6217, three certificates (66,927 area leaves, least margin 2.60×10^{-5}; RUN-048) and three more (115,511 area leaves, least 2.66×10^{-5}; RUN-049) hold in exact arithmetic and take the root's open cells from 68 to 61.

9.3 Numerical only

- **Geometric closure of the atlas** is a floating-point exact-core computation under the package's pad (hull -10^{-10}): 51 of 77 cells close by lift depth 36, the only complete common budget (RUN-029); continuing past 40, cells 31, 27 and 7 close at depths 42, 43 and 45, and 66 of 77 are geometrically complete (RUN-035, RUN-040).
- **The barrier at root 6217.** Its constants — $\theta^\star \approx 0.454045$, the hull area 0.8349861105592, the triple contact to 80 digits — reproduce in high-precision floating point, not in interval enclosures (RUN-048, RUN-056, RUN-057, RUN-060, RUN-061).
- **First-order criticality of the rotation at the barrier** holds numerically: the one-sided rate along the re-minimized profile is proportional to ρ (5.3×10^{-5} at $\rho = 0.001$); neither the finite active-contact model nor the document's pressure split shows it (RUN-058).
- **Gate 6217-A on the $\rho_5 = 0$ slice** holds numerically by 2.05×10^{-4} (RUN-048, RUN-056); the profile $m(\rho) \geq m(0)$ on $[0, \pi/120]$ holds at every sample (RUN-061).

9.4 Refuted or corrected, where it matters

- "All other regions have already been eliminated by finite certificates" is false: 44.15 % of root 6217's box lies outside the 200 subroots and was closed only in a search whose trees are not shipped, and the $\rho_5 = 0$ arc side is covered by no certificate (RUN-048, RUN-049).
- "Gate 6217-A: CLOSED" is not established: the $\rho_5 = 0$ reduction drops the F3 constraint (RUN-048).
- The 195-root shape ledger omits the 14.6 % of the box closed for A_3^+ only (RUN-058).
- First-order criticality is derived from a pressure split that is not a subgradient: in the finite model the least first-order change is +0.0293, not 0; $(g_3, 0)$ falls short of a subgradient by 0.0663 over 72 directions (RUN-057, RUN-058, RUN-061).
- The germ stiffness at root 127 is 0.2366, not 0.4543; the ratio of the two germs' stiffnesses is 1.92, not 1.00078; and the "unstable" class of the isolation rule cannot occur (RUN-066, RUN-067).
- Evenness of the rotation profile gives $m'(0^+) = -m'(0^-)$, not $m'(0) = 0$ (RUN-061).
- The second-order slivers of the blow-up open at $\pi/2$, $\theta^\star$ and $2\pi - \theta^\star$, not at $\pi \mp \theta^\star$ (RUN-060).
- The least audited slack of the full atlas is 8.47×10^{-11}, not 2.8982×10^{-9} (RUN-029).

9.5 Open

- **The headline bound.** $a_{\mathrm{Leb}} \geq 0.835$ is not certified (RUN-035, RUN-039, RUN-040).
- **The main atlas.** Eleven cells are not geometrically complete; at depth 40 fourteen are still open under a common budget (RUN-030); cell 33 keeps 78,507 open boxes at depth 42 (RUN-040); ten residual frontiers are not recovered (RUN-040). The programme's MPFR replays and dominance audits are recorded, not recomputed (RUN-036 to RUN-040).
- **The F3 branch.** The two erosion certificates close 644 of its 22,516 roots and do not address the rest (RUN-047); roots 8001, 8015 and 8016, and the parts of roots 6217 and 3766 outside their certified subroots, are closed only in search logs whose trees are not shipped (RUN-047). Root 6217 keeps 61 open cells, the arc side of its $\rho_5 = 0$ slice, Gate 6217-B on $0 < \rho_5 \leq \pi/120$, and the 44 % closed by search only (RUN-048, RUN-049); root 3766 and the F5 branch are untouched (RUN-048, RUN-049).

9.6 Targets

Closure in this programme is target-specific, and does not carry from one target to another (RUN-056, RUN-061). Three targets appear:

target	value	where it stands
headline	$167/200 = 0.835$	not certified; the root-6217 barrier, a configuration of the witness family used there, has hull area $0.8349861105592 < 0.835$ (numerical)
tight reference	0.8349865	has 3.894×10^{-7} of room, not the stated 6.5×10^{-8} (RUN-049); cell 131's candidate is the only one below it
A_3^+	0.834781191	the ceiling of the shape ledger at root 6217 (RUN-056); the barrier lies 2.049196×10^{-4} above it (RUN-048)

All three lie above the published lower bound 0.8344 (Mishra, cited, not verified here) and below the published upper bound 0.8440935944 (Gibbs, cited, not verified here).

10. Open problems

The ledger of §9 leaves these open, in roughly the order the programme would have to meet them.

1. **Close the atlas.** Eleven of the 77 necessity cells are not geometrically complete; ten residual frontiers are not recovered, and cell 33 keeps 78,507 open boxes at depth 42 (RUN-040). Every closure then has to be re-certified in exact arithmetic, as cells 47, 48, 68, 75 and 76 were (§4.6).
2. **Certify what is now only geometry.** 66 cells are complete in floating-point exact-core runs under a 10^{-10} pad (RUN-035, RUN-040); a certificate needs either a derived error term or an enclosure, cell by cell.
3. **The F3 branch.** Of 22,516 roots, 644 are closed by certificate (RUN-047). The parts closed only in search logs — roots 8001, 8015 and 8016, and the regions of 6217 and 3766 outside their certified subroots — need certificates whose trees are kept (RUN-047, RUN-049); root 3766 and the F5 branch are untouched (RUN-048, RUN-049).
4. **Root 6217.** Sixty-one open cells (RUN-049); the arc side of the $\rho_5 = 0$ slice (RUN-048); Gate 6217-B on $0 < \rho_5 \leq \pi/120$ (RUN-048); and the profile inequality $m(\rho) \geq m(0)$ on $[0, \pi/120]$, which holds at every sample and is not proved (RUN-061).

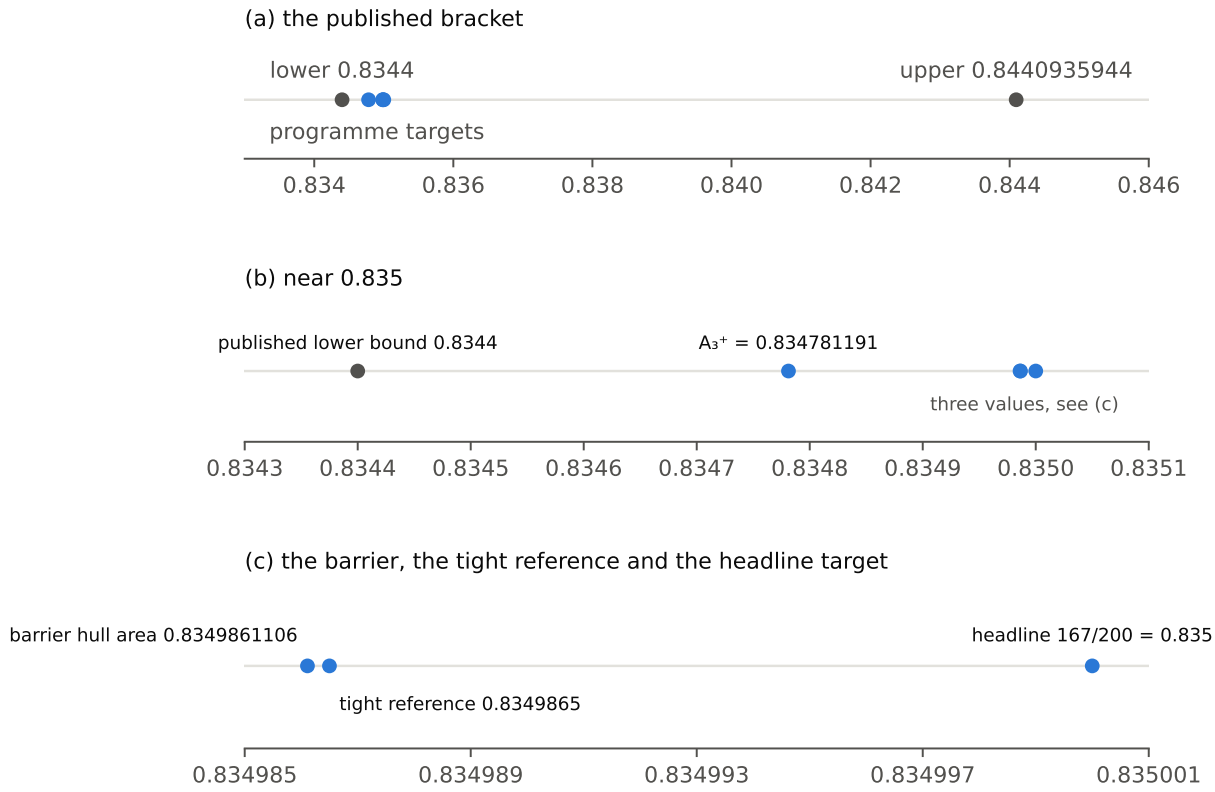


Figure 3: The published bracket and the programme's targets. (a) The bracket $0.8344 \le a_{\mathrm{Leb}} \le 0.8440935944$ (both bounds cited, not verified here). (b) Near 0.835: A_3^+ and the cluster. (c) The root-6217 barrier's hull area (numerical), the tight reference and the headline target. Values from RUN-002, RUN-035, RUN-048, RUN-049, RUN-056.

5. **The barrier's stationarity.** A dual certificate that means stationarity at the barrier would need the far-branch constraint with its own multiplier and edge-turning moves (RUN-056); the translation KKT condition holds numerically (least feasible rate 0.0354) but not by the argument given (RUN-057).
6. **Which target.** Closure here is target-specific (§9.6). Below the headline 167/200, the targets that remain meaningful near root 6217 must be fixed before any certificate is built for them — and the programme's own canonical target is not written down in its archive (RUN-043).

11. 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 the 68 runs, as they stood at the pinned commit, are published with SHA-256 sums and a deterministic archive at https://amral.evemisslab.com/lebesgue/recompute/ (English: https://amral.evemisslab.com/en/lebesgue/recompute/), one page per document of the programme. Every load-bearing number in this paper occurs in one of those gate logs or in `data/paper-figures.json` (Appendix C), which adds them up, counts them, or quotes them from the programme's documents with the line; the inventory (Appendix A) gives each item with a quotation that occurs in its report.

What re-running needs. The programme's documents and packages are two archives, kept byte for byte:

archive	size (bytes)	SHA-256
`AMRAL_Lebesgue_Finite_Closure_Round.zip`	57,576,727	`af850b6d09e9c77e0cf027451e571b465e0b`
`AMRAL_LUC_FC_Round_37_Package_v0.1_corrected.zip`	29,948,807	`942a79b10add0c981e125f48eacd23ec2e11`

Together they hold 1,459 leaf files (401.6 MB, nested archives included); the input manifest `data/SOURCE_MANIFEST.v1.json` gives the size, SHA-256 and git blob id of both archives 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 with mpmath 1.3, numpy 2.4 and sympy 1.14; 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/run030_round27_interval_replay.py
python code/run030_round27_interval_replay.py --drill
```

and likewise for every `code/runNNN_*.py`. 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` from the gate logs, the inventory, the input manifest and the archive, and is correct on the same terms. The figures of this paper are redrawn from the gate logs, as they stood at the pinned commit, by `scripts/figures.py` beside the manuscript; 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.

12. 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 and the methodologies' applications to this problem, on the programme Neo.K set out. Every statement this paper attributes to "the programme" or to "a document" is from those documents; none is the recomputation's invention.

Neo.K set out the research programme: the finite-closure strategy and its target, and the four methodologies it runs on — DLMVC, UESFCM, LESR and BCODR — which the first layer of the AMRAL site records as Neo.K's own.

懸衡 (Xuán Héng) — the independent recomputation; model: Claude, Opus series — did the recomputation and wrote this paper: every `COMPUTED_HERE` figure rebuilt with independently written code, never by running the programme's scripts; every certificate this paper calls certified re-verified in its own integer arithmetic; every theorem-level step read against its source; the mutation drills of §8.3; the corrections of §9.4; the inventory, its strength audit and the figures.

見證 (Witness), who is not an author of this paper, curated the programme's documents into the first layer of the AMRAL site, as written. That the 42 documents there are byte-identical to the archived originals is the recomputation's finding, not the curator's statement: RUN-001 matched each by content hash against every leaf of the two archives.

Priority, stated plainly. No result in this paper is claimed as new mathematics. The bounds and theorems it cites are other people's. What is offered as this line's own is narrower: an independent recomputation of a finite-closure programme, the strength at which each of its results stands, and the corrections the recomputation found.

References

1. P. Gibbs, *An upper bound for Lebesgue's covering problem*, https://arxiv.org/abs/1810.10089 (2018).
2. U. Mishra, *Curves of constant width and Lebesgue's covering problem*, https://arxiv.org/abs/2608.30538 (v2, 2026).
3. 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).
4. N. Xie, *A certified lower bound for Lebesgue's universal cover problem*, https://arxiv.org/abs/2606.04458 (v4, 2026).
5. Z. Deng, *Closed escape path of smallest diameter in forest: dual formulation of Lebesgue's universal covering problem*, https://arxiv.org/abs/2609.32931 (2026).

References [1]–[3] are the programme's own baseline, checked against arXiv on 2026-09-24 (RUN-002); [4] and [5] and the absence of a newer bound were checked against arXiv on 2026-10-18. Searches covered arXiv only.

Appendices

The appendices ship with the manuscript and the evidence; they are part of this paper.

- **Appendix A — the inventory.** `INVENTORY.md` and `data/inventory.json` beside this manuscript: every item of the 68 reports, the strength at which it stands after the audit (an item the audit moved says so), and a quotation that occurs in its report.
- **Appendix B — the corrections.** The corrections table of `INVENTORY.md`: each row the source's value against the recomputed one.
- **Appendix C — the figures this paper adds up, counts or quotes.** `data/paper-figures.json` at the pinned commit, written by `code/paper_figures.py`. It names every field of the gate logs it reads and asserts what a sum relies on (every leaf certified, no failure). The figures quoted from the programme's documents are kept under their own key, with the document, its SHA-256 and the line, and are not counted as recomputed.
- **Appendix D — the input manifest.** `data/SOURCE_MANIFEST.v1.json` at the pinned commit: the size, SHA-256 and git blob id of both archives, and the SHA-256 of every one of their leaves.