Exact values and certified lower bounds for the no-three-in-line problem in the cube

Aleksei Kudriashov (Alex Komang)
Nusa Dua Studio, Nusa Dua, Bali, Indonesia
ORCID: 0009-0006-8885-5338

3 September 2026 — version 1.5
doi:10.5281/zenodo.22019279

Abstract

Let $a(n)$ be the largest number of points of the grid $\{0, 1, \ldots, n-1\}^3$ no three of which are collinear. Pór and Wood proved $a(n) = \Theta(n^2)$, the lower bound by the construction $\{(x, y, x^2+y^2 \bmod p)\}$, which has no three collinear precisely when $p \equiv 3 \pmod 4$ (Lemma 4 of [1]); no exact values appear to have been recorded, and the sequence was not in the OEIS when this work began (it is now A399138, approved August 2026). We determine

$$a(1), \ldots, a(6) = 1,\ 8,\ 16,\ 28,\ 40,\ 64,$$

give certified lower bounds $a(7) \geq 73$, $a(8) \geq 94$, $a(9) \geq 116$, $a(10) \geq 138$, $a(11) \geq 164$ (the bounds for $n \geq 8$, new in version 1.5, come from exact optimisation inside the symmetry classes of the cube group), and prove $a(p) \geq p^2$ for every prime p by an elementary construction from an anisotropic binary quadratic form.

The method is a certified decision procedure rather than a search: a line carrying at most two lattice points forbids nothing, so admissibility is exactly "at most two chosen points on every line with three or more lattice points", which is a propositional formula on n^3 variables. Unsatisfiability is answered with a DRAT certificate that an independent checker verifies, and the lower bounds are witnesses re-verified by exhaustive integer arithmetic over every triple. We are explicit about the three separate links this requires — that the formula expresses the problem, that it is unsatisfiable, and that the symmetry pruning discards nothing — because a certificate attests only to the formula it was given, and an error in the other two links would pass through it unseen.

Along the way the optima turn out to share a layer structure that fails at exactly one of the values computed: for $n = 2, 3, 4, 6$ the two outer layers carry the planar maximum $2n$ and the inner ones $2n-2$, giving $2n^2 - 2n + 4$; at $n = 5$ two outer layers provably cannot both carry $2n$ at any total, and at $n = 7$ the corresponding profile is impossible as well. We have no explanation for this.

1 Definitions and elementary bounds

Points of $G_n := \{0, \ldots, n-1\}^3$; a set is *lawful* if no three of its points are collinear in $\mathbb{R}^3$ (equivalently, all 2×2 minors of the difference matrix vanish for no triple). Each of the $3n^2$ axis-parallel lines carries at most two points of a lawful set, and each point lies on three of them, so $|S| \leq 2n^2$; this is the trivial upper bound, attained at $n = 2$.

Proposition 1. *For every prime p, $a(p) \geq p^2$.*

Proof. Let $Q(x, y) = x^2 - dy^2$ with d a quadratic non-residue modulo p, so that $Q(a, b) \equiv 0$ only for $a \equiv b \equiv 0$ (an anisotropic binary form; one exists for every p). Put $S = \{(x, y, Q(x, y) \bmod p) : 0 \leq x, y < p\}$, of size p^2. Suppose three of its points are collinear. Their projections to the

(x, y)-plane are distinct (the third coordinate is a function of the first two), hence lie on a line with primitive integer direction (a, b), and the three points are $(x_0+t_i a,\ y_0+t_i b,\ \cdot\,)$ with distinct integers t_i, $|t_i| < p$. Along this line the third coordinate satisfies $z \equiv Q(x_0 + ta, y_0 + tb) = Q(a, b)\, t^2 + \ell(t) \pmod p$ with ℓ affine, whereas collinearity in $\mathbb{R}^3$ forces z to be an affine function of t. A quadratic congruence agreeing with an affine one at three points distinct modulo p has vanishing leading coefficient, so $Q(a, b) \equiv 0$ and therefore $(a, b) = (0, 0)$, a contradiction. $\square$

Neither $x^2 + y^2$ (anisotropic iff $p \equiv 3 \bmod 4$) nor $x^2 + xy + y^2$ (iff $p \equiv 2 \bmod 3$) works for all p; the choice $x^2 - dy^2$ does. Proposition 1 is thereby the completion, to every odd prime, of Lemma 4 of Pór and Wood [1], whose construction is the case $Q = x^2 + y^2$ and works precisely for $p \equiv 3 \pmod 4$; we verified their dichotomy directly (6 and 286 collinear triples at $p = 5$ and $p = 13$, none at $p = 7$, 11). For $p = 2$ the statement is trivial: no line meets $\{0, 1\}^3$ in three points, so $a(2) = 8 > 4 = p^2$. The bound of Proposition 1 is not sharp: $a(5) \geq 40 > 25$ and $a(7) \geq 73 > 49$. As a guard against a slip in transcribing the construction rather than in the proof, we also built S explicitly for $p = 3, 5, 7, 11, 13$ and tested *every* one of its triples by integer cross product — 790 244 triples at $p = 13$ — finding none collinear.

2 Method: a certified decision procedure

A line carrying at most two lattice points of $[n]^3$ forbids nothing, so

$$S \subseteq [n]^3 \text{ is admissible} \iff |S \cap L| \leq 2 \text{ for every line } L \text{ with } \geq 3 \text{ lattice points.} \tag{1}$$

This turns the question "is there an admissible set of size M?" into a Boolean satisfiability problem with n^3 variables — 125 for $n = 5$, 216 for $n = 6$ — one at-most-two cardinality constraint per rich line, and one global at-least-M constraint. The global counter is a totalizer truncated at $2n^2$; the truncation is legitimate because each of the n^2 axis-parallel lines carries at most two points, so $a(n) \leq 2n^2$. We add the 47 lexicographic leader constraints of the cube group, which preserve satisfiability because the property in (1) is invariant under that group.

We adopted this route after measuring the alternative. A prefix-split exhaustive search on 42 cores closed 5 of 204 prefixes in two hours at $n = 5$, extrapolating to roughly 88 core-days; the same question is decided by a SAT solver in minutes. Speed, however, is not the reason. An exhaustive search establishes completeness by *its own* covering argument, which only its authors' code can check; an unsatisfiability answer comes with a DRAT certificate that an independent, widely used checker verifies without trusting our solver or our code at all.

Three separate links. It is worth being explicit about what is and is not certified, because the three claims below rest on different foundations.

1. *The formula expresses the problem.* Verified two independent ways: by comparing the set of triples forbidden by the rich-line constraints against the set with vanishing cross product, and by forcing each collinear triple as an assumption and requiring unsatisfiability. At $n = 4$ and $n = 5$ there are 376 and 1858 collinear triples; none was left unforbidden and none was forbidden spuriously. A DRAT checker cannot verify this link — it knows nothing of cubes and lines — and this is exactly where an encoding bug would hide.

2. *The formula is unsatisfiable.* Verified by `drat-trim` on the certificate emitted by `kissat`.

3. *A witness is admissible.* Verified by exhaustive integer arithmetic over all $\binom{|S|}{3}$ triples, by a program sharing no code with the search.

Links 1 and 2 give the upper bound; links 1 and 3 give the lower bound.

Removing link 1 for small n. The two checks above verify that the encoding is faithful. A stronger thing is possible: to compute the same values without an encoding at all, so that link

1 does not need to be trusted because it is not used. We wrote a plain exhaustive traversal that places points in the cube and tests triples by integer cross product — the literal definition — with no propositional formula, no solver and no certificate. Two points kill the line through them, so every cell carries a counter of dead lines covering it and a cell is available while that counter is zero; a branch is abandoned when the points placed plus the cells still available cannot exceed the best found. Run from a bar of zero, so that no value can be fed in and echoed back:

n	result	nodes
1	1 (exhausted)	2
2	8 (exhausted)	37
3	16 (exhausted)	49 812

These agree with the certified values and share nothing with the certified route but the definition of the prohibition itself. Our own full enumeration at $n = 3$ visits all 3 813 884 lawful subsets (the empty one included) and finds exactly *two* optima of size 16 — a chiral pair, each with stabiliser of order 24 in the full cube group, both with layer profile $[6, 4, 6]$ along all three axes. At $n = 4$ a traversal from a bar of 27 *exhausted* its tree (9 012 859 649 nodes, 3180 s), proving $a(4) = 28$ with no encoding at all — a bar is harmless under exhaustion, since it prunes only branches unable to exceed it. The count at $n = 4$ is now ours as well: a 32-shard traversal (sharded by first cell, run to completion on every shard, 13 215 017 061 nodes in all) finds exactly 14 optima of size 28, in three classes under the full cube group with stabiliser orders 8, 8, 24 $(6 + 6 + 2 = 14)$ — in agreement with an independent external recount. Every optimum at $n \leq 4$ therefore has the layer profile $[2n, (2n-2)^{n-2}, 2n]$: the empirical observation of Section 5 is a fact up to $n = 4$.

Correction. An earlier version of this paragraph quoted two numbers received from the first solver's message and never passed through the verification bench of this very note: "2 220 075 subsets" and "the count 10 960". Both are withdrawn, and both origins have since been reconstructed: 2 220 075 is exactly $\binom{27}{8}$ — the enumeration of all 8-subsets of 3^3 that verifies the uniqueness of the 8-point optimum of the *no-four-coplanar* problem (A280538 gives 1 there) — and $10\,960 = 48 \cdot 232 - 176$ is the number of labelled optima of that same neighbouring problem at $n = 4$, where A280538 gives 232 classes. Two numbers from the wrong problem, quoted into the section about removing trust; recorded, not smoothed.

We are explicit about how far this extends. From a bar of zero the traversal reaches 28 at $n = 4$ without exhausting in an hour $(8.85 \cdot 10^9$ nodes); from a bar of 27 it exhausts $(9.01 \cdot 10^9$ nodes, 53 minutes), so $a(4) = 28$ is established twice over as well. The honest division: for $n \leq 4$ both the value and the bound are encoding-free in addition to certified; for $n \geq 5$ the encoding is checked but used.

Calibration. Before computing anything new we recomputed, as a calibration, the two smallest nontrivial values $a(3) = 16$ and $a(4) = 28$ — there is no prior published source for them that we could find, so the calibration is against our own earlier runs, not against the literature — each with unsatisfiability at $M + 1$ in under a second, independently on two machines. For instances that resist a single solver call we split into independent pieces by fixing the contents of the first columns; the split is complete because three points of a column are collinear, so subsets of size at most two exhaust it. The aggregator refuses a verdict unless every piece of the manifest is present and unsatisfiable — a missing piece and an unsatisfiable one must never look alike.

3 Exact values

n	$a(n)$	$a(n)/n^2$	how established
1	1	—	trivial
2	8	2.00	exhaustive; meets the trivial bound $2n^2$
3	16	1.78	exhaustive, and certified
4	28	1.75	exhaustive, and certified
5	40	1.60	certified
6	64	1.78	certified

The value $a(6) = 64$ appears to be new; a search of the OEIS for $1, 8, 16, 28, 40$ and for $8, 16, 28, 40, 64$ returns nothing, and the sequence itself is not in the database.

For $n \leq 4$ the values were also produced by two independent exhaustive programs whose node counts differ by a factor of about 72 ($17/3855/4.21 \cdot 10^7$ against $17/20068/3.03 \cdot 10^9$ for $n = 2, 3, 4$), so their agreement is not an artefact of a shared design.

For $n = 5$ and $n = 6$ the decisive question — does a set of 41, respectively 65, points exist? — was split into independent pieces by fixing the contents of the first two z-columns, and the heaviest pieces were split once more on the third column. Both splits are complete because three points of a column are collinear, so subsets of size at most two exhaust it. The final tally, reproduced in `logs/no3_3d/FINAL_VERDICT.txt`:

	$n = 5$, $M = 41$	$n = 6$, $M = 65$
top-level pieces	256	484
closed directly	221	450
closed through all their continuations	35	34
continuations required per piece	16	22
satisfiable pieces found	0	0
pieces left undecided	0	0

The aggregator builds the expected set of pieces from the manifest rather than from the results, treats a missing piece and an undecided one as failures, and distinguishes three outcomes rather than two: only an explicit unsatisfiable answer closes a piece, a satisfiable one anywhere would destroy the claim, and a killed or timed-out run counts as *no information* — it neither closes nor refutes. Eleven such uninformative lines occur in the $n = 6$ record, from runs interrupted when work was moved between machines; each of the pieces concerned is closed by an explicit unsatisfiable answer from another run.

The lower bounds are witnesses, checked by exhaustive integer arithmetic over all $\binom{|S|}{3}$ triples by a program sharing no code with the search: 40 points at $n = 5$ (9880 triples) and 64 points at $n = 6$ (41664 triples). For $n = 6$ two witnesses were obtained independently, by a SAT solver and by a constraint solver in decision form; they turned out to be the same set, which the stabiliser explains — it has order 24, so the configuration has only two distinct images under the 48 symmetries of the cube.

As a further check that is not part of the proof, a constraint solver was given the $n = 6$, $M = 65$ question in one piece for an hour and returned no answer. That is not evidence of impossibility — it is the solver saying it does not know — but it found no counterexample either, while at $M = 64$ it produces a solution in seconds.

4 Structure of the optima, and an anomaly at $n = 5$

Writing the layer profile of a set as the number of its points in each of the n layers perpendicular to an axis, every optimum we have found has the *same* profile along all three axes:

n	$a(n)$	layer profile (identical for all three axes)
2	8	$4, 4$
3	16	$6, 4, 6$
4	28	$8, 6, 6, 8$
5	40	$8, 8, 8, 8, 8$
6	64	$12, 10, 10, 10, 10, 12$

A layer is an $n \times n$ grid with no three collinear, so it holds at most $2n$ points. For $n = 2, 3, 4, 6$ the two outer layers attain that planar maximum and the inner ones hold $2n - 2$, giving $2 \cdot 2n + (n-2)(2n-2) = 2n^2 - 2n + 4$ — which reproduces $8, 16, 28$ and 64 exactly.

At $n = 5$ the pattern fails, and it fails for a structural reason rather than by accident. Prescribing the profile as a hard constraint and deciding feasibility shows that

profile at $n = 5$	verdict
$10, 8, 8, 8, 10$ (total 44, the value the pattern predicts)	infeasible
$10, 9, 8, 9, 10$; $10, 8, 9, 8, 10$; $10, 10, 10, 10, 10$; $9, 9, 9, 9, 9$	infeasible
$10, 8, 8, 8, 8$ (total 42) and $10, 7, 8, 7, 10$ (total 42)	infeasible
$\mathbf{10, 6, 8, 6, 10}$ (total 40)	**infeasible**
$8, 8, 8, 8, 8$ (total 40)	feasible
$n = 6$: $12, 10, 10, 10, 10, 12$ (total 64)	feasible
$n = 7$: $14, 12, 12, 12, 12, 12, 14$ (total 88)	infeasible

A caveat on how these verdicts are to be read. Prescribing the profile along all three axes at once is a restriction, so an infeasibility obtained that way would exclude only symmetric configurations, not that number of points. Every infeasibility above was therefore re-derived with the profile prescribed along a *single* axis, which excludes every configuration with that layer distribution. The satisfiable rows need no such care: a solution is a solution. The fourth line is the sharpest: at $n = 5$ two outer layers cannot both attain the planar maximum $2n$ at *any* total, so the optimum is forced onto a uniform profile. The same computation at $n = 6$ finds the analogous profile realisable, and at $n = 7$ it is again impossible. The pattern $[2n, (2n-2)^{n-2}, 2n]$ is thus realisable at $n = 2, 3, 4, 6$ and impossible at $n = 5, 7$ — the obstruction is not an artefact of one case, and it is not an obstruction to the idea of maximal outer layers as such.

The obstruction can be priced exactly. A layer is a planar no-three-in-line configuration and so carries at most $2n$ points; requiring *both* outer layers to attain that planar maximum, and maximising the total, gives

n	both outer layers	maximum total	$a(n)$
4	8	28	28 — costs nothing
5	10	**39**	**40 — costs exactly one point**
6	12	64	64 — costs nothing

So the anomaly is not that a fitted formula occasionally fails. It is that two planar-optimal configurations placed in the outer planes are compatible with the spatial optimum at $n = 4$ and $n = 6$ and incompatible at $n = 5$, and incompatible by a single point.

We do not have a proof of why. A count of the planar configurations attaining $2n$ points suggests rigidity: the 4×4 grid admits 11 such configurations in 4 classes under the square's symmetries, the 5×5 grid admits 32 in only 5 classes, and the 6×6 grid admits 50 in 11 classes. We offer this as an observation accompanied by its data, not as an argument. One natural mechanism can be ruled out: the midpoint of a segment joining the two outer layers is a lattice point when n is odd, and must stay empty, but an exact count shows the middle layer still retains up to 16 free cells — above its own planar ceiling of 10 — so that constraint does not bind.

5 Certified lower bounds for $n = 7, \ldots, 11$

n	$a(n) \geq$	$/n^2$	source of the witness
7	73	1.49	fourteen inequivalent configurations known; the first was invariant under every permuta
8	94	1.47	two inequivalent configurations, each invariant under a subgroup of order 4 (version 1.4
9	116	1.43	two inequivalent configurations: stabilisers of order 4 and of order 12; the latter is optir
10	138	1.38	full stabiliser of order 12; optimal within the order-12 class
11	164	1.36	stabiliser of order 12; the bound inside the class is not closed

The bounds for $n = 8, \ldots, 11$ (version 1.5, 2–3 September 2026) come from a sweep over all 33 conjugacy classes of subgroups H of the symmetry group of the cube (order 48): in each class the maximum of $|S|$ over H-invariant admissible sets is computed by exact optimisation (CP-SAT over the H-orbits of cells, with "at most two chosen cells on every grid line") under a time limit of 25–40 minutes per class. Where the solver closes the bound the value is the exact maximum of the class (this happened for the order-12 class at $n = 9$ and $n = 10$); elsewhere it is a lower bound inside the class. Every configuration was verified by two independent programs over all $\binom{|S|}{3}$ triples, its stabiliser computed under the 48 symmetries, and its local optimality certified: no point can be replaced (all are rigid, min $\kappa \geq 2$ in the sense of the disjoint-killers lemma of the companion repository `lemma-atelier`), and no exchange removing at most four points (73, 94, 116) or three points (138, 164) improves it. The configurations, the per-class results and the verifier are archived separately [6].

The rest of this section is the version-1.4 account of $n = 7, 8$, kept as data. Each configuration is re-verified by testing all $\binom{|S|}{3}$ triples with integer cross products. These are lower bounds only; we make no claim about how far they are from the truth, and we expect them to be weak, since a randomised search is a poor instrument at this size. The $n = 7$ bound illustrates the point: it stood at 71 from randomised search and rose to 73 within the hour once the search was restricted to configurations invariant under a cyclic permutation of the axes. Restricting to a subgroup collapses the cells into orbits and shrinks the search by the order of the group; it is legitimate for lower bounds only, since a configuration found is a configuration, while failure to find one among the symmetric ones says nothing about the rest. The size of the group has an optimum in the middle rather than at either end, and the optimum is a balance: the group must be large enough to make the problem solvable and small enough that good configurations still fall inside the symmetric stratum. At $n = 8$:

order of the group	best found
2	82, 84
3	88
4	88, 90
6	**92, 92, 93**
8	80, 88
12	88
24	88 (proved optimal within its stratum)
48	80 (proved optimal within its stratum)

(The full sweep of version 1.5, 40 minutes per class, later gave 94 in two classes of order 4, so the order-6 value 93 of this table was not the maximum of its own neighbourhood of classes either.) At $|G| = 48$ only 20 orbits remain, the solver settles the stratum in seconds, and the answer is 80, because genuinely good configurations do not have that much symmetry. Note also that the order alone does not decide it: two different groups of order 8 gave 80 and 88. We record this as data. An earlier draft of this note stated instead that a larger group simply gives a better result — a pattern read off the first four rows, which the remaining four contradict. (An earlier draft of this note quoted 72 and 90 here. Those numbers came from a run whose witness was never saved, and when we came to attach witnesses to them the journals we could reproduce gave 71 and 89. We report the values we can exhibit.) Two indications. First, the layer structure of the smaller optima suggests $2n^2 - 2n + 4 = 88$ at $n = 7$ — and although we have shown that the corresponding *profile* is impossible, that leaves the value itself untouched. Second, at $n = 6$ the same randomised search reached 64, which the certified computation then confirmed to be optimal, so the instrument is not hopeless either; we simply do not know which case $n = 7$ resembles.

For the record we also measured what the abandoned route would have cost, though the measurement deserves a caveat. Knuth's random-probe estimator, calibrated against the exact tree at $n = 4$, put an unaided exhaustive proof at $n = 6$ at some $4.6 \cdot 10^{22}$ nodes, and that figure is what sent us to a decision procedure. We later found the plain estimator unreliable on trees of this family — on a sibling problem the mean of twenty thousand probes was 89% a single probe, with the median five orders of magnitude below the mean — and replaced it by a stratified version that agrees with the exact tree at $n = 4$ to within 0.1%. The $4.6 \cdot 10^{22}$ was produced by the plain estimator and has not been recomputed with the stratified one, so it should be read as an order of magnitude and nothing finer. It says nothing about the difficulty of the problem in any case, only about the difficulty of one way of attacking it.

6 Reproducibility, and what is *not* claimed

All programs, journals, witnesses, manifests and the verification protocol are public: `https://github.com/iwasborninbali/saturation`. The protocol `docs/VERIFICATION.md` is written for a reader who trusts none of our code and gives the commands in order.

The following are the limits of what is established here, stated so that a reader need not infer them.

- At $n = 5$ the upper bound no longer rests on one solver: the same 256-piece split was closed independently by Glucose, of the MiniSat lineage rather than kissat's, with every piece unsatisfiable and none satisfiable, the last taking 2785 seconds and 14.4 million conflicts. Two solvers of unrelated descent, two independently written aggregators, two independent semantic checks and two independently verified witnesses now stand behind $a(5) = 40$. Worth recording separately: the *monolithic* instance defeated Glucose for 10,000 seconds while its pieces fell in minutes, so the case split is a property of the method rather than of one solver. **At $n = 6$ this has not been done** — there are 484 pieces and they are heavier — and the reader should not carry $n = 5$ over to it.

- The upper bound for $n = 6$ rests on one SAT solver. The instances that a single call decides carry DRAT certificates verified by `drat-trim`, which removes the dependence on the solver; the instances that had to be split do not carry stored certificates, because the certificates run to tens of gigabytes. They are regenerated deterministically on demand, and the case split, its completeness and its aggregation are each verified separately.

- We have no proof of the upper bounds independent of the propositional encoding. Our own exhaustive enumerator does not reach $n = 5$: it closed 24 of 4096 prefixes at 15–33 billion nodes each, with the full traversal on the order of 10^{14} nodes. It corroborates the lower bound only.

- The verification of the symmetry pruning shows that our *implementation* discards nothing; it is not a proof of the underlying lex-leader argument, which is standard but which we did not reprove.
- The stratum sweep (version 1.5) proves nothing about $a(n)$: a class maximum is an upper bound for that class only, and the classes are enriched in maxima but need not contain them — at $n = 6$ every class with a non-trivial symmetry stays below 64 except those attained by the certified optimum itself.
- The layer-structure anomaly at $n = 5$ and $n = 7$ is a computation, not an explanation. A count of the planar configurations attaining $2n$ points suggests rigidity (11 configurations in 4 classes for 4×4; 32 in only 5 classes for 5×5; 50 in 11 classes for 6×6), and one natural mechanism — the lattice midpoint between two outer layers — is ruled out by an exact count. Beyond that we do not know.

References

[1] A. Pór and D. R. Wood, No-Three-in-Line-in-3D, *Algorithmica* 47 (2007) 481–488; preliminary version in Proc. Graph Drawing 2004, Lecture Notes in Comput. Sci. 3383, Springer.

[2] A. Biere, K. Fazekas, M. Fleury and M. Heisinger, CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020, Proc. SAT Competition 2020, University of Helsinki, 2020.

[3] N. Wetzler, M. J. H. Heule and W. A. Hunt, Jr., DRAT-trim: efficient checking and trimming using expressive clausal proofs, Proc. SAT 2014, Lecture Notes in Comput. Sci. 8561, Springer, 2014.

[4] G. Audemard and L. Simon, Predicting learnt clauses quality in modern SAT solvers, Proc. IJCAI 2009.

[5] OEIS Foundation Inc., The On-Line Encyclopedia of Integer Sequences, `https://oeis.org`, sequence A399138.

[6] A. Kudriashov, Witness configurations and lower bounds for the no-three-in-line problem in the $n \times n \times n$ grid (OEIS A399138), Zenodo, 2026, `https://doi.org/10.5281/zenodo.22271375`; source `https://github.com/iwasborninbali/a399138-witnesses`.

AI/tool-use disclosure

The searches, the programs and the text were produced by autonomous AI agents (Anthropic Claude) working under the direction of the author of record, who posed the question, chose what to compute, decided what to claim and takes responsibility for the content. Two agents worked with deliberately different program designs and audited each other; several errors reported in the repository were caught that way and not by their author, including one — an unverified symmetry pruning — that no certificate, witness check or coverage check would have detected.

The work is arranged so that this should not matter. Every claim above is meant to be checkable by a third party without trusting the agents or the author: certificates by `drat-trim`, witnesses by brute-force integer arithmetic, coverage by a mechanical aggregator that refuses a verdict on a missing or undecided piece, and the encoding by comparison against the direct collinearity criterion. Where a claim is not checkable in that way, it is listed in the preceding section.