Nineteen certified configurations, eighteen independent bounds for the no-four-coplanar problem in the cube, and what they do not establish

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

3 September 2026 — version 3.4

Abstract

Let $a(n)$ be the largest number of points of $\{0, \dots, n-1\}^3$ no four of which are coplanar (OEIS A280537). Earlier versions of this note claimed that published data stop at $a(8) = 20$. *That claim was false*, and this version exists to correct it: a January 2017 comment on the OEIS entry itself records lower bounds up to $a(17) \geq 42$ [1], and a 2016 public programming contest ran exactly this problem on the first twenty-five primes, with per-size bests recoverable from its final report [2] — 28 at $n = 11$, 32 at 13, 42 at 17, 46 at 19, 54 at 23, 66 at 29, down to ratio $1.94\,n$ at $n = 97$. Three rows of our table ($n = 14, 15, 17$) are below the 2017 bounds and four prime rows ($n = 17, 19, 23, 29$) are below the 2016 contest bests; Section 2 gives the full comparison.

What remains, stated against that prior art. Nineteen configurations, each verified by two programs sharing no code, with all coordinates public — the prior bounds come as numbers without accessible witnesses. Four improvements over the known bounds *and their monotone closure* under $a(n+1) \geq a(n)$: $a(12) \geq 31$ against the recorded 30, and $a(21) \geq 47$, $a(22) \geq 49$, $a(27) \geq 56$ against the 46, 46, 54 implied by the contest values at $n = 19$ and $n = 23$. The remaining composite rows (n = 18, 20, 24, 25) fall *below* that closure and are public witnesses, not records: version 3.1 called them “apparently first published”, and the monotone comparison that removes them is due to Hugo Pfoertner. And two proved statements above $n = 8$, both about the cyclically invariant subspace and both exhausted by two independent implementations: its maximum is 23 at $n = 9$ and 26 at $n = 10$.

Version 3.3 adds a search over all 33 symmetry classes of the cube group (Section 9): no bound changes, but the best known values at n = 7, 8, 9 are attained by three pairwise inequivalent configurations each (stabilisers of orders 1, 2, 3 at n = 7 and 8), all rigid and mutually unreachable by small exchanges, and four kinds of symmetry are shown to be incompatible with the problem.

No upper bound above $n = 8$ is known to us beyond the trivial $a(n) \leq 3n$; every number in the table is a lower bound only, and the earlier failure to check the field before claiming novelty is recorded in Section 2 rather than erased.

1 Reformulation

Four points are coplanar exactly when they lie in a common plane, and a plane carrying at most three lattice points forbids nothing. Hence

$$S \subseteq [n]^3 \text{ is admissible } \iff |S \cap P| \leq 3 \text{ for every plane } P \text{ with } \geq 4 \text{ lattice points.}$$

Two consequences are used throughout. Three collinear points are forbidden, since with any fourth point they are coplanar; and each of the n planes $x = \text{const}$ carries at most three points, whence

Proposition 1. *$a(n) \leq 3n$ for every n.*

This is the only upper bound above $n = 8$ that we can prove.

2 Prior art, and what earlier versions of this note missed

Two sources predate this note and were missed by every version of it up to 2.9; both sit in plain sight on the OEIS page this note cites.

The 2017 comment. The entry A280537 carries, since January 2017: "Currently (January 2017) known lower bounds for the next terms are $a(9) \geq 23$, $a(10) \geq 26$, $a(11) \geq 28$, $a(12) \geq 30$, $a(13) \geq 32$, $a(14) \geq 35$, $a(15) \geq 36$, $a(16) \geq 38$, $a(17) \geq 42$" [1]. The same entry states that terms up to $a(6)$ were found by exhaustive search and that "$a(7)$ and $a(8)$ are based on extensive numerical evidence" — which also corrects an error repeated below in Section 5 by earlier versions.

The 2016 contest. Al Zimmermann's programming contest *Non-Coplanar Points* (March–June 2016) posed exactly this problem for the first twenty-five primes [2]. The final report lists, for each size, the best score achieved outside the first 48 hours; the score is the square of the number of points, so the per-size bests are recoverable exactly:

n	best	n	best	n	best	n	best	n	best
2	5	13	32	31	71	53	110	73	145
3	8	17	42	37	81	59	121	79	157
5	13	19	46	41	89	61	125	83	164
7	18	23	54	43	92	67	137	89	173
11	28	29	66	47	100	71	142	97	188

(The full per-size decode from the final report, verified against the exact values $5, 8, 13, 18$ at $n = 2, 3, 5, 7$. Where this note overlaps: ties at 11, 13; below by 4, 3, 4, 7 at 17, 19, 23, 29.)

The contest reached ratio $2.67n$ at $n = 3$ falling to $1.94n$ at $n = 97$; whether that decline reflects the problem or the budget is precisely the question Section 7 argues search cannot settle — but the numbers themselves are prior art and stand.

The comparison for every row of this note.

n	this note	prior (source)	verdict
9	23	23 (2017)	tie, witness now public
10	26	26 (2017)	tie, witness now public
11	28	28 (2017; 2016)	tie, witness now public
12	31	30 (2017)	**improved by 1**
13	32	32 (2017; 2016)	tie, witness now public
14	34	35 (2017)	*below prior art*
15	35	36 (2017)	*below prior art*
16	38	38 (2017)	tie, witness now public
17	38	42 (2017; 2016)	*below prior art by 4*
18	41	42 (monotone from $a(17)$)	*below the monotone consequence*
19	43	46 (2016)	*below prior art*
20	45	46 (monotone from $a(19)$)	*below the monotone consequence*
21	47	46 (monotone from $a(19)$)	**improved by 1**
22	49	46 (monotone from $a(19)$)	**improved by 3**
23	50	54 (2016)	*below prior art*
24	52	54 (monotone from $a(23)$)	*below the monotone consequence*
25	53	54 (monotone from $a(23)$)	*below the monotone consequence*
27	56	54 (monotone from $a(23)$)	**improved by 2**
29	59	66 (2016)	*below prior art by 7*

Values marked "monotone" are the closure of the published bounds under $a(n+1) \geq a(n)$ (an n-cube configuration is also an $(n+1)$-cube configuration). Version 3.1 of this table compared against the published rows only and marked the composite sizes "apparently first published"; the monotone comparison is due to Hugo Pfoertner. The 2017 and 2016 numbers come without accessible configurations — the contest server verified scores, but the point sets themselves were not published; every number of ours comes with a public configuration, verified twice by programs sharing no code. *Certified public witnesses* against *score-only records* is the honest division of value, and it is why the rows below prior art stay in the table: they carry the only accessible configurations at their sizes that we know of.

One more thing the comparison measures. Our search above $n = 16$ was confined to the cyclically invariant subspace, and the gap to the 2016 records is 4, 3, 4, 7 exactly on the primes where they exist: an external measurement of what that confinement costs at growing n, consistent with Section 8 — enrichment is not containment, and here its price has a number.

How this was missed is worth recording, because the mechanism is general: the OEIS DATA line does stop at $a(8) = 20$, and we read the line and not the page. The comment with nine further bounds sat directly under it, and the contest is linked from the same entry. An external review pointed both out.

3 The bounds

n	lower bound	$3n$	gap	quadruples checked
9	23	27	4	8 855
10	26	30	4	14 950
11	28	33	5	20 475
12	31	36	5	31 465
13	32	39	7	35 960
14	34	42	8	46 376
15	35	45	10	52 360
16	38	48	10	73 815
17	38	51	13	73 815
18	41	54	13	101 270
19	43	57	14	123 410
20	45	60	15	148 995
21	47	63	16	178 365
22	49	66	17	211 876
23	50	69	19	230 300
24	52	72	20	270 725
25	53	75	22	292 825
26	53	78	25	292 825
27	56	81	25	367 290
28	56	84	28	367 290
29	59	87	28	455 126

Twenty-one rows, but only nineteen configurations and eighteen independent statements. The rows for $n = 26$ and $n = 28$ carry no configuration of their own: they follow from $n = 25$ and $n = 27$ by monotonicity alone, and are printed only to keep the table contiguous. The row for $n = 17$ does have its own configuration, but no longer carries its own information, since $a(17) \geq a(16) \geq 38$ follows by monotonicity. This has happened three times as the neighbouring row was improved, and it is the reason to read the table as a set of bounds rather than of values — a row is worth something only while it exceeds its left neighbour.

The configurations themselves are in the repository cited below. The gap to Proposition 1 grows with n; whether this reflects the problem or the reach of the search we cannot say, and Section 7 explains why we believe the question is not decidable by search.

4 Verification

Each configuration was checked twice, by programs written independently and testing different statements. The first computes the 3×3 determinant of every quadruple in exact integer arithmetic. The second takes every triple, forms the normal of the plane through it, and tests the scalar product against every remaining point; it also checks that the points are distinct and lie inside the cube, which the first test does not ask. Before each run the verifier is given a deliberately corrupted witness and must reject it; a check that cannot fail confirms nothing. All runs reported zero violations.

5 What is not established

No configuration proves an upper bound. A witness certifies "at least" and can never certify "at most". The values $a(5) = 13$ and $a(6) = 16$ rest on exhausted search trees; $a(7) = 18$

and $a(8) = 20$ rest, per the OEIS entry, on "extensive numerical evidence" only — an earlier version of this note wrongly promoted them to exhausted; above $n = 8$ we have Proposition 1 and nothing else. In particular the numbers in the table are not claimed to be maxima, and we expect several of them to be improved.

6 A bound of ours that these witnesses refuted

Setting $\log \binom{N}{m}$ equal to the expected number of coplanar quadruples in a random m-subset, computed *as if the constraints were independent*, gives a threshold that landed on exactly $2n+5$ for $n = 5, \ldots, 12$. Against the values then available its slack ran $2, 1, 1, 1, 0, 0$: monotonically to zero, which reads as a bound becoming tight. It is not a bound. The witness for $a(11) \geq 28$ exceeds $2n + 5 = 27$.

The failure has a mechanism. The estimate treats the forbidden quadruples as independent; they are positively correlated, since a set that has already avoided one plane avoids the next more easily. A first-moment threshold therefore *understates*, and we verified the sign on the companion estimate as well: the quadruple-independence threshold lies below the certified lower bound at every n from 5 to 12, with the deficit $-1, -2, -3, -3, -4, -4, -5, -4$: growing to $n = 11$, receding by one at the last step. What looked like a bracket around the truth was one biased estimator evaluated at two coarsenings, both leaning the same way.

We record the reading error separately, because it is not specific to this problem: a slack that converges to zero does not distinguish a tight estimate from one about to be crossed, and the two cannot be told apart from inside the model.

7 Effort, and why the growth rate is not measurable by search

The ratio of certified bound to n decreases over the range of the table. We do not read this as a property of the problem. Giving each seed an equal time budget on one core (eight seeds, fifteen seconds of uniform random greedy restarts each, the cell drawn uniformly among the alive) and recording how many seeds reach the OEIS value, we measure 8/8 at $n = 5$ and 0/8 at $n = 6, 7, 8$ — at $n = 6$ none of roughly 160 000 restarts per seed reaches 16. *Correction:* versions up to 2.9 quoted 8/8, 3/8, 2/8, 1/8 here; those four fractions were measured on the *planar* no-three-in-line problem and migrated into this note without a cube re-measurement. The corrected cliff is steeper, which strengthens the point but does not excuse the migration. Equal time is therefore not equal effort, and a curve obtained under a fixed budget records the budget. Deciding whether $a(n)/n$ tends to 3 or to something smaller would need a budget growing exponentially in n.

A second contamination is worth naming. In a search whose target doubles as its stopping condition, asking for K and receiving K is not evidence about K. Several values in an earlier draft of this table were produced that way, agreed with a prediction, and were withdrawn when the target was removed and the same program returned more.

8 The symmetric subspace

All configurations above were found by searching the subspace invariant under a cyclic permutation of coordinates. That subspace is enormously enriched in maxima, and enriched more as n grows. In the two-dimensional analogue, where we hold complete material — every maximum configuration for $n \leq 11$, reproduced against A000755 — the proportion of maxima carrying a symmetry of the square, divided by the same proportion in a null model matched on row and column counts, runs

$$\times 8,\ \times 56,\ \times 125,\ \times 769,\ \times 3729,\ \times 73\,000 \quad (n = 6, \ldots, 11),$$

the last figure resting on seven events in four million and therefore stated to one significant figure only. The raw proportion *falls* across the same range, from 36% to 13%; the ratio to the null model rises.

Enrichment is not containment. At $n = 6$ in three dimensions the cyclically invariant subspace contains no maximum at all: its exhausted maximum is 15 while $a(6) = 16$. Exhaustive traversal of that subspace at $n = 9$ gives exactly 23, which bounds the subspace and says nothing about $a(9)$ beyond the witness.

One consequence deserves emphasis. A sample drawn from a region enriched by three or four orders of magnitude cannot support statements about the structure of solutions in general. Structural claims in this note are made only where complete material is available.

9 Symmetry strata (versions 3.3–3.4)

This version adds a second search route and three structural facts, and changes no bound in the table.

Route. Instead of one symmetric subspace, take every conjugacy class of subgroups H of the symmetry group of the cube (order 48; 33 classes), and in each class maximise $|S|$ over H-invariant sets by exact optimisation: a CP-SAT model over the H-orbits of cells with the constraints "at most two points on every grid line" and "at most three on every lattice plane carrying at least four cells", the plane list computed exactly (12 453, 73 056, 269 895, 965 826, 2 688 585, 7 241 856 planes for $n = 5, \ldots, 10$). A lazy version of the model, which adds planes only when a solution violates them, does not converge (at $n = 6$ it stalls at 12 after two hundred rounds) — the full list is necessary.

Four kinds of symmetry are incompatible with the problem outright. If H contains the central inversion or a rotatory reflection of order 6, an H-invariant admissible set has at most 3 points (two antipodal pairs are coplanar); if H contains a reflection, at most 5 (the segments x–σx are parallel, so any two such pairs are coplanar, and the mirror carries at most three points); if H contains a rotation of order 4, the set lies on the axis. Only nine of the 33 classes remain: the trivial group, two classes of half-turns, the 3-fold rotation, two Klein four-groups, the cyclic group of a rotatory reflection of order 4, the dihedral group of order 6 and the tetrahedral rotation group. (The elementary proofs, with machine checks and a Lean formalisation of the inversion case, are in the companion repository `lemma-atelier`, lemma 003.)

Calibration. Within the strata the known values are reached at $n = 5, 7, 8$ (13, 18, 20) by the classes containing a 3-fold rotation or a half-turn; at $n = 6$ the best symmetric value is $15 = a(6) - 1$, in agreement with Section 8. At $n = 9, 10$ the 3-fold stratum reproduces 23 and 26 — the subspace maxima already exhausted by the first route — with a one-hour budget, and nothing above them. The trivial stratum (no symmetry imposed) reached 17 at $n = 7$ and 19 at $n = 8$ within 900 seconds and proves nothing.

Structural facts, every configuration checked by the two verifiers of Section 5 and its class computed under the 48 symmetries: (i) at $n = 7$ there are at least three pairwise inequivalent 18-point configurations, with stabilisers of orders 1, 2 and 3; at $n = 8$ three inequivalent 20-point configurations (orders 1, 2, 3); at $n = 9$ three inequivalent 23-point configurations (all of order 3). The order-2 classes lie outside the cyclically invariant subspace of Section 8. (ii) All nine are rigid — no point can be replaced by another — and no two are connected by an exchange of at most two points (at most three points for $n = 7$); under the best alignment the classes differ in at least 80% of their points (16 of 20 for one pair at $n = 8$, more for the others). (iii) Every stratum at $n = 5, \ldots, 8$ was solved to optimality or to the time limit; the per-stratum table and the configurations are in the repository.

What this does not establish: the strata give lower bounds inside symmetry classes only, and the strata at $n \geq 11$ (best known 28 and 31) have not been exhausted.

References

[1] OEIS Foundation Inc., Sequence A280537 (M. Scheuermann, 2017), with the January 2017 comment recording lower bounds to $a(17) \geq 42$ and the sourcing of $a(7)$, $a(8)$; `https://oeis.org/A280537`. See also A280538 (numbers of optimal solutions) and the references there: E. Pegg Jr. (Wolfram Demonstrations; Mathematics Stack Exchange), T. Sillke (rec.puzzles, 1992), W. Möhres (exhaustive search for $n = 6$, private communication, 2016).

[2] Al Zimmermann's Programming Contests, *Non-Coplanar Points*, March–June 2016; problem statement and final report at `http://azspcs.com/Contest/Tetrahedra`. Scores are squares of set sizes; per-size bests decoded accordingly (winner T. Foden).

AI/tool-use disclosure

The searches, the programs and the text were produced by autonomous AI agents (Anthropic Claude) 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 independently: one searched, the other verified and did not read the searcher's code, so that the verification would not inherit its blind spots. Verification protocol, programs, journals and witnesses: `https://github.com/iwasborninbali/saturation`.