

The direction spectrum of no-three-in-line solutions: a line model derives its shape, not its scale

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

3 September 2026 — version 1.1

Abstract

A *solution* of the no-three-in-line problem is a set of $2n$ points of the $n \times n$ grid with no three collinear. For a primitive direction v let $c_v(n)$ be the mean number of point pairs of a solution whose difference is a positive multiple of v , divided by n (averaged over the images of v under the symmetries of the square). Measured on Flammenkamp’s database of all known solutions, the short directions are nearly constant in n ($c_{(1,1)} = 0.731$, $c_{(1,2)} = 0.564$) and depleted against a null model with fixed row and column sums (ratio 0.58 and 0.75), while long directions are slightly enriched and grow with n . We show that the *shape* of this spectrum — the ratios between directions and their order — follows from a single rule with no fitted parameter: distribute $2n$ points over the lines of one direction with at most two points per line, the weight of a line of L cells carrying x points being $\binom{L}{x}$. The expected number of doubly occupied lines is given exactly by a generating function; one global renormalisation to the exact number of non-axial pairs of a solution, $2n^2 - 3n$, fixes the scale. The model reproduces fifteen measured constants within 12% (model/data 0.89–1.05 at $n = 20$, 0.88–1.02 at $n \approx 30$) and their order at matched n up to one adjacent pair; seven of them were predicted before being measured by a second agent who had not seen the model. What the model does *not* do is stated with equal care: the scale is imposed, not derived — the model’s own total falls 17% short of $2n^2 - 3n$, and the shortfall it redistributes is exactly the deficit seen in the measured directions; the growth of long directions with n is a property of the solutions that the model only partly tracks. A multi-channel prior-art search has now been carried out (Section 5): we found no published or web-posted measurement of pairs per direction on actual solutions, no κ -like constant, and no per-direction decomposition of the Guy–Kelly count; the nearest objects are listed and delimited there, and the search’s own limits are stated.

1 Definitions and data

Let $[n] = \{0, \dots, n-1\}$. A set $S \subseteq [n]^2$ is *admissible* if no three of its points are collinear and a *solution* if moreover $|S| = 2n$, the maximum possible since every row carries at most two points. For $v = (a, b)$ with $\gcd(a, b) = 1$ and $(a, b) \neq (0, 0)$, a *pair of direction* v in S is an unordered pair $\{p, q\} \subseteq S$ with $q - p = kv$ for some $k \geq 1$ (so every pair has exactly one primitive direction up to sign). For a solution, every line carries at most two points, hence the number of pairs of direction v equals the number of lines of direction v carrying exactly two points. Define

$$c_v(n) = \frac{1}{n} \mathbb{E}[\#\{\text{pairs of direction } v\}],$$

the expectation over the uniform distribution on the solutions of size n in the sample at hand, averaged over the class $\{(a, b), (b, a), (a, -b), (b, -a)\}$ of v under the symmetries of the square.

Since every row and every column of a solution carries exactly two points, $c_{(1,0)} = c_{(0,1)} = 1$ identically, and the total over all non-axial classes is fixed:

$$\sum_{v \text{ non-axial}} n c_v(n) \cdot |\text{class}(v)| = \binom{2n}{2} - 2n = 2n^2 - 3n. \quad (1)$$

Data. We use Achim Flammenkamp’s database of all known solutions (state of 31 August 2026). It is complete for $n \leq 20$ (for $n = 19$: 32 577 solutions; $n = 20$: 118 057) and contains only solutions with a non-trivial symmetry for $21 \leq n \leq 57$. The measurements were made by a second agent with code written independently of the model below; the values quoted as “mixed n ” average over all 391 812 solutions with weight one. Ensemble measurements on this database have been published since 1997 (Flammenkamp’s frequency maps of cell occupation [2], continued by Prellberg in 2026 for $n = 56, 57$); the present note continues that genre with a different observable.

class	mixed n	$n = 19\text{--}20$	$n = 29\text{--}31$	$n = 32\text{--}57$
(1, 1)	0.7306	0.7308	0.7327	0.7325
(1, 2)	0.5633	0.5638	0.5635	0.5548
(1, 3)	0.4545	0.4545	0.4586	0.4526
(2, 3)	0.4095	0.4097	0.4121	0.4070
(1, 4)	0.3744	0.3741	0.3807	0.3793
(1, 5)	0.3073	0.3074	0.3228	0.3261
(2, 5)	0.2846	0.2842	0.3023	0.3049
(3, 4)	0.3043	0.3039	0.3175	0.3186
(1, 6)	0.2618	0.2492	0.2687	0.2800
(1, 7)	0.2233	0.2071	0.2340	0.2461
(2, 7)	0.2081	0.1889	0.2208	0.2335
(3, 5)	0.2701	0.2609	0.2758	0.2819
(4, 5)	0.2390	0.2243	0.2486	0.2575
(1, 8)	0.1859	0.1692	0.1954	0.2160
(3, 7)	0.1961	0.1791	0.2069	0.2209

Short directions are constant in n to 1%; long directions grow with n (by 23% for (3, 7) and 28% for (1, 8) between $n \approx 20$ and $n \geq 32$); the sizes $21 \leq n \leq 28$ (54 604 solutions) are omitted from the table. A “constant of the database” for a long direction is therefore an average over a mixture of sizes, and the model must be compared with the data *at the same n* . Symmetry does not explain the growth: at $n = 19$ and 20, where the database is complete, the spectrum of the solutions with a half-turn symmetry (592 and 675 of them) differs from that of all solutions by at most 4%, and for the long directions in the *downward* direction ($c_{(1,7)}$: 0.1986 against 0.2067 at $n = 19$).

The null model. Under the uniform distribution on all configurations with exactly two points in every row and column (the correct null model: a solution is such a configuration), the expected number of pairs of direction v is the number of cell pairs of direction v times the probability $2(2n - 3)/(n(n - 1)^2)$ that a given pair of cells in distinct rows and columns is occupied (the group $S_n \times S_n$ acts transitively on such pairs). At $n = 20$ this gives $c_{(1,1)}^{\text{null}} = 1.27$, $c_{(1,2)}^{\text{null}} = 0.75$, ..., so the solutions are depleted in the short directions (0.58, 0.75, 0.89, 0.94 for (1, 1), (1, 2), (1, 3), (2, 3)) and slightly enriched in the long ones (1.01–1.11). This is the pattern to be explained.

2 The line model

Fix a direction v . The cells of $[n]^2$ are partitioned into the lines of direction v (a cell without a neighbour in direction v is a line of length 1); let the lines have lengths $L_1, \dots, L_r$.

Theorem 1. *Let $2 \leq m \leq \sum_i \min(L_i, 2)$ and let S be uniformly distributed on the m -subsets of the cells that contain at most two cells of every line. Then the number of such subsets is $Z_m = [t^m] \prod_{i=1}^r (1 + L_i t + \binom{L_i}{2} t^2)$ and*

$$\mathbb{E}[\#\{i : |S \cap \ell_i| = 2\}] = \frac{1}{Z_m} \sum_{i=1}^r \binom{L_i}{2} [t^{m-2}] \prod_{j \neq i} (1 + L_j t + \binom{L_j}{2} t^2).$$

Proof. A subset is admissible for v exactly when it takes $x_i \in \{0, 1, 2\}$ cells on line i ; there are $\binom{L_i}{x_i}$ ways to choose them, independently across lines, so the admissible m -subsets are counted by the coefficient of t^m in the product. The expectation is $\sum_i P(x_i = 2)$, and the admissible m -subsets with $x_i = 2$ are counted by $\binom{L_i}{2}$ times the coefficient of t^{m-2} in the product over $j \neq i$. $\square$

This is the canonical ensemble of hard particles in r independent boxes of capacity two with degeneracies $1, L_i, \binom{L_i}{2}$ (equivalently, a uniform placement conditioned on “at most two per box”); the identity is elementary. We did not find an established name for it; the nearest umbrella term in the refereed literature is *restricted occupancy* (Freund 1956 [6]; the unweighted equal-box case is an exercise in Comtet), and the corresponding relaxation appears as “capacity-two constraints” in Prellberg’s linear-programming bound for a parity variant [3]. What is not elementary is that it describes the solutions of the no-three-in-line problem.

The model. Put $m = 2n$ and let E_v be the expectation of Theorem 1 for the lines of direction v . The model does not know that the pairs of a solution are counted by (1): summed over all non-axial primitive directions, $\sum_v E_v = 615.3$ at $n = 20$ against the true 740, a shortfall of 17%. We therefore impose (1) by one global factor, $Z_n = (2n^2 - 3n) / \sum_v E_v$ ($Z_{20} = 1.203$, $Z_{30} = 1.163$, $Z_{40} = 1.139$), and set $c_v^{\text{model}}(n) = Z_n E_v / n$, averaged over the class of v . The factor is not fitted to the measured spectrum — it is the same for every direction — but it is not derived either. Everything the model predicts is therefore a statement about *ratios* between directions.

3 Results

class	model 20	data 19–20	ratio	model 30	data 29–31	ratio
(1, 1)	0.7637	0.7308	1.045	0.7441	0.7327	1.016
(1, 2)	0.5629	0.5638	0.998	0.5541	0.5635	0.983
(1, 3)	0.4327	0.4545	0.952	0.4316	0.4586	0.941
(2, 3)	0.3766	0.4097	0.919	0.3792	0.4121	0.920
(1, 4)	0.3428	0.3741	0.916	0.3492	0.3807	0.917
(1, 5)	0.2781	0.3074	0.905	0.2890	0.3228	0.895
(2, 5)	0.2544	0.2842	0.895	0.2665	0.3023	0.882
(3, 4)	0.2737	0.3039	0.901	0.2839	0.3175	0.894
(1, 6)	0.2322	0.2492	0.932	0.2440	0.2687	0.908
(1, 7)	0.1931	0.2071	0.932	0.2105	0.2340	0.900
(2, 7)	0.1800	0.1889	0.953	0.1978	0.2208	0.896
(3, 5)	0.2309	0.2609	0.885	0.2443	0.2758	0.886
(4, 5)	0.2076	0.2243	0.926	0.2223	0.2486	0.894
(1, 8)	0.1657	0.1692	0.979	0.1826	0.1954	0.934
(3, 7)	0.1669	0.1791	0.932	0.1852	0.2069	0.895

At $n = 20$ the model’s order of the fifteen classes coincides with the data’s except for the adjacent pair $(1, 6)$, $(3, 5)$ (which differ by 3% in the data); the model is below the data in every class but $(1, 1)$, by 5–12%, and at $n \approx 30$ by 6–12% in every class but $(1, 1)$ and $(1, 2)$.

Blind test. The first eight classes were known to us when the model was built. The remaining seven were *not*: their model values were written into a dated ledger (3 September 2026, 16:05 WITA) before a second agent, who had not seen the model or the ledger, measured them on the database with independent code; all seven fell within 2–10% of the prediction (model/data at mixed n : 0.93, 0.95, 0.95, 0.90, 0.93, 0.98, 0.94) and in the predicted order up to pairs differing by less than 3%. At matched $n = 20$ the agreement for the long directions improves (model/data 0.93 for both $(1, 7)$ and $(3, 7)$ instead of 0.87 and 0.85 against the mixture).

4 What fails, and what is not claimed

A first-order pair model fails. The natural first attempt — a pair of direction v survives with probability $(1 - \rho)^L$, ρ the conditional occupancy of a further cell on its line and L the number of such cells, then renormalise — over-depletes the short directions by a factor of two ($c_{(1,1)} = 0.40$ against 0.73) and puts $(1, 2)$ above $(1, 1)$. The exclusions along a long line are far from independent; the line model succeeds because it treats each line as a unit.

The scale is imposed. The model’s total falls 17% short of (1). After renormalisation the fifteen measured classes sum to 17.1 (model) against 18.6 (data) out of the fixed total $2n - 3 = 37$ at $n = 20$; the model’s deficit in the measured directions is exactly its excess in the tail of rare, steep directions, and no second renormalisation can move it. Of the shortfall, at most a third can be attributed to the row and column sums that the model ignores (a free $2n$ -subset has 705.7 non-axial pairs in expectation, the fixed-marginal model 740); the rest is the model’s treatment of one direction at a time, whereas a solution is constrained in all directions simultaneously.

Drift in n . The model’s values drift by 3–8% between $n = 20$ and 40 (upward for long directions), in the same direction as the data but not by the same amount; the constancy of the short directions in the data is reproduced only approximately ($c_{(1,1)}^{\text{model}}$: 0.764, 0.744, 0.731 at $n = 20, 30, 40$). We found neither a measurement nor an explanation of the growth of the long directions in the literature; the nearest theoretical statement, Remark 1.7 of Ghosal et al. [4], predicts for $k = 2$ an equal contribution of every dyadic scale of the direction length — that is, no growth.

One inversion. At $n = 20$ the model orders $(1, 6)$ above $(3, 5)$ by 0.6%; the data order them the other way by 4.7% (0.2492 against 0.2609) — the model’s largest error, –11.5%, is on $(3, 5)$.

Not claimed. No theorem about the ensemble of solutions is proved here; Theorem 1 is a statement about a model. No derivation of the scale, and no explanation of the growth of the long directions with n , is offered.

5 Prior art (searched 3 September 2026)

A multi-channel search (direct retrieval of the primary sources; full-text search inside arXiv, zbMATH, OEIS, Google Books and Open Library; API queries to Crossref, OpenAlex, Semantic Scholar, GitHub and the Wayback Machine) found no published measurement of pairs per direction, or of diagonal occupation, on actual no-three-in-line solutions; no κ -like constant; no per-direction decomposition of the Guy–Kelly count; and no model of the kind of Section 2. The nearest objects, none of which is prior art for the results above, are: Guy and Kelly’s expected number of collinear triples of a random $2n$ -subset [1], a sum over primitive directions that

is collapsed through Euler’s function and never split by direction; Flammenkamp’s frequency maps of cell occupation on the same database [2] and his “superficial explanation” of the shading cost of a pair of points, aggregated over symmetry classes; Prellberg’s linear-programming upper bound with capacity-two constraints on rows, columns and the two diagonal families for a parity variant [3] — a relaxation, not an ensemble, and without the weights $\binom{L}{x}$; Simkin’s occupation of diagonals by direction for n -queens (capacity one), with an explicit independence assumption across directions [5]; and a web-posted, unrefereed analysis of July 2026 [7] that measured central directions of symmetric solutions against a random null model and reached the same qualitative conclusion (short directions depleted). In view of the last item the qualitative observation is not claimed as new here; what is new is the quantitative spectrum over all pairs, its fifteen constants at matched n , and the model. *Limits of the search*: free-text web search was not available during it, and several primary sources were not read (Ungar 1982, Jamison 1984–86, Kaplansky–Riordan 1946, Riordan 1958, Hall–Jackson–Sudbery–Wild 1975, Flammenkamp’s JCTA papers of 1992 and 1998, the restricted-occupancy papers of Freund), so the negative result is stated with that caveat; the full record of channels, verdicts and confidence is public in the repository.

References

- [1] R. K. Guy and P. A. Kelly, The no-three-in-line problem, *Canad. Math. Bull.* 11 (1968) 527–531.
- [2] A. Flammenkamp, Frequency distributions of points of no-three-in-line configurations, <https://wwwhomes.uni-bielefeld.de/achim/no3in/density.html> (1997–2026).
- [3] T. Prellberg, No-three-in-line sets on the checkerboard grid, arXiv:2605.09215 (May 2026).
- [4] A. Ghosal, R. Goenka, A. Grebennikov, P. Keevash, M. Kwan and H. T. Pham, No- $(k+1)$ -in-line problem for $k \geq 3$, arXiv:2607.05255 (July 2026).
- [5] M. Simkin, The number of n -queens configurations, *Adv. Math.* 427 (2023) 109127.
- [6] J. E. Freund, Restricted occupancy theory — a generalization of Pascal’s triangle, *Amer. Math. Monthly* 63 (1956) 20–27.
- [7] J. Du, no3inline-rigidity (GitHub repository `aujurd22/no3inline-rigidity`, July 2026), unrefereed.

6 Reproducibility

The measurement code, the model code with an executable statement of Theorem 1 checked against brute force for $n \leq 5$, the tests written by a second agent, the two adversarial reviews (one of which found the false claim “the data are constant in n ” in an earlier draft and forced the comparison at matched n), the dated ledger of predictions, and the prior-art search of Section 5 (`docs/research/deep_research_14_answer.md`) are public: <https://github.com/iwasborninbali/lemma-atelier> (lemma 005, need 008) and <https://github.com/iwasborninbali/saturation>. The database is Flammenkamp’s (<https://wwwhomes.uni-bielefeld.de/achim/no3in/readme.html>).

AI/tool-use disclosure

The model, 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 built the model, the other measured the database with its own code and did not read the model before the blind test — and a third, adversarial pass reviewed the claims before publication.