Finding most likely Delaunay triangulations of imprecise points is NP-hard

Pedro M. M. de Castro
Centro de Informática, Universidade Federal de Pernambuco
Recife, Pernambuco, Brazil
pmmc@cin.ufpe.br

November 23, 2026

hexagon:2611.00951v1

Abstract

For n independent planar Gaussian sites, we seek the most probable Delaunay triangulation, specified by the vertex labels of its triangles. The search is NP-hard even for rational centres in general position and a fixed common isotropic covariance. For four such centres, six triangulations can be more probable than the centre triangulation. For sufficiently large n, at least $(n-1)!/2$ triangulations can outrank it. In one dimension, sorting the centres of independent Gaussians with common variance gives a Delaunay path with at least half the maximum probability in $O(n \log n)$ comparisons. In the plane, however, finding a triangulation with even half the maximum probability is NP-hard for independent standard Gaussian displacements. This hardness holds for every fixed fraction between zero and one, so a relative polynomial-time approximation scheme would imply $P = NP$. Both hardness proofs encode a planar cubic graph so that every qualifying triangulation returned by one call yields a maximum independent set of that graph. For Gaussian input, sampling finds, with probability at least $1 - \delta$, a triangulation whose probability is within ε of the maximum. Its running time is polynomial in the input bit length, $1/\varepsilon$ and $\log(1/\delta)$. Exact and constant-factor search for Delaunay triangulations remain NP-hard for fixed bounded rotation-invariant planar densities for which the logarithm of one plus the displacement length has a finite positive-order moment, and for fixed invertible linear images of these laws. This includes every fixed positive definite Gaussian covariance, uniform-ellipse and Cauchy noise, and laws without a positive moment of the displacement length. Bounded projection densities also suffice for some singular planar laws. For the convex hull, finding a most probable set of extreme labels is NP-hard for site-dependent isotropic Gaussian variances. Its complexity for a common variance remains open.

AI-use disclosure. OpenAI's ChatGPT and Codex, and Anthropic's Claude, were used as research tools to assist the author with literature retrieval, mathematical exploration, symbolic and numerical checks, illustrations and draft preparation and revision. Antigravity assisted with editorial review. The author formulated the research questions and mathematical framework, developed and refined the results and arguments with this assistance, and critically reviewed the manuscript throughout. Responsibility for the mathematical statements and the final manuscript rests entirely with the author.

Keywords. Delaunay triangulation; uncertain points; modal structure; Gaussian perturbation; approximation algorithms; NP-hardness.

Mathematics Subject Classification (2020). 68Q17; 68U05; 60D05; 68W25.

1 Introduction

For a planar set of at least three points in general position, the Delaunay triangulation divides its convex hull into triangles whose circumcircles enclose no other point. Suppose that each measured site is represented by a planar Gaussian centred at its reported position (Fig. 1). The joint density is largest when every site is at its centre, so the Delaunay triangulation of the centres is a natural representative. Is it also the most probable combinatorial structure? Its probability is the chance that the sampled sites give exactly the same labelled faces. Even for centres in general position, the configurations realizing one other set of labelled faces can have greater total probability than those realizing the centre triangulation.

More precisely, let $X_i \sim N_2(\mu_i, \sigma^2 I_2)$ be independent, with a common scale $\sigma > 0$. A configuration is in general position, or generic, if its points are distinct, no three are collinear, and no four are cocircular. Gaussian samples satisfy these conditions with probability one, so their Delaunay triangulation is almost surely unique. Its complete labelled type, or simply its type, lists the bounded triangular faces by their vertex labels. This list also determines the convex hull. We seek a type T maximizing

$$p_\mu(T) = \mathbb{P}\{\mathrm{DT}(X_1, \ldots, X_n) = T\}, \qquad p_* = \max_T p_\mu(T).$$

A type attaining p_* is called a mode.

1.1 Can the centres identify a likely triangulation?

The centre triangulation is easy to compute and reliable under small noise. Every generic configuration has a positive structural tolerance: perturbations below this tolerance preserve its Delaunay type [2, pp. 221, 223–224]. This stability underlies relocation filters [8]. Bounded uncertainty within the tolerance gives the centre type probability one. For Gaussian noise, its probability tends to one as the scale tends to zero. Does the centre type remain most probable at an arbitrary noise scale?

A counterexample needs only four sites. Take independent Gaussian points with common standard deviation 20 and rational centres

$$\mu_A = (-1, -1), \qquad \mu_B = (1, -1), \qquad \mu_C = (0, 1), \qquad \mu_D = (0, 0).$$

The centres are generic. Since D lies inside ABC, their triangulation has faces ABD, ACD and BCD. Compare it with the triangulation whose hull has cyclic order $ABCD$ and whose diagonal is AC (Fig. 2). Writing p_{centre} and p_{diagonal} for the probabilities of these two specific face lists, respectively, we obtain

$$8.35\% \le p_{\mathrm{centre}} \le 9.20\% \quad < \quad 10.39\% \le p_{\mathrm{diagonal}} \le 11.24\%.$$

All six labelled types with a four-vertex hull have probabilities in the second interval. Each stays close to the value they share when the centres coincide (Proposition 5). Dividing every centre by 20 gives the same type distribution under unit-variance noise. With four sites on the line, however, sorting the centres gives a most probable labelled Delaunay path under independent Gaussian displacements of any common positive variance (Section 9.3).

In this planar example, at least six labelled triangulations are more probable than the triangulation of the centres. With more sites, the number that outrank it can grow factorially. To see where so many competitors come from, start with coincident centres. Exchangeability gives the same probability to every relabelling of a given type. Combining this symmetry with the rarity of fully convex Gaussian samples yields, for every sufficiently large n, rational generic centres in convex position whose triangulation is outranked by at least $(n-1)!/2$ types. For the explicit centres $\mu_i = (i, i^2)/(n!)^3$, their triangulation has probability at most

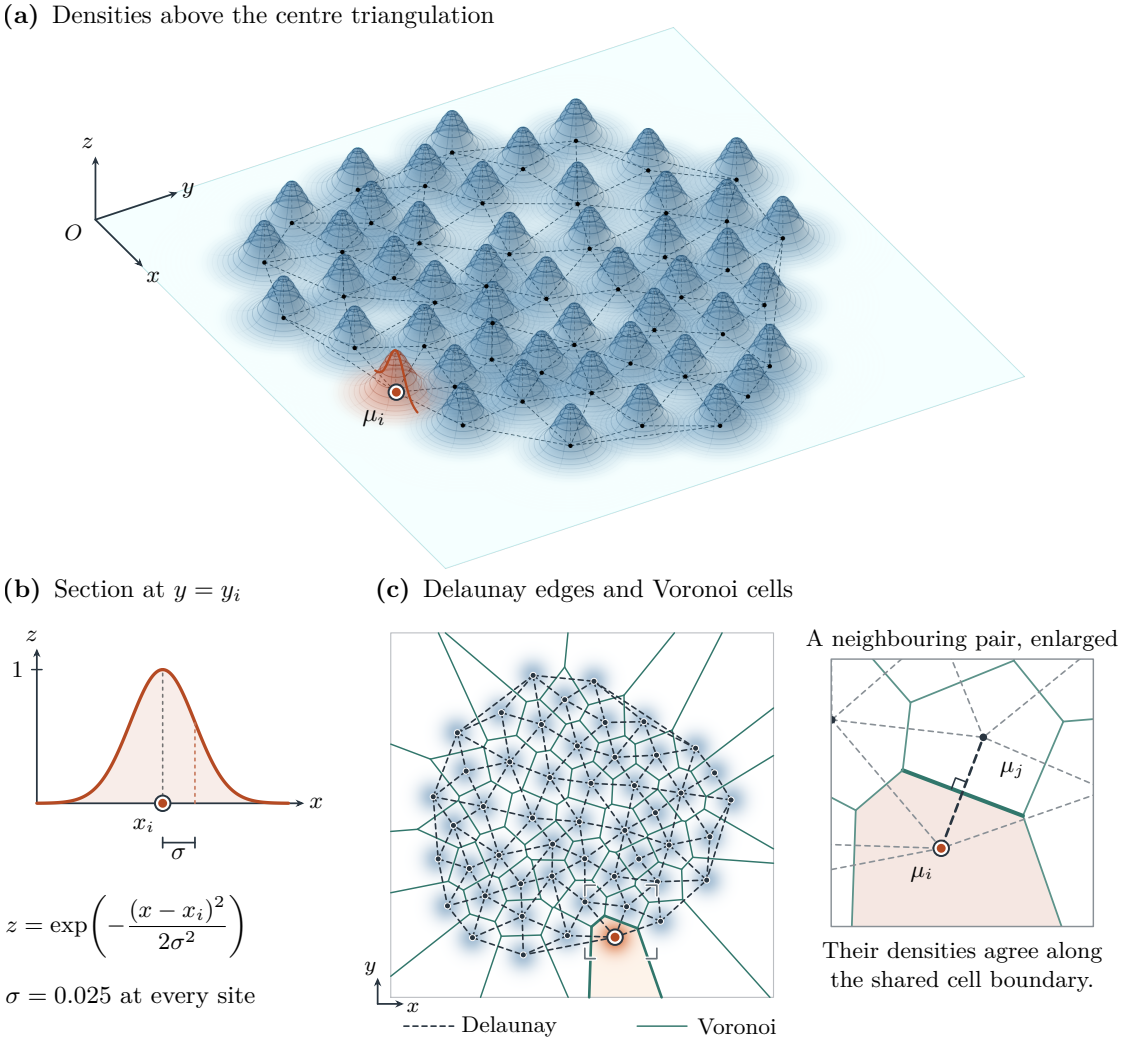

Figure 1: Fifty Gaussian sites with common standard deviation $\sigma = 0.025$. In (a), each bell shows a site density scaled to peak height one. Horizontal circles and four equally spaced vertical sections trace each surface. The ring marks the same centre $\mu_i = (x_i, y_i)$ in all three views. In (b), the orange curve shows the section highlighted in (a). In (c), equal widths make the closest centre the one with the greatest density. The shaded Voronoi cell shows where the marked site attains that maximum. The detail pairs one shared cell boundary with the Delaunay edge joining the two centres. Fig. 3 shows how sampling from these same sites can change the triangulation.

$\exp(-n \ln n + O(n))$ times the modal probability. Rational centres arbitrarily close to these have a unique mode of probability below $3/(n-1)!$ (Proposition 6, Fig. 5). An algorithm choosing only among sampled types and the centre type therefore needs factorially many samples for a fixed positive chance of finding the mode. This leaves open the possibility of efficient algorithms that construct unobserved types, motivating our complexity results.

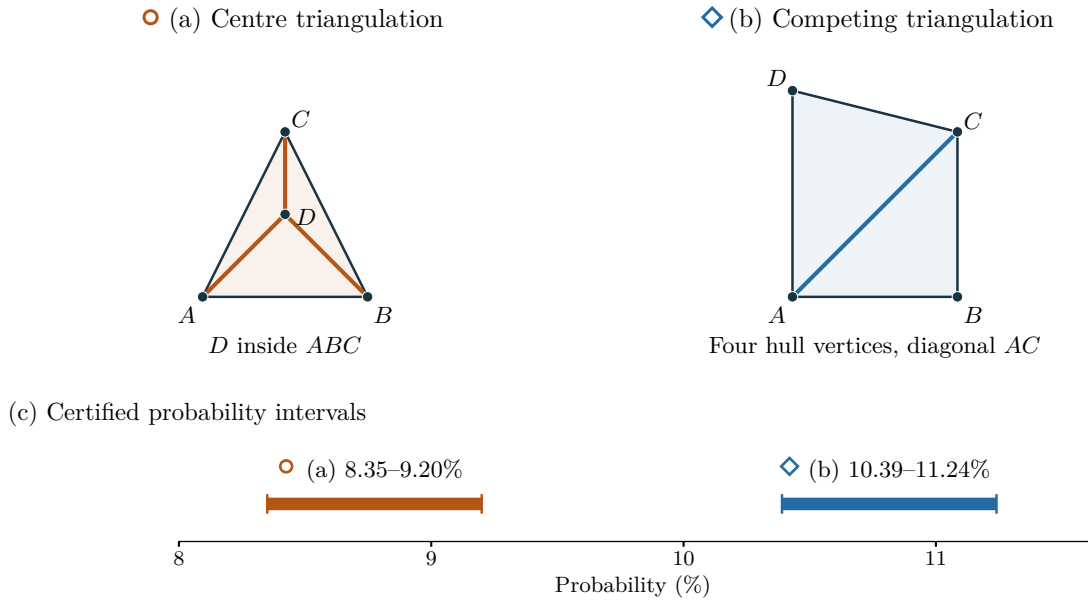

Figure 2: Four Gaussian centres whose Delaunay triangulation is less probable than a competitor. In (a), D lies inside ABC at the given centres. In (b), a configuration with D outside ABC illustrates the competing triangulation with diagonal AC. In (c), each interval bounds the probability of the corresponding labelled triangulation when the Gaussian sites are centred as in (a). The bounds in Proposition 5 separate the two probabilities.

1.2 Results and their scope

Our results distinguish exact modal search, relative approximation and additive sampling (Table 1). This distinction explains how polynomial-time additive guarantees coexist with the hardness of finding a mode or a type attaining a fixed fraction of its probability.

Theorem 1. *Given rational centres $\mu_1, \dots, \mu_n \in \mathbb{Q}^2$ in general position, finding any complete labelled Delaunay triangulation of maximum probability for independent $X_i \sim N_2(\mu_i, I_2)$ is NP-hard.*

The reduction constructs Gaussian centres in polynomial time from a planar cubic graph, which can be drawn without crossings and has three edges at every vertex. From any most probable triangulation, it recovers in polynomial time a maximum independent set, a largest set of pairwise nonadjacent vertices. One call to a solver is therefore sufficient, regardless of how ties are resolved. The general reduction encodes several parameter choices at once. For Gaussian displacements one explicit choice suffices (Proposition 17). Section 8.2 completes the proof of Theorem 1. Sections 5–8 also prove the more general Theorem 3.

On the line, sorting the centres of independent Gaussians with common variance finds a most probable left-to-right order [29, Theorem 1]. This takes $\Theta(n \log n)$ comparisons in the worst case. Each labelled Delaunay path corresponds to two opposite orders. Its probability is the sum of their probabilities. Consequently, the path obtained by sorting has at least half the maximum path probability (Proposition 22). Sorting is exact for any number of sites when the multiset of means is symmetric about a point. To the best of our knowledge, it remains open whether finding a most probable unoriented Delaunay path is NP-hard for independent Gaussian sites with a common variance and arbitrary n.

In the plane, however, finding a triangulation with even half the maximum probability is NP-hard. Theorem 21 establishes this hardness for every fixed fraction $c \in (0, 1)$, including

under independent standard Gaussian displacements. A relative PTAS would therefore imply $P = NP$, since it would achieve every fixed factor below one in polynomial time.

These planar hardness results extend to fixed laws with a bounded radial density and a logarithmic tail bound (condition (L), Definition 2). This bound allows the probability of a very large displacement to decay as slowly as an inverse power of the logarithm of its length. Section 9 proves hardness even for a factor of $2^{-n^{1-\epsilon}}$, for each fixed $\epsilon \in (0, 1)$, where n counts sites. Section 10 extends the reductions to fixed invertible linear images, including Gaussian displacements with any fixed positive definite covariance.

For planar Gaussian sites, sampling can find a triangulation whose probability comes within a chosen distance ε of the maximum p_*. For example, if the maximum is 12%, an additive error of one percentage point allows any triangulation with probability at least 11%. The algorithm of Theorem 8 meets this guarantee with probability at least $1 - \delta$. Each site may have its own covariance matrix, supplied through a rational invertible 2×2 factor. The algorithm runs in time polynomial in the input length, $1/\varepsilon$ and $\log(1/\delta)$. This additive guarantee becomes uninformative when $p_* \leq \varepsilon$. A supplied positive lower bound on p_* permits a relative guarantee. An additional lower bound on the relative gap between a unique mode and its competitors permits exact recovery, with the costs stated in Table 1.

In the reductions for triangulations, we use the same displacement law for every instance. The input then consists of the rational site centres in general position. Any finite positive moment of the displacement length suffices for (L). Even some laws with no such moment satisfy (L), as the examples in Section 11 show.

Table 1: Summary of the results. Here L is the binary input length and p_* the modal probability. Types are planar except in the row marked 1D. The hardness reductions for triangulations use independent displacements from a fixed law, rational generic centres and one solver call. The sampling algorithms succeed with probability at least $1 - \delta$.

Task	Noise law	Result	Location
Most probable type	standard Gaussian	NP-hard	Thm. 1
Most probable type	bounded radial density satisfying (L)	NP-hard	Thm. 3
Type with $p \geq c\,p_*$, fixed $c \in (0, 1)$	bounded radial density satisfying (L)	NP-hard; no PTAS unless $P = NP$	Thm. 21
Both tasks above	fixed linear image of a rotation-invariant law with bounded projection density and (L)	NP-hard	Thm. 23
Path with $p \geq p_*/2$ (1D)	Gaussian, common variance	deterministic; $O(n \log n)$ comparisons	Prop. 22
Type with $p \geq p_* - \varepsilon$	Gaussian, rational centres and invertible rational 2×2 covariance factors	time $\mathrm{poly}(\mathsf{L}, 1/\varepsilon, \log(1/\delta))$	Thm. 8
Type with $p \geq (1 - \varepsilon)p_*$	same Gaussian input, given rational $0 < \underline{p} \leq p_*$	time $\mathrm{poly}(\mathsf{L}, 1/\varepsilon, 1/\underline{p}, \log(1/\delta))$	Prop. 9
The unique Delaunay mode	same Gaussian input, given inverse-polynomial lower bounds on p_* and the relative gap	randomized polynomial time	Prop. 10
Most probable hull	isotropic Gaussian, individual scales	NP-hard	Thm. 24

1.3 Probabilistic geometry and related work

Two closely related lines of work motivate our study: modal geometric structures and the distribution of Delaunay types under random perturbations. Their models clarify what must be proved for independent Gaussian sites.

Suri, Verbeek and Yıldız [33, Theorems 3 and 4] prove NP-hardness and strong relative inapproximability for the most likely planar convex hull in their multipoint model, where each point has finitely many possible locations and may also be absent. Their reduction encodes clause satisfaction by enclosing designated points inside the hull [33, Fig. 3 and Lemmas 8 and 9]. Their conclusion explicitly proposes analogous questions for Delaunay triangulations and Voronoi diagrams [33, Section 5]. Theorem 1 addresses this direction using one fixed Gaussian law translated to rational generic centres. The choice of model matters: for fixed locations with independent presence probabilities, their planar hull problem admits an $O(n^3)$ algorithm [33, Theorem 1].

The complete labelled Voronoi diagram has the same type probabilities by duality. In general position, each Voronoi vertex records the three site labels of one Delaunay triangle. Conversely, triangles sharing a pair of labels give adjacent Voronoi vertices, and a pair occurring in just one triangle gives an unbounded ray. This correspondence and its inverse take polynomial time, so both exact and relative hardness carry over to the incidence structure of the whole labelled diagram.

One might try to transfer hull hardness by asking for a most probable Delaunay type and reading off its hull. This would require every most probable type of the same instance to have a most probable hull. Under a common distribution of configurations, the probability of a hull type is the sum of the probabilities of all complete triangulation types with that hull, and maximizing one term does not guarantee a maximum for the sum. The models differ as well. In the multipoint model, the output polygon has specified vertex locations [33, Appendix H], an event of probability zero under Gaussian displacements, so we group realizations by their extreme labels. Imitating the finitely many locations of a point by narrow Gaussian components would give each site a mixture, whereas our Delaunay reduction uses a single isotropic Gaussian per site with a common covariance. For the hull itself, Theorem 24 gives a separate reduction from Subset Sum when the Gaussian variances may differ. It applies both to the set of extreme labels and to their cyclic order.

Khanimov and Sharir [22, Section 2, Fig. 2.2] study Delaunay triangulations of grids and regular polygons under small Gaussian perturbations. They use linearized incircle tests to identify the most and least probable classes for the 3×3 grid and compare the resulting probabilities with simulations. Bandyopadhyay and Snoeyink [3, Section 5] express the probability that a tetrahedron is Delaunay under isotropic Gaussian displacements as an integral over circumspheres, evaluated by Monte Carlo. Leipold et al. [23] derive probabilities for Voronoi-cell topologies in planar lattices under infinitesimal perturbations with a radially symmetric joint law. These studies describe local structures or triangulations near degenerate configurations. We maximize probability over complete labelled types for arbitrary rational generic centres.

A related question asks for a region that covers many sampled sites. For n sites with normal location distributions, Atalay, Friedler and Xu [37, Definition V.1 and Theorem V.2] compute the convex hull of the mean positions and output the convex hull of the confidence regions of its vertices. If each local test certifying the hull of the means also holds for the random positions with probability at least ϕ, the output region contains at least ϕn of the sites in expectation. Theorem 24 asks which set of labels is most likely to occur on the hull. Results on geometric features include reliable incircle tests for disks [10] and rectangles [21], guaranteed Voronoi diagrams and common Delaunay edges [11], first-order stability tests allowing dependent locations [17], and almost-Delaunay simplices under bounded

perturbations [3]. The most likely Voronoi diagram of Suri and Verbeek [32] identifies, at each query location, the most probable nearest neighbour among fixed sites with independent presence probabilities. Reconstruction methods preprocess bounded uncertainty regions to accelerate triangulation after coordinates are revealed [25, 9, 7].

A related approximation question arises in Bayesian networks: given observed variables, find values of the hidden variables whose conditional probability is within a fixed factor of the maximum. Abdelbar and Hedetniemi [1] proved hardness for this task. For independent copies, the probability ratio for a joint outcome is the product of the individual ratios. This principle appears in the inference reduction of Park and Darwiche [30, Theorem 8 and Appendix A] and the uncertain-hull construction of Suri et al. [33, Lemma 7]. For triangulations, the challenge is to keep this product when all copies belong to one geometric output (Section 9).

1.4 Proof ideas and classical tools

We encode an entire planar cubic graph in one configuration of sites, which we call a block. Each vertex becomes a cycle of quadrilaterals. Each quadrilateral can choose one of two diagonals, recorded as a sign. Shared sites make matching signs more probable along a cycle. An additional quadrilateral between the cycles of two adjacent vertices, called a bridge, favours opposite signs more weakly. A small shift at one site per cycle sets the reward for selecting a vertex. In the resulting score, every selected vertex earns the same reward and each edge joining two selected vertices costs more than that reward. Removing either endpoint of such an edge improves the score. Among independent sets, the score increases with cardinality. Section 7 bounds the approximation error tightly enough to preserve each improvement. Thus every most probable vector of first-order diagonal signs identifies a maximum independent set.

For Gaussian displacements, or more generally for a fixed law in Theorem 3 with a finite second moment, one block suffices (Proposition 17). A reader interested only in Theorem 1 can follow Sections 5–7, this proposition and the proof in Section 8.2. For a general law covered by the theorem, the suitable parameters depend on quantities that the reduction does not compute, so we place several blocks with different parameter choices in the same input. At least one choice is suitable. Given any most probable triangulation, we read the diagonals chosen in each block. We retain the vertex set of a block when its diagonals give a common sign on each cycle and the positive cycles form an independent set. We return a largest retained set. Section 8 proves that at least one of them is maximum.

Could a solver find a triangulation with even half the maximum probability? To rule this out, we repeat the construction. If none of the independent copies yields a maximum independent set, each copy incurs a relative probability loss. The probability ratios multiply. When the copies are joined in one geometric triangulation, a site with a large displacement can destroy some local diagonal choices. Our construction limits how many blocks one such site can affect. We retain the probability factors from the remaining blocks and show that every output meeting the required fraction still decodes to a maximum independent set (Section 9).

For approximation, we sample realizations and record their Delaunay types, using the sampling principle of Löffler and Phillips [24]. Jørgensen et al. [18, Theorem 4] approximate the cumulative distribution of a scalar geometric statistic using finite supports. Here, rounding each Gaussian coordinate to one of q rational quantile representatives changes the distribution of complete types by at most $O(n^4/q)$ in total variation (Section 4). The Dvoretzky–Kiefer–Wolfowitz inequality with Massart's constant [27] controls empirical mode selection.

Sections 10–12 treat elliptically contoured noise, examples at the boundary of the tail condition and the convex hull. Appendix A gives the coordinates, the normalization of the incircle derivative and the completion points of the construction, and Appendix B the rational parameters of the hull reduction.

2 The model and the tail condition

The Delaunay reduction must place its sites far enough apart to control the effect of large displacements. We describe which laws permit the required distances to be encoded with polynomially many bits.

Fix a probability law ν on $\mathbb{R}^2$, let $Z_1, \dots, Z_n$ be independent with law ν, and set $X_i = \mu_i + Z_i$, so that μ_i identifies site i even when ν has no mean. For the planar densities considered below, independent sites are in general position almost surely. Section 10 establishes the same property for the wider class used there. There are finitely many types, and their probabilities $p_\mu(T) = \mathbb{P}\{\mathrm{DT}(X_1, \dots, X_n) = T\}$ therefore have a maximum. Any type attaining this maximum is an acceptable output. Fig. 3 shows three realizations of the sites of Fig. 1.

We measure input length in bits, using the binary numerators and denominators of the rational centre coordinates. For hull search (Section 12), the input also gives one rational scale per site. For the sampling algorithms of Section 4, the input also includes the rational covariance factors.

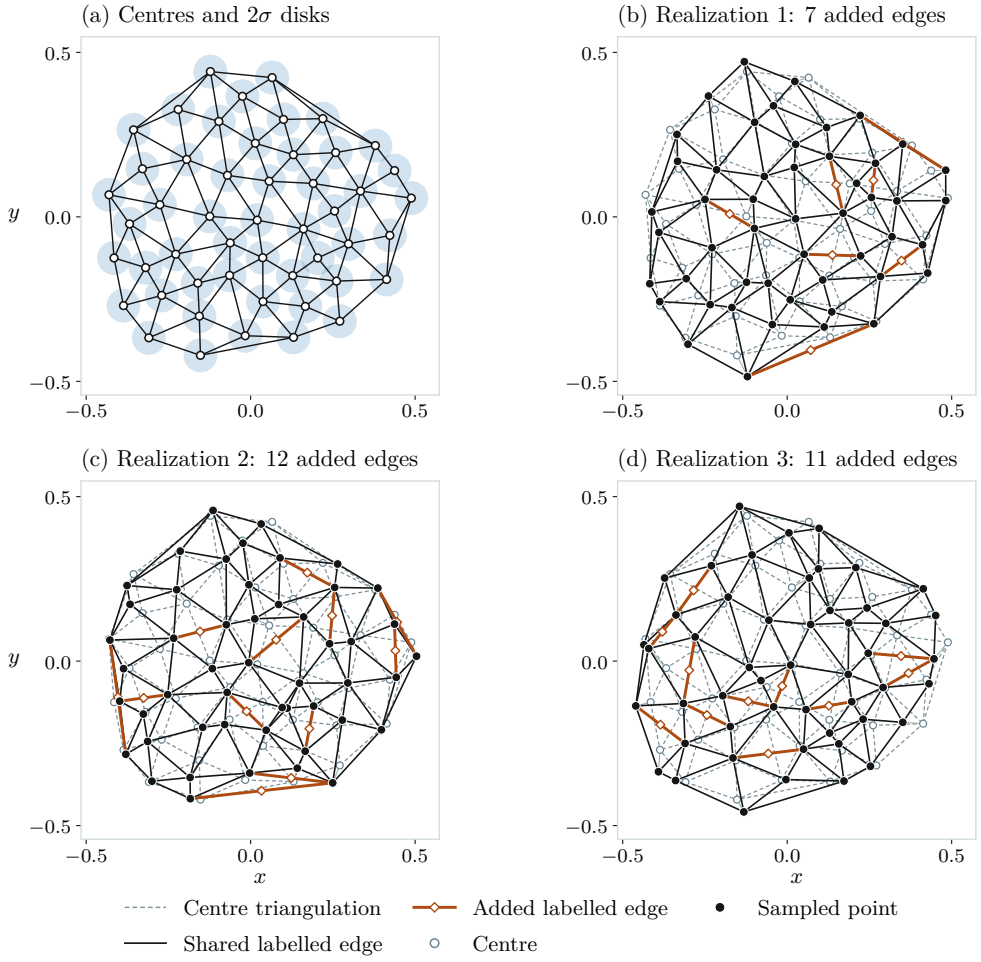

Figure 3: The same fifty labelled sites before and after Gaussian perturbation with $\sigma = 0.025$. In (a), hollow markers show the centres, with blue disks of radius 2σ and the centre triangulation. In (b)–(d), solid markers and thin displacement segments show three independent full realizations at the same scale. The centre triangulation appears as a fine dashed reference beneath each realization. Dark solid edges retain label pairs from (a), orange edges with a small diamond connect new label pairs, and the counts give the number of label pairs absent from (a).

Definition 2. Let Z have law ν on $\mathbb{R}^2$, and put $T_\nu(R) = \mathbb{P}\{\|Z\| > R\}$ for every real $R \geq 0$. We say that ν satisfies condition *(L)* if there are a fixed integer $C_L \geq 1$ and a fixed rational $\gamma > 0$ such that

$$T_\nu(2^s) \leq C_L s^{-\gamma} \qquad \text{for every integer } s \geq 1. \tag{L}$$

Theorem 3. *Fix a probability law ν on $\mathbb{R}^2$ with Lebesgue density $\rho(z) = \varphi(\|z\|)$, where $0 \leq \rho \leq M_\rho < \infty$, satisfying* (L)*. Given rational centres $\mu_1, \dots, \mu_n$ in general position, finding any most probable complete labelled Delaunay triangulation of the sites $X_i = \mu_i + Z_i$, with independent displacements Z_i of law ν, is NP-hard. Construction and decoding take polynomial time, and the reduction makes one modal oracle call.*

Sections 5–8 prove Theorem 3. Once the law ν is fixed, finitely many rational constants depending on it suffice for a single reduction that handles every planar cubic input graph. The algorithm uses those constants directly. It makes one query, and every answer, including any tied mode, decodes to a maximum independent set.

It remains open whether one algorithm could take a description of the law as an additional input. We also leave open whether comparing the probabilities of two given types exactly is in NP, meaning that every yes answer has a polynomially verifiable certificate.

Proposition 4 characterizes (L) through a moment of the logarithm of the displacement length and through a geometric property. With probability at least $1/2$, one disk whose radius has polynomial bit length contains all n displacements. In the reduction, Section 8 uses (L) directly to choose radii at the higher confidence levels needed to compare triangulation probabilities.

Proposition 4. *Condition* (L) *holds if and only if $\mathbb{E}[(\log(1 + \|Z\|))^a] < \infty$ for some $a > 0$. Equivalently, there is a polynomial p such that for every $n \geq 1$ the independent displacements $Z_1, \dots, Z_n$ with law ν satisfy*

$$\mathbb{P}\Big\{ \max_{1 \leq i \leq n} \|Z_i\| \leq 2^{p(n)} \Big\} \geq \tfrac{1}{2}.$$

Proof. We express both conditions through the tail of the displacement length. Put $W = \log_2(1 + \|Z\|)$. Under (L), for $u \geq 4$, the inequalities $2^u - 1 \geq 2^{u-1}$ and monotonicity give

$$\mathbb{P}(W > u) \leq T_\nu(2^{\lfloor u-1 \rfloor}) \leq C_L \lfloor u - 1 \rfloor^{-\gamma} \leq C_L 2^\gamma u^{-\gamma}.$$

Choose $0 < a < \gamma$. The tail identity $\mathbb{E}W^a = a \int_0^\infty u^{a-1} \mathbb{P}(W > u)\, du$ proves finiteness. Conversely, if $\mathbb{E}W^a < \infty$, choose a rational $0 < \gamma < a$. Markov's inequality gives $T_\nu(2^s) \leq \mathbb{P}(W > s) \leq \mathbb{E}W^a\, s^{-a} \leq C_L s^{-\gamma}$ for every integer $s \geq 1$, after increasing the integer C_L. Changing the logarithm base multiplies the moment by a constant.

For confinement, (L) permits $R_n = 2^{s_n}$ with $s_n = \lceil (2C_L n)^{1/\gamma} \rceil$. An integer polynomial bounds this exponent, and a union bound gives the required probability. Conversely, independence and confinement give $(1 - T_\nu(2^{p(n)}))^n \geq 1/2$, hence $T_\nu(2^{p(n)}) \leq (\ln 2)/n$. Choose constants $C \geq 1$ and an integer $k \geq 1$ with $p(n) \leq Cn^k$. For large s, take $n = \lfloor (s/C)^{1/k} \rfloor$. Then $p(n) \leq s$, so monotonicity gives $T_\nu(2^s) = O(s^{-1/k})$. Enlarging the constant for small s proves (L). □

3 How far can the centres be from the mode?

Choosing the Delaunay type of the centres can fail even at four generic sites. For large n, factorially many types can be more probable. We first certify the four-site comparison, then combine relabelling symmetry with the rarity of fully convex Gaussian samples.

3.1 An explicit example with four centres

We begin with the possible outputs on four labels, then certify the example of Section 1.1. Four generic labelled sites have ten complete types: six have a four-vertex hull and two bounded faces, and four have a three-vertex hull and three bounded faces. Kabluchko and Panzo [19, Theorem 3.2] compute probabilities of combinatorial hull types for $d+2$ i.i.d. standard Gaussian points in $\mathbb{R}^d$. In the planar case, we further distinguish the complete labelled Delaunay types. Before certifying the counterexample, we look at a configuration that behaves differently. Fig. 4 shows how the frequencies of the ten types change with the noise scale for a separate, nearly square configuration, with centres $A=(0,0)$, $B=(1,0)$, $C=(1,1.03)$, $D=(0,1)$ and Delaunay diagonal BD. We generated 4 000 000 realizations at each of four scales. Only the two diagonal alternatives occurred at the two smaller scales, and all ten types at the two larger ones. At $\sigma=0.4$, about 29% of the realizations changed the order or size of the hull. The centre type was the most frequent at every sampled scale, with its lead over the next type falling from 71 to 1.2 percentage points.

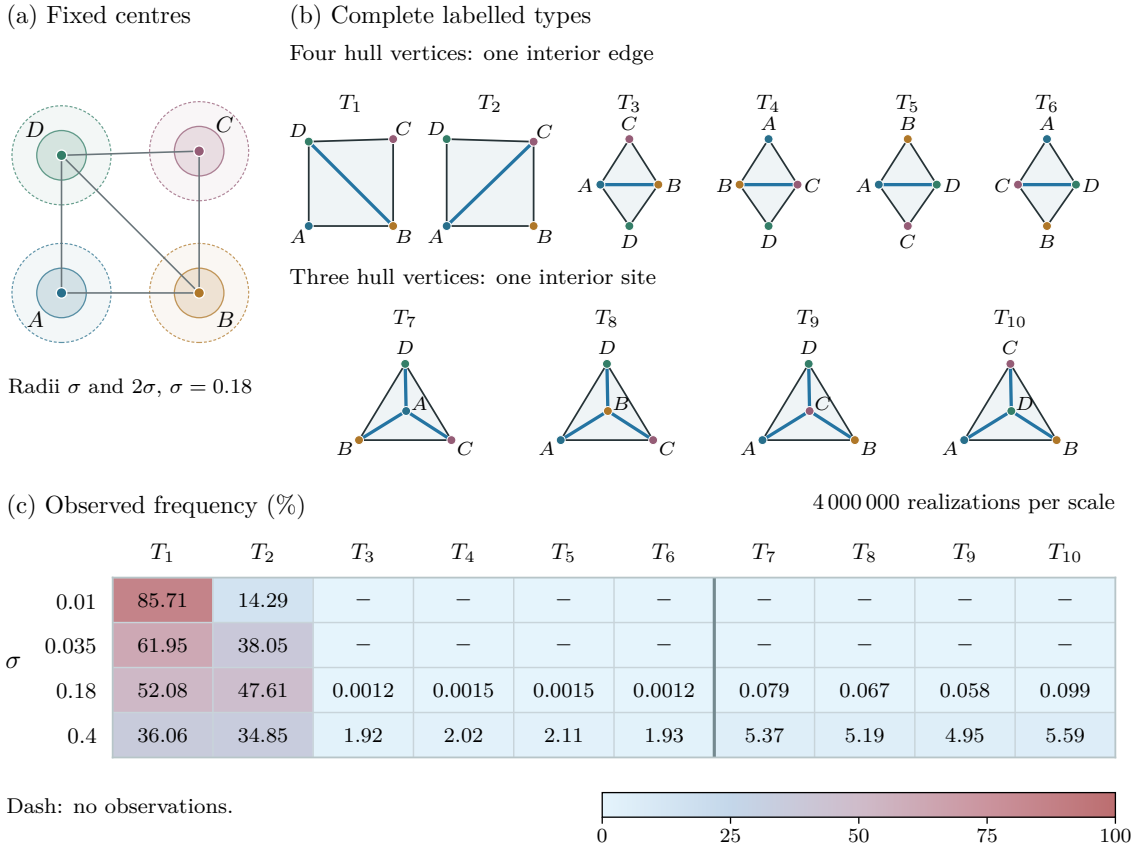

σ	T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9	T_{10}
0.01	85.71	14.29	–	–	–	–	–	–	–	–
0.035	61.95	38.05	–	–	–	–	–	–	–	–
0.18	52.08	47.61	0.0012	0.0015	0.0015	0.0012	0.079	0.067	0.058	0.099
0.4	36.06	34.85	1.92	2.02	2.11	1.93	5.37	5.19	4.95	5.59

Figure 4: The ten complete types on four labels. In (a), disks show radii σ and 2σ at $\sigma=0.18$, and the edges show the centre type T_1. In (b), canonical drawings identify the face lists. Their coordinates are chosen for legibility. In (c), each column gives the percentages in 4 000 000 realizations. Low frequencies are light blue and high frequencies dark red. The printed entries give the percentages independently of colour. Fig. 2 compares T_{10} and T_2 at a different set of centres.

Under independent Gaussian sampling, Hoeffding's inequality [14, Theorem 2] and a union bound over the 40 frequencies give simultaneous error below 0.0961 percentage point at confidence 0.95. This is less than half of each observed lead, so the centre type ranks first at all four scales at that confidence level. Every type has positive probability because a nonempty open set of configurations realizes it. A separate 95% bound puts every unobserved

probability below $\ln(800)/N < 1.7\cdot10^{-6}$, since a type of probability p is absent in $N = 4\,000\,000$ independent draws with probability at most e^{-Np}.

These bounds assume ideal sampling. The seeded binary64 generator has no certified bound on its sampling or rounding error. CGAL [36] computes each face list, and a second classifier checks it with exact determinants on the same binary64 coordinates. For rare types, we also computed marginal Wilson intervals [15, p. 210] from the batch counts.

Proposition 5. *Let X_A, X_B, X_C, X_D be independent planar Gaussian points with common standard deviation 20 and centres $(-1,-1)$, $(1,-1)$, $(0,1)$, $(0,0)$, respectively. The centres are in general position and their Delaunay type is $T_{10} = \{ABD, ACD, BCD\}$. Each complete type with a three-vertex hull has probability in $[0.0835, 0.0920]$, whereas each complete type with a four-vertex hull has probability in $[0.1039, 0.1124]$. In particular, at least six complete types are strictly more probable than T_{10}.*

Proof. We first identify the centre type. We then compute the probabilities for coincident centres and bound how much separating the centres can change any of them. The orientation determinants of ABC, ABD, ACD and BCD are 4, 2, -1 and 1, and the incircle determinant of $ABCD$ is 6, so the centres are generic. Since $\mu_D = \mu_A/4 + \mu_B/4 + \mu_C/2$, the point D lies strictly inside ABC and the centre type is T_{10}.

At coincident centres the four sites are independent and identically distributed. Write $p_0(T)$ for the probability of type T in this configuration. For the probability that all four sites are hull vertices, Blatter [13, p. 837] proves $P_{\text{convex}} = (6/\pi)\arcsin(1/3)$ and credits the original proof to Maehara [12]. In every other generic realization, one point lies inside the triangle of the others. Relabelling acts transitively on the six convex types and on the four interior-point types, so each convex type has probability q and each interior-point type probability r, where

$$q = \frac{P_{\text{convex}}}{6} = \frac{\arcsin(1/3)}{\pi}, \qquad r = \frac{1 - P_{\text{convex}}}{4} = \frac{1}{4} - \frac{3q}{2}.$$

To separate the centres, subtract their centroid $\bar{\mu} = (0, -1/4)$ from all four points and divide by 20. This preserves types and gives the joint law $Q_+ = N_8(v, I_8)$ with

$$v = ((\mu_i - \bar{\mu})/20)_{i=A,B,C,D} \in \mathbb{R}^8, \qquad a = \|v\|^2 = \frac{19}{1600}.$$

For Gaussian sites with common standard deviation σ, joint laws with mean vectors μ and μ' have relative entropy $\|\mu - \mu'\|^2/(2\sigma^2)$, so Pinsker's inequality [4, Theorem 4.19] gives

$$|p_\mu(T) - p_{\mu'}(T)| \le \frac{\|\mu - \mu'\|}{2\sigma} \qquad \text{for every type } T. \tag{1}$$

Applied against coincident centres, this gives the bound $\sqrt{a}/2 > 0.054$, which exceeds $q - r$ and cannot separate the two probabilities. A simultaneous half-turn $z \mapsto -z$ preserves each type event, so it has the same probability under Q_+, under $Q_- = N_8(-v, I_8)$ and under their mixture $Q_s = (Q_+ + Q_-)/2$. Averaging cancels the term linear in v in the likelihood ratio ψ relative to the coincident-centre law $Q_0 = N_8(0, I_8)$. Since $v \cdot z \sim N(0, a)$ under Q_0,

$$\psi(z) = e^{-a/2}\cosh(v \cdot z), \qquad \mathbb{E}_0\psi^2 = e^{-a}\frac{1 + e^{2a}}{2} = \cosh a.$$

This gives a common bound $\eta = \frac{1}{2}\sqrt{\cosh a - 1} = O(a)$ on the change in each type probability. For two laws P, Q, their total variation distance $d_{\text{TV}}(P, Q) = \sup_E |P(E) - Q(E)|$ is the largest probability change over measurable events. Since $\mathbb{E}_0\psi = 1$, Cauchy–Schwarz gives

$$|p_\mu(T) - p_0(T)| \le d_{\text{TV}}(Q_s, Q_0) = \frac{1}{2}\mathbb{E}_0|\psi - 1| \le \frac{1}{2}\sqrt{\mathbb{E}_0\psi^2 - 1} = \eta.$$

Thus a four-vertex hull type has probability between $q-\eta$ and $q+\eta$. For a three-vertex hull type the same bounds hold with r in place of q. Substituting $0.10817 < q < 0.10818$, $r = 1/4 - 3q/2$ and $\eta < 0.0042$ gives the intervals in the proposition after rounding the lower endpoints down and the upper endpoints up. The rational certificate bounds the series for $\arcsin(1/3)$ and $\cosh(19/1600) - 1$ and uses Machin's formula for π. □

The bounds establish six strict comparisons, with the identity of a mode still undetermined. These comparisons persist under small changes of the centres: if centres $\mu' \in \mathbb{R}^8$ with standard deviation 20 satisfy $\|\mu - \mu'\| < 0.23$ and keep μ'_D strictly inside $\mu'_A\mu'_B\mu'_C$, they remain generic with type T_{10}, and (1) with $\sigma = 20$ changes every type probability by at most $\|\mu - \mu'\|/40 < 0.0058$. Each convex type then has probability above 0.0981 and each interior-point type below 0.0978.

3.2 Factorially many more probable types

Relabelling symmetry supplies many competitors to a fully convex centre type. We preserve their advantage by separating coincident centres along a generic convex configuration (Fig. 5), and a final perturbation makes the mode unique.

Proposition 6. *For all sufficiently large n, take the rational centres $\mu_i = (i, i^2)/(n!)^3$, $1 \le i \le n$. They are in generic convex position. For independent $X_i \sim N_2(\mu_i, I_2)$, at least $(n-1)!/2$ types are strictly more probable than the centre triangulation. Its probability is at most $\exp(-n \ln n + O(n))$ times the modal probability. Arbitrarily close rational centres retain these bounds and can be chosen to have a unique modal type with*

$$p_* = \max_T p_\mu(T) < \frac{3}{(n-1)!}.$$

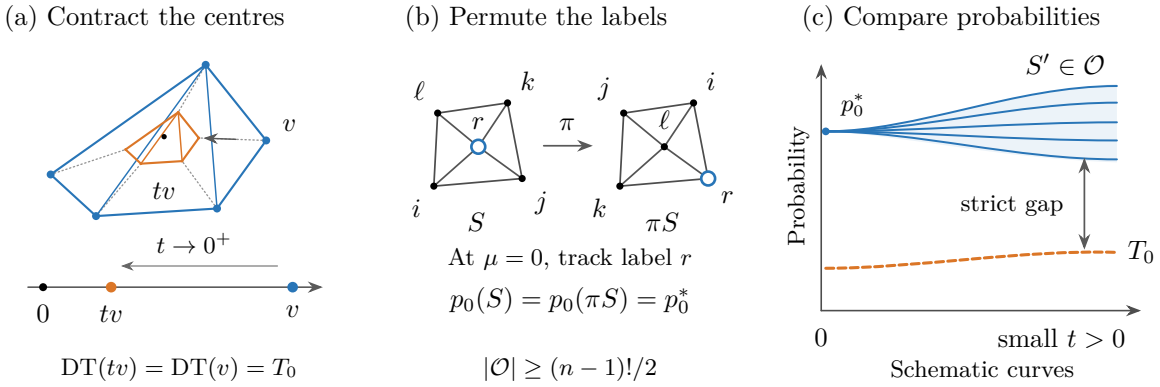

Figure 5: Why the centre type can have factorially poor rank. In (a), contracting a generic convex configuration v along its ray in $\mathbb{R}^{2n}$ preserves T_0 at a fixed noise scale. In (b), π permutes the labels of a modal type S. At coincident centres, every type in its relabelling orbit $\mathcal{O}$ has the same probability. In (c), for sufficiently large n, every type in $\mathcal{O}$ is strictly more probable than the fully convex type T_0, and the advantage persists for small $t > 0$. The small triangulations illustrate the two operations. The probability curves are schematic, and the rank conclusion applies for sufficiently large n.

Proof. For each complete labelled type T, let $p_0(T)$ be the probability that n independent planar points with law $N_2(0, I_2)$ have type T, and set $p_0^* = \max_T p_0(T)$. We first bound the number of labelled types obtained by relabelling one type. A relabelling that fixes a realizable type T is an automorphism of its complex of bounded faces. It preserves the boundary cycle,

formed by the edges lying in only one bounded face, and acts on it as an element of the dihedral group of order $2h \le 2n$, where h is its length. This element determines the automorphism: the images of the ends of a boundary edge determine the image of the unique face on that edge, hence of its third vertex. Likewise, known images of the vertices of a face determine the image of the face across each of its interior edges, hence of its third vertex, and these steps reach every face through the connected face-adjacency graph of the triangulated disk. At most $2n$ of the $n!$ relabellings fix T, so its orbit has at least $n!/(2n)$ elements. Their events are disjoint and have equal probability by exchangeability. Consequently,

$$|\operatorname{Orb}(T)| \ge \frac{(n-1)!}{2}, \qquad p_0(T) \le \frac{2}{(n-1)!}. \tag{2}$$

We next need a lower bound on p_0^*, obtained by counting all realizable types. Place a new vertex 0 in the outer face and join it to each boundary vertex. This triangulates the sphere. The resulting graph is simple and maximal planar on $n+1 \ge 4$ vertices, hence 3-connected. Whitney's theorem determines its facial cycles [5, Theorem 11], or equivalently its spherical embedding up to reflection [6, Theorem 1.2]. The faces that omit 0 recover the type. This construction maps distinct types to distinct labelled planar graphs on $n+1$ vertices, whose number is $(n+1)!\, e^{O(n)} \le n!\, C^n$ for an absolute constant $C > 1$ [16, Theorem 1]. Hence there are at most $n!\, C^n$ realizable types, and since their probabilities sum to one,

$$p_0^* \ge \frac{1}{n!\, C^n}. \tag{3}$$

Let c_n be the probability that n independent points with law $N_2(0, I_2)$ are all vertices of their convex hull. Order them from left to right. The abscissas determine this ordering and are independent of all ordinates, so the reordered ordinates are still independent $N(0,1)$ variables. Now consider a realization in convex position. Going from left to right, the upper hull rises to its highest vertex and then falls. The lower hull falls to its lowest vertex and then rises. Assign the leftmost and rightmost vertices to the upper chain, and include each turning vertex in just one of its two runs. This partitions the ordinates, still in left-to-right order, into four possibly empty monotone subsequences [35, Section 2]. There are at most 4^n assignments of points to the four runs. For any fixed assignment with run sizes $m_1, \dots, m_4$, the four runs use disjoint sets of ordinates. By independence and symmetry of the ranks, the probability that each run follows its prescribed order is $1/(m_1! \cdots m_4!) \le 4^n/n!$. The inequality follows from the multinomial theorem. A union bound gives

$$c_n \le \frac{16^n}{n!}. \tag{4}$$

If all vertices of a type T_0 lie on the boundary, its whole orbit lies in the convex-position event, and (2) gives $p_0(T_0) \le 2c_n/(n-1)!$. By (3) and (4), $p_0(T_0) < p_0^*$ as soon as $2n(16C)^n < n!$, which holds for all sufficiently large n. Fix such an n and a modal type S of p_0. By (2), the orbit $\mathcal{O}$ of S has at least $(n-1)!/2$ types, each of probability p_0^* and thus more probable than every fully convex type.

For the centres, take $v_i = (i, i^2)$, $1 \le i \le n$. Tangents to the parabola expose every point, and no line meets three of them. A circle $x^2 + y^2 + a_1 x + a_2 y + a_3 = 0$ through four of them would make their four distinct positive abscissas the four roots of $x^4 + (1+a_2)x^2 + a_1 x + a_3$, which sum to zero. Hence v is generic and in convex position, $T_0 = \mathrm{DT}(v)$ is fully convex, and $\mathrm{DT}(tv) = T_0$ for every $t > 0$.

Choose $t = (n!)^{-3}$. The coordinates of tv take $O(n \log n)$ bits each. By (1), moving the centres from zero to tv changes each type probability by at most $t\|v\|/2 \le n^3/(n!)^3$. This is $o(1/(n!C^n))$, so each type in $\mathcal{O}$ still has probability at least $1/(2n!C^n)$ for large n. For the

centre type, (2) and (4) instead give

$$p_{tv}(T_0) \le \frac{2 \cdot 16^n}{n!(n-1)!} + \frac{n^3}{(n!)^3} < \frac{3 \cdot 16^n}{n!(n-1)!}.$$

Dividing by the lower bound for any member of $\mathcal{O}$ gives a ratio below $6(16C)^n/(n-1)! = \exp(-n \ln n + O(n))$. Also, every type has probability at most $2/(n-1)! + n^3/(n!)^3 < 3/(n-1)!$ for large n. At tv every orientation and incircle determinant is nonzero, so sufficiently small changes of the centres preserve the hull and its triangulation T_0. By (1) every type probability is continuous in the centres. The bounds just proved are strict, the competing types have positive probabilities, and there are finitely many types for this n. Hence one neighbourhood preserves all the comparisons and both strict upper bounds. Within it we can now break the ties.

It remains to choose centres with a unique mode. The function $p_\mu(T)$ is the convolution of the indicator of the type region with a nondegenerate Gaussian density. It is real analytic in μ: for complex means, the integrand is holomorphic and, on each compact set of means, is bounded by a constant times $e^{-\|x\|^2/4}$. For distinct realizable types T and T', the difference $p_\mu(T) - p_\mu(T')$ is not identically zero: take a generic realization w of T and set $\mu = \lambda w$. Dividing the sites by λ preserves their type and reduces the noise scale to $1/\lambda$. As $\lambda \to \infty$, the probability of T tends to one, so the difference tends to one as well.

Each tie set, the zero set of a nonzero real-analytic function, is therefore closed with empty interior, and so is their finite union. Its complement contains rational points arbitrarily close to tv in this neighbourhood. At that point all realizable types have distinct probabilities, so the mode is unique. □

On the instances of Proposition 6, the unique mode differs from the centre type. A rule returning the centre type or one of N independently sampled types can therefore find the mode only if it was sampled. Its success probability is at most

$$1 - (1 - p_*)^N \le N p_* < \frac{3N}{(n-1)!}.$$

The explicit centres give the rank and probability bounds in polynomial bit length. The further perturbation proves existence of a unique mode without specifying its bit length.

4 Additive approximation and sampling

The sampling algorithm returns the most frequent type among its samples. To implement it with finite bits, we round the coordinates of the standard Gaussian vectors that generate the displacements to rational values and bound the probability that rounding changes the type. This gives an additive guarantee. Lower bounds on the modal probability and on its gap from the competitors will then give relative approximation and exact recovery.

Throughout this section, site i is $X_i = \mu_i + A_i Z_i$, where the centre $\mu_i \in \mathbb{Q}^2$ and the invertible matrix $A_i \in \mathbb{Q}^{2\times 2}$ are part of the input. The vectors $Z_1, \dots, Z_n$ are independent, each with two independent $N(0,1)$ coordinates. Thus $X_i \sim N_2(\mu_i, A_i A_i^{\mathsf{T}})$, allowing each site its own, possibly anisotropic, covariance. The Gaussian reduction uses $A_i = \sigma I_2$ with the same rational $\sigma > 0$ at every site. The input length L includes the entries of these matrices and any supplied accuracy, confidence and probability bounds. Let $\mathcal{T}_n$ be the finite set of complete labelled Delaunay types realizable by n planar points with labels $1, \dots, n$. For each $T \in \mathcal{T}_n$, let $p(T)$ be its probability under these Gaussian sites, and set $p_* = \max_{T \in \mathcal{T}_n} p(T)$. We assume $n \ge 4$, since $\mathcal{T}_n$ has one element for $n \le 3$.

4.1 A finite distribution with controlled error

To round the coordinates of the Z_i, we divide the standard normal line into intervals of equal probability and choose one rational value per interval (Fig. 6). Rounding to these values can produce degeneracies. We test all $\binom{n}{3}$ orientation and $\binom{n}{4}$ incircle determinants. Coincident sites also make an orientation determinant vanish. If any test vanishes, we output the symbol $\perp$. Otherwise we output the complete type. A *valid* type is an element of $\mathcal{T}_n$. The extra symbol $\perp$ has probability zero under the original distribution p. For any rounded output distribution $\widetilde{p}$ on this common alphabet, the supremum over events defining total variation can equivalently be written as

$$d_{\mathrm{TV}}(p, \widetilde{p}) = \frac{1}{2} \sum_{T \in \mathcal{T}_n \cup \{\perp\}} |p(T) - \widetilde{p}(T)|.$$

Lemma 7. *For an integer $q \geq 2$, partition the standard normal line into q consecutive intervals of probability $1/q$, choose a rational representative in the interior of each, and replace each coordinate of each Z_i by the representative of its interval. If $\widetilde{p}_q$ is the resulting output distribution, then*

$$d_{\mathrm{TV}}(p, \widetilde{p}_q) \leq \frac{n^4}{q}.$$

The representatives can be generated individually in time polynomial in $\log q$, using polynomially many bits.

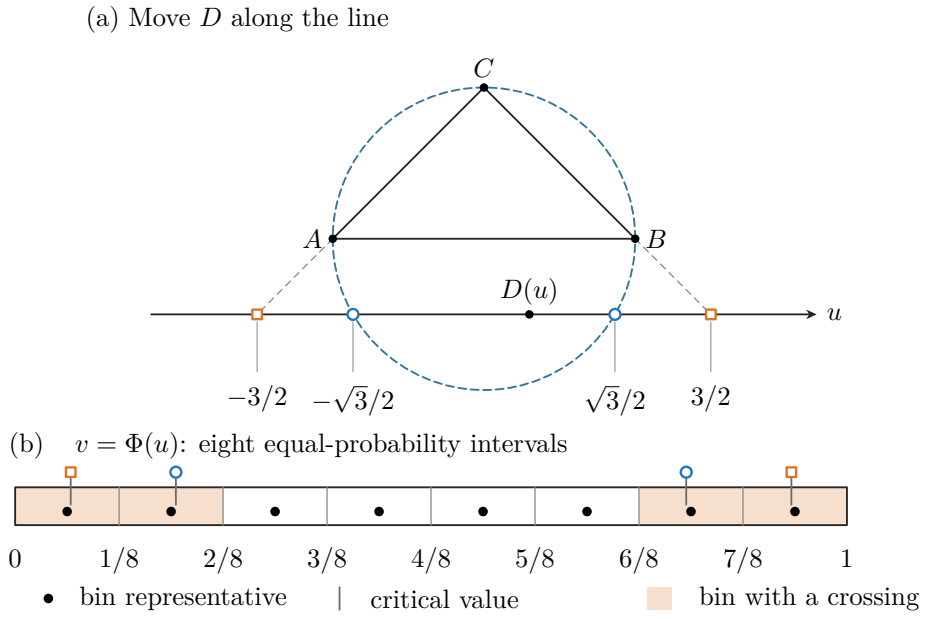

Figure 6: Quantile coupling on a one-dimensional slice of configuration space. In (a), all coordinates except the abscissa u of D are fixed, with $A = (-1, 0)$, $B = (1, 0)$, $C = (0, 1)$ and $D(u) = (u, -1/2)$. Orientation changes (squares) occur at $u = \pm 3/2$, where D crosses the lines AC and BC, and incircle changes (circles) occur at $u = \pm\sqrt{3}/2$, where D crosses the circumcircle of ABC. In (b), the centre of D has abscissa 0, so u is the standard normal noise of this coordinate, and the standard normal cumulative distribution function $v = \Phi(u)$ maps the four critical values into bins $1, 2, 7, 8$. Coupling a draw to its bin representative preserves the complete type in each unshaded bin, so on this slice rounding can change the type only in a shaded bin. The proof of Lemma 7 applies this count to each coordinate in turn and adds the resulting bounds.

Proof. Round the $2n$ coordinates of $Z_1, \ldots, Z_n$ one at a time in a fixed order, using the same sample throughout. If the original and final triangulations differ, some rounding step changed the output. It therefore suffices to add up the probabilities of a change at each step. The coordinates start out independent, and each rounded value depends only on its own original coordinate. Conditional on the current values of all other coordinates, the coordinate about to be rounded still has the standard normal law. In particular, each of its q intervals still has probability $1/q$.

With those other coordinates fixed, varying the remaining coordinate of Z_i moves X_i along a straight line in the direction of the corresponding column of A_i (Fig. 6(a)). Every determinant involving only the other sites is constant along the line. If any orientation or incircle determinant vanishes identically, the algorithm reports $\perp$ both before and after this rounding step. Testing every determinant catches even degeneracies that a triangulation algorithm may never encounter. If no determinant vanishes identically, each pair of fixed sites defines a line that the moving site can meet at most once, and each triple defines a circle that it can meet at most twice. The corresponding determinants are polynomials of degree at most one and two in the moving coordinate. Thus at most $\binom{n-1}{2} + 2\binom{n-1}{3} \le n^3/2$ values are critical. On each interval between critical values, including the two outer intervals, every test keeps its sign and the triangulation stays the same.

The sampled coordinate and its representative lie in the interior of the same quantile interval almost surely. Rounding can change the output only if this interior contains a critical value (Fig. 6(b)). At most $n^3/2$ intervals can do so. The conditional probability of a change is therefore at most $n^3/(2q)$. Summing over $2n$ steps bounds the probability that the coupled outputs differ, and hence their total variation distance, by n^4/q, uniformly in the centres.

Given j, we still need to construct a rational number inside its interval. Let Φ be the cumulative distribution function of the standard normal. On the probability axis in Fig. 6(b), the jth interval runs from $(j-1)/q$ to j/q, with midpoint $v_j = (j-\frac{1}{2})/q$. The midpoint is $1/(2q)$ from either endpoint on the probability axis. Thus any rational r_j with $|\Phi(r_j) - v_j| < 1/(4q)$ lies strictly inside its quantile interval on the real line, even for the unbounded first and last intervals.

We locate such a point by bisection on the real line. For $K = \lceil \log_2 q \rceil + 2$ the Gaussian Chernoff bound gives $\Phi(-K) = 1 - \Phi(K) \le e^{-K^2/2} < 1/(8q)$, so $[-K, K]$ contains the quantile $\Phi^{-1}(v_j)$. At the midpoint r of the current interval, compute a rational $\widetilde{\Phi}$ with $|\widetilde{\Phi} - \Phi(r)| \le 1/(16q)$. If $|\widetilde{\Phi} - v_j| \le 1/(8q)$, return r. Indeed, $|\Phi(r) - v_j| \le 3/(16q) < 1/(4q)$. Otherwise $\Phi(r) - v_j$ has the sign of $\widetilde{\Phi} - v_j$, which shows on which side of r the quantile lies, and we keep that half. The test must succeed once the search interval has length below $1/(8q)$. Its midpoint is then within $1/(16q)$ of the quantile, and $\Phi' < 1$ gives $|\Phi(r) - v_j| < 1/(16q)$, hence $|\widetilde{\Phi} - v_j| < 1/(8q)$. Starting from length $2K$, this takes $O(\log q)$ halvings, and every midpoint is a dyadic rational with $O(\log q)$ bits.

To compute $\Phi(r)$, we integrate the standard normal density from 0 to r and add $1/2$. Expand e^{-w} at zero through degree k and substitute $w = x^2/2$. The resulting polynomial in x has degree $2k$. Every derivative of e^{-w} has magnitude at most one for $w \ge 0$, so the pointwise Taylor error is at most $w^{k+1}/(k+1)!$. Between 0 and r, whatever the sign of r, we have $w \le K^2/2$ along an interval of length $|r| \le K$. Integrating the pointwise bound therefore gives truncation error at most $K(K^2/2)^{k+1}/(k+1)!$. Since $K = O(\log q)$, taking k a sufficiently large constant multiple of K^2 and using a certified rational approximation to $1/\sqrt{2\pi}$ gives the required accuracy $1/(16q)$ with a number of bits polynomial in $\log q$. □

4.2 A polynomial additive search algorithm

The additive algorithm samples $\widetilde{p}_q$ and returns a most frequent valid type. It stores only observed types. To compare that output with an unobserved mode, we need a probability

bound that holds simultaneously for every type, regardless of the size of the support.

Theorem 8. *Given independent Gaussian sites specified by rational centres and invertible rational 2×2 covariance factors, and rational $0 < \varepsilon, \delta < 1$, a randomized algorithm returns a valid type $\widehat{T}$ such that*

$$\mathbb{P}\{p(\widehat{T}) \geq p_* - \varepsilon\} \geq 1 - \delta.$$

It uses $O(\varepsilon^{-2} \log(2/\delta))$ independent discrete realizations and runs in time polynomial in $n, \mathsf{L}, 1/\varepsilon, \log(1/\delta)$.

Algorithm 1 Additive approximation for Gaussian input

Require: Rational centres, invertible rational covariance factors A_i, and rational $0 < \varepsilon, \delta < 1$
Ensure: $p(\widehat{T}) \geq p_* - \varepsilon$ with probability at least $1 - \delta$

1: Choose the smallest power of two $q \geq 4n^4/\varepsilon$.
2: Set $N = \lceil 32\varepsilon^{-2} \rceil \lceil \log_2(2/\delta) \rceil$.
3: **for** $j = 1, \ldots, N$ **do**
4: Sample $2n$ uniform quantile indices, generate their rational representatives and form the rational sites.
5: Evaluate all orientation and incircle predicates exactly. Record $\perp$ if one vanishes, and the canonical type otherwise.
6: **end for**
7: Return a most frequent valid type. If none was observed, return the Delaunay type of (i, i^2), $i = 1, \ldots, n$.

Proof. Algorithm 1 chooses q so that rounding changes the probability of any type by at most $\varepsilon/4$ (Lemma 7). To control the sampling error, fix any ordering of the discrete outputs, including $\perp$. By the Dvoretzky–Kiefer–Wolfowitz inequality with Massart's constant [27], with probability at least $1 - 2\exp(-N\varepsilon^2/32)$ every empirical cumulative frequency is within $\varepsilon/8$ of its population value. This bound holds for discrete distributions as well: applying their inverse distribution function to uniform draws gives the same cumulative frequencies at the corresponding probability thresholds. The algorithm takes $N \geq 32\varepsilon^{-2} \ln(2/\delta)$, since $\log_2 x \geq \ln x$ for $x \geq 1$. The bound thus holds simultaneously at all thresholds with confidence at least $1 - \delta$.

Subtract the cumulative frequency just before a type from the cumulative frequency through it. Each of the two cumulative errors is at most $\varepsilon/8$, so the resulting frequency $\widehat{p}(T)$ differs from $\widetilde{p}_q(T)$ by at most $\varepsilon/4$. Its denominator counts all N draws, including $\perp$. Adding the rounding error gives $|p(T) - \widehat{p}(T)| \leq \varepsilon/2$ for every valid type on this same event.

The algorithm returns a valid type $\widehat{T}$ with a largest count. In particular, it occurs at least as often as any true mode T_*, whose count is zero if it was never sampled. Applying the error bound first to T_* and then to $\widehat{T}$ gives

$$p(T_*) \leq \widehat{p}(T_*) + \varepsilon/2 \leq \widehat{p}(\widehat{T}) + \varepsilon/2 \leq p(\widehat{T}) + \varepsilon.$$

If every draw is $\perp$, all valid types have empirical frequency zero. The fallback is therefore also an empirical maximizer. It is valid because the points (i, i^2) are generic, as shown in the proof of Proposition 6. Each draw takes $2n \log_2 q$ random bits. By Lemma 7, we generate its rational representatives in polynomial time with polynomial bit length. Exact determinant evaluation and sorted lists of labelled faces then let us compute and count the observed types in polynomial time.

Without sampling, enumerating the q^{2n} grid realizations for $q = \lceil 2n^4/\varepsilon \rceil$ gives $\widetilde{p}_q$ exactly, and a most probable valid type of $\widetilde{p}_q$, with the same fallback, loses at most $2d_{\mathrm{TV}}(p, \widetilde{p}_q) \leq \varepsilon$. This deterministic variant takes time $(n/\varepsilon)^{O(n)} \operatorname{poly}(\mathsf{L})$. □

4.3 Relative guarantees and favourable instances

When $\varepsilon \geq p_*$, the additive guarantee allows every valid type. To obtain a useful relative guarantee, we supply a positive lower bound on p_* and choose a smaller additive tolerance. If a unique mode is also separated from its competitors by a known positive gap, a tolerance smaller than that gap forces exact recovery.

Proposition 9. *If a positive rational lower bound $\underline{p} \leq p_*$ is supplied, Gaussian input admits randomized relative search with*

$$\mathbb{P}\{p(\widehat{T}) \geq (1-\varepsilon)p_*\} \geq 1-\delta$$

in time polynomial in $n, \mathsf{L}, 1/\varepsilon, 1/\underline{p}, \log(1/\delta)$. For arbitrary Gaussian input there is such an algorithm with bit cost

$$2^{O(n \log n)} \operatorname{poly}(\mathsf{L}, 1/\varepsilon, \log(1/\delta)).$$

Proof. With tolerance $\varepsilon\underline{p}$, the algorithm of Theorem 8 returns a type satisfying $p(\widehat{T}) \geq p_* - \varepsilon\underline{p} \geq (1-\varepsilon)p_*$ in time polynomial in $1/(\varepsilon\underline{p})$ and the other parameters. For the second assertion, the type count used in (3) gives $p_* \geq 1/(n!C^n)$ for a fixed integer $C > 1$, since every realization is almost surely generic. Using this lower bound for $\underline{p}$ gives the stated $2^{O(n \log n)}$ factor. □

Proposition 10. *Suppose the mode T_* is unique, with probability $p_* = p(T_*)$, gap $\Delta = p_* - \max_{T \neq T_*} p(T) > 0$, and relative gap $\gamma = \Delta/p_*$. The empirical mode of N independent ideal Gaussian realizations equals T_* with probability at least $1-\delta$ whenever*

$$N \geq \frac{20}{p_* \gamma^2} \ln \frac{2|\mathcal{T}_n|}{\delta}.$$

For Gaussian input with known inverse-polynomial lower bounds on p_* and γ, exact modal recovery also has a randomized polynomial-time implementation in bits.

Proof. Each empirical frequency $\widehat{p}(T)$ averages N independent indicators with variance at most p_*. Bernstein's inequality [4, Section 2.7, inequality (2.10)] at $t = \Delta/3$ gives

$$\mathbb{P}\{|\widehat{p}(T) - p(T)| \geq \Delta/3\} \leq 2\exp\left(-\frac{N\Delta^2}{18p_* + 2\Delta}\right) \leq 2\exp\left(-\frac{N\Delta^2}{20p_*}\right),$$

since $\Delta \leq p_*$. As $\Delta^2/p_* = p_*\gamma^2$, a union bound over $\mathcal{T}_n$ shows that, for the stated N, all these errors are below $\Delta/3$ with probability at least $1-\delta$. Then T_* has strictly largest empirical frequency. For the finite-bit implementation, let $\underline{p}$ and $\underline{\gamma}$ be the supplied positive lower bounds on the modal probability p_* and its relative gap γ, respectively. Their product is at most $p_*\gamma = \Delta$. Use half that product as the additive tolerance in Theorem 8. Every other type has probability at most $p_* - \Delta$, so only T_* can meet the guarantee. □

The coupling of Lemma 7 uses independent Gaussian coordinates, and the algorithms need efficient finite sampling. The reductions for Delaunay triangulations require neither independence between coordinates nor an algorithm for evaluating the density or sampling efficiently. The sampling guarantees therefore do not automatically apply to every law covered by those theorems.

5 Decision cells in a complete subdivision

The input to the reduction is a planar graph $G = (V, E)$ in which every vertex has degree three. Our aim is to recover a maximum independent set of G from any most probable Delaunay triangulation in the constructed instance. To encode a binary choice, we place four sites on an empty circle. Their convex quadrilateral, which we call a *decision cell*, can split along either diagonal when the sites move. The construction has to control which choices occur together and ensure that they determine the entire triangulation, including its hull.

5.1 From graph vertices to cycles of circles

We represent a vertex v by a cycle $\mathcal{C}_v$ of tangent unit circles. Each circle carries one decision cell, and two neighbouring cells share the site at their point of tangency. Its displacement can affect the diagonal in both cells. We arrange the cells so that agreement around the cycle is favoured. The common choice will indicate whether v belongs to the independent set. To connect two vertices joined by an edge, we add one more circle, called its *bridge*, touching both cycles. Fig. 7 shows the three bridges incident to one vertex.

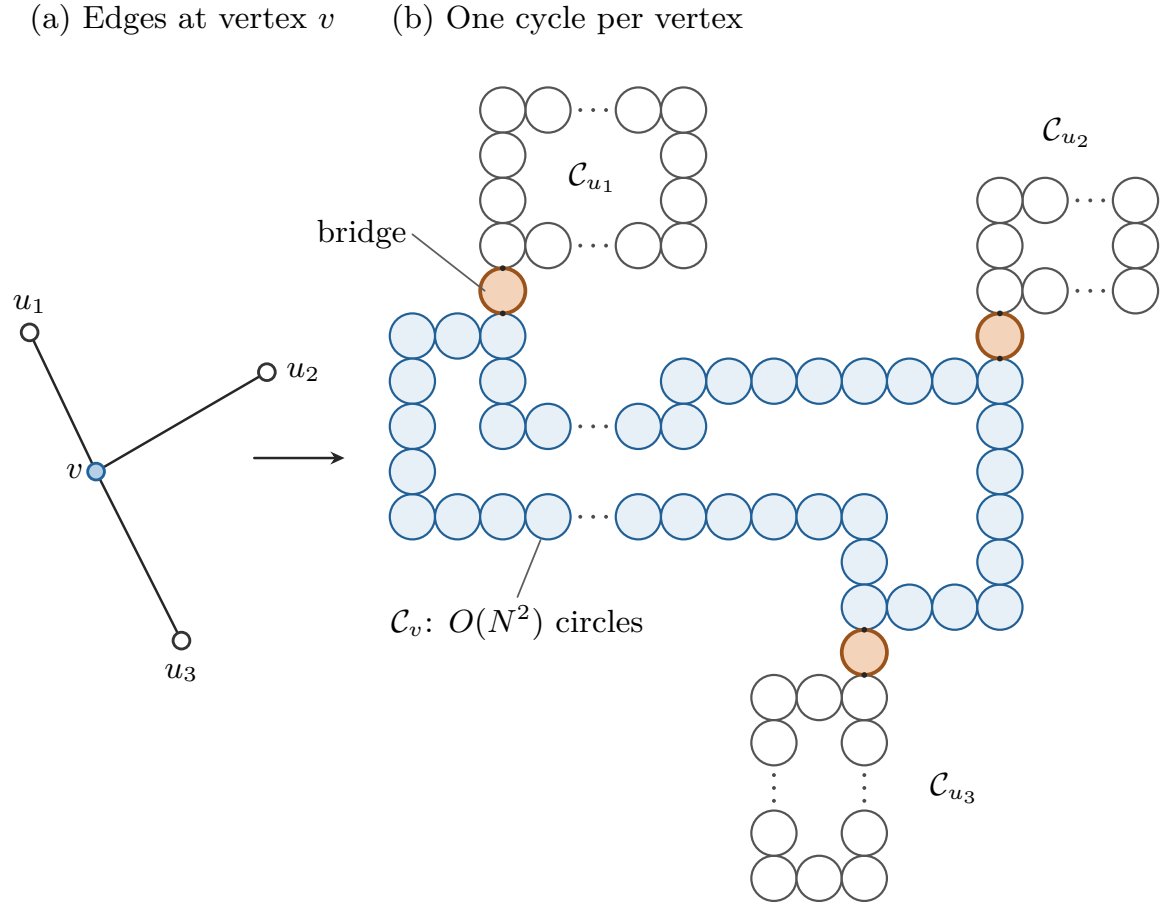

Figure 7: In (a), v is a vertex of the input graph $G = (V, E)$, and the three edges incident to v are shown. In (b), the blue cycle represents v. Each edge vu_i supplies an orange bridge tangent to both endpoint cycles. Matching dotted gaps on opposite sides of each cycle stand for equally long runs of omitted circles. With $N = |V|$, the construction of Appendix A.1 uses at least 80 unit circles per cycle. The $O(N^2)$ bound applies to the whole cycle, whose length depends on its three incident edges.

Lemma 11. *Given a planar cubic graph $G = (V, E)$ with $N = |V|$, one can construct in polynomial time $m = O(N^3)$ unit circles with distinct centres in $2\mathbb{Z}^2$. The circles consist of one induced tangency cycle $\mathcal{C}_v$ for each vertex v, together with one bridge for each edge uv joining vertices u and v. The bridge of uv touches exactly one circle of $\mathcal{C}_u$ and one of $\mathcal{C}_v$. Each cycle turns through a right angle at its incident circle. The two points of tangency are opposite ends of a diameter of the bridge. All other tangencies are between consecutive circles of a single cycle. Each cycle circle touches at most one bridge. In each cycle we designate a circle that touches no bridge. Its two neighbours touch it at opposite ends of a diameter. All disk interiors are disjoint, and no point of the plane lies on three of the circles.*

Appendix A.1 gives the coordinates and the proof. We start with a bar-visibility drawing [34,

Section 4, Theorem 4], in which horizontal bars represent vertices and vertical lines of sight represent edges. Thicken each bar and extend three arms towards its neighbours, leaving one lattice point between facing arms for their bridge. Double the coordinates of that point and of each lattice point on the outline of each bar and its arms. These are the centres of the bridge circles and the cycle circles. The cycles can thus have different lengths. Each outline has a constant number of straight sides, each of length $O(N^2)$ on the grid. Summing over the N vertices and adding the $3N/2$ bridges gives the cubic upper bound.

We now place the four sites of a cell. On a cycle, travel counterclockwise around the outline and consider a circle q with centre o_q. We enter q at its contact $\mathbf{B}_q$ with the preceding circle and leave at its contact $\mathbf{A}_q$ with the next. Where the cycle turns through a right angle, take $\mathbf{C}_q$ at the other end of the diameter from $\mathbf{B}_q$. If a bridge touches the corner, this is its contact with q (Appendix A.1). Along a straight side, take $\mathbf{C}_q$ at the cardinal point facing out of the polygon, to the right as we travel. Fig. 8(a) and (b) show these two shapes.

The last site, $\mathbf{D}_q(t)$, belongs only to q. Place it close to $\mathbf{A}_q$ on the arc towards $\mathbf{C}_q$ that avoids $\mathbf{B}_q$. With origin o_q and $\mathbf{A}_q$ due east, its coordinates are

$$\mathbf{D}_q(t) = \left(\frac{1 - t^2}{1 + t^2}, \varsigma_q \frac{2t}{1 + t^2} \right).$$

Here the sign ς_q selects the arc, and the rational parameter $0 < t < 1/4$ is common to all cycle cells. The four sites occur in cyclic order $\mathbf{A}_q, \mathbf{B}_q, \mathbf{C}_q, \mathbf{D}_q$. The two possible diagonals join $\mathbf{A}_q$ to $\mathbf{C}_q$ and $\mathbf{B}_q$ to $\mathbf{D}_q$. A site shared by two cells carries one label and one random displacement, as in Fig. 8(c). Distinct sites have independent displacements with the fixed law, including the sites later added to complete the subdivision.

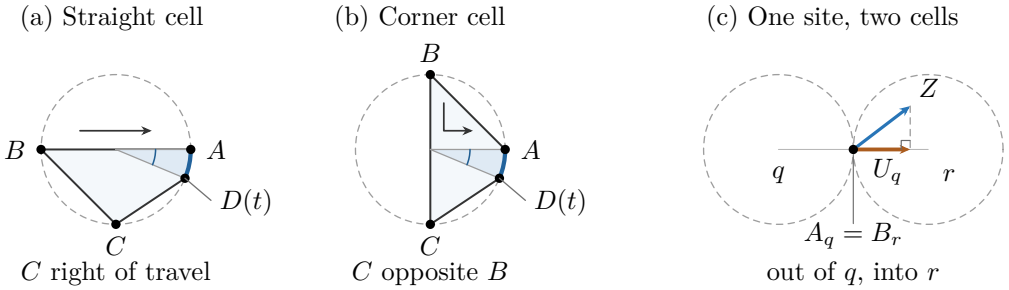

Figure 8: The four sites of a cell, drawn with $t = 1/5$. In (a), travel runs from the predecessor contact B towards the successor contact A. The unused cardinal point on the right is C. In (b), the turn at the cell places C opposite B. In both cases the marked arc from A towards C places $D(t)$ at angle $2 \arctan t$ from A. In (c), the same site $A_q = B_r$ belongs to consecutive cells. The displacement Z (blue) has radial component U_q along the outward normal of q. The orange arrow shows this component pointing into r.

For a decision cell to be a face of the Delaunay subdivision, its closed circumdisk must contain exactly its four vertices, all on the boundary. To check this, take any other circle, with centre $o_q + 2\mathbf{u}$ for some $\mathbf{u} \in \mathbb{Z}^2 \setminus \{0\}$, and a site $o_q + 2\mathbf{u} + \mathbf{v}$ on it, where $\|\mathbf{v}\| = 1$. Its squared distance to o_q is

$$1 + 4(\|\mathbf{u}\|^2 + \langle \mathbf{v}, \mathbf{u} \rangle) \geq 1.$$

Indeed, $\langle \mathbf{v}, \mathbf{u} \rangle \geq -\|\mathbf{u}\| \geq -\|\mathbf{u}\|^2$. Equality requires $\mathbf{v} = -\mathbf{u}$ with $\|\mathbf{u}\| = 1$, which gives the tangency point of the two unit circles, and this site already belongs to both cells. Every other site of another cell is thus strictly outside the closed circumdisk of q.

5.2 How a displacement selects a diagonal

At the base configuration the four sites are cocircular. To compare the two diagonals when the sites move, we compute the linear change of the incircle determinant. The size of the change depends on which vertex moves. The weights of the four vertices explain how to favour one diagonal and how a shared site affects the choices in two cells.

Fix a cell q and write A, B, C, D for its four sites. Its diagonals AC and BD meet in their interiors. Equating the two convex combinations for their intersection and scaling the coefficient of D to -1 gives

$$a_q\mathbf{A}_q - b_q\mathbf{B}_q + c_q\mathbf{C}_q - \mathbf{D}_q = 0, \qquad a_q - b_q + c_q - 1 = 0. \tag{5}$$

The intersection is interior, so a_q, b_q, c_q are positive. Their magnitudes also have a direct geometric interpretation. The weight at each vertex is the area of the triangle formed by the other three, divided by the area of ABC. For instance, a_q is the area of BCD divided by that of ABC, and b_q uses ACD in the numerator (Appendix A.2). Substitution of the coordinates gives Table 2. Reflecting the cell exchanges north and south and reverses ς_q while preserving the weights.

Table 2: Exact coefficients of the affine dependence (5) for a decision cell rotated so that **A** is at east. N, W, S denote the other cardinal points. Entries in parentheses give the reflected case, which has the same coefficients. The first row gives the corner cycle cells, in which **B** is a quarter turn from **A** and **C** is opposite to **B**. The second row gives the straight cells.

B	**C**	ς	$(1+t^2)a$	$(1+t^2)b$	$(1+t^2)c$
N (S)	S (N)	-1 (1)	$1-t^2$	$t(1-t)$	$t(1+t)$
W	S (N)	-1 (1)	$1-t$	$t(1-t)$	$2t$

Table 2 gives $a_q = 1 + O(t)$, $c_q = O(t)$ and $b_q = t(1 + O(t))$. For a corner cell the weight of $\mathbf{C}_q$ is $c(t) = t(1+t)/(1+t^2)$. In particular, the two cycle cells at the ends of any bridge have this same weight.

The exact test of a cell is a normalized incircle determinant, denoted Γ_q. Take its rows in the order $(\mathbf{P}_1, \ldots, \mathbf{P}_4) = (\mathbf{A}_q, \mathbf{B}_q, \mathbf{C}_q, \mathbf{D}_q)$, with coordinates relative to o_q. A site with coordinates (x_i, y_i) has row $(x_i, y_i, x_i^2 + y_i^2, 1)$. Multiply this determinant by $\iota_q/4$, where ι_q is the sign of orient$(\mathbf{A}_q, \mathbf{B}_q, \mathbf{C}_q)$ at the base configuration, to obtain Γ_q. The cyclic order may run clockwise or counterclockwise. With this normalization, a small generic perturbation gives $\Gamma_q > 0$ when $\mathbf{D}_q$ lies inside the circle through the other three sites. The selected diagonal is then $\mathbf{B}_q\mathbf{D}_q$, as in Fig. 9(b). A negative Γ_q selects $\mathbf{A}_q\mathbf{C}_q$.

At the base, let $\mathbf{n}_i = \overrightarrow{o_q\mathbf{P}_i}$ be the outward unit normal at vertex i. For small displacements $h\mathbf{z}_i$, the coefficient of h in the change of Γ_q is

$$a_q\langle \mathbf{n}_1, \mathbf{z}_1\rangle - b_q\langle \mathbf{n}_2, \mathbf{z}_2\rangle + c_q\langle \mathbf{n}_3, \mathbf{z}_3\rangle - \langle \mathbf{n}_4, \mathbf{z}_4\rangle. \tag{6}$$

Appendix A.2 derives this formula from the determinant. Each inner product measures how far the corresponding vertex moves along its own outward radius, in units of h. A displacement tangent to the circle has zero inner product with that radius and can affect the test only at second or higher order. Thus the weights in (5) give the response of the test to radial motion, with opposite signs at consecutive vertices.

In the base coordinates the circles have radius one. At site i the random displacement is σZ_i, where the Z_i are independent with law ν. Section 8.2 chooses the common scale $\sigma > 0$. All coordinates are then divided by σ, so the final instance has displacements Z_i. In units of σ, write V_q and Q_q for the signed radial components at $\mathbf{C}_q$ and $\mathbf{D}_q$, respectively, taking the

outward direction from o_q as positive at each site. Fig. 9(a) shows these projections for one draw, in which the projection points inward at $\mathbf{C}_q$ and outward at $\mathbf{D}_q$. By (6) the two sites contribute $c_qV_q - Q_q$ to the linear test.

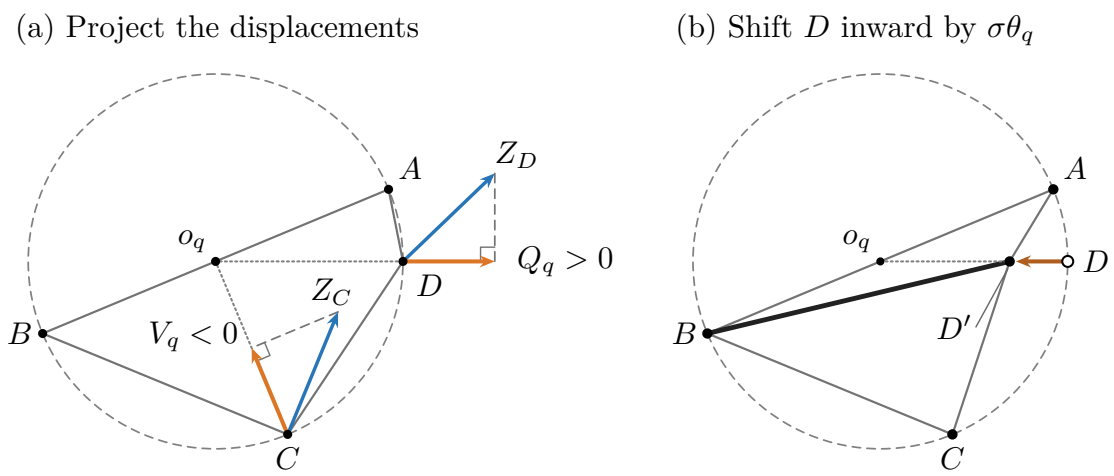

Figure 9: Both drawings show the same cell q, with sites labelled A, B, C, D. It is the straight cell of Fig. 8(a), rotated so that the radius of D is horizontal. In (a), the blue arrows are the site displacements Z_C, Z_D. Measured outward in units of σ, the orange projections have signed lengths V_q at C and Q_q at D. In (b), A, B, C stay fixed. Shifting the centre from D (open) to D' (filled) puts it inside their circle. The bold segment is the resulting Delaunay diagonal in the absence of random displacements.

To favour the diagonal $\mathbf{B}_q\mathbf{D}_q$, move the nominal centre of $\mathbf{D}_q$ inwards by $\sigma\theta_q$, with $\theta_q > 0$ small, as in Fig. 9(b). For the same random draw, the total signed radial displacement from the original position $\mathbf{D}_q$, in units of σ, is now $Q_q - \theta_q$. The weight -1 at this site turns the inward shift into a gain of $+\theta_q$ in the linear test. We call θ_q the cell's *field*. For either sign of the field the shift is $-\sigma\theta_q(\mathbf{D}_q - o_q)$.

The remaining terms come from the two contacts along the cycle. At $\mathbf{A}_q = \mathbf{B}_r$ in Fig. 8(c), let U_q be the radial component of the displacement measured outwards from o_q. Measured from the centre o_r of the next cell, the same displacement has radial component $-U_q$. Thus (6) gives a_qU_q in the first test and $(-b_r)(-U_q) = b_rU_q$ in the second. Moving the contact towards o_r increases both tests to first order. This common contribution is the source of agreement along a cycle. At the entry contact of q, shared with its predecessor q^-, the same calculation gives $b_qU_{q^-}$. Combining the two contacts of q, its other two sites and its field gives the linearized test

$$\Lambda_q = a_qU_q + b_qU_{q^-} + c_qV_q - Q_q + \theta_q. \tag{7}$$

Once a cycle has a common sign, how can an edge affect the choices on its endpoint cycles? We use the bridge to favour opposite signs by sharing one displacement with each cycle. The bridge uses the same four-site test with a separate rational parameter $0 < t' \le t$. Fig. 10(a) shows its two contacts and its two unshared sites. In the lattice construction, the contacts are the south and north points of the bridge x, which we call $\mathbf{B}_x$ and $\mathbf{C}_x$, respectively. Place $\mathbf{A}_x$ at its unused west point and $\mathbf{D}_x(t')$ on the arc from $\mathbf{A}_x$ towards $\mathbf{C}_x$ that avoids $\mathbf{B}_x$. A half-turn identifies the bridge with the corner row of Table 2, giving weights

$$(a', b', c') = \frac{1}{1+t'^2}(1 - t'^2,\ t'(1-t'),\ t'(1+t')).$$

At its two unshared sites $\mathbf{A}_x, \mathbf{D}_x$, write U_x, Q_x for the outward radial components in the bridge. Its lower contact is $\mathbf{C}_{q_B(x)}$ on a cycle cell $q_B(x)$, and its upper contact is $\mathbf{C}_{q_C(x)}$ on a cell $q_C(x)$

of the other cycle. The bridge normal is opposite the cycle normal at each contact. In (6) the lower site has weight $-b'$, so reversing the normal gives $+b'V_{q_B(x)}$. The upper site has weight $+c'$, giving $-c'V_{q_C(x)}$. Thus the linearized bridge test is

$$\Lambda_x = a'U_x + b'V_{q_B(x)} - c'V_{q_C(x)} - Q_x. \tag{8}$$

Both shared projections have weight $+c(t)$ in their respective cycle tests. The bridge therefore responds in the same direction as the lower cell and in the opposite direction to the upper cell at their respective contacts (Fig. 11(c)). The bridge is itself one of the m cells, and the sign of its test is recorded with the signs of the cycle cells. At zero field the bridge sign has positive correlation w_B with the lower endpoint sign and negative correlation $-w_C$ with the upper one (Section 6). When the endpoint signs differ, one choice of bridge sign gains at both contacts (Fig. 10(b)). Choosing t' smaller weakens the two bridge contacts without changing the cycle cells. Across the construction, radial components measured once per distinct site are independent. By rotational invariance each has the law of U, the projection of a random displacement with law ν onto a fixed unit vector.

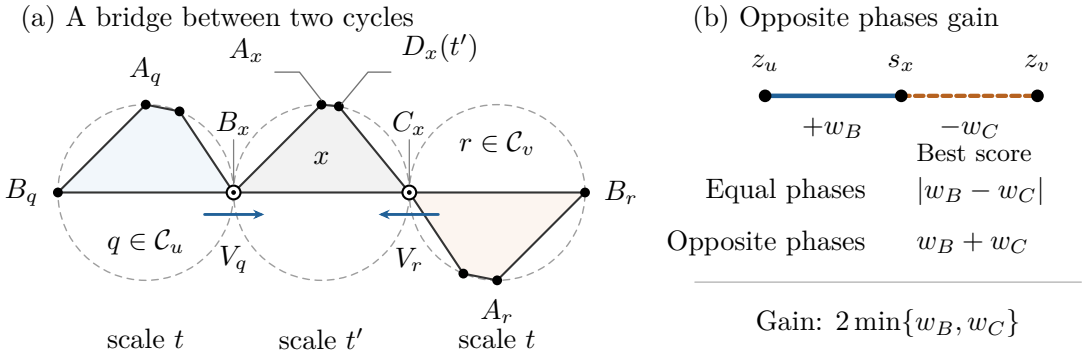

Figure 10: In (a), the bridge shares the two sites B_x, C_x with the endpoint cycles. The drawing shows the lattice after a clockwise quarter-turn, with the lower contact B_x on the left and the upper contact C_x on the right. Only the bridge uses A_x, D_x. Each shared site is opposite the predecessor contact in its cycle. The arrows point out of the cycle cells $q = q_B(x)$ and $r = q_C(x)$ and give the positive directions for V_q, V_r in (8). In (b), z_u, z_v are the common signs on the respective cycles, and s_x is the bridge sign. With zero fields, w_B, w_C are the magnitudes of the two sign correlations in (15). Choosing the better bridge sign gives $|w_B z_u - w_C z_v|$. Opposite phases gain $2\min\{w_B, w_C\}$ over equal phases. The drawing uses $t = 1/5$ on the cycles and $t' = 1/10$ on the bridge.

If the tests on an entire cycle have the same sign, call that sign its *phase*. A positive phase selects the vertex and a negative phase leaves it out. We put a field only on the designated circle in each cycle. These fields and the bridge choices will together favour maximum independent sets (Section 7). A *block* is one copy of all m cells for G, with one pair of parameters (t, t') and one assignment of fields θ. Write $p^{\theta}_{t,t'}(s)$ for the probability that all m linear tests have the signs $s \in \{-1, 1\}^m$. Fig. 11(a) and (b) contrast the two local diagonal choices for one cell, and Fig. 12 shows how to read the selected vertices from the diagonals in a pair of cycles.

5.3 From local diagonals to the complete triangulation

The modal oracle returns a complete labelled triangulation. If a face outside the decision cells has several possible diagonals, the probability of one sign vector can be spread over several triangulations. The shaded faces in Fig. 12 illustrate the problem. We insert sites to triangulate these extra faces and put three more sites around the construction to fix the convex hull.

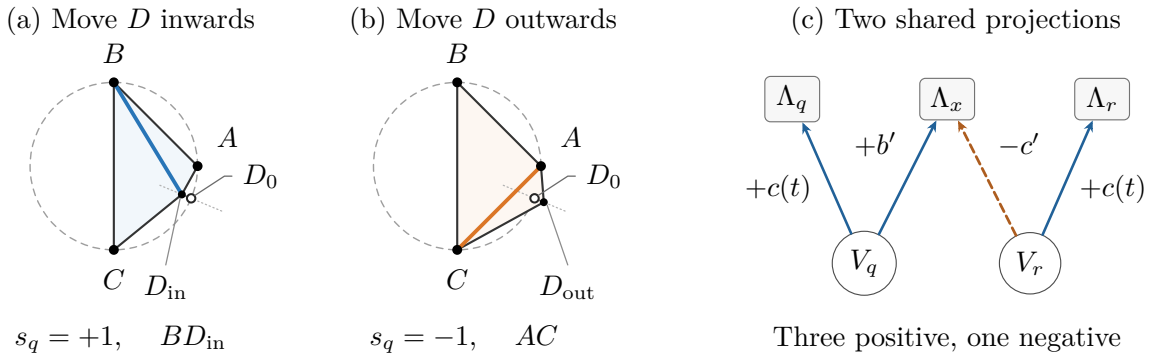

Figure 11: Diagonal decisions and shared randomness. In (a) and (b), $D_0 = D(1/5)$ is the unperturbed position. A radial displacement of 12% gives D_{in} or D_{out} and selects the blue diagonal BD_{in} or the orange diagonal AC. In (c), each arrow gives the coefficient of a shared projection in a linear test. The solid arrows carry positive coefficients. The right-hand projection V_r comes from the upper contact in the lattice. The dashed arrow gives its negative coefficient in (8).

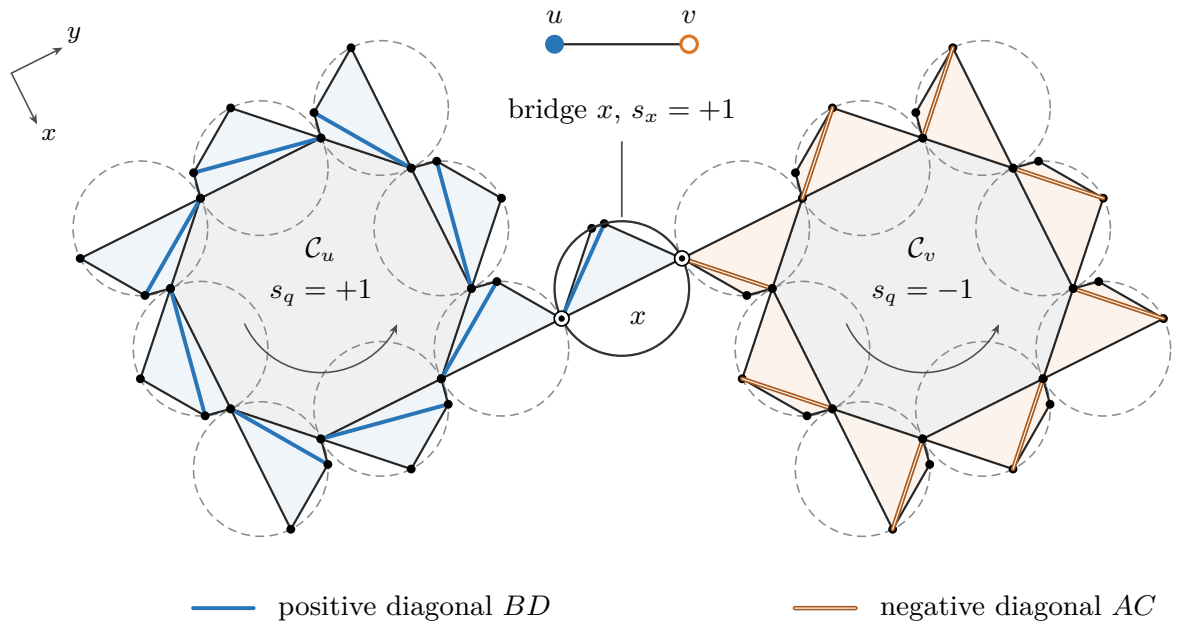

Figure 12: Two vertex cycles linked by one bridge circle. Each schematic cycle is shortened to eight cells, with $t = 1/5$. The lattice construction uses at least 80 cells per cycle. The bridge cell x uses $t' = 1/10$. Blue diagonals select the positive phase on the left, and orange diagonals select the negative phase on the right. The bridge takes its favourable positive sign for this phase pair. Rings mark its two shared contacts. The small graph records the edge uv. The curved arrows give the counterclockwise orientation of the cycles, and the axes give the lattice directions of the rotated drawing. The shaded octagons are nontriangular faces outside the decision cells, which Lemma 12 triangulates with additional sites. The drawing shows the cocircular base before fields and generic perturbations.

Lemma 12. *Consider any finite collection of these decision cells, including separated copies built with different parameters. In polynomial time one can add rational sites outside their closed circumdisks so that the convex hull is triangular, each closed circumdisk still contains exactly the four vertices of its cell, and the decision cells are the only nontriangular Delaunay faces. With M cells the final number of sites satisfies $n \leq 12M + 4$. All coordinates have polynomial bit length in the input coordinates and M.*

Proof. First enclose the cells in a triangle of three additional sites, called guards. Then put one more site inside each unwanted nontriangular face, outside all the closed circumdisks of the cells. At each insertion we choose a grid point distinct from all earlier sites, off the lines through pairs of them and off the circles through noncollinear triples of them. Appendix A.3 gives rational grids on which such points can be found in polynomial time.

Choose one guard in each of three rational boxes whose placement guarantees that the guard triangle encloses all closed circumdisks in its interior (Fig. 13(a)). The $n_0 \leq 4M + 3$ sites and guards now have a triangular hull. To find all faces of their Delaunay subdivision, test the circumcircles of all noncollinear triples for emptiness and group together triples on the same empty circle.

Consider an unwanted face F. If the closed circumdisks covered its interior, the connected open set $\operatorname{int} F$ would lie inside one of them. Indeed, their interiors are disjoint and their tangencies lie on the boundary of their union. All vertices of F would then belong to that closed disk. Its only sites are the four vertices of one decision cell, so F would be that cell. Hence some open disk inside F is disjoint from the closed circumdisks. Fix one such disk with rational centre and radius for each unwanted face before inserting any points. Their bit lengths then depend only on the original subdivision (Appendix A.3).

After the insertions, no guard or inserted site can belong to a face with four or more vertices. The last chosen vertex of such a face would lie on the circle through three earlier vertices, which the grid choice excludes. Any remaining nontriangular face therefore has only original vertices. Its open circumdisk is empty in the final subdivision and was already empty before the insertions. All the original sites on its boundary are still present, so they formed the same face before insertion. In an unwanted face, however, we inserted a site strictly inside

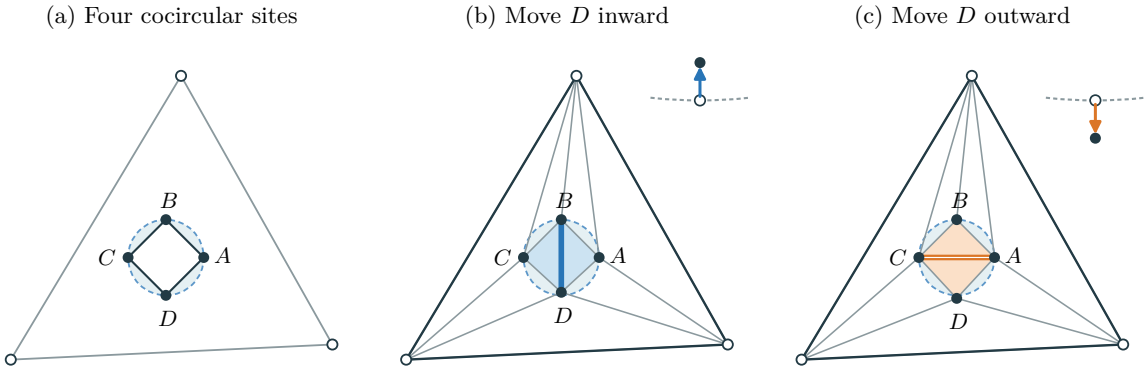

Figure 13: Completing a configuration with one decision cell. In (a), the unit circumcircle passes through the four vertices of the cell, and three guards form a triangular hull. In (b) and (c), moving D radially inwards or outwards by 8% selects BD or AC. In each enlarged view, the open point shows the original position of D and the filled point shows its displaced position. The two complete triangulations share seven faces and differ by the diagonal of the quadrilateral. For legibility, the four sites are drawn at the cardinal points. The construction places D close to A, as in Fig. 12.

its circumdisk. Thus every remaining nontriangular face is a decision cell. The guards still form the hull, and the closed circumdisk of each cell still contains exactly its four chosen sites.

To count the insertions, triangulate the original subdivision with the guards in place. Its n_0 sites and three hull vertices give $2n_0-5$ bounded triangles. Each unwanted face contains at least one of these triangles, so at most $2n_0-5$ sites were inserted. Thus $n \le 3n_0-5 \le 12M+4$. $\square$

Fig. 13(b) and (c) show how the chosen diagonal determines the complete triangulation for one cell. We now give the inequalities that preserve this correspondence when the sites move.

Let f_3 count the triangular faces just after completion and before perturbation. Each of the M decision cells contributes two more triangles once a diagonal is chosen. Euler's formula for n sites with three on the hull gives

$$f_3 + 2M = 2n - 5. \tag{9}$$

The count remains valid for every generic realization of the displaced sites that satisfies the following conditions. All nonguard sites lie strictly inside the guard triangle, and the f_3 triangles keep their orientations and their empty circumcircles. In each decision cell the four triples keep their orientations, so the cell stays convex with the same cyclic order. For each triple, every site outside the cell remains strictly outside its circumcircle. All four circumcircles coincide at the base and can separate when the vertices move, so this exteriority condition is imposed on each of them. These inequalities are all strict at the base, and each fixes the sign of an orientation or incircle determinant.

In any generic realization satisfying these inequalities, each Γ_q selects exactly one local diagonal and with it two triangles of the cell. For each of these triangles, the test puts the remaining vertex of the cell outside its circumcircle. The exteriority conditions exclude all sites outside the cell. The two selected triangles are therefore Delaunay. Together with the f_3 fixed triangles, the selected triangles already account for all $2n-5$ bounded faces of the perturbed triangulation, by (9) and the fixed triangular hull. The sign vector determines every face. Conversely, one reads the sign in cell q by checking whether the labelled edge $\mathbf{B}_q\mathbf{D}_q$ is present, so distinct sign vectors give distinct complete types. Section 8 bounds the probability that a geometric inequality fails or that an exact diagonal test differs in sign from its linear approximation.

6 Bounded responses and exact weights

This section and the next study one block through the signs of its m linear tests. Section 8 carries the conclusions to complete triangulations. The two parameters t and t' separate the tasks of the construction. Agreement along a cycle has weight comparable to $\mathcal{A}(t)$, a bounded average of the noise. Each connection through a bridge has weight at most a constant times $\mathcal{A}(t')$. Choosing the second average smaller will therefore make a nonconstant cycle unprofitable. We express the weights in terms of the signs of individual cells and pairs of cells, and bound their sizes. The proofs use one property of the noise, a bounded density of its projection. The tail condition (L) enters through (13), which bounds how many candidate scales the reduction must list.

Let g and F be the density and distribution function of a unit projection U. By radial symmetry this law is symmetric and independent of the direction. Integrating the planar density along a line and substituting $u = x^2 + y^2$ gives

$$g(x) = \int_{x^2}^{\infty} \frac{\varphi(\sqrt{u})}{\sqrt{u - x^2}}\, du, \qquad \int_0^{\infty} \varphi(\sqrt{u})\, du = 1/\pi. \tag{10}$$

Since $0 \le \varphi \le M_\rho$, this integral is largest when the mass $1/\pi$ sits at height M_ρ on an interval of length $1/(\pi M_\rho)$ next to the kernel singularity. Hence g is bounded, and

$$\|g\|_\infty \le 2\sqrt{M_\rho/\pi}, \qquad |F(x) - F(y)| \le \|g\|_\infty |x - y|, \qquad F(-x) = 1 - F(x). \tag{11}$$

Boundedness of g is the only property of the density used in this section and the next. The difference of two independent projections has density $h = g * g$. Since $g \in L^1 \cap L^\infty$, the function h is uniformly continuous, and the Cauchy–Schwarz inequality gives $0 \le h \le h(0)$. We write $J = h(0) = \int g^2 > 0$.

For $Y = |U|$, the bounded average

$$\mathcal{A}(t) = \mathbb{E} \min(tY, 1)$$

measures the expected size of a response whose argument is multiplied by t. For the tests covered by (18), fixing one projection of magnitude y and averaging over the others changes the probability of a positive sign from $1/2$ by an amount comparable to $\min(ty, 1)$ when its coefficient is of order t. The average $\mathcal{A}(t)$ tends to zero with t, even when Y has no finite positive moment. With $a_0 = \mathbb{E} \min(Y, 1) > 0$, we have

$$a_0 t \le \mathcal{A}(t) \le 1 \qquad (0 < t \le 1), \qquad \mathcal{A}(2t) \le 2\mathcal{A}(t). \tag{12}$$

Under (L), splitting the expectation at $Y = 2^{j/2}$ gives

$$\mathcal{A}(2^{-j}) \le 2^{-j/2} + C_L \lfloor j/2 \rfloor^{-\gamma} \le C_A j^{-\gamma} \qquad (j \ge 3), \tag{13}$$

for a fixed integer C_A. The first term bounds the contribution below the split. The second bounds its complement using $|U| \le \|Z\|$. The reduction will use this bound without evaluating $\mathcal{A}$.

The interaction of two cells through one site involves responses at both scales. Define

$$\widehat{w}(t, t') = \mathbb{E}[\min(tY, 1) \min(t'Y, 1)].$$

Fig. 14 shows how the smaller scale limits this interaction. We shall write $\mathcal{A} = \mathcal{A}(t)$, $\mathcal{A}' = \mathcal{A}(t')$ and $\widehat{w} = \widehat{w}(t, t')$ when the parameters are fixed.

The estimates below use the coefficients of the tests only through fixed ranges. By Table 2, for $0 < t' \le t < 1/4$ and $C_0 = 2$,

$$\begin{gathered} |a_q - 1| \le C_0 t, \qquad \frac{t}{C_0} \le b_q \le C_0 t, \qquad 0 \le c_q \le C_0 t, \\ |a' - 1| \le C_0 t', \qquad \frac{t'}{C_0} \le b', c' \le C_0 t', \end{gathered} \tag{14}$$

and $c_q = c(t) \ge t/C_0$ at the two corners touched by a bridge. By Lemma 11 and the choices of Section 5, all bridges have the same coefficients, the two endpoints of every bridge are corners with the same coefficients, and all designated cells are straight cells with the same coefficients. The proofs of this section and the next use the coefficients only through (14) and these congruences, so their constants depend only on ν and C_0.

For a cycle cell q, first average over the three projections U_q, V_q, Q_q of its test (7), holding the predecessor's contribution fixed at x:

$$\widetilde{L}_{q,t}(x) = \mathbb{P}(a_q U + c_q V - Q + x > 0) - \tfrac{1}{2},$$

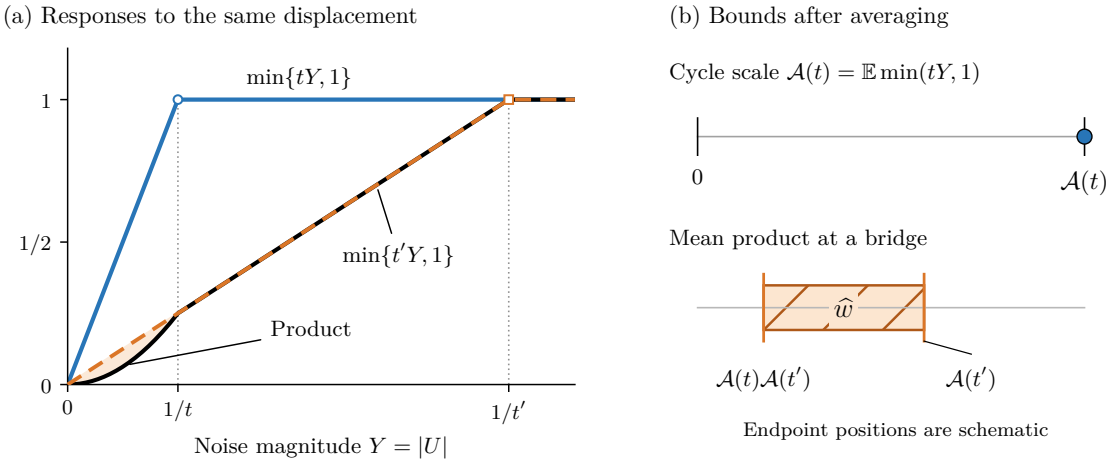

Figure 14: Two scales separate cycle agreement from contacts between cycles. In (a), $t = 1/5$ and $t' = 1/20$. The black curve is the product of the bounded responses, which saturate at $Y = 1/t$ and $Y = 1/t'$. In (b), the orange interval represents the bounds $\mathcal{A}(t)\mathcal{A}(t') \le \widehat{w} \le \mathcal{A}(t') \le \mathcal{A}(t)$. Its endpoint positions are schematic and use no particular noise law. The bridge interaction is of order $\widehat{w}$, while agreement inside a cycle remains of order $\mathcal{A}(t)$.

where U, V, Q are independent unit projections. This function is odd and nondecreasing. The predecessor contributes $x = b_q U_{q^-}$. With zero fields for the first three quantities, the cycle weight, the two bridge weights and the bias of a designated cell d under a field θ are

$$\begin{aligned} \kappa_q(t) &= 4\mathbb{E}[(F(a_{q^-}U) - \tfrac{1}{2})\widetilde{L}_{q,t}(b_q U)], \\ w_B(t,t') &= \mathbb{E}[\operatorname{sgn}\Lambda_x \operatorname{sgn}\Lambda_{q_B(x)}], \qquad w_C(t,t') = -\mathbb{E}[\operatorname{sgn}\Lambda_x \operatorname{sgn}\Lambda_{q_C(x)}], \\ \mathcal{M}(\theta) &= \mathbb{E}[\operatorname{sgn}\Lambda_d] = \operatorname{sgn}(\theta)\,\mathbb{P}(|\Lambda_d^0| \le |\theta|). \end{aligned} \tag{15}$$

To interpret κ_q, keep only $a_{q^-}U_{q^-} - Q_{q^-}$ in the predecessor test. At zero field, the sign of this part and the sign of Λ_q are conditionally independent given U_{q^-}. Their conditional means are twice the two factors in the first expectation, so κ_q is their sign correlation. Section 7.1 splits the conditional mean of each cycle sign into a function of its exit projection and a correction. At zero field, pairing the correction at q with the main part at q^- gives κ_q in (22). In (21), w_B and $-w_C$ are the exact sign correlations of bridge x with its endpoints $q_B(x)$ and $q_C(x)$, respectively. The bias is the mean sign of a designated cell whose field is θ. Here Λ_d^0 is its test with zero field, which has a symmetric continuous law. Since the bridges are congruent and the designated cells have equal coefficients, these three quantities do not depend on x or d.

Both bridge weights are positive: the sign correlation is positive at the endpoint $q_B(x)$ and negative at $q_C(x)$, like the coefficients $+b'$ and $-c'$ in (8). Indeed, given the projection $V = V_{q_B(x)}$ shared by a bridge and its endpoint $q = q_B(x)$, their decisions are independent, and their conditional probabilities of a positive sign differ from $1/2$ by responses of the same kind as $\widetilde{L}_{q,t}$. With $L_i(y) = \mathbb{P}(W_i < y) - \frac{1}{2}$,

$$w_B = 4\mathbb{E}[L_1(c(t)Y)L_2(b'Y)], \qquad w_C = 4\mathbb{E}[L_1(c(t)Y)L_3(c'Y)], \tag{16}$$

where $W_1 = Q - a_q U - b_q U'$ for an endpoint q, $W_2 = Q - a'U + c'V'$ and $W_3 = Q - a'U - b'V'$, all variables being independent unit projections. Each L_i is odd and nondecreasing. At $q_C(x)$ the shared projection has opposite signs in the two tests, which gives the negative correlation $-w_C$.

For Gaussian unit projections the tests are jointly Gaussian. Two centred Gaussian variables of correlation r can be represented, after rescaling, by projections of a standard

planar Gaussian onto unit vectors at angle $\arccos r$. The projections have opposite signs on two sectors of that angle. The direction is uniform, so their sign product has expectation $1 - 2\arccos(r)/\pi = (2/\pi)\arcsin r$ [20, Section 27]. Let $v_q = a_q^2 + b_q^2 + c_q^2 + 1$ be the variance of a cycle test and $v' = a'^2 + b'^2 + c'^2 + 1$ that of a bridge test. Then

$$w_B = \frac{2}{\pi}\arcsin\frac{c(t)\,b'}{\sqrt{v_q v'}}, \qquad w_C = \frac{2}{\pi}\arcsin\frac{c(t)\,c'}{\sqrt{v_q v'}}, \qquad \mathcal{M}(\theta) = \operatorname{erf}\Big(\frac{\theta}{\sqrt{2v_d}}\Big),$$

with q an endpoint, and $\kappa_q(t) = t/\pi + O(t^2)$. Both bridge weights are $tt'/\pi + O(tt'(t+t'))$, and $\mathcal{M}(\theta)$ is $\theta/\sqrt{\pi}$ up to a factor $1 + O(t + \theta^2)$. The next two lemmas give the corresponding bounds for every law in the theorem. The first makes every response to a small argument comparable to that argument. It is stated for a family of variables because the coefficients of the tests vary with the cell and with the scales t and t'.

Lemma 13. *For $0 < \varepsilon \le 1/2$, let $\mathcal{W}_\varepsilon$ be the set of random variables $W = Q - aU - \gamma_1 V_1 - \gamma_2 V_2$, where Q, U, V_1, V_2 are independent unit projections, $|a - 1| \le \varepsilon$ and $|\gamma_1|, |\gamma_2| \le \varepsilon$. Each W is symmetric with a density $k_W \le \|g\|_\infty$, and $\sup_{W \in \mathcal{W}_\varepsilon} \|k_W - h\|_\infty \to 0$ as $\varepsilon \to 0^+$. Consequently there are $\varepsilon_1 \in (0, 1/2]$ and $y_1 \in (0, 1]$, depending only on ν, such that*

$$\tfrac{11}{12} J \le k_W(y) \le \tfrac{13}{12} J \qquad (|y| \le y_1,\ W \in \mathcal{W}_{\varepsilon_1}). \tag{17}$$

Proof. We compare each density to the density of the difference of two independent projections, first under dilation and then under translation. The variable W is the sum of Q and an independent symmetric variable, so it is symmetric and $k_W(y) = \mathbb{E}g(y + aU + \gamma_1 V_1 + \gamma_2 V_2) \le \|g\|_\infty$ is a density of W. Let g_a be the density of aU and $k_a = g * g_a$, a density of $Q - aU$. It is uniformly continuous, as the convolution of an integrable and a bounded function, and so is $k_W = \mathbb{E}k_a(\cdot + \gamma_1 V_1 + \gamma_2 V_2)$. We always use these continuous versions. Translations and dilations act continuously on $L^1(\mathbb{R})$. Thus $\omega_1(r) = \sup_{|u| \le r} \|g(\cdot + u) - g\|_1$ tends to zero with r, and $\|g_a - g\|_1$ tends to zero as $a \to 1$. Young's inequality gives $\|k_a - h\|_\infty \le \|g\|_\infty \|g_a - g\|_1$ and, for $a \ge 1/2$, $|k_a(y) - k_a(y')| \le \|g_a\|_\infty\, \omega_1(|y - y'|) \le 2\|g\|_\infty\, \omega_1(|y - y'|)$, since $g_a(z) = a^{-1} g(z/a)$. Put $S = \gamma_1 V_1 + \gamma_2 V_2$. Then $k_W(y) = \mathbb{E}k_a(y + S)$, and for every $r > 0$

$$|k_W(y) - k_a(y)| \le 2\|g\|_\infty\, \omega_1(r) + 2\|g\|_\infty \mathbb{P}(|S| > r) \le 2\|g\|_\infty\, \omega_1(r) + 4\|g\|_\infty \mathbb{P}(\varepsilon Y > r/2).$$

Given $\eta > 0$, choose r and then ε so that each of the two terms and $\|g\|_\infty \|g_a - g\|_1$ is at most $\eta/3$ for $|a - 1| \le \varepsilon$. This proves the uniform convergence. For (17), take $\eta = J/24$ and use the continuity of h at zero. □

The lemma is qualitative: no rate of convergence is needed. Let $W \in \mathcal{W}_{\varepsilon_1}$ and $L_W(y) = \mathbb{P}(W < y) - \frac{1}{2} = \int_0^y k_W$. By (17) and $k_W \le \|g\|_\infty$, for $y \ge 0$ and a coefficient c with $t/C_0 \le c \le C_0 t$,

$$c_\nu \min(ty, 1) \le L_W(cy) \le C_\nu \min(ty, 1), \tag{18}$$

with $c_\nu = \frac{11}{12} J \min(y_1, 1/C_0)$ and $C_\nu = \max(1, C_0 \|g\|_\infty)$. The same holds with t' for coefficients in $[t'/C_0, C_0 t']$. From now on, c_ν, C_ν denote these two constants.

Lemma 14. *There is $t_1 \in (0, 1/4)$, depending only on ν and C_0, such that for $0 < t' \le t \le t_1$ and every cycle cell q,*

$$\mathcal{A}\mathcal{A}' \le \widehat{w} \le \mathcal{A}', \qquad c_w \widehat{w} \le w_B, w_C \le C_w \widehat{w}, \qquad c_\kappa \mathcal{A} \le \kappa_q \le 2C_\nu \mathcal{A}, \tag{19}$$

with $c_w = 4c_\nu^2$, $C_w = 4C_\nu^2$ and $c_\kappa = c_\nu/2$. Moreover $\mathcal{M}$ is odd and nondecreasing, $|\mathcal{M}(\theta)| \le 2\|g\|_\infty |\theta|$ for every θ, and

$$\tfrac{11}{6} J\theta \le \mathcal{M}(\theta) \le \tfrac{13}{6} J\theta \quad (0 \le \theta \le y_1), \qquad \mathcal{M}(\tfrac{5}{4}\theta) \le \tfrac{3}{2}\mathcal{M}(\theta) \quad (0 \le \theta \le \tfrac{4}{5} y_1). \tag{20}$$

Proof. We first compare $\widehat{w}$ with $\mathcal{A}\mathcal{A}'$, then bound both bridge weights by the same product of truncated responses, and finally obtain the lower bound needed for agreement within a cycle from the sign of a covariance. Two nondecreasing functions of the same random variable have nonnegative covariance: for an independent copy Y', the product $(f_1(Y)-f_1(Y'))(f_2(Y)-f_2(Y'))$ is nonnegative. Taking its expectation with $f_1(y)=\min(ty,1)$ and $f_2(y)=\min(t'y,1)$ gives $\widehat{w}\ge\mathcal{A}\mathcal{A}'$. The upper bound follows from $\min(tY,1)\le 1$.

Let $t_1=\varepsilon_1/C_0$, reduced if necessary so that $\mathbb{P}(|Q|\le a|U|)\ge 1/4$ whenever $|a-1|\le C_0t_1$. This is possible because $\mathbb{P}(|Q|\le|U|)=1/2$ for two independent projections and the probability is continuous in a. By (14), every variable W_i of (16), the variable $Q-a_qU-c_qV$ behind $\widetilde{L}_{q,t}$, and $-\Lambda_d^0=Q-a_dU-b_dU'-c_dV$ belong to $\mathcal{W}_{\varepsilon_1}$.

The bounds for w_B and w_C follow from (16) and (18). For κ_q, both factors in (15) are odd in U, so we may replace U by Y, and both are then nondecreasing functions of Y. The upper bound uses $|F-\frac12|\le\frac12$ and (18). The covariance inequality gives

$$\kappa_q\ge 4\big(\mathbb{E}F(a_{q^-}Y)-\tfrac12\big)\,\mathbb{E}\widetilde{L}_{q,t}(b_qY).$$

Since Q is symmetric, $\mathbb{E}F(aY)-\frac12=\mathbb{P}(0\le Q\le aY)=\frac12\mathbb{P}(|Q|\le aY)\ge\frac18$, and the second factor is at least $c_\nu\mathcal{A}$ by (18). Finally, let k_d be the density of $-\Lambda_d^0$. Then $\mathcal{M}(\theta)=\int_{-\theta}^{\theta}k_d$ for $\theta\ge0$, which is nondecreasing and at most $2\|g\|_\infty\theta$, and (17) gives (20). The last inequality holds because $\frac54\cdot\frac{13}{11}<\frac32$. □

7 From local decisions to most probable triangulations

We still consider one block and the probability $p^\theta_{t,t'}(s)$ of a vector s of signs of its linear tests. We first show how to increase the probability of a sign vector whose signs vary around a cycle. We then show that, in every most probable vector, the positive cycles encode a maximum independent set. We keep the probability within each cycle as one factor. Two additional terms per bridge describe dependence between cycles, and the fields adjust the reward for a positive phase. This follows the construction: cycles represent vertices and bridges represent edges.

For a vertex v, let $\pi_v(s)$ be $2^{|\mathcal{C}_v|}$ times the marginal probability of the signs on $\mathcal{C}_v$, including the field of its designated cell. It depends only on those signs, on t and on that field. $\pi_v^0(s)$ denotes its value when the field is zero. The notation q^- denotes the predecessor of a cycle cell q. $q_B(x)$ and $q_C(x)$ are the two cells incident to bridge x, as in (8). We abbreviate $\mathcal{A}=\mathcal{A}(t)$, $\mathcal{A}'=\mathcal{A}(t')$ and $\widehat{w}=\widehat{w}(t,t')$. The constants in O_ν may depend on the law and on the fixed coefficient bound C_0 in (14). Here $C_0=2$.

Lemma 15. *Suppose $0<t'\le t<1/4$, the fields vanish outside the designated cells of the cycles, and $|\theta_q|\le\mathcal{A}$. If $m\mathcal{A}$ is below a sufficiently small constant depending only on ν, then, uniformly over all sign vectors and these fields,*

$$2^m p^\theta_{t,t'}(s)=\prod_{v\in V(G)}\pi_v(s)+w_B\sum_x s_xs_{q_B(x)}-w_C\sum_x s_xs_{q_C(x)}+O_\nu(m^2\mathcal{A}\,\widehat{w}).\tag{21}$$

For each cycle, with zero field,

$$\pi_v^0(s)=1+\sum_{q\in\mathcal{C}_v}\kappa_qs_qs_{q^-}+O_\nu(|\mathcal{C}_v|^2\mathcal{A}^2),\tag{22}$$

and reversing all its signs leaves $\pi_v^0(s)$ unchanged. A field θ at the designated cell changes $\pi_v(s)$ by at most $6\|g\|_\infty|\theta|$, and the two constant sign vectors of the cycle satisfy

$$\pi_v(+)-\pi_v(-)=2\mathcal{M}(\theta)+O_\nu(|\mathcal{C}_v|\,\mathcal{A}\,|\theta|).\tag{23}$$

The proof is given in Section 7.1. Equation (22) measures the cost of each sign change within a cycle. Equation (21) then compares whole blocks while retaining the exact cycle factors, and (23) measures how a field separates the two phases of a cycle. The weights w_B, w_C and the bias $\mathcal{M}$ are exact, so no approximation of these three quantities enters the comparison. In particular, $\kappa_q = O_\nu(\mathcal{A})$ implies $\pi_v = 1 + O_\nu(|\mathcal{C}_v|\mathcal{A} + |\theta|)$. The product over any subset of cycles is therefore $1 + O_\nu(m\mathcal{A})$ for the fields used below. This controls the factors left unchanged when we repair one cycle, that is, when we replace the signs on a nonconstant cycle by one common sign.

The smaller scale makes this repair possible with a fixed margin in units of the relevant weights. We require $\widehat{w} \le c_1\mathcal{A}$, where $c_1 \in (0, 1]$ is a fixed rational constant depending only on ν and C_0, specified at the start of the proof of Lemma 16. This holds in particular when $\mathcal{A}' \le c_1\mathcal{A}$. By Lemma 14, the cycle gain is at least a constant times $\mathcal{A}$, whereas its three bridge links and its field cost at most a constant times $\widehat{w}$. Fig. 15 shows which signs and terms change in the repair.

Once every cycle has a constant sign, we must choose which cycles have positive phase. Put $W = \min(w_B, w_C)$. Without a field, the positive phase of a cycle would gain more from its three bridges than an edge between two positive phases costs, as the objective (25) below shows when $\mathcal{M}(\lambda)$ is replaced by zero. We therefore put the negative field $-\lambda$ on the designated cell of every cycle, with

$$\frac{7}{5}W \le \mathcal{M}(\lambda) \le \frac{13}{5}W. \tag{24}$$

The reduction in Section 8.1 lists rational scale pairs and fields. At least one pair satisfies $\widehat{w} \le c_1\mathcal{A}$, and at least one of its listed fields satisfies (24).

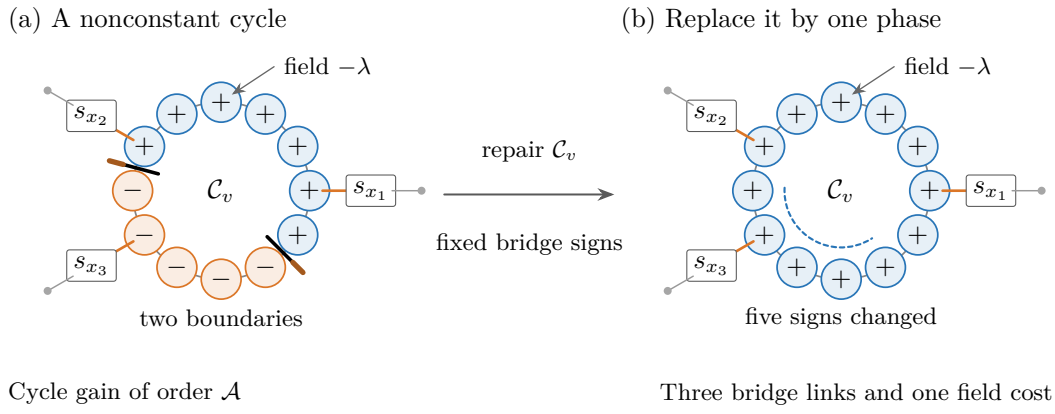

Figure 15: Repairing one cycle. Each round node represents a cell sign, and the three boxes represent its incident bridges. In (a), the two short strokes cross the boundaries between opposite signs. In (b), the dotted arc marks the five signs that change in this example. The bridge signs and all other cycle signs stay fixed. The cycle gains a term of order $\mathcal{A}$, while the three incident bridge links and the designated field cost at most a constant times $\widehat{w}$. Under the hypotheses of Lemma 16, the cycle-factor gain dominates these costs and the remainder. The drawing represents sign dependencies.

For fixed phases z_u, z_v, the best sign s_x of their bridge contributes $|w_B z_u - w_C z_v|$. Opposite phases give a value larger by exactly $2W$ than equal phases (Fig. 10(b)). If S contains the vertices of positive phase, cubicity gives $3|S| - 2|E(S)|$ edges leaving S, where $E(S)$ contains its internal edges. By (23), the field $-\lambda$ lowers the factor of each positive cycle by $2\mathcal{M}(\lambda)$ relative to a negative one, up to a smaller error. Up to a constant, the optimized bridge signs thus contribute $2W(3|S| - 2|E(S)|)$ and the fields $-2\mathcal{M}(\lambda)|S|$. The leading contribution in

(21) that depends on S is therefore $4W\Psi(S)$, where

$$\Psi(S) = \alpha|S| - |E(S)|, \qquad \alpha = \frac{3}{2} - \frac{\mathcal{M}(\lambda)}{2W}, \qquad \frac{1}{5} \le \alpha \le \frac{4}{5}. \tag{25}$$

Each selected vertex receives reward α, and each edge between selected vertices costs one. Thus a selected vertex with a selected neighbour can be profitably removed. Among independent sets, the reward increases with cardinality. Fig. 16 illustrates this familiar independent-set objective [26, Section 4.2].

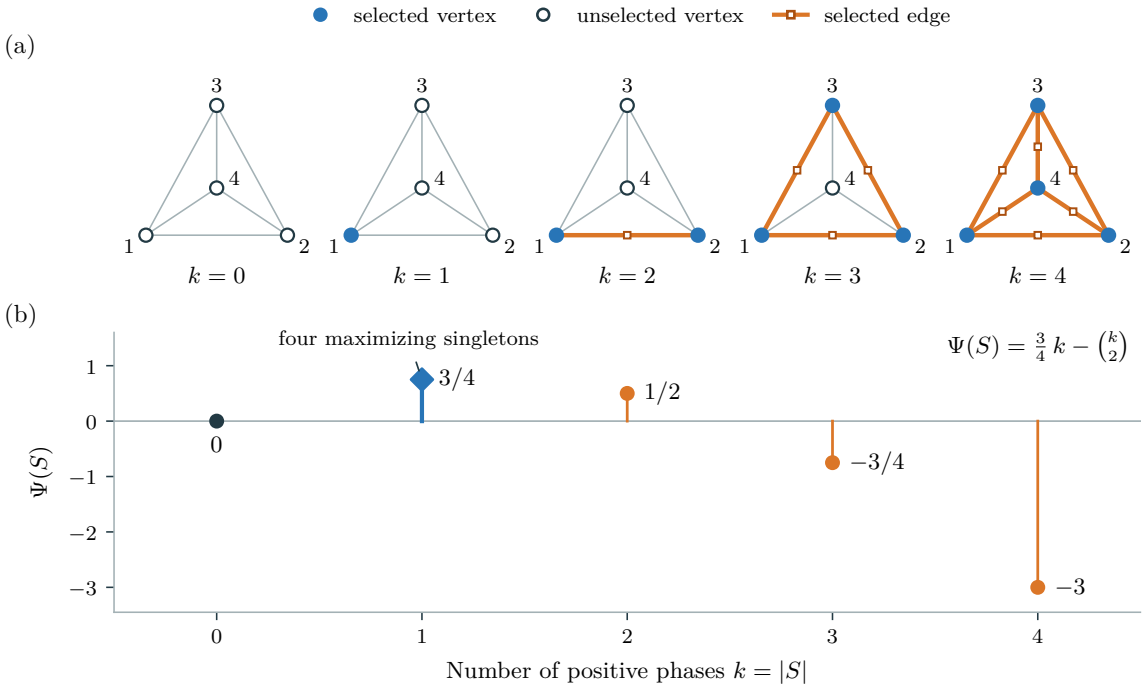

Figure 16: The independent-set objective on K_4. In (a), filled vertices are selected, and highlighted edges have both endpoints selected. In (b), $\alpha = 3/4$ gives $\Psi(S) = (3/4)k - \binom{k}{2}$, maximized by the four singleton sets. Blue marks the maximum, orange marks sets containing an edge, and black marks the empty set. This example takes $\mathcal{M}(\lambda) = 3W/2$, within (24), after each bridge sign has been optimized.

Lemma 16. *Suppose $t' \le t \le t_1$, $\widehat{w}(t,t') \le c_1\mathcal{A}(t)$, the fields satisfy (24), and $m^2\mathcal{A}(t)$ is below a sufficiently small constant depending only on ν. Call a sign vector correct if it is constant on every cycle and its positive cycles form a maximum independent set of G. Bridge signs are unrestricted. Every mode of $p^{\theta}_{t,t'}$ is correct. For some fixed rational $0 < d_\nu \le 1/2$, every incorrect vector s satisfies*

$$p^{\theta}_{t,t'}(s) \le (1 - d_\nu\mathcal{A}(t)\mathcal{A}(t')) \max_r p^{\theta}_{t,t'}(r). \tag{26}$$

Its probability is also at least $2d_\nu\mathcal{A}(t)\mathcal{A}(t')2^{-m}$ below a mode.

Proof. Put

$$C_\lambda = \frac{78\,C_w}{55\,J},$$

and let c_1 be a fixed rational with

$$0 < c_1 \le \min\left(1,\ \frac{1}{C_\lambda},\ \frac{c_\kappa}{6C_w + 12\|g\|_\infty C_\lambda}\right).$$

The entry $1/C_\lambda$ ensures that the field is at most $\mathcal{A}$, as Lemma 15 requires. The last entry bounds the combined loss from the field and the three bridge links when we repair a cycle. The hypothesis makes $m\mathcal{A}$ small. We first bound the field. By (24) and Lemma 14, $\mathcal{M}(\lambda) \le \frac{13}{5} C_w \widehat{w}$. Since $\widehat{w} \le \mathcal{A}$ is small, this is less than $\mathcal{M}(y_1)$, so $\lambda < y_1$ and (20) gives

$$\lambda \le C_\lambda \widehat{w} \le C_\lambda c_1 \mathcal{A} \le \mathcal{A}.$$

Thus Lemma 15 applies, every cycle factor is $1 + O_\nu(|\mathcal{C}_v|\mathcal{A} + \widehat{w})$, and the product of the factors of any set of cycles is $1 + O_\nu(m\mathcal{A})$.

Suppose first that s is nonconstant on $\mathcal{C}_v$. Its number of sign changes around this cycle is even and at least two. Give every cell of $\mathcal{C}_v$ one common sign, leaving all other signs unchanged, to obtain s'. Let $\kappa_{\min}$ be the least cycle weight κ_q on this cycle. Each sign change at a cell q lowers the sum in (22) by $2\kappa_q$, so this equation increases π_v^0 by at least $4\kappa_{\min} - O_\nu(|\mathcal{C}_v|^2\mathcal{A}^2)$. Since the field changes each of $\pi_v(s')$ and $\pi_v(s)$ by at most $6\|g\|_\infty\lambda$, we obtain $\pi_v(s') - \pi_v(s) \ge 4\kappa_{\min} - 12\|g\|_\infty\lambda - O_\nu(m^2\mathcal{A}^2)$. Only three bridge links change. Their total cost and the field loss satisfy

$$6\max(w_B, w_C) + 12\|g\|_\infty\lambda \le (6C_w + 12\|g\|_\infty C_\lambda)\widehat{w} \le c_\kappa\mathcal{A},$$

by the choice of c_1. Since $\kappa_{\min} \ge c_\kappa\mathcal{A}$, the gain on the repaired cycle is positive for sufficiently small $m^2\mathcal{A}$. The product of the other cycle factors is then at least $7/8$. Multiplying the cycle gain by this factor retains at least $(7/2)\kappa_{\min}$ of the leading term. Subtracting the field and bridge losses above and absorbing the remainder of (21) into $O_\nu(m^2\mathcal{A}^2)$ gives a net gain of at least $(5/2)c_\kappa\mathcal{A} - O_\nu(m^2\mathcal{A}^2)$. Thus

$$2^m\left[p^\theta_{t,t'}(s') - p^\theta_{t,t'}(s)\right] \ge 2c_\kappa\mathcal{A}.$$

The smaller bridge scale thus makes either constant sign adequate for this comparison.

Now suppose every cycle is constant, with phases z_v. Let $\bar{\pi}_v$ be the mean of $\pi_v(+)$ and $\pi_v(-)$. Equation (23) with $\theta = -\lambda$ gives directly

$$\pi_v(z_v) = \bar{\pi}_v - z_v\mathcal{M}(\lambda) + O_\nu(|\mathcal{C}_v|\mathcal{A}\lambda).$$

Multiplying over the vertices,

$$\prod_v \pi_v(z_v) = \prod_v \bar{\pi}_v - \mathcal{M}(\lambda)\sum_v z_v + O_\nu(m^2\mathcal{A}\,\widehat{w}).$$

The product of the means is $1 + O_\nu(m\mathcal{A})$ and does not depend on the phases. The error collects $m\mathcal{A}\cdot N\lambda$, $m\mathcal{A}\lambda$ and $(N\lambda)^2$, all $O_\nu(m^2\mathcal{A}\widehat{w})$ since $N \le m$ and $\lambda \le \min(C_\lambda\widehat{w}, \mathcal{A})$. Since $\sum_v z_v = 2|S| - N$, maximizing the bridge terms for fixed phases gives $4W\Psi(S)$ up to an additive constant and the stated error, with Ψ as in (25). If the positive set S contains an edge, deleting a selected vertex with a selected neighbour gains at least $1 - \alpha \ge 1/5$ in Ψ. If S is independent but not maximum, replacing it by a maximum independent set gains at least $\alpha \ge 1/5$. Choose optimal bridge signs for the new phases. Whatever the original bridge signs were, the main terms improve by at least $4W/5$, because their original value was at most the optimized value for S. The errors total at most $2W/5$ for sufficiently small $m^2\mathcal{A}$, using $W \ge c_w\widehat{w}$. The normalized gain is therefore at least

$$\frac{2}{5}W \ge \frac{2c_w}{5}\widehat{w} \ge \frac{2c_w}{5}\mathcal{A}\mathcal{A}'.$$

Choose a fixed rational $0 < d_\nu \le \min(1/2, c_\kappa, c_w/5)$. Both cases improve every incorrect vector by at least $2d_\nu\mathcal{A}\mathcal{A}'2^{-m}$, so every mode is correct and the stated absolute gap holds. Finally, (21) and $\widehat{w} \le \mathcal{A}$ give $2^m p^\theta_{t,t'}(s) \le 2$ for every vector under the same smallness condition. Dividing the absolute gap by the modal probability gives the relative gap (26). □

7.1 Deriving the probability estimates

We expand the probability of a sign vector in the correlations of the signs of sets of cells. The sets without a bridge sum to the product of the cycle factors, and the two endpoint pairs of each bridge have the correlations w_B and $-w_C$ of (15). Every other set that contains a bridge is an error term. To bound these terms, we condition on the projections U_q and V_q of the cycle cells and write the conditional mean of the sign of each cycle cell as a main part plus a small correction. The main part depends on one projection and has mean zero. By symmetry, a product vanishes when a main part occurs without the correction of the next cell, or a bridge without the correction of either endpoint. The remaining products are small.

Proof of Lemma 15. Sign correlations. For a set L of cells, write $s_L = \prod_{q\in L} s_q$ and $\varrho_L = \mathbb{E}\prod_{q\in L}\operatorname{sgn}\Lambda_q$, both equal to one when L is empty. Each test is almost surely nonzero, and the indicator of the event $\operatorname{sgn}\Lambda_q = s_q$ is then $(1 + s_q \operatorname{sgn}\Lambda_q)/2$. Multiplying these indicators over the m cells and taking the expectation gives

$$2^m p^{\theta}_{t,t'}(s) = \sum_L s_L \varrho_L, \tag{27}$$

where L ranges over all sets of cells. The same identity on one cycle gives $\pi_v(s) = \sum_{L\subseteq\mathcal{C}_v} s_L\varrho_L$. Cells of different cycles have no site in common, so $\varrho_L = \prod_v \varrho_{L\cap\mathcal{C}_v}$ for every set L without bridges. These sets therefore contribute exactly $\prod_v \pi_v(s)$ to (27), with the fields included. A bridge x and its endpoints carry no field, so (15) gives $\varrho_L = w_B$ for $L = \{x, q_B(x)\}$ and $\varrho_L = -w_C$ for $L = \{x, q_C(x)\}$. These pairs give the two sums over bridges in (21). It remains to bound the other sets that contain a bridge and to examine one cycle.

Main parts and corrections. Condition on the projections U_q and V_q of all cycle cells. Every other projection (namely Q_q on each cycle cell and U_x, Q_x on each bridge) is used by one cell only. The signs are therefore conditionally independent, and ϱ_L is the mean of the product over L of the conditional means of the signs, which we denote by χ_q. A cycle cell has $\chi_q = 2F(a_qU_q + b_qU_{q^-} + c_qV_q + \theta_q) - 1$. We write $\chi_q = \xi_q + \Delta_q$, where the main part $\xi_q = 2F(a_qU_q) - 1$ is an odd function of U_q alone with $|\xi_q| \le 1$, and Δ_q is the correction. For a bridge x with endpoints $q = q_B(x)$ and $r = q_C(x)$, averaging the two unshared projections in (8) gives $\chi_x = 2\mathbb{P}(Q - a'U < b'V_q - c'V_r) - 1$, an odd function of (V_q, V_r). The variable $Q - a'U$ has a density at most $\|g\|_\infty$, so its distribution function satisfies the Lipschitz bound in (11) as well. Put $u(y) = \min(ty, 1)$ and $u'(y) = \min(t'y, 1)$. With the constant C_ν of (18), this Lipschitz bound, the trivial bound $|F(x) - F(y)| \le 1$ and the ranges (14) give

$$|\Delta_q| \le M_q = 2C_\nu\big[u(|U_{q^-}|) + u(|V_q|)\big] + 2\|g\|_\infty|\theta_q|, \qquad |\chi_x| \le M_x = 2C_\nu\big[u'(|V_q|) + u'(|V_r|)\big].$$

Expanding $\chi_q = \xi_q + \Delta_q$ for the cycle cells of every L turns (27) into a sum of terms

$$T(K,H,X) = \mathbb{E}\Big[\prod_{q\in K}\Delta_q \prod_{q\in H}\xi_q \prod_{x\in X}\chi_x\Big],$$

each multiplied by $s_{K\cup H\cup X}$. Here K and H are disjoint sets of cycle cells, represented by their corrections and by their main parts, respectively, and X is a set of bridges. Most terms vanish by one of three symmetries. (a) If a cell q of H is the predecessor of no cell of K, the projection U_q occurs only in the factor ξ_q, whose mean is zero. (b) If a bridge of X has no endpoint in K, its two shared projections occur only in χ_x, which changes sign when both projections are negated together. (c) Negating all projections reverses every main part, every bridge factor χ_x and every correction of a cell without field. A term with an odd number of factors, none of them the correction of a cell with a field, is therefore zero.

In a term not cancelled by (a), every cell of H is the predecessor of a cell of K, so each K admits at most $2^{|K|}$ sets H. Bounding the main parts by one leaves the mean of the product

of the majorants M_q, $q \in K$, and M_x, $x \in X$. Majorants of different cycle cells depend on disjoint projections, because M_q uses only U_{q^-} and V_q and $q \mapsto q^-$ is bijective on each cycle, while M_x shares projections only with the majorants of the endpoints of x. Each cycle cell meets at most one bridge (Lemma 11). The mean therefore factors over the bridges of X, each taken with its endpoints in K, and the remaining cells of K. For a cycle cell q, and for a bridge x taken with one endpoint q or with both endpoints q, r,

$$\mathbb{E}M_q \le 6C_\nu\mathcal{A}, \qquad \mathbb{E}[M_xM_q] \le 16C_\nu^2\widehat{w}, \qquad \mathbb{E}[M_xM_qM_r] \le 64C_\nu^3\widehat{w}\mathcal{A}. \tag{28}$$

The first bound uses $|\theta_q| \le \mathcal{A}$ and $\|g\|_\infty \le C_\nu$. The endpoints carry no field. Expanding the products, the second mean is $4C_\nu^2(\widehat{w} + 3\mathcal{A}\mathcal{A}')$ and the third is $32C_\nu^3\mathcal{A}(\widehat{w} + \mathcal{A}\mathcal{A}')$, where $\mathcal{A}\mathcal{A}' \le \widehat{w}$ by the covariance inequality in the proof of Lemma 14. By hypothesis $m\mathcal{A}$ is small, so $4C_\nu\mathcal{A} \le 1$ and $|\mathcal{C}_v|\mathcal{A}$ is small for every cycle.

One cycle without field. Consider the terms with $K, H \subseteq \mathcal{C}_v$ and $X = \emptyset$, at zero field. With their signs $s_{K\cup H}$, they sum to $\pi_v^0(s)$. By (a), the only term with $K = \emptyset$ is the empty one, equal to one. For $K = \{q\}$, the term with $H = \emptyset$ vanishes by (c). The term with $H = \{q^-\}$ is $\mathbb{E}[\Delta_q\xi_{q^-}]$. Averaging Δ_q over U_q and V_q leaves $2\widetilde{L}_{q,t}(b_qU_{q^-})$, while $\xi_{q^-} = 2F(a_{q^-}U_{q^-}) - 1$, so this term equals κ_q by (15). A set K with l cells can be chosen in at most $|\mathcal{C}_v|^l/l!$ ways and admits at most 2^l sets H, and each of its terms has absolute value at most $(6C_\nu\mathcal{A})^l$. The terms with $l \ge 2$ therefore total at most $\sum_{l\ge2}(12C_\nu|\mathcal{C}_v|\mathcal{A})^l/l!$, which is $O_\nu(|\mathcal{C}_v|^2\mathcal{A}^2)$. This proves (22). Negating every projection of the cycle reverses all its zero-field tests, so reversing all its signs leaves $\pi_v^0(s)$ unchanged.

Sets containing a bridge. By (b), a term with $X \ne \emptyset$ vanishes unless K contains an endpoint of every bridge of X. Fix X and add the absolute values of its terms over K and H, letting each cell of K choose whether its predecessor is in H. By the factorization above, the sum is at most a product with one factor for each bridge of X and one for every other cycle cell. The factor of a bridge x with endpoints q, r adds the three ways in which K can meet its endpoints,

$$2\mathbb{E}[M_xM_q] + 2\mathbb{E}[M_xM_r] + 4\mathbb{E}[M_xM_qM_r] \le 128C_\nu^2\widehat{w},$$

by (28) and $4C_\nu\mathcal{A} \le 1$. The factor of any other cycle cell q' is $1 + 2\mathbb{E}M_{q'} \le 1 + 12C_\nu\mathcal{A}$, which covers its absence from K and its presence with or without its predecessor in H. The terms of X therefore total at most $(128C_\nu^2\widehat{w})^{|X|}e^{12C_\nu m\mathcal{A}}$. There are at most $m^j/j!$ sets X with j bridges. The terms with at least two bridges thus total $O_\nu(m^2\widehat{w}^2)$, which is $O_\nu(m^2\mathcal{A}\widehat{w})$ because $\widehat{w} \le \mathcal{A}' \le \mathcal{A}$.

Now let $X = \{x\}$. The terms in which K contains a cell other than the endpoints of x total at most $128C_\nu^2\widehat{w}\,(e^{12C_\nu m\mathcal{A}} - 1) = O_\nu(m\mathcal{A}\widehat{w})$, by the same count. When K consists of both endpoints, each of the at most four terms is at most $64C_\nu^3\widehat{w}\mathcal{A}$. There remains $K = \{q\}$ with q an endpoint of x. The term with $H = \emptyset$ is the only nonzero term of ϱ_L for $L = \{x, q\}$, already counted as w_B or $-w_C$. The term with $H = \{q^-\}$ has three factors. Main parts involve no field, even at a designated cell, so this term vanishes by (c). Summing over at most m bridges, the terms with one bridge other than the two sums in (21) total $O_\nu(m^2\mathcal{A}\widehat{w})$. This proves (21).

The field of one cycle. Let d be the designated cell of $\mathcal{C}_v$, with field θ, and let $\mathcal{E}$ be the event $|\Lambda_d^0| \le |\theta|$. Compute the signs with and without the field from the same projections. The two sign vectors differ at most at d, and only on $\mathcal{E}$. Given all projections other than Q_d, the signs of the other cells are determined and $\mathcal{E}$ has conditional probability at most $2\|g\|_\infty|\theta|$. The event that the remaining cells agree with s outside d is the union of the events of the two zero-field vectors matching s outside d. By the count for (22), each of those two vectors has probability at most $2^{-|\mathcal{C}_v|}e^{12C_\nu|\mathcal{C}_v|\mathcal{A}} \le \frac{3}{2}\,2^{-|\mathcal{C}_v|}$, so their union has probability at most $3 \cdot 2^{-|\mathcal{C}_v|}$. Multiplying by $2^{|\mathcal{C}_v|} \cdot 2\|g\|_\infty|\theta|$ shows that the field changes $\pi_v(s)$ by at most $6\|g\|_\infty|\theta|$.

For the two constant vectors, negating every projection of the cycle shows that $\pi_v(-)$ is the value of $\pi_v(+)$ under the field $-\theta$. Under the fields θ and $-\theta$, the events that all signs

are positive differ, up to a null event, exactly when $\mathcal{E}$ holds and the other signs are positive. Expanding the indicator of these $|\mathcal{C}_v| - 1$ signs as in (27),

$$\pi_v(+) - \pi_v(-) = 2\,\mathrm{sgn}(\theta) \sum_{L \subseteq \mathcal{C}_v \setminus \{d\}} \mathbb{E}\Big[\mathbf{1}_{\mathcal{E}} \prod_{q \in L} \mathrm{sgn}\,\Lambda_q\Big].$$

The empty set contributes $2\,\mathrm{sgn}(\theta)\mathbb{P}(\mathcal{E}) = 2\mathcal{M}(\theta)$, by (15). For the other sets, condition again on the projections U_q and V_q. The event $\mathcal{E}$ and the signs of the cells of L are conditionally independent. Let η be the conditional probability of $\mathcal{E}$. It is at most $2\|g\|_\infty|\theta|$, depends only on U_d, U_{d^-} and V_d, and is unchanged when all projections are negated. Each set L thus gives the terms $T(K, H, \emptyset)$ with the extra factor η inside the expectation. Since $d \notin L$, we have $d \notin H$, so rule (a) still applies to every cell of H other than d^-, whose projection U_{d^-} occurs in η. The term with $K = \emptyset$ and $H = \{d^-\}$ vanishes, because the negation of all projections preserves η and reverses ξ_{d^-}. A nonzero term thus has $K \neq \emptyset$, at most $2^{|K|+1}$ sets H and absolute value at most $2\|g\|_\infty|\theta|(6C_\nu\mathcal{A})^{|K|}$. Summing as for (22) gives $O_\nu(|\mathcal{C}_v|\,\mathcal{A}\,|\theta|)$, which proves (23). □

8 A polynomial construction and one oracle call

We now turn the probability comparisons of Section 7 into the reduction of Theorem 3. Given a planar cubic graph, we construct rational generic centres such that every most probable complete labelled triangulation yields a maximum independent set. The construction uses one solver call and polynomially many bits, while the noise law remains fixed.

For noise with a finite second moment, Proposition 17 gives one explicit block. For the general class, we enumerate a list of candidate blocks. At least one meets the scale and field bounds specified below. We call such a block *suitable*. The decoder reads every block and returns the largest independent set obtained, so it works without identifying which block is suitable. We then separate the blocks, complete their union, and choose a small geometric noise scale. A point that moves far can change the surrounding triangulation, including its hull. The comparison below covers these complete outputs by counting how many decision cells lose their local faces. Each distant point affects at most four cells in at most two blocks. Independence across the other blocks preserves the probability advantage needed for decoding.

Algorithm 2 summarizes the construction. We write m for the cell count of one block, including its bridges, B for the number of candidates, $M = Bm$ for their total cell count, and n for the site count after completion. Condition (L) makes both the candidate list and the coordinate bit lengths polynomial.

Algorithm 2 One-call reduction for a fixed law satisfying (L)

Require: Planar cubic graph G; constants fixed for ν
Ensure: A maximum independent set, using a modal solver once

1: Construct the vertex cycles and bridges. Determine their total cell count m.
2: Choose the thresholds and scale bound in (29).
3: Build a block for every pair $3 \le j \le j' \le j_{\max}$ and field in (30).
4: Separate and complete the blocks jointly. Record the base configuration and its size.
5: Compute the relative gap β in (32) and choose the tail and sign-error bounds in (37).
6: Choose the noise scale and generic displacements from Section 8.2. Form the centres (34).
7: Query the modal solver once for a complete labelled triangulation.
8: Read every block from its cell faces, using the decoder of Section 8.3. Return the largest retained independent set.

8.1 Choosing the candidate blocks

Lemma 16 needs two scales with $\widehat{w}(t,t') \le c_1\mathcal{A}(t)$, for which $\mathcal{A}(t') \le c_1\mathcal{A}(t)$ suffices, fields in (24), and a small value of $m^2\mathcal{A}(t)$. We obtain the separation by taking the first scales at which $\mathcal{A}$ falls below two thresholds, a and $c_1a/2$, defined next. The reduction lists all possible pairs up to a computable bound and a grid of fields for each pair. Recall the fixed rational c_1 from Section 7, and set

$$\begin{aligned} \ell &= \lceil \log_2(m+1) \rceil, \qquad a = 2^{-K_\nu \ell}, \\ j_{\max} &= \min\{ j \ge 3 : C_A j^{-\gamma} \le c_1 a/2 \}. \end{aligned} \tag{29}$$

The minimum is taken over integers, and C_A is the integer in (13). Choose the fixed integer K_ν first so that $a < a_0/4$. By (12), $\mathcal{A}(1/4) \ge a_0/4$, so the first crossing of a has index at least three. The decay bound (13) guarantees that both thresholds have been crossed by $j_{\max}$.

Let j and j' be the first indices at which $\mathcal{A}(2^{-j}) \le a$ and $\mathcal{A}(2^{-j'}) \le c_1a/2$, respectively (Fig. 17(b), p. 42). Since $\mathcal{A}(2t) \le 2\mathcal{A}(t)$, these first crossings give

$$a/2 < \mathcal{A}(t) \le a, \qquad c_1a/4 < \mathcal{A}(t') \le c_1a/2, \qquad t = 2^{-j}, \quad t' = 2^{-j'}.$$

In particular, $j \le j'$ and $\mathcal{A}(t') \le c_1\mathcal{A}(t)$. At this pair,

$$m^2\mathcal{A}(t) < 2^{-\ell(K_\nu - 2)}.$$

One fixed large K_ν makes this smaller than every constant required by Lemma 16, uniformly in $\ell \ge 1$. Increasing it also puts $t \le a/a_0$ below all fixed small-scale thresholds, including t_1 of Lemma 14 and $2\|g\|_\infty t < c_w a_0^2$ for the field grid below. The requirements on K_ν in the rest of this subsection have the same form: a quantity bounded by a fixed multiple of 2^{-K_ν} must lie below a positive threshold depending only on ν. The bound $j_{\max} = O_\nu(m^{K_\nu/\gamma})$ is polynomial for every fixed $\gamma > 0$. To compute $j_{\max}$, raise the defining inequality to the denominator of the fixed rational exponent γ and clear denominators. This gives an integer comparison for each candidate j.

List every pair $3 \le j \le j' \le j_{\max}$. At each pair $t = 2^{-j}$, $t' = 2^{-j'}$, list fields

$$\lambda_k = t^2 t' (5/4)^k, \qquad k = 0, 1, \dots, \tag{30}$$

up to the first value at least $1/4$. Put field $-\lambda_k$ only on the designated cell of each cycle. All bridge fields vanish.

At the pair located above, one of these fields satisfies (24), although the reduction evaluates neither $W = \min(w_B, w_C)$ nor $\mathcal{M}$. The first grid value has a bias below the interval $[\frac{7}{5}W, \frac{13}{5}W]$ of (24) and the last one a bias above it. Between them, one step of the grid multiplies a small bias by at most $3/2$, less than the ratio $13/7$ of the two ends of the interval. The first bias to reach the lower endpoint therefore stays below the upper endpoint. Indeed, Lemma 14 and (12) give $W \ge c_w\mathcal{A}\mathcal{A}' \ge c_w a_0^2 tt'$. Thus $\mathcal{M}(\lambda_0)/W \le 2\|g\|_\infty t/(c_w a_0^2) < 1$, below the target $7/5$. The last grid value is at least $1/4$, and $\mathcal{M}(\min(y_1, 1/4))$ exceeds $13W/5$ for large K_ν, because $W \le C_w\mathcal{A}' \le C_w c_1 a/2$. Let $k \ge 1$ be the first index with $\mathcal{M}(\lambda_k) \ge 7W/5$. Then $\mathcal{M}(\lambda_{k-1}) < 7W/5$. For large K_ν this is below $\mathcal{M}(4y_1/5)$, which forces $\lambda_{k-1} \le 4y_1/5$, and (20) gives

$$\frac{7}{5}W \le \mathcal{M}(\lambda_k) \le \frac{3}{2}\mathcal{M}(\lambda_{k-1}) < \frac{21}{10}W < \frac{13}{5}W,$$

as required by (24). The grid has $O(j + j' + 1)$ elements, all below $5/16$, of bit length

$$O(j + j' + 1). \tag{31}$$

The complete list contains $B = O(j_{\max}^3)$ blocks. A listed block is suitable when $a/2 < \mathcal{A}(t) \le a$, $c_1a/4 < \mathcal{A}(t') \le c_1a/2$, and its field satisfies (24). The first crossings and the grid argument

prove that such a block exists. Since $m = O(N^3)$, the list is polynomial in the graph size. Its construction evaluates neither $\mathcal{A}$ nor the bridge weights.

We use a computable lower bound for the relative gap of Lemma 16. Write $p_b(s)$ for the probability of sign vector s in a block and $p_b^* = \max_s p_b(s)$. In a suitable block the first-crossing bounds give $\mathcal{A}\mathcal{A}' > c_1 a^2/8$. Therefore every incorrect vector satisfies

$$p_b(s) \le (1 - d_\nu \mathcal{A}\mathcal{A}')p_b^* \le (1-\beta)p_b^*, \qquad \beta = \frac{d_\nu c_1 a^2}{8}. \tag{32}$$

Here d_ν is the fixed rational constant of Lemma 16, so β is rational and computable. We have $0 < \beta \le 1/2$ and $\beta^{-1} = O_\nu(a^{-2})$, polynomial in m. Bridge signs are included in each probability. Correctness requires constant cycle signs encoding a maximum independent set, with no restriction on the bridge signs.

When $\mathbb{E}\|Z\|^2 < \infty$, the list is unnecessary. The leading behaviour of the weights is then known in advance, and one block with explicit parameters satisfies the hypotheses of Lemma 16. Its two scales can even coincide, because $\widehat{w}(t,t)$ is then of order t^2 while $\mathcal{A}(t)$ is of order t.

Proposition 17. *Suppose $\mathbb{E}U^2 < \infty$, and fix a rational κ_0 with $\frac{14}{5} J\,\mathbb{E}U^2 < \kappa_0 < \frac{26}{5} J\,\mathbb{E}U^2$. There is an integer K, depending only on ν and κ_0, such that for every m and $\ell = \lceil \log_2(m+1) \rceil$ the block with*

$$t = t' = 2^{-K\ell}, \qquad \lambda = \kappa_0 t^2$$

satisfies the hypotheses of Lemma 16. For this block the comparison (32) *holds with $\beta = d_\nu a_0^2\, 2^{-2K\ell}$, after replacing a_0 by a rational lower bound. For the standard Gaussian, $J\,\mathbb{E}U^2 = 1/(2\sqrt{\pi})$, and one may take $\kappa_0 = 1$, that is, $\lambda = t^2$.*

Proof. We check the cycle and bridge scales and then compute their limiting ratio to the field. First, $\widehat{w}(t,t) \le t^2 \mathbb{E}U^2 \le c_1 a_0 t \le c_1 \mathcal{A}(t)$ for $t \le c_1 a_0/\mathbb{E}U^2$. Second, $m^2 \mathcal{A}(t) \le m^2 t\, \mathbb{E}|U| \le 2^{(2-K)\ell} \mathbb{E}|U|$, which is small for large K. Third, let $t' = t \to 0^+$. The densities of the variables W_i in (16) converge uniformly to h by Lemma 13, so $L_i(cy)/t \to Jy$ for each $y \ge 0$ whenever $c/t \to 1$, while $L_i(cy)/t \le C_\nu y$ by (18). The ratios $c(t)/t$, b'/t and c'/t tend to one and $\mathbb{E}Y^2 < \infty$, so dominated convergence gives $w_B/t^2 \to 4J^2\mathbb{E}U^2$ and $w_C/t^2 \to 4J^2\mathbb{E}U^2$. Likewise $\mathcal{M}(\kappa_0 t^2)/t^2 \to 2J\kappa_0$. Hence $\mathcal{M}(\lambda)/W \to \kappa_0/(2J\,\mathbb{E}U^2)$, which lies strictly between $7/5$ and $13/5$, so (24) holds for small t. A large K puts $t \le 2^{-K}$ below t_1 and below the thresholds of these three conditions. The gap follows from Lemma 16 and $\mathcal{A}(t)\mathcal{A}(t') = \mathcal{A}(t)^2 \ge a_0^2 t^2$. For the standard Gaussian, $J = 1/(2\sqrt{\pi})$ and $\mathbb{E}U^2 = 1$, and the interval for κ_0 contains $[0.79, 1.46]$. □

With this block in place of the candidate list, steps 2 and 3 of Algorithm 2 become one explicit construction with $B = 1$. The rest of the reduction uses a suitable block only through its relative gap, so it applies with this value of β, and the decoder reads a single block. This gives a direct route to Theorem 1. For laws with infinite variance, we retain the candidate list.

8.2 Local geometric stability

So far the signs were those of the linear tests Λ_q. When the four displaced sites of a cell retain their convex cyclic order, the exact incircle test Γ_q selects their local diagonal. We compare its *geometric sign* with the *linear sign* of Λ_q using the same displacements Z_i. This subsection makes the two signs agree outside events of small probability. We choose a radius R and call a site *large* when its displacement Z_i has norm greater than R. All other sites stay sufficiently close to their base positions to preserve the strict orientation and incircle certificates involving only those sites. A separate threshold τ controls sign errors near zero. Keeping these two sources of error apart lets us treat large sites individually without requiring all sites to remain close at once. All realizations below are generic outside a null event, which we discard.

Translate the blocks by even integer vectors into boxes at mutual distance at least four. Their circle centres remain in $2\mathbb{Z}^2$, and sites in different blocks have independent displacements. Complete the union by Lemma 12, obtaining a base configuration P with $n \le 12M + 4$ sites and polynomial coordinate bit length. This completion includes the hull and all surrounding faces. Its added sites belong to no decision cell.

Let $0 < r_* \le 1$ be a rational lower bound for the absolute values of all strict orientation and incircle determinants listed after (9). We call these determinants the *certificates*. Let $C_P \ge 1$ bound the sums of the absolute values of the nonconstant Taylor coefficients of the certificates and of the normalized tests Γ_q, at P. Their degree is at most four. If the relevant sites move by vectors of norm at most $u \le 1$, every certificate changes by at most $C_P u$, and the part of Γ_q beyond its linear term has absolute value at most $C_P u^2$.

For now $R \ge 1$ and $0 < \tau < 1$ are free parameters. The proof of Theorem 1 fixes them for Gaussian noise, and (37) does so in general. Let σ be the largest dyadic power with

$$\sigma \le \frac{\min(r_*/400, \tau)}{2C_P(R+2)^2}. \tag{33}$$

Put $d_{\mathbf{D}_q} = -\theta_q(\mathbf{D}_q - o_q)$ and $d_i = 0$ elsewhere. The radius is one, so $\|d_{\mathbf{D}_q}\| = |\theta_q| < 1/2$ for every listed field and for the single-block choice with K sufficiently large. Choose rational displacements ζ_i, site by site, on a grid with $n^3 + 1$ points per side in the open square $\|\zeta\|_\infty < \sigma\tau/32$. Choose ζ_i so that $P_i + \sigma d_i + \zeta_i$ avoids the earlier centres in these coordinates, their pairwise lines and their triple circumcircles. Appendix A.3 bounds the excluded grid points, so this makes the centres generic in polynomial time. Define

$$\mu_i = \frac{P_i + \sigma d_i + \zeta_i}{\sigma}, \qquad X_i = \mu_i + Z_i. \tag{34}$$

In the coordinates of P, the realization is $\sigma X_i = P_i + \Delta_i$, where $\Delta_i = \sigma(d_i + Z_i) + \zeta_i$. The homothety leaves the triangulation unchanged and makes the noise scale one in the final instance.

For every nonlarge site, $\|\Delta_i\| < \sigma(R+2)$. The choice of σ gives

$$\sigma(R+2) \le \frac{r_*}{800C_P} \le \frac{1}{800}, \qquad C_P\sigma(R+2) \le r_*/2, \qquad C_P\sigma(R+2)^2 \le \tau/2.$$

Consequently all strict certificates involving only nonlarge sites remain valid. If the four sites of a cell are nonlarge, its normalized test satisfies

$$|\Gamma_q(P + \Delta)/\sigma - \Lambda_q| < \tau.$$

Indeed, the nonlinear part is at most $\tau/2$. The generic displacements contribute less than $3\sqrt{2}\tau/32 < \tau/2$, because the four gradient norms sum to $a_q + b_q + c_q + 1 = 2 + 2b_q \le 3$, by (5), (6) and Table 2.

Let $\mathcal{N}_\tau$ be the event that $|\Lambda_q| \le \tau$ for some cell. The projection at the site $\mathbf{D}_q$, used only by cell q, has coefficient -1 and density at most the fixed rational $L_g \ge \max(1, \|g\|_\infty)$. Conditioning on the displacements of all other sites bounds the probability of $|\Lambda_q| \le \tau$ by $2L_g\tau$ for each cell, and a union bound over the M cells gives

$$\mathbb{P}(\mathcal{N}_\tau) \le 2L_g M\tau.$$

Outside this event, every cell with four nonlarge sites has the same geometric and linear decision sign.

Proof of Theorem 1. Use the single Gaussian block of Proposition 17, so $M = m$. Each linear sign vector s selects two triangles per cell. Together with the f_3 fixed triangles they give a complete type $T(s)$, by the count following (9). The chosen diagonals distinguish these types. A linear mode has probability $p_b^* \geq 2^{-m}$, and every incorrect vector has probability at most $(1-\beta)p_b^*$.

Choose R and τ so that

$$ne^{-R^2/2} + 2L_g m\tau \leq \beta 2^{-m}/4.$$

Since $\mathbb{P}(\|Z_i\| > R) = e^{-R^2/2}$, an integer $R = O(m+\log(n/\beta))$ and a dyadic τ with polynomially many bits suffice. Except when some site is large or $\mathcal{N}_\tau$ occurs, every strict certificate holds and every geometric sign agrees with its linear sign. Thus the geometric type equals $T(s)$, including its hull, except on an event of probability at most $\beta p_b^*/4$.

For a linear mode s, this gives geometric probability at least $(1-\beta/4)p_b^*$ for $T(s)$. For an incorrect vector, the corresponding upper bound is $(1-3\beta/4)p_b^*$. Any type outside the image of T occurs only on the exceptional event. Consequently every geometric mode yields a correct vector, even in a tie. Its cycle diagonals give a maximum independent set. Each certificate and normalized test involves at most four sites and has degree at most four. Its rational coefficients have polynomial bit length, as shown in Section 8.4. With these bounds, (33) and (34) give rational centres of polynomial bit length. The construction and decoding take polynomial time and use one modal query. Maximum independent set is NP-hard on 2-connected planar cubic graphs [28, Theorem 4.1(a)]. □

Under (L) alone, that global comparison can require too many bits for R. It needs all sites to be nonlarge except on an event of probability at most $\beta 2^{-M}/4$, and under (L) a tail this small can need an exponent of R that is exponential in M (Section 11). Section 8.4 instead makes the probability of a large displacement at most $\beta/(256n)$ per site. For a specified set of large sites, independence supplies one tail probability per site. We will show that each site can obscure only a bounded number of cell decisions, so the number of possibilities grows by at most a constant factor per site.

Lemma 18. *Suppose the four sites of a decision cell are nonlarge. The circumdisk through any three of them is contained in the open disk of radius* $11/10$ *about the base centre of the cell. After all large sites are deleted, the cell has its two selected triangular faces in the Delaunay triangulation. If neither face occurs in the complete triangulation before deletion, some large site lies within distance* $11/10$ *of this centre.*

Proof. We control the circumcircle after perturbation by bounding its determinant outside a slightly larger disk. Move the cell centre to the origin and write $u = \sigma(R+2) \leq r_*/800$. The three perturbed vertices have norms within u of one, and their orientation has absolute value at least $r_*/2$, with its original sign. For a point $x = (x_1, x_2)$, subtracting the final column from the squared-norm column of the incircle determinant of these three vertices and x expresses it as

$$-(\|x\|^2 - 1)\,\mathrm{orient}(P_1 + \Delta_1, P_2 + \Delta_2, P_3 + \Delta_3) + e(x),$$
$$|e(x)| \leq 20u(\|x\| + 1).$$

For completeness, each of the first three entries of that column has absolute value at most $2u + u^2$. Its cofactor has absolute value at most $2(1+u)(|x_1| + |x_2|) + 2(1+u)^2$. Summing the three products and using $|x_1| + |x_2| \leq \sqrt{2}\,\|x\|$ gives the displayed bound for $u \leq 1/800$. If $\|x\| \geq 11/10$, the main term has absolute value at least $r_*(\|x\| + 1)/20$, while the error is at most half that value. The determinant cannot vanish there. Thus the circumcircle, and by convexity its disk, lies within radius $11/10$.

After deleting large sites, the preserved certificates keep the four vertices convex and every other remaining site outside all four triple circumcircles. The geometric decision selects two empty triangles, which are Delaunay faces. Reinserting sites can remove an empty triangle only by placing a site inside its circumdisk. If neither selected face remains in the full triangulation, a reinserted large site therefore lies in the stated disk. □

8.3 Comparing every complete output

We now compare probabilities while allowing large sites anywhere in the configuration. Fix a complete labelled triangulation T and a set I of site labels, and consider the event that the triangulation is T and that the large sites are exactly those of I. The proof uses I to identify the cells whose signs can be read from T in every realization of this event. The decoder defined after Lemma 19 uses the returned type and the known construction. A cell is *intact* if none of its four labels is in I. Among intact cells, read those for which T contains a triangle on three of the four labels. Choose the first such triangle in label order. It selects one of the two base diagonals, and hence a sign. If another such face selects the other sign, the event under consideration is empty: after deleting I, both faces would persist, contradicting the unique triangulation of the intact quadrilateral.

Let k_b be the number of cells of block b whose signs cannot be read this way. Suppose the event determined by T, I is nonempty. Each label in I belongs to at most two decision cells, both in its original block. Lemma 18 accounts for the other unread cells: a large site must enter their disks of radius $11/10$. A point lies in at most two such disks (Fig. 17(a)). Indeed, three centres within this distance would be pairwise less than $22/10$ apart. In $2\mathbb{Z}^2$, each pair would then have distance two, and three lattice points cannot form an equilateral triangle. The boxes of different blocks are at distance at least four, so the two affected cells belong to one block. For every such nonempty event,

$$\sum_b k_b \le 4|I|, \qquad \#\{b : k_b > 0\} \le 2|I|.$$

Both bounds are multiples of $|I|$ and do not involve the number of blocks. Thus $|I|$ large sites leave at most $2|I|$ blocks undetermined. This is what allows the independent repetitions of Section 9 to use a truncation radius with polynomially many bits. Below, $q_R = T_\nu(R)$ denotes the probability that a given site is large.

Lemma 19. *For a fixed complete type T and fixed set I, let the determined blocks be those with $k_b = 0$, and let $s_b(T)$ be the sign vectors read from T. Then*

$$\begin{aligned}\mathbb{P}\{\mathrm{DT}(\sigma X) = T,\ \{i : \|Z_i\| > R\} = I,\ \mathcal{N}_\tau^c\}\\ \le (16q_R)^{|I|} \prod_{k_b=0} p_b(s_b(T)) \prod_{k_b>0} p_b^*.\end{aligned} \tag{35}$$

If the event on the left is nonempty, at most $2|I|$ blocks are undetermined.

Proof. The inequality is immediate for an empty event. Otherwise, every face of T whose vertices are outside I remains a Delaunay face after deleting I, since its circumcircle remains empty. Thus each sign read above agrees with the geometric decision of the intact cell, and with its linear test outside $\mathcal{N}_\tau$. All these tests depend only on displacements of sites outside I, because the cells that are read are intact. The event on the left of (35) is therefore contained in the intersection of two independent events: the tests of the read cells have the signs read from T, and $\|Z_i\| > R$ for every $i \in I$. Blocks share no site, so the probability of the first event factors over the blocks. In block b, the partial vector read from T has at most 2^{k_b} completions and hence probability at most $2^{k_b} p_b^*$. When $k_b = 0$, its probability is exactly $p_b(s_b(T))$. The factor $q_R^{|I|}$ and the inequality $\sum_b k_b \le 4|I|$ prove (35). □

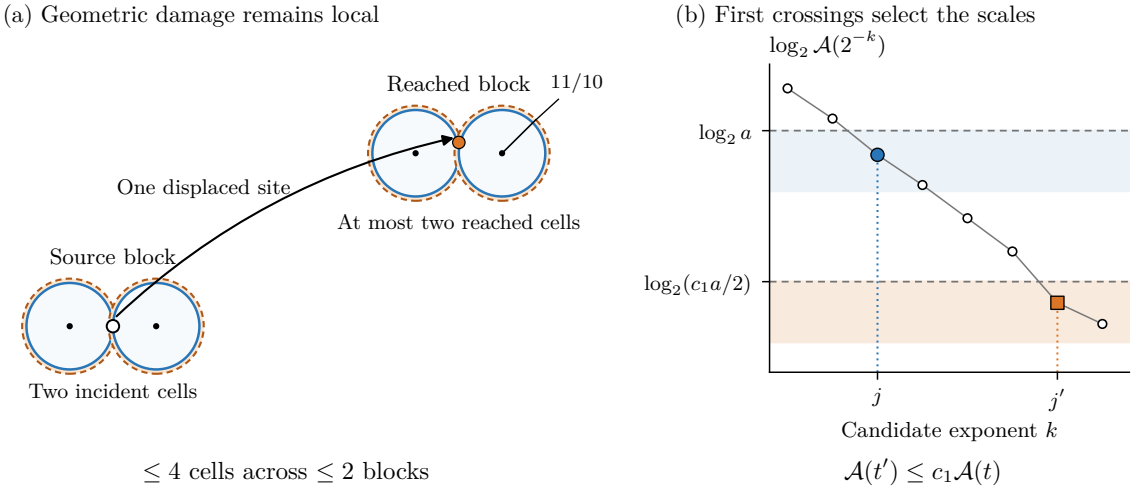

Figure 17: Two local controls used in the reduction. In (a), the hollow and filled points show a site's original and displaced positions. Blue circles have radius 1. Dashed orange circles have radius 11/10. The site can affect at most two incident cells at its source and at most two cells near its new position, hence at most four cells in two blocks. In (b), the first crossings of a and $c_1 a/2$ choose the cycle and bridge scales. The shaded strips each span a factor two, since halving the scale changes $\mathcal{A}$ by at most that factor. Thus $\mathcal{A}(t') \le c_1 \mathcal{A}(t)$. The scale plot is schematic.

The decoder uses only the returned faces and their labels. It checks every cell, including the bridges. A block is read when each cell has at least one triangle on three of its four labels and all such triangles in that cell select the same base diagonal. Only cycle signs determine the selected vertex set. The decoder discards blocks with nonconstant cycles or a selected set that is not independent in G, and returns a largest remaining set, with a fixed rule for ties. If none remains, it returns the empty set. This rule takes polynomial time and always returns an independent set. Every determined block in a nonempty event of Lemma 19 is read by this rule. In particular, a correct determined block gives a maximum independent set.

The following comparison applies both to one candidate list and to repeated lists. Write $\bar{p}(T) = p_\mu(T)$ and $P_* = \prod_b p_b^*$, the modal probability in the full linear model. Independence and the 2^M possible sign vectors give $P_* \ge 2^{-M}$.

Proposition 20. *Suppose the configuration contains r independent copies of the candidate list, each with a suitable block satisfying* (32)*, where $0 < \beta \le 1/2$. For every complete type T whose decoder fails to return a maximum independent set,*

$$\begin{aligned} \bar{p}(T) &\le e^{64nq_R}(1-\beta)^r P_* + 2L_g M\tau, \\ \max_{T'} \bar{p}(T') &\ge (1-4nq_R)P_* - 2L_g M\tau. \end{aligned} \tag{36}$$

Here M, n count all cells and sites in the repeated configuration.

Proof. Fix a complete type T for which decoding fails and a set I for which the event in (35) is nonempty. At least $r - 2|I|$ chosen suitable blocks are determined. The decoder reads every determined block, so each must be incorrect for decoding to fail. By (35), the probability of this event is at most

$$(16q_R)^{|I|}(1-\beta)^{r-2|I|}P_* \le (64q_R)^{|I|}(1-\beta)^r P_*.$$

The inequality uses $(1-\beta)^{-2} \le 4$. If $r - 2|I| < 0$, the factor $(1-\beta)^{r-2|I|}$ exceeds one and still gives an upper bound. Sum over all subsets I, use $(1+64q_R)^n \le e^{64nq_R}$, and add the probability

of $\mathcal{N}_\tau$. This proves the first inequality for each complete output separately, including outputs with a different hull.

For the lower bound, fix a full linear vector attaining P_*. For any site i, omit the at most two cell tests that use its displacement. The remaining event is independent of $\|Z_i\| > R$, and its probability is a sum of at most four full-vector probabilities, each at most P_*. Thus the chosen vector and a large site i have joint probability at most $4q_R P_*$. A union bound shows that the vector occurs with all sites nonlarge and outside $\mathcal{N}_\tau$ with probability at least $(1-4nq_R)P_* - 2L_g M\tau$. On this event all certificates and decision signs agree. The completion argument following (9) then gives one complete type, proving the lower bound. □

8.4 Choosing the precision and decoding a mode

For an exact mode, one copy of the candidate list suffices. Proposition 20 compares every failed complete output with a type obtained from a linear maximizer. We choose the loss from large displacements and the probability of sign errors to be small fractions of their relative separation:

$$\begin{aligned} s_R &= \min\{s \in \mathbb{Z}_{\geq 1} : C_L s^{-\gamma} \leq \beta/(256n)\}, \\ R &= 2^{s_R}, \qquad \tau = \frac{\beta\, 2^{-M}}{16 L_g M}. \end{aligned} \tag{37}$$

By (L), $q_R \leq \beta/(256n)$. The value s_R is found by integer power comparisons, as for $j_{\max}$, and is $O_\nu((n/\beta)^{1/\gamma})$. In particular, the exponent needed for R is polynomial. The radius itself is never enumerated.

All coordinates have polynomial bit length. To see this, let B_P bound the bit lengths in the completed base configuration. There are polynomially many certificates and normalized tests, each of degree at most four and involving at most four sites. Exact evaluation at the base gives r_*, and expansion in the displacements gives C_P, both in polynomial time. A nonzero value of such a determinant is a rational whose denominator has $O(B_P)$ bits, and its Taylor coefficients have $O(B_P)$ bits as well. Hence $\log(1/r_*) + \log C_P = O(B_P + \log n)$, and (33) gives

$$\log_2(1/\sigma) = O_\nu(B_P + \log n + s_R + \log_2(1/\tau)).$$

Here $\log_2(1/\tau) = M + O_\nu(\log M + \log(1/\beta))$, and all terms are polynomial in N. The grids for ζ_i, the fields in (31) and the homothety (34) preserve this bound. The constants $C_L, C_A, K_\nu, c_1, L_g, d_\nu$, and the rational exponent γ, are fixed in the reduction for the law ν. No probability or quantile is computed.

Proof of Theorem 3. Maximum independent set is NP-hard on 2-connected planar cubic graphs [28, Theorem 4.1(a)]. Apply the construction above to such a graph. Since $P_* \geq 2^{-M}$, the sign-error term in (36) is at most $\beta P_*/8$. For every failed complete output, the upper bound is therefore

$$\bar{p}(T) \leq (e^{\beta/4}(1-\beta) + \beta/8)P_* \leq (1 - 5\beta/8)P_*.$$

The last inequality follows from $e^{\beta/4} \leq (1-\beta/4)^{-1}$, which gives $e^{\beta/4}(1-\beta) \leq 1 - 3\beta/4$. The lower bound in (36) is $(1 - 9\beta/64)P_*$, strictly larger. Hence every global geometric mode, including each tied mode, decodes to a maximum independent set. The comparison uses complete labelled types and covers every possible hull. Construction and decoding take polynomial time, and the solver is called once. □

9 A barrier to relative approximation

The same construction also excludes relative approximation. A suitable block has the relative gap β from (32), whose reciprocal is polynomial in the graph size. Independent repetitions

multiply the probability losses of failed outputs. Proposition 20 retains this multiplication for complete geometric triangulations, because k large sites can leave at most $2k$ blocks undetermined.

Theorem 21. *Fix a bounded rotationally invariant planar density satisfying* (L). *For every fixed $c \in (0,1)$, it is NP-hard to return a complete labelled type T satisfying*

$$p_\mu(T) \geq c \max_{T'} p_\mu(T').$$

The reduction uses one approximate-solver call and takes polynomial time. The sites of the constructed instance have rational generic centres μ_i and positions $X_i = \mu_i + Z_i$, where the displacements Z_i are independent draws from the fixed law. Consequently, unless $P = NP$, no polynomial-time algorithm achieves any fixed factor $c \in (0,1)$. In particular there is no polynomial-time relative approximation scheme. A randomized polynomial algorithm meeting this guarantee with probability at least $2/3$ would imply $NP = RP$.

Section 9.2 proves the theorem. The final comparison also permits factors that decrease with the number of sites. We give the quantitative consequence after the proof.

9.1 The gap and the number of repetitions

The gap $\beta = d_\nu c_1 a^2/8$ of (32) compares each incorrect vector with the modal probability of its suitable block. In Lemma 16, dividing the absolute loss by the modal probability cancels the factor 2^{-m}. With $a = 2^{-K_\nu \ell}$, we have

$$B = O_\nu(m^{3K_\nu/\gamma}), \qquad \beta^{-1} = O_\nu(m^{2K_\nu}).$$

Both exponents are fixed for the law. The reduction repeats the entire list, so it can use this gap without finding a suitable block. When the decoder fails, no block that it reads is correct, and each copy contributes a relative loss unless a large site makes its suitable block undetermined.

For the Gaussian law there is a shorter route to Theorem 21. Repeat the single block of Proposition 17 $r = \lceil (\lceil \log_2(1/c) \rceil + 4)/\beta \rceil$ times, with $M' = rm$ cells and, after completion, $n' \leq 12M' + 4$ sites. In the linear model each failed output has probability at most $(1-\beta)^r P_* < cP_*/8$, where $P_* \geq 2^{-M'}$. As in the proof of Theorem 1, choose R, τ with $n' e^{-R^2/2} + 2L_g M' \tau \leq c2^{-M'}/8$. The probability of each complete type changes by at most $cP_*/8$. Thus every failed type has probability below $cP_*/4$, while a linear mode retains at least $(1 - c/8)P_*$. The precision is polynomial in M' and $\log(1/c)$, including for the decreasing factors below.

9.2 Replication and complete geometric outputs

We choose enough repetitions to make the linear loss smaller than the requested factor. We then bound the geometric correction by Proposition 20. The tail condition is used with an inverse-polynomial tail probability, while the exponentially smaller sign-error tolerance affects only a coordinate bit length.

Proof of Theorem 21. It suffices to treat rational c, since achieving a larger factor also achieves c. Build the candidate list of Section 8, put $h = \lceil \log_2(1/c) \rceil + 4$, and place

$$r = \lceil h/\beta \rceil$$

independent copies of the entire list in separate boxes (Fig. 18(a)). Complete their union once, with $M' = rBm$ cells and $n' \leq 12M' + 4$ sites. Every copy contains a suitable block. Use the local geometric construction of Section 8.2, now with

$$s_R = \min\{s \geq 1 : C_L s^{-\gamma} \leq 1/(64n')\}, \qquad R = 2^{s_R}, \qquad \tau = \frac{c\,2^{-M'}}{16 L_g M'}.$$

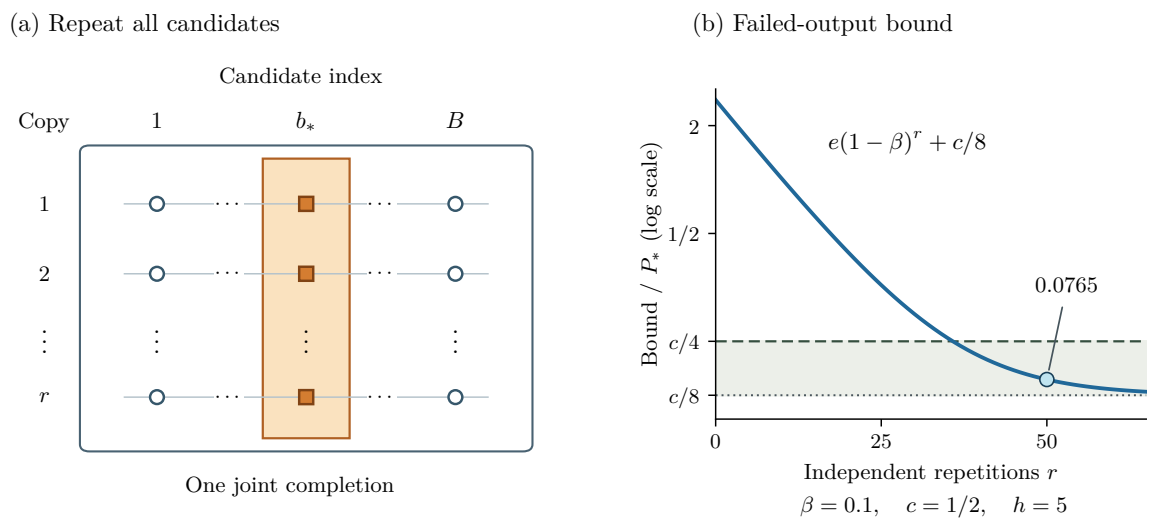

Figure 18: Independent repetition of the full candidate list. In (a), each row contains all B candidate blocks, and each column uses one parameter choice. The orange column marks a suitable choice b_*, whose identity the reduction does not need. The outer frame represents one joint geometric completion. In (b), the curve is an analytic upper bound on failed-output probabilities divided by P_*, on a logarithmic vertical axis. For $\beta = 0.1$ and $c = 1/2$, take $h = 5$ and $r = \lceil h/\beta \rceil = 50$. The bound is $e\,0.9^{50} + c/8 \approx 0.0765 < c/4 = 0.125$.

Thus $64n'q_R \le 1$ and $2L_gM'\tau \le cP_*/8$, where $P_* \ge 2^{-M'}$ is the maximum probability in the full linear model. Proposition 20 bounds the probability of each failed complete geometric output by

$$\bar{p}(T) \le (e(1-\beta)^r + c/8)P_* \le (e^{1-h} + c/8)P_* < \frac{c}{4}P_*. \tag{38}$$

Indeed, $h \ge \ln(1/c) + 4$, so $e^{1-h} \le ce^{-3} < c/8$. Fig. 18(b) plots the first bound in (38). The lower bound in the same proposition gives

$$\bar{p}_* = \max_{T'} \bar{p}(T') \ge (1 - 1/16 - c/8)P_* \ge \frac{13}{16}P_*.$$

It follows that every output with $\bar{p}(T) \ge c\bar{p}_*$ makes the decoder return a maximum independent set. This applies to every qualifying complete type, regardless of its hull or of ties.

The size and precision remain polynomial. The estimates of Section 9.1 give $n' = O_\nu(hm^{1+3K_\nu/\gamma+2K_\nu})$, and $m = O(N^3)$. The exponent $s_R = O_\nu((n')^{1/\gamma})$ and the bit length $\log_2(1/\tau) = M' + O_\nu(\log M' + \log(1/c))$ are polynomial as well. Section 8.4 then bounds the bits in the final centres. One call to the approximate solver would solve maximum independent set on planar cubic graphs [28, Theorem 4.1(a)]. In particular, a relative PTAS used at factor $1/2$ would imply $P = NP$.

For the randomized claim, take an instance (G, k_0) of the decision version of maximum independent set, run the solver once and reject malformed outputs. Apply the decoder to any well-formed output, verifying that every retained set is independent. Accept (G, k_0) exactly when the returned set has size at least k_0. An input with no independent set of size k_0 is never accepted. Every yes-instance is accepted whenever the output qualifies, with probability at least $2/3$. Thus the NP-complete decision problem is in $RP \subseteq NP$, giving $NP = RP$. □

Factors that decrease with the number of sites. The reduction uses polynomially many copies whenever the value of h is polynomial in the number N of graph vertices and the rational factor c has a polynomial-length encoding. Fix an integer $k \ge 1$, take $c = 2^{-N^k}$ and $h = N^k + 4$,

and use the same radius and tolerance. For a constant a_ν depending only on the law, the completed instance has $n' \le N^{k+a_\nu}$ sites for all sufficiently large N. Smaller graphs are solved directly. For example, any fixed integer $a_\nu > 3 + 9K_\nu/\gamma + 6K_\nu$ suffices. Given $\epsilon \in (0,1)$, choose $k \ge a_\nu(1-\epsilon)/\epsilon$. Then

$$(n')^{1-\epsilon} \le N^{(k+a_\nu)(1-\epsilon)} \le N^k, \qquad 2^{-(n')^{1-\epsilon}} \ge c.$$

Hence, under the hypotheses of Theorem 21 and for every fixed $\epsilon \in (0,1)$, it is NP-hard to return a type of probability at least $2^{-n^{1-\epsilon}}$ times the maximum, where n is the number of sites. The precision remains polynomial because the tail probability is bounded in terms of n', and the additional N^k bits enter only through τ.

9.3 Delaunay paths on the line

On the line, the Delaunay triangulation joins consecutive points into a labelled path. We obtain a constant relative guarantee by first finding a most probable left-to-right order. Each path has two orientations, which accounts for the factor $1/2$ in the comparison with the planar barrier.

For independent Gaussian sites $X_1, \dots, X_n$ of positive variance, write $p(\pi) = \mathbb{P}\{X_{\pi(1)} < \cdots < X_{\pi(n)}\}$ for each permutation π of their labels. Let $[\pi]$ denote its path and π^{rev} the reversed order. For $n \ge 2$, the path has probability $P([\pi]) = p(\pi) + p(\pi^{\mathrm{rev}})$. Set $P_* = \max_\pi P([\pi])$.

Proposition 22. *Let $X_i \sim N(\mu_i, \sigma^2)$, $1 \le i \le n$, be independent, with $n \ge 2$ and a common $\sigma > 0$. If π_0 orders the means nondecreasingly, then*

$$P([\pi_0]) \ge \frac{1}{2} P_*.$$

Equality $P([\pi_0]) = P_*$ holds when $n = 4$, and for every n when the multiset of means is symmetric about a point. The path $[\pi_0]$ can be constructed deterministically with $O(n \log n)$ comparisons, and in polynomial bit time for rational means.

Proof. Palmer, Hill and Scheding [29, Theorem 1 and Appendix I] prove that sorting the means gives an optimal oriented order. Rank-order probabilities were studied earlier by Savage [31]. We recall the short density comparison to derive the path guarantee. For the Gaussian variables considered here, if $\mu_i \le \mu_j$, the corresponding normal densities satisfy

$$\frac{f_i(x) f_j(y)}{f_j(x) f_i(y)} = \exp\!\left(\frac{(\mu_j - \mu_i)(y - x)}{\sigma^2} \right) \ge 1 \qquad (x < y).$$

Swapping an adjacent inversion and integrating over the ordered configurations therefore cannot decrease the order probability. Repeating the swaps gives $p(\pi) \le p(\pi_0)$ for every π, with ties between means allowed. For distinct means, every corrected inversion gives a strict increase, so finding an oriented mode determines their sorted order. Each path combines two orders, so

$$P_* \le 2p(\pi_0) \le 2P([\pi_0]).$$

Sorting the rational means and joining consecutive labels uses the stated number of comparisons and polynomially many bit operations. The algorithm evaluates no probabilities.

For the exact case, translate and rescale to unit variance and centred, sorted means $d_1 \le \cdots \le d_n$. Write $c = (d_{\pi(1)}, \dots, d_{\pi(n)})$ and let Z be the increasing order statistics of n independent standard normals. The Gaussian likelihood ratio on the ordered region gives $p(\pi) = e^{-\|d\|^2/2} \mathbb{E} e^{c \cdot Z} / n!$. Since Z and $-Z^{\mathrm{rev}}$ have the same law, adding the reverse order gives

$$P([\pi]) = \frac{2 e^{-\|d\|^2/2}}{n!} \, \mathbb{E} \cosh(c \cdot Z).$$

For symmetric means $d^{\mathrm{rev}} = -d$. The rearrangement inequality then gives $-d \cdot Z \le c \cdot Z \le d \cdot Z$. The evenness of cosh and its monotonicity on $[0, \infty)$ prove optimality.

For four sites the sorted path is optimal even for asymmetric means. For an ordered vector z, a single term $\cosh(c \cdot z)$ can exceed $\cosh(d \cdot z)$. We compare each term together with the term for the reverse order. By the symmetry of Z, averaging the two preserves the expectation in the path probability. Write

$$x = c \cdot z, \qquad y = c^{\mathrm{rev}} \cdot z, \qquad a = d \cdot z, \qquad b = -d^{\mathrm{rev}} \cdot z,$$

where the superscript denotes reversal. By the rearrangement inequality, every permutation c' of d satisfies $-b \le c' \cdot z \le a$. In particular $x, y \in [-b, a]$. The mean of $c' \cdot z$ over all permutations is zero, because $\sum_i d_i = 0$. Hence $a \ge 0$ and $b \ge 0$. Since the entries of c sum to zero, $c_1 + c_4 = \frac{1}{2}[(c_1 - c_2) + (c_4 - c_3)]$, and each difference of two entries of d has absolute value at most $d_4 - d_1$. Thus $|c_1 + c_4| \le d_4 - d_1$. Using $c_2 + c_3 = -(c_1 + c_4)$ and $2z_1 \le z_2 + z_3 \le 2z_4$, we get

$$|x + y| = |c_1 + c_4|\,|z_1 + z_4 - z_2 - z_3| \le (d_4 - d_1)(z_4 - z_1) \le a + b,$$

where the last step uses $a+b = (d_4-d_1)(z_4-z_1)+(d_3-d_2)(z_3-z_2)$. Together with $|x-y| \le a+b$ (since $x, y \in [-b, a]$), this gives $\max(|x|, |y|) \le \max(a, b)$ and $|x| + |y| = \max(|x+y|, |x-y|) \le a + b$. Order the two pairs so that $|x| \ge |y|$ and $a \ge b$. If $|y| \le b$, both entries are already bounded by their counterparts. Otherwise transfer the excess $|y| - b$ from the smaller entry to the larger. Convexity can only increase the sum of their hyperbolic cosines, and the larger entry becomes $|x| + |y| - b \le a$. In both cases $\cosh x + \cosh y \le \cosh a + \cosh b$. Finally, Z and $-Z^{\mathrm{rev}}$ have the same distribution, so taking expectations yields $P([\pi]) \le P([\mathrm{id}])$, as required. □

10 Elliptically contoured noise

The Delaunay triangulation is not invariant under linear maps, so a fixed anisotropic law cannot be reduced to an isotropic one by a change of coordinates. Both hardness theorems nevertheless extend to fixed invertible linear images of the rotation-invariant laws specified below. We keep the Euclidean construction and let the anisotropy change only the coefficients of the linear tests. Sections 6 and 7 use these coefficients only through the ranges (14). The same argument weakens the assumption on the density: only the projection of the rotation-invariant law needs a bounded density.

Let Y be a random vector in $\mathbb{R}^2$ with a rotation-invariant law satisfying (L), whose unit projection U has a bounded density g. Let M be a fixed real invertible 2×2 matrix, with singular values $0 < s_{\min} \le s_{\max}$ and $\varkappa = s_{\max}/s_{\min}$, and let ν be the law of MY. For a unit vector n,

$$\langle n, MY \rangle = \langle M^{\mathsf{T}} n, Y \rangle \overset{d}{=} s(n)\, U, \qquad s(n) = \|M^{\mathsf{T}} n\|, \tag{39}$$

and $s_{\min} \le s(n) \le s_{\max}$, $s(-n) = s(n)$, $|s(n) - s(n')| \le s_{\max}\|n - n'\|$. Fig. 19 illustrates this directional scaling while keeping the Euclidean cell fixed. Every Gaussian law $N_2(0, \Sigma)$ with a fixed positive definite Σ has this form. So do the laws of Theorem 3, by (11), and some laws with unbounded radial density, such as $\varphi(r) = \frac{3}{4\pi} r^{-1/2} \mathbf{1}_{\{r \le 1\}}$, whose projection density is bounded.

Theorem 23. *For every law ν of this form, the conclusions of Theorems 3 and 21 hold, including the factor $2^{-n^{1-\epsilon}}$ of Section 9. The matrix M and the law of Y are fixed. Only the rational generic centres form the input.*

Proof. We retain the Euclidean construction of Section 5. After checking general position and the symmetries of the special cells, we rescale the linear tests and show that the block estimates and the reduction still apply.

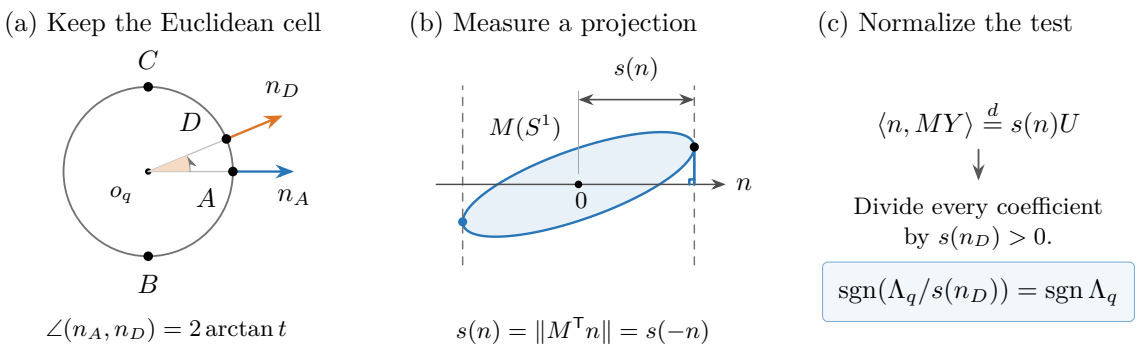

Figure 19: Directional scales of elliptic noise. In (a), the circumcircle of the decision cell is unchanged. The normals at A and D differ by $2\arctan t$. In (b), the ellipse is the image of the unit circle under M. Its projection onto n has half-width $s(n)$, marked by the dashed supporting lines. This scale is meaningful even for displacements with unbounded support. Opposite normals have the same scale. In (c), division by the positive scale at D preserves the sign of the decision test.

General position. Write $Y = R(\cos\Theta, \sin\Theta)$ with Θ uniform and independent of R. Bounded projection density precludes an atom of R at zero. An atom at $r > 0$ would put mass of order $\sqrt{h/r}$ in the projection interval $[r - h, r]$ as $h \to 0^+$, by the angular width of that arc. The density bound gives mass at most $h\|g\|_\infty$. Thus the radius is atomless. The inverse image of a line or circle under an invertible affine map is a line or ellipse. It meets a circle of radius R in finitely many points except at a possible single radius where the curves coincide. Conditioning on R gives probability zero in either case. Conditioning successively on independent sites now rules out coincident points, collinear triples and cocircular quadruples.

Homogeneity. Among the symmetries of the lattice construction, the law ν is invariant under the point reflection $z \mapsto -z$, and its translates are permuted by translations. By the choices of Section 5 and Appendix A.1, all bridges are translates of each other, all lower endpoints of bridges are translates of each other, each upper endpoint is the image of a lower one under a point reflection, and all designated cells are translates of each other.

Tests. The first variation (6) is a geometric identity, valid for arbitrary displacement vectors. A site with unit normal n in a cell and noise MY_i contributes $\langle n, MY_i\rangle = s(n)\langle M^\mathsf{T} n/s(n), Y_i\rangle$, and the second factor is a unit projection of Y_i. Distinct sites give independent projections. A site shared by two cells has opposite normals in them, hence the same scale and opposite scalars, as in Section 5. For a site $\mathbf{P}$ of a cell, write $s(\mathbf{P})$ for (39) at its unit normal in that cell. Dividing the test of each cell by $s(\mathbf{D}_q) > 0$ preserves its sign and gives tests of the form (7) and (8) in independent copies of U, with coefficients

$$\tilde{a}_q = a_q \frac{s(\mathbf{A}_q)}{s(\mathbf{D}_q)}, \qquad \tilde{b}_q = b_q \frac{s(\mathbf{B}_q)}{s(\mathbf{D}_q)}, \qquad \tilde{c}_q = c_q \frac{s(\mathbf{C}_q)}{s(\mathbf{D}_q)},$$

similarly for bridges, and fields $\theta_q / s(\mathbf{D}_q)$. Since $\mathbf{D}_q$ is at angle $2\arctan t \le 2t$ from $\mathbf{A}_q$, we have $|s(\mathbf{D}_q)/s(\mathbf{A}_q) - 1| \le 2\varkappa t \le \frac{1}{4}$ when $t \le 1/(8\varkappa)$. With (14) for $C_0 = 2$, this gives $|\tilde{a}_q - 1| \le 6\varkappa t$, $t/(3\varkappa) \le \tilde{b}_q \le 3\varkappa t$ and $\tilde{c}_q \le 3\varkappa t$, with $\tilde{c}_q \ge t/(2\varkappa)$ at the endpoints of bridges, and the same bounds with t' for bridges. Thus the normalized tests satisfy (14) with $C_0 = 6\varkappa$. Translations preserve all normals and a point reflection reverses them, which preserves s. By the homogeneity above, the normalized coefficients are therefore the same for all bridges, for their endpoints, and for all designated cells. In particular $s_\mathrm{d} = s(\mathbf{D}_d)$ is the same for every designated cell d.

Block. Lemmas 14, 15 and 16 apply to the normalized tests for $t \le \min(t_1, 1/(8\varkappa))$. Their constants, in particular t_1, c_1 and d_ν, are those given by $C_0 = 6\varkappa$. They depend

on the law of Y and on $\varkappa$. The field $-\lambda$ of the construction acts as $-\lambda/s_{\mathrm{d}}$, so (24) reads $\frac{7}{5}W \le \mathcal{M}(\lambda/s_{\mathrm{d}}) \le \frac{13}{5}W$.

Reduction. In Section 8, the quantities $\mathcal{A}$, $\widehat{w}$ and a_0 are those of U, and (13) holds because $|U| \le \|Y\|$. A site is large when its noise MY_i has norm greater than R. Since $\|MY\| \le s_{\max}\|Y\|$, the law ν satisfies (L) with the same γ and a larger C_L. The normalized tests give the block gap. The geometric comparison with Γ_q/σ uses the original linear test and its tolerance τ. The projection of the noise at $\mathbf{D}_q$ on its normal has density at most $\|g\|_\infty/s_{\min}$, which replaces $\|g\|_\infty$ in the choice of L_g. In Section 8.1, the grid argument applies to $\mathcal{M}(\lambda_k/s_{\mathrm{d}})$ once the threshold on t is lowered to $2\|g\|_\infty t < c_w a_0^2 s_{\min}$. The first value then satisfies $\mathcal{M}(\lambda_0/s_{\mathrm{d}}) \le 2\|g\|_\infty t^2 t'/s_{\min} < W$, and, for large K_ν, the last exceeds $13W/5$ because $\mathcal{M}(\lambda/s_{\mathrm{d}}) \ge \mathcal{M}(\min(y_1, 1/(4s_{\max})))$ for $\lambda \ge 1/4$. All other steps of Sections 8 and 9 use the block only through its probabilities $p_b(s)$, the gap β, the independence of distinct sites and L_g. They apply without change. For listed blocks outside the suitable range, the reduction uses only the bound $p_b(s) \le p_b^*$. □

The projection assumption even admits singular planar laws. For example, choose $R = 1+X$, where X has the uniform middle-thirds Cantor distribution, independently of the angle. These radii form a set of length zero, so their union of circles has area zero by polar integration. Yet the projected density is bounded. Indeed, for $\alpha = \log 2/\log 3 > 1/2$, an interval of length $h \le 1$ has radial mass $O(h^\alpha)$: at the level where $3^{-j-1} < h \le 3^{-j}$, it meets at most two level-j intervals of mass 2^{-j}. At abscissa u, the conditional density for $r > |u|$ is $1/(\pi\sqrt{(r-|u|)(r+|u|)})$. As $r \ge 1$, the contribution from $2^{-k-1} < r - |u| \le 2^{-k}$ is $O(2^{-k(\alpha-1/2)})$, uniformly in u. These terms are summable for $k \ge 0$, and the kernel is bounded on $r - |u| > 1$. Bounded support gives (L).

The reduction still contains finitely many rational constants, now depending on M and on the law of Y, and it evaluates no density, probability or quantile. Laws that are symmetric under the point reflection but not elliptically contoured have projections whose laws depend on the direction in a more general way. We do not treat them.

The single-block simplification of Proposition 17 also holds when $\mathbb{E}\|Y\|^2 < \infty$. Write e_1, e_2 for the horizontal and vertical unit vectors. Choose a fixed rational κ_0 in

$$\frac{14}{5}J\mathbb{E}U^2\frac{s(e_2)^2}{s(e_1)} < \kappa_0 < \frac{26}{5}J\mathbb{E}U^2\frac{s(e_2)^2}{s(e_1)},$$

and set $t = t' = 2^{-K\ell}$, $\lambda = \kappa_0 t^2$. In the construction of Section 5 and Appendix A.1, the two bridge contacts lie on a vertical diameter. The normals at $\mathbf{A}$ in bridges, endpoint cells and designated cells are horizontal. The dominated-convergence calculation of Proposition 17 therefore gives

$$\frac{w_B}{t^2}, \frac{w_C}{t^2} \longrightarrow 4J^2\mathbb{E}U^2\left(\frac{s(e_2)}{s(e_1)}\right)^2, \qquad \frac{\mathcal{M}(\lambda/s_{\mathrm{d}})}{t^2} \longrightarrow \frac{2J\kappa_0}{s(e_1)}.$$

These limits put the field in (24). The bounds $\widehat{w}(t,t) \le t^2\mathbb{E}U^2$ and $\mathcal{A}(t) \le t\mathbb{E}|U|$ give the other two conditions exactly as before. A sufficiently large fixed integer K thus gives one block with the same inverse-polynomial relative gap. The constants depend on the fixed law and M. Their choice requires no computation of M from the input.

11 Examples and the boundary of the argument

The tail assumption admits bounded noise, Gaussian and Cauchy noise, and laws with much larger displacements. We compare these examples through the number of bits needed to make

a large displacement have probability at most $1/(64n)$, the truncation level used in the product comparison. Fig. 20 compares the required radii.

Every finite positive moment of $\|Z\|$ implies every positive moment of $\log(1+\|Z\|)$, so Proposition 4 gives (L). The standard planar Gaussian has $T_\nu(R) = e^{-R^2/2}$, and a uniform law on a disk has compact support. The isotropic planar Cauchy law

$$\rho(z) = \frac{1}{2\pi(1+\|z\|^2)^{3/2}}, \qquad T_\nu(R) = (1+R^2)^{-1/2},$$

has finite moments of order $0 < a < 1$. Its finite moments of order below one imply (L), so both hardness theorems apply. Scaling any of these fixed laws preserves their hypotheses.

More slowly decaying tails also qualify. For $0 < \beta < 1$, the bounded radial density

$$\rho_\beta(z) = \begin{cases} (1-e^{-1})/\pi, & \|z\| \le 1, \\ \dfrac{\beta}{2\pi\|z\|^2}(\ln\|z\|)^{\beta-1}e^{-(\ln\|z\|)^\beta}, & \|z\| \ge e, \\ 0, & 1 < \|z\| < e \end{cases}$$

integrates to one and has $T_\nu(R) = e^{-(\ln R)^\beta}$ for $R \ge e$. It satisfies (L) for every $\gamma > 0$. Every positive moment of $\|Z\|$ diverges, since

$$\mathbb{E}\|Z\|^a \ge a\int_e^\infty R^{a-1}T_\nu(R)\,dR = a\int_1^\infty e^{au-u^\beta}\,du = \infty \qquad (a > 0).$$

The theorems thus cover laws for which no positive power of the displacement is integrable.

The bounded radial family

$$\rho_\eta(z) = \frac{\eta\,\mathbf{1}_{\{\|z\|\ge e\}}}{2\pi\|z\|^2(\ln\|z\|)^{\eta+1}}, \qquad T_\nu(R) = (\ln R)^{-\eta} \quad (R \ge e) \tag{40}$$

shows the further reach of the product comparison. For every $\eta > 0$, choose a rational $0 < \gamma \le \eta$ and enlarge C_L. Then (L) holds. A tail bound of $1/(64n)$ needs a radius $R = \exp((64n)^{1/\eta})$, whose binary length is polynomial in n. Requiring tail 2^{-s} instead needs $\ln R \ge 2^{s/\eta}$, so its bit length is exponential in s. Keeping the probabilities of unaffected blocks is what permits the weaker truncation level in our reduction.

Condition (L) is exactly what allows all displacements to remain within a radius with polynomially many bits with a fixed positive probability (Proposition 4). Our transfer uses this confinement event to bound the modal probability from below. The complexity of modal Delaunay triangulation outside (L) remains open.

For a fixed uniform law on a disk, Theorem 21 gives worst-case hardness as the centres vary. A fixed generic configuration also has the stable regime described in Section 1.1: moving each centre by less than its structural tolerance preserves its Delaunay type. Multiplying all centres by a common positive factor multiplies this tolerance by that factor. Once the tolerance exceeds a bound on the noise norm, every realization has the type of the centres, with probability one.

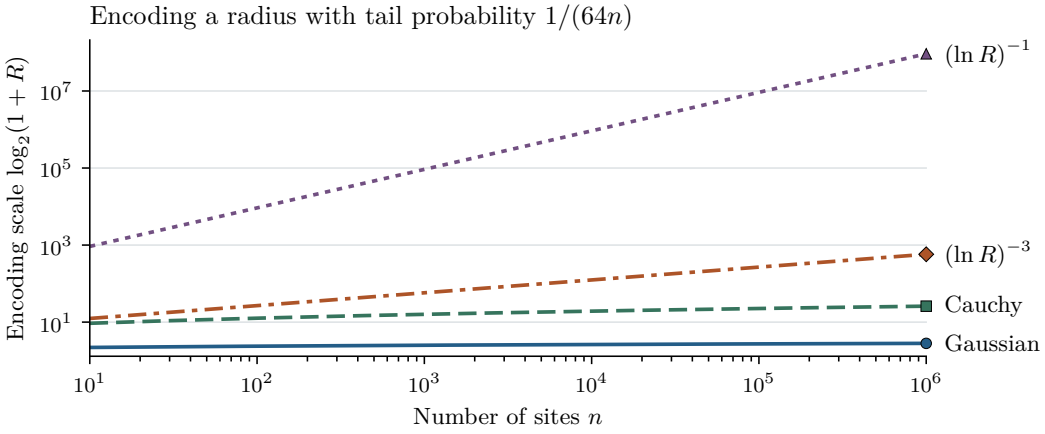

Figure 20: The radius needed for tail probability $1/(64n)$. The ordinate $\log_2(1+R)$ measures its encoding length, and both axes are logarithmic. The two curves labelled by their tails $(\ln R)^{-\eta}$ and the Gaussian and Cauchy examples all satisfy (L). The bridge construction covers all four laws, including the tail with exponent $\eta = 1$.

12 The most probable convex hull

For the hull, each realization records which sites are extreme. One can also ask for their cyclic order around the boundary, up to rotation and reversal. Both descriptions discard the interior diagonals that encode a graph in the Delaunay reduction. Here the choice will be visible on the boundary itself.

Theorem 24. *For independent planar Gaussian sites with rational centres and covariance matrices $s_i^2 I_2$, where each $s_i > 0$ is rational, finding a most probable set of extreme labels is NP-hard. The same holds for the cyclic order of the extreme labels. Each search problem admits a polynomial-time reduction from* Subset Sum *with one solver call, with any tie breaking.*

Proof. Apart from the standard normal distribution function Φ, the notation of this proof and of Appendix B is independent of the earlier sections. Take positive integers $a_1, \dots, a_m$ and a target $1 \le T \le A = \sum_i a_i$, given in binary [38, problem 18 and its reduction from Exact Cover]. We assign one item to each of m consecutive edges of a convex polygon K. A candidate site near the midpoint of its edge can land outside the supporting line (Fig. 21(a)). In this case, its label joins the hull and its two neighbours enclose a triangular cap with the old edge. A further independent site P, called the probe, stays off the boundary whenever it falls in the enlarged polygon. Thus opening a cap incurs a probability cost and gives the probe more room. We balance these effects so that the most probable hulls select items of total weight T.

Let $B = \mathbb{P}(P \in K)$, with $B \ge 7/8$, and choose a small $\lambda > 0$ with $\lambda A \le 1/8$. The probe mass of a cap is the probability that the probe falls in it, for a given position of its candidate. The ideal probability p_i of opening cap i and the mean of its probe mass, conditional on opening, are chosen to satisfy

$$\frac{p_i}{1-p_i} = e^{-c\lambda a_i}, \qquad c = \frac{1}{B+\lambda T}, \qquad \text{mean probe mass of cap } i = \lambda a_i.$$

Call a hull *canonical* for S when its extreme labels are the polygon vertices and precisely the candidates indexed by S, in boundary order when order is requested. For $x = \lambda \sum_{i \in S} a_i$, independence gives opening probability $\prod_i (1-p_i) e^{-cx}$. For fixed candidate positions, add the

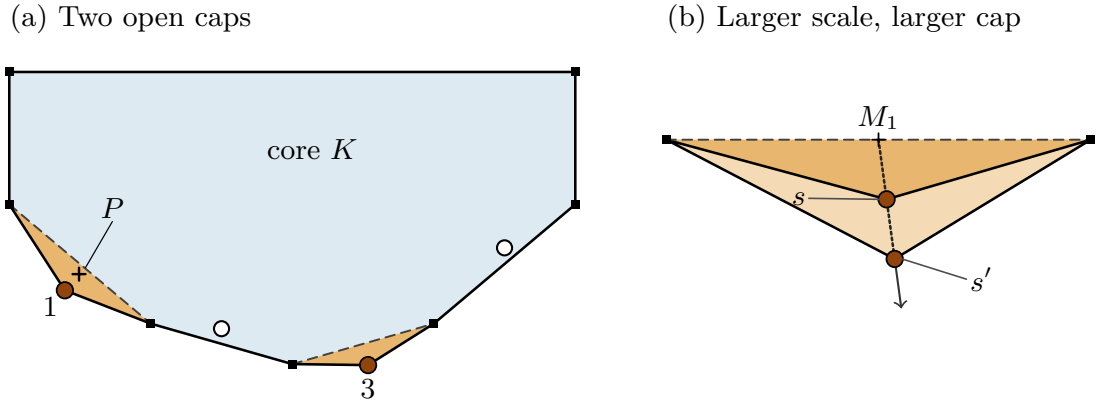

Figure 21: Encoding items on the boundary. In (a), filled candidate circles open disjoint caps on edges 1 and 3. The open circles lie inside the core and the probe P is inside a cap. In (b), increasing the scale from s to s' moves the apex farther along the same ray from the edge midpoint M_1, so the caps are nested. The dashed line is the original edge. Distances perpendicular to the edges are exaggerated.

probe masses of the open caps to B. On the event of separated caps in Appendix B, this sum is the probability that the probe lies in the enlarged polygon. Conditional on the opening pattern alone, each open cap has mean mass λa_i in the ideal model, so the sum has mean $B + x$. Multiplying by the opening probability gives the ideal score

$$\left(\prod_i (1 - p_i)\right) F(x), \qquad F(x) = (B + x)e^{-cx}, \qquad F'(x) = ce^{-cx}(\lambda T - x). \tag{H}$$

The score peaks at the target (Fig. 22). Throughout the range of x, $ce^{-cx} \geq 1/4$. Integrating $|F'|$ between x and the peak gives a loss of at least $(x - \lambda T)^2/8$. Since the weights are integers and $p_i \leq 1/2$, a target subset in a yes instance beats every non-target subset by at least $G = 2^{-m}\lambda^2/8$ in ideal probability. Its ideal probability is at least $2^{-m}B$, since $F(\lambda T) \geq F(0) = B$.

To realize these probabilities, first keep the polygon vertices fixed. Use lower vertices $v_j = 20(j, j(j - m))$, $0 \leq j \leq m$, and top vertices $(0, 20m^2), (20m, 20m^2)$. All four corners of the rectangle $[0, 20m] \times [0, 20m^2]$ are vertices of K. By convexity, K contains the rectangle and hence the disk of radius $10m$ about $o = (10m, 10m^2)$. The probe $P \sim N_2(o, I_2)$ therefore satisfies $1 - B \leq e^{-50m^2} < 2^{-m-5}$. At the midpoint M_i of each lower edge, with outward unit normal u_i, place an independent candidate

$$X_i = M_i + s_i(d_i u_i + Z_i), \qquad Z_i \sim N_2(0, I_2).$$

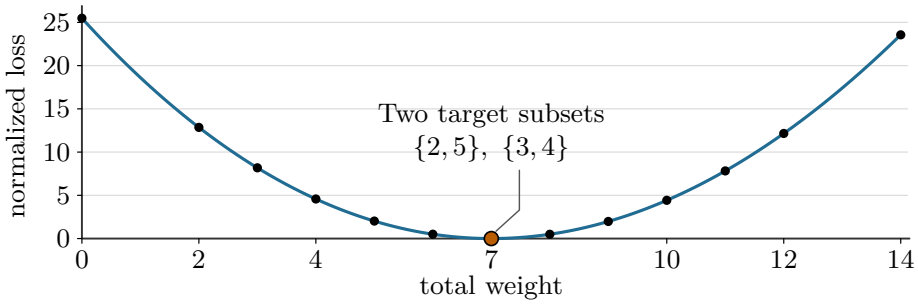

Figure 22: Why the score selects a sum. For weights $2, 3, 4, 5$ and target 7, the dots show the loss from the ideal maximum at each attainable total, in units of $F(\lambda T)(\lambda/(B + \lambda T))^2$. The choices $\{2, 5\}$ and $\{3, 4\}$ attain the same maximum. The curve evaluates (H) with $B = 7/8$ and $\lambda = 1/128$ for illustration.

Its probability of crossing the edge is $\Phi(d_i)$, independent of its scale s_i. After setting d_i we can therefore vary s_i to tune the mass of the cap. For a fixed outward draw, the apex travels on a ray from M_i, so the cap grows by inclusion (Fig. 21(b)).

Appendix B gives polynomial-bit rational input parameters for which every canonical probability differs from (H) by less than $G/4$, and the total probability of all other hulls is less than 2^{-m-4}. The construction also replaces every fixed polygon vertex by a Gaussian site of positive variance. These bounds hold for both hull descriptions. In a yes instance, a canonical hull for a target subset has probability at least $2^{-m}B - G/4$. This probability is more than the total probability of all noncanonical hulls, and it is at least $G/2$ more than the probability of any canonical hull for a non-target subset. Hence *every* mode is at a target subset, even when several subsets attain T. The decoder reads the candidate labels and checks their sum. In a no instance no subset passes that test. This proves the one-call reduction. □

The probe, the candidates and the polygon vertices use different Gaussian scales. The complexity of exact labelled hull search for independent planar Gaussian sites with one common isotropic variance remains open.

13 Conclusion

The Delaunay triangulation of the centres can be less probable than factorially many other types. Finding a most probable type is NP-hard already for standard Gaussian noise, as is finding one within any fixed factor of the maximum. For Gaussian input, a sampling algorithm finds with high probability a type whose probability is within a chosen additive error of the maximum. A supplied lower bound on the modal probability permits relative approximation, and an additional one on the relative gap of a unique mode permits exact recovery. The running times are polynomial also in the reciprocals of the error and of these bounds.

The Delaunay reductions encode graph vertices by cycles of local diagonal choices. Agreement along each cycle and opposite phases at bridges make the positive cycles describe a maximum independent set. The bridge interactions act at a smaller probability scale than agreement along a cycle, including for heavy tails. A comparison that retains the factors of unaffected blocks transfers this mechanism to complete triangulations under (L), including logarithmic tails of every positive order and laws with no positive moment of the noise magnitude. It also preserves the strong relative inapproximability obtained by replication. Section 10 extends both results to every fixed invertible linear image of a rotation-invariant law with bounded projection density, under (L), including Gaussians with a fixed positive definite covariance. Further questions concern heavier tails outside (L), anisotropic laws that are not elliptically contoured, and geometric conditions guaranteeing a large modal probability or separation. The latter would make the sampling guarantees effective on recognizable classes of uncertain point sets.

Determining the exact complexity of Delaunay paths on the line is a natural next question. For independent Gaussians with common variance, our 1/2-approximation uses an optimal oriented order. Exact path optimization involves the sum of the probabilities of two opposite orders. Sorting is exact for four sites and for any number of sites with a symmetric multiset of means. With unequal variances, the density comparison underlying the sorting argument also fails in general. For planar convex hulls, Theorem 24 uses individual variances to tune the contribution of each boundary label. Whether the most probable labelled hull is NP-hard to find when all sites have the same isotropic Gaussian variance remains open.

A Auxiliary geometric constructions

The construction in Section 5 uses explicit coordinates, a normalized incircle derivative and generic rational completion points. We give these details here to verify the geometric clearances and polynomial bit lengths.

A.1 Coordinates for the cycles and bridges

A thickened bar-visibility representation supplies each cycle. A one-point gap between the arms of an edge supplies its bridge while keeping every unintended pair of circles apart.

Proof of Lemma 11. We turn a visibility drawing into cycles by thickening the bars into rectangles and running arms along the edge corridors. The facing corners of two arms leave room for a single bridge. Since G is cubic, $N \geq 4$. Add edges to obtain a maximal planar graph, which is 3-connected. The ε-visibility theorem of Tamassia and Tollis [34, Section 4, Theorem 4], which requires 2-connectivity, gives horizontal bars at distinct heights and an open vertical band of lines of sight for every edge. Keep the bands for the edges of G, and replace endpoint abscissas and heights by their integer ranks. This coordinatewise monotone change preserves the bands and makes their endpoints integers.

Number the edges $e = 1, \ldots, E$. Choose an integer j_e with $(j_e, j_e + 1)$ in the band of e, multiply bar coordinates by $100(E + 1)$, and put $x_e = 100(E + 1)j_e + 100e$. Distinct bar heights, endpoints and chosen abscissas are at least 100 apart. Every bar strictly between the endpoints of e misses its open band and is at horizontal distance at least 100 from x_e.

For $e = uv$ with heights $h_u < h_v$, set $y_e = (h_u + h_v)/2$, an integer. Thicken each bar by ten in the ℓ_∞ metric. Add to the polygon for u the arm $[x_e - 20, x_e] \times [h_u + 10, y_e]$, and to that for v the arm $[x_e, x_e + 20] \times [y_e + 2, h_v - 10]$. The two facing corners are $q_e = (x_e, y_e)$ and $r_e = (x_e, y_e + 2)$. Rectangles have height twenty, arms width twenty and length at least thirty-eight. The separations of 100 keep all other pieces of different vertices at least ten apart and keep nonincident sides of each polygon at least twenty apart.

Each rectangle with its attached arms is an orthogonal polygon homeomorphic to a disk. Two boundary lattice points are at distance one exactly when they are consecutive along the boundary: nonincident sides are separated, and the nearest points on two incident sides other than their common corner are at distance $\sqrt{2}$. The boundary points therefore form an induced cycle. The only boundary points at distance one from $\beta_e = (x_e, y_e + 1)$ are q_e, r_e. Its other two lattice neighbours have ordinate $y_e + 1$, absent from the two arms, and the other pieces are separated by at least ten. Different β_e's are separated by their abscissas. When every boundary is traversed counterclockwise, the predecessor of q_e is the lattice point below it and the predecessor of r_e is the one above it, while β_e lies above q_e and below r_e. At both corners the bridge contact is therefore opposite to the predecessor contact. The successor of q_e is the lattice point on its left and the successor of r_e is the one on its right. Hence the cells at the corners q_e, for the different edges, are translates of each other, and the cell at r_e is the image of the cell at q_e under a point reflection. The two lattice neighbours of β_e with ordinate $y_e + 1$ carry no circle, so the west point of every bridge is unused and all bridges are translates of each other.

Double all coordinates, place a unit circle at each boundary point and add the unit circle centred at $2\beta_e$ for each edge. The preceding distances give exactly the tangencies in the statement, with the two bridge contacts opposite. Each cycle has at least eighty cells and only three bridge incidences. If the thickened bar of a vertex is $[\ell - 10, r + 10] \times [h - 10, h + 10]$, its designated cell is the one at the lattice point $(\ell, h + 10)$ of the upper side. Every arm is at horizontal distance at least eighty from this point, so its two cycle neighbours lie on that side and its contacts with them are opposite. The counterclockwise orientation runs westwards

along the upper side, so every designated cell has its successor contact at its west point, its predecessor contact at its east point and $\mathbf{C}$ at its north point. The designated cells of all cycles, with their four sites, are therefore translates of each other. Every pair of distinct centres has distance at least two. Their open disks are disjoint. Three such centres cannot have a common unit-distance point: each pair would then have distance exactly two, and three points of a unit circle cannot be pairwise at distance two. Thus no site belongs to three circles. The $O(N)$ pieces lie in a box of side $O(N^2)$, and their boundaries contain $O(N^3)$ lattice points. Including the $E = 3N/2$ bridges, $m = O(N^3)$, and the construction takes polynomial time. □

A.2 Normalization of the incircle derivative

The normalization of (6) follows directly in the coordinates of Section 5, with origin at o_q. Displace each site from its base position $\mathbf{n}_i$ to $\mathbf{n}_i + h\mathbf{z}_i$. Subtracting the last column of the incircle determinant from its third leaves the entries $\|\mathbf{n}_i + h\mathbf{z}_i\|^2 - 1$, all zero at $h = 0$. The cofactors of the other columns therefore vanish there, and the derivative of the third entry at $h = 0$ is the scalar product $2\langle \mathbf{n}_i, \mathbf{z}_i \rangle$. Thus, after multiplication by ι_q, the determinant has first variation $2\sum_i K_i \langle \mathbf{n}_i, \mathbf{z}_i \rangle$, where K_i is the cofactor in the third column multiplied by ι_q.

To identify these cofactors, replace the third column by any other column. The resulting determinant is zero, so $\sum_i K_i \mathbf{n}_i = 0$ and $\sum_i K_i = 0$. The affine dependence of four distinct concyclic points is unique up to scale. Equation (5) therefore makes K proportional to $(a_q, -b_q, c_q, -1)$. The fourth cofactor is $K_4 = -\iota_q \operatorname{orient}(\mathbf{A}_q, \mathbf{B}_q, \mathbf{C}_q) = -2$, since three distinct cardinal points of the unit circle have absolute orientation determinant two. Hence $K = 2(a_q, -b_q, c_q, -1)$. Dividing the first variation by four gives (6), with the same signs and scale.

A.3 Rational insertion points for completion

For completion, we must choose rational points inside open regions, with polynomial bit length and no new collinearities or cocircularities involving earlier points. A horizontal sweep gives an open disk with rational centre and radius in $\operatorname{int} F$ outside the closed disks for each unwanted face F of Lemma 12. The critical ordinates of this open set, at vertices, circle extrema and pairwise intersections of sides and circles, are algebraic of bounded degree and polynomial bit length. Root separation lets us choose a horizontal line at a rational height between two consecutive critical ordinates, meeting the open set. On this line we choose a rational point in the set. The strict defining inequalities at that point give an open disk with dyadic centre and radius of polynomial bit length.

Grid count. If S has at most $\bar{n}$ points and $k = \bar{n}^3 + 1$, every k-by-k square grid has a point outside S, off the lines through two and the circles through three noncollinear points of S. Indeed, a line contains at most k grid points and a circle at most $2k$, so at most $\bar{n} + k\binom{\bar{n}}{2} + 2k\binom{\bar{n}}{3} \le k\bar{n}^3 < k^2$ grid points are excluded. With $\bar{n} = 12M + 4$, bounding the number of points throughout Lemma 12, grids fixed in the guard boxes and disks supply the guards and inserted points. Section 8 uses the same count with $\bar{n} = n$.

B Rational Gaussian parameters for the hull reduction

We bound the error in (H) uniformly over all subsets. There are $N = 2m + 4$ sites. The following choices leave room both to calibrate the cap masses and to give the polygon vertices positive variance:

$$U = 100(m+1)^2, \quad k = 16\big(U^2 + m + \lceil \log_2(A+1) \rceil + 32\big), \quad \lambda = 2^{-k}, \quad \beta = \frac{2^{-m}\lambda^2}{2^{16}N^2}.$$

Here U bounds the diameter of K and the distance to o of every point within one unit of K. All four parameters have polynomial binary length.

Keeping the caps separate. Put $\rho = 1/(100U)$. Among the polygon vertices and lower edge midpoints, a collinear triple must lie on a single edge by convexity. Every other orientation determinant is a nonzero integer. Moving the midpoints by at most ρ changes any determinant by at most $4U\rho + 4\rho^2 < 1/2$. In the event that every candidate lies within ρ of its midpoint, all inward candidates lie in K. In boundary order, the vertices and outward candidates have strictly positive orientations for every ordered triple in counterclockwise order, including the triple on each opened edge. Their hull is therefore K together with the triangles of the active caps, whose interiors are disjoint. For $|d_i| \le 1/4$ and $s_i \le s_{\max} = \rho/(16k)$, the probability of any candidate moving farther than ρ from its midpoint is at most $m \exp(-(16k-1)^2/2) < \beta/N$.

Tuning one cap. Choose rational $d_i \in [-1/4, 1/4]$ so that $|\Phi(d_i) - p_i| \le \beta$. The bounds $c \le 8/7$ and $\lambda A \le 1/8$ put the odds $e^{-c\lambda a_i}$ in $[e^{-1/7}, 1]$, so $\Phi(-1/4) < p_i \le 1/2 = \Phi(0)$. Denote by $w_i(s)$ the expected probe mass of the triangle conditional on opening. For a fixed outward draw of normal height $h = d_i + Z_i \cdot u_i > 0$, increasing s by Δs adds area $\ell_i h\, \Delta s/2$, where ℓ_i is the edge length. The probe density is at most $1/(2\pi)$ and $\mathbb{E}[h \mid h > 0] < 2$. Thus $w_i(0) = 0$ by continuity and w_i is nondecreasing and U-Lipschitz. With probability greater than $1/100$, the standard normal draw has normal component in $[1, 2]$ and tangential component in $[-1, 1]$. On this event the height is at least $s/2$ and the probe density throughout the triangle exceeds 2^{-U^2-4}. Since $\ell_i \ge 1$,

$$w_i(s_{\max}) \ge s_{\max} 2^{-U^2-14} > \lambda A.$$

Bisection gives rational s_i with $|w_i(s_i) - \lambda a_i| \le \beta$. Use function values accurate to $\beta/8$, stopping if the error interval contains the target or the scale interval has width at most $\beta/(2U)$. This takes polynomially many steps, and $s_i \ge \lambda/(2U)$ since $U s_i \ge w_i(s_i) \ge \lambda - \beta \ge \lambda/2$.

Evaluating the cap mass. Use the edge tangent e_i, the standard normal densities φ, φ_2, and coordinates h, v normal and tangential to the edge. Conditional on the candidate coordinates, integrate the probe density along segments from base to apex. This gives

$$\Phi(d_i) w_i(s) = \ell_i s \int h(1-t)\varphi(h - d_i)\varphi(v) \cdot \varphi_2\big(M_i - o + (1-t)\ell_i y e_i + ts(hu_i + ve_i)\big)\, dh\, dv\, dt\, dy.$$

Here $h \ge 0$, $v \in \mathbb{R}$, $t \in [0, 1]$ and $y \in [-1/2, 1/2]$. The base coordinate is y, the fraction of the way to the apex is t, and $\ell_i s h(1-t)$ is the area Jacobian. Division by $\Phi(d_i) > 1/3$ costs a constant number of precision bits.

To evaluate $w_i(s)$ within 2^{-q} in time polynomial in U and q, truncate $h, |v|$ at $10(q+10)$. The discarded integral is bounded by the Gaussian tail since every triangle has probe mass at most one. On rescaling to the unit box, the integrand is a polynomial times e^{-Q}, where Q has fixed degree and $0 \le Q \le Q_{\max} = O((U+q+10)^2)$. The polynomial prefactor has polynomial magnitude. The Taylor remainder through degree D is at most $Q_{\max}^{D+1}/(D+1)!$, so $D = O(Q_{\max} + q + \log(U+q+10))$ suffices. Four variables and fixed degree of Q give only polynomially many monomials, each integrable exactly over the box. To see the bit bound, round the Gaussian normalizations and the square roots in the unit vectors to dyadic intervals. The sum of absolute coefficients of each power through D is at most $\mathrm{poly}(U+q+10)^D$, so $O(D \log(U+q+10) + q)$ bits suffice in the coefficients. Rational polynomial integration then uses polynomially many bits. Triangulating K evaluates B by the same Taylor bound in two variables. The one-dimensional integral for Φ and the exponential in these odds have bounded arguments. These evaluations and bisection require $q = O(\log(1/\beta) + \log N)$.

From cap masses to hull probabilities. With fixed vertices, multiply the indicator of the chosen opening pattern by the sum of the probe indicators for K and for each chosen triangle.

Its expectation is the opening probability times $B + \sum_{i \in S} w_i(s_i)$, by independence. When the candidates stay within ρ, a hull with the canonical set of labels has only one cyclic order, obtained by inserting its active candidates along the edges of K. The random variable thus agrees with the canonical hull indicator for either output on this event. On the exceptional event its magnitude is at most $m + 1$. The error from that event is at most $(m + 2)\beta/N$. Replacing the opening probabilities and conditional masses by their ideal values costs at most $4m\beta$, by telescoping the opening factors and using $B + \sum w_i < 2$. A noncanonical hull requires an exceptional draw or $P \notin K$, with combined probability at most $\beta/N + 1 - B$.

Rational centres and positive variances. Round each vector $d_i u_i$ to a rational vector within β/N before multiplying by s_i. By Pinsker's inequality for a Gaussian shift, each normalized shift of length at most β/N costs at most $\beta/(2N)$ in total variation. Summing over candidates costs less than β. It remains to move the fixed vertices without changing the hull except on an event of probability less than β.

Set $s_* = \min(1, s_1, \dots, s_m)$, $b = \beta/(100N^3)$ and $\delta = s_*^2 b^2$. In any triple containing a random site, draw that site last. The other two lie at distance at least $s_* b$ except with probability at most $b^2/2$, by bounding the Gaussian density on a disk. Two fixed vertices are already at least one unit apart. Given a separated pair, its orientation determinant with the last site is normal with standard deviation at least $s_*^2 b$, so it lies within δ of zero with probability at most b. The all-fixed determinants have magnitude at least one. A union bound over all triples leaves probability less than $\beta/10$ of any small determinant.

Replace each fixed vertex v by $v + \tau Z$ with independent standard planar Gaussian Z. Put $J = 10(U + \lceil \log_2(N/b) \rceil + 10)$. Except with probability less than $\beta/10$, all old positions and all new standard draws have norm at most J, by the planar Gaussian tail. Choose the largest dyadic power $\tau \le \delta/(32J^2)$. Moving the vertices changes each determinant by at most $(8\tau + 4\tau^2)J^2 < \delta$, preserving every orientation sign outside these exceptional events. Since $s_* \ge \lambda/(2U)$, the binary length of τ is polynomial. The errors from calibrating the candidates and perturbing the fixed vertices total less than $(4m+3)\beta$, far below $G/4 = 2^{11} N^2 \beta$. Including both perturbations, every canonical probability therefore has error less than $G/4$, and the total noncanonical mass is at most $1 - B + \beta/N + 2\beta < 2^{-m-4}$.

Data and code availability

The computational materials include the code, centres, noise scales, seeds and batch counts for the four-site experiment, the rational certificate for Proposition 5, and the inputs and code for the illustrations and finite geometric checks, including the convex hull examples for Theorem 24. The author will release these materials in a public repository once the preprint is public on both arXiv and Hexagon. The SHA-256 checksum of the release archive is `0441748a74848264fff6ddd48c18c7cd2de803bfa4558166fff37d731248958f`

References

[1] A. M. Abdelbar and S. M. Hedetniemi. Approximating MAPs for belief networks is NP-hard and other theorems, *Artificial Intelligence* 102(1) (1998), 21–38. doi:10.1016/S0004-3702(98)00043-5.

[2] M. Abellanas, F. Hurtado and P. A. Ramos. Structural tolerance and Delaunay triangulation, *Information Processing Letters* 71(5–6) (1999), 221–227. doi:10.1016/S0020-0190(99)00107-6.

[3] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: nearest neighbor relations for imprecise points, in *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2004), 410–419. `https://dl.acm.org/doi/10.5555/982792.982851`.

[4] S. Boucheron, G. Lugosi and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. doi:10.1093/acprof:oso/9780199535255.001.0001.

[5] H. Whitney. Congruent graphs and the connectivity of graphs, *American Journal of Mathematics* 54(1) (1932), 150–168. doi:10.2307/2371086.

[6] G. Brinkmann. A simple and elementary proof of Whitney's unique embedding theorem, *Ars Mathematica Contemporanea* 20(2) (2021), 195–197. doi:10.26493/1855-3974.2334.331.

[7] K. Buchin, M. Löffler, P. Morin and W. Mulzer. Preprocessing imprecise points for Delaunay triangulation: simplified and extended, *Algorithmica* 61(3) (2011), 674–693. doi:10.1007/s00453-010-9430-0.

[8] P. M. M. de Castro, J. Tournois, P. Alliez and O. Devillers. Filtering relocations on a Delaunay triangulation, *Computer Graphics Forum* 28(5) (2009), 1465–1474. doi:10.1111/j.1467-8659.2009.01523.x.

[9] O. Devillers. Delaunay triangulation of imprecise points, preprocess and actually get a fast query time, *Journal of Computational Geometry* 2(1) (2011), 30–45. doi:10.20382/v2i1a3.

[10] J. S. Ely and A. P. Leclerc. Correct Delaunay triangulation in the presence of inexact inputs and arithmetic, *Reliable Computing* 6(1) (2000), 23–38. doi:10.1023/A:1009977923779.

[11] W. Evans and J. Sember. Guaranteed Voronoi diagrams of uncertain sites, in *Proceedings of the 20th Canadian Conference on Computational Geometry* (CCCG 2008), 207–210. `https://www.cccg.ca/proceedings/2008/paper50.pdf`. Extended version: Technical Report TR-2008-17, University of British Columbia, 2008, `https://www.cs.ubc.ca/tr/2008/tr-2008-17`.

[12] H. Maehara. On Sylvester's problem (in Japanese), *Proceedings of the Institute of Statistical Mathematics, Tokyo* 25 (1978), 81–86. `https://www.ism.ac.jp/editsec/toukei/ihou.html`.

[13] C. Blatter. Four shots for a convex quadrilateral, *American Mathematical Monthly* 115(9) (2008), 837–843. doi:10.1080/00029890.2008.11920598.

[14] W. Hoeffding. Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58(301) (1963), 13–30. doi:10.1080/01621459.1963.10500830.

[15] E. B. Wilson. Probable inference, the law of succession, and statistical inference, *Journal of the American Statistical Association* 22(158) (1927), 209–212. doi:10.1080/01621459.1927.10502953.

[16] O. Giménez and M. Noy. Asymptotic enumeration and limit laws of planar graphs, *Journal of the American Mathematical Society* 22(2) (2009), 309–329. doi:10.1090/S0894-0347-08-00624-3.

[17] R. Gitik and L. Joskowicz. Voronoi diagram and Delaunay triangulation with independent and dependent geometric uncertainties, *International Journal of Computational Geometry & Applications* 31(2–3) (2021), 75–121. doi:10.1142/S0218195921500059.

[18] A. Jørgensen, M. Löffler and J. M. Phillips. Geometric computations on indecisive points, in *Algorithms and Data Structures – WADS 2011*, LNCS 6844, Springer, 2011, 536–547. doi:10.1007/978-3-642-22300-6_45.

[19] Z. Kabluchko and H. Panzo. A refinement of the Sylvester problem: probabilities of combinatorial types, *Discrete & Computational Geometry* 76 (2026), 1177–1211. doi:10.1007/s00454-026-00820-2.

[20] W. F. Sheppard. On the application of the theory of error to cases of normal distribution and normal correlation, *Philosophical Transactions of the Royal Society of London, Series A* 192 (1899), 101–167. doi:10.1098/rsta.1899.0003.

[21] A. A. Khanban and A. Edalat. Computing Delaunay triangulation with imprecise input data, in *Proceedings of the 15th Canadian Conference on Computational Geometry* (CCCG 2003), 94–97. `https://www.cccg.ca/proceedings/2003/38.pdf`.

[22] M. Khanimov and M. Sharir. Delaunay triangulations of degenerate point sets, preprint, arXiv:1510.04608, 2015. arxiv.org/abs/1510.04608.

[23] H. Leipold, E. A. Lazar, K. A. Brakke and D. J. Srolovitz. Statistical topology of perturbed two-dimensional lattices, *Journal of Statistical Mechanics: Theory and Experiment* (2016), 043103. doi:10.1088/1742-5468/2016/04/043103.

[24] M. Löffler and J. M. Phillips. Shape fitting on point sets with probability distributions, in *Algorithms – ESA 2009*, LNCS 5757, Springer, 2009, 313–324. doi:10.1007/978-3-642-04128-0_29.

[25] M. Löffler and J. Snoeyink. Delaunay triangulation of imprecise points in linear time after preprocessing, *Computational Geometry* 43(3) (2010), 234–242. doi:10.1016/j.comgeo.2008.12.007.

[26] A. Lucas. Ising formulations of many NP problems, *Frontiers in Physics* 2 (2014), article 5. doi:10.3389/fphy.2014.00005.

[27] P. Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality, *Annals of Probability* 18(3) (1990), 1269–1283. doi:10.1214/aop/1176990746.

[28] B. Mohar. Face covers and the genus problem for apex graphs, *Journal of Combinatorial Theory, Series B* 82(1) (2001), 102–117. doi:10.1006/jctb.2000.2026.

[29] A. W. Palmer, A. J. Hill and S. J. Scheding. Modelling resource contention in multi-robot task allocation problems with uncertain timing, in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, 3693–3700. doi:10.1109/ICRA.2018.8460981. Author's accepted version: `https://arxiv.org/abs/1607.04358v3`.

[30] J. D. Park and A. Darwiche. Complexity results and approximation strategies for MAP explanations, *Journal of Artificial Intelligence Research* 21 (2004), 101–133. doi:10.1613/jair.1236.

[31] I. R. Savage. Contributions to the theory of rank order statistics: the "trend" case, *Annals of Mathematical Statistics* 28(4) (1957), 968–977. doi:10.1214/aoms/1177706797.

[32] S. Suri and K. Verbeek. On the most likely Voronoi diagram and nearest neighbor searching, *International Journal of Computational Geometry & Applications* 26(3–4) (2016), 151–166. doi:10.1142/S0218195916600025.

[33] S. Suri, K. Verbeek and H. Yıldız. On the most likely convex hull of uncertain points, in *Algorithms – ESA 2013*, LNCS 8125, Springer, 2013, 791–802.

doi:10.1007/978-3-642-40450-4_67. Author's version with appendices: https://sites.cs.ucsb.edu/~suri/psdir/esa13.pdf.

[34] R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs, *Discrete & Computational Geometry* 1(4) (1986), 321–341. doi:10.1007/BF02187705.

[35] H. R. Tiwary. On the largest convex subsets in Minkowski sums, *Information Processing Letters* 114(8) (2014), 405–407. doi:10.1016/j.ipl.2014.02.013. Author's preprint: https://kam.mff.cuni.cz/~hansraj/publications/minkowskisubset.pdf.

[36] M. Yvinec. 2D Triangulations, in *CGAL User and Reference Manual*, version 6.1.1, 2026. https://doc.cgal.org/6.1.1/Triangulation_2/index.html.

[37] F. B. Atalay, S. A. Friedler and D. Xu. Convex hull for probabilistic points, *29th SIBGRAPI Conference on Graphics, Patterns and Images* (2016), 48–55. doi:10.1109/SIBGRAPI.2016.016.

[38] R. M. Karp. Reducibility among combinatorial problems, in *Complexity of Computer Computations*, Plenum Press (1972), 85–103. doi:10.1007/978-1-4684-2001-2_9.