

Radius Functions on Poisson–Delaunay Mosaics and Related Complexes Experimentally*

H. Edelsbrunner, A. Nikitenko, K. Ölsböck, and P. Synak

IST Austria (Institute of Science and Technology Austria), Am Campus 1,
3400 Klosterneuburg, Austria,
herbert.edelsbrunner, anton.nikitenko, katharina.oelsboeck, peter.synak@ist.ac.at

Abstract

Discrete Morse theory has recently led to new developments in the theory of random geometric complexes. This article surveys the methods and results obtained with this new approach, and discusses some of its shortcomings. It uses simulations to illustrate the results and to form conjectures, getting numerical estimates for combinatorial, topological, and geometric properties of weighted and unweighted Delaunay mosaics, their dual Voronoi tessellations, and the Alpha and Wrap complexes contained in the mosaics.

2010 AMS Mathematics Subject Classification: 60D05; 60-04; 65D18; 68U05.

Keywords and phrases: Voronoi tessellations, Delaunay mosaics, Alpha complexes, Wrap complexes; slices, weights; discrete Morse theory, critical simplices, intervals; stochastic geometry, Poisson point processes, expectations; simulation.

1 Introduction

Natural phenomena are often characterized by spatial decompositions reflecting local proximity. Indeed, such phenomena arise in different disciplines of science and beyond, so that a variety of names were established all referring to the same geometric model: *Voronoi diagrams*, *Dirichlet tessellations*, *Wigner–Seitz cells*, *Thiessen polygons*, *Brillouin zones* etc.; see [3]. The basic version is defined for a locally finite set, $X \subseteq \mathbb{R}^n$, and assigns to each $x \in X$ the region of points that are at least as close to x as to any other point in X . We refer to the collection of such regions as the *Voronoi tessellation* of X . Assuming the points are in general position, the nerve of the Voronoi tessellation is a simplicial complex in \mathbb{R}^n , which we refer to as the *Delaunay mosaic* of X . Beyond the basic version, we limit ourselves to the weighted case, in which the squared Euclidean distance is replaced by the power distance to a point. The resulting decomposition is often referred to as *power diagram* or *Laguerre tessellation* but we will call it a *weighted Voronoi tessellation*. To introduce randomness, we use a stationary Poisson point process in Euclidean space, and we refer to the resulting random geometric structures as *Poisson–Voronoi tessellations* and their dual *Poisson–Delaunay mosaics*. While the former have more interesting geometry, the latter are more convenient to work with when combinatorial and topological aspects are in the focus. We will be interested in such aspects of the Alpha and Wrap complexes of Delaunay mosaics.

Prior work and contributions. The systematic investigation of random Voronoi tessellations and Delaunay mosaics was initiated by Miles’ extensive study of the 2-dimensional

* This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreements No 78818 Alpha and No 638176). It is also partially supported by the DFG Collaborative Research Center TRR 109, ‘Discretization in Geometry and Dynamics’, through grant no. I02979-N35 of the Austrian Science Fund (FWF).

case [39], but some results were already known to Meijering [38]. While Miles settled a wide variety of stochastic questions in \mathbb{R}^2 , few theoretical results beyond two dimensions were known prior to [26]. An exception is the expected number of Voronoi vertices or, equivalently the top-dimensional simplices in the Delaunay mosaic, although many integral expressions and relations for geometric characteristics like intrinsic volumes and numbers of simplices were available; see [11, 12, 42], and [46, Chapter 10] for a general survey. In this context, we also mention [6], where general Gamma-type results for distributions of various associated quantities were obtained. Parallel to the purely mathematical interest, the study of random tessellations in \mathbb{R}^3 is motivated by questions in material science, and a wealth of primarily experimental findings on 3-dimensional Poisson–Voronoi tessellations can be found in [35, 37]. Random weighted Voronoi tessellations were studied in [33, 34], with the weights following their own distribution, such as uniform or normal. Alternatively, we may construct a weighted Voronoi tessellation as a slice of a higher-dimensional unweighted Voronoi tessellation, and this construction was briefly considered in [42].

In this paper we continue the study of the questions pioneered in the already mentioned works of Miles [39, 40], and we consider them from the viewpoint of discrete topology. When we construct a Delaunay mosaic incrementally, then each simplex acquires topological significance. To formalize this idea, we may consider the (generalized) discrete Morse function that encodes the family of Alpha complexes contained in the mosaic [5]. We are motivated to shift to this view by the widespread use of persistence diagrams in topological data analysis [13, 19]. A first step in this analysis turns the data into a filtration of complexes, and the most common types are the Čech, Vietoris–Rips, and Alpha complexes. The stochastic properties of the first two applied to a Poisson point process have been investigated by Kahle [30, 31], by Bobrowski and coauthors [7, 9], by Decreusefond et al. [15], and recently by us [43]. The stochastic properties of Alpha complexes have come into focus recently [23, 24, 25, 26], and some of these findings will be surveyed in this paper. In addition, we will present experimental data to illustrate the theoretical results but also to motivate further studies aimed at shedding light on observed but mathematically not yet understood stochastic phenomena.

Approach. There is a subtle but important difference between the conventional approach to stochastic geometry and the approach taken in this paper. To explain the difference, consider the problem of counting the simplices in a Poisson–Delaunay mosaic in \mathbb{R}^n , possibly differentiating between simplices of different dimensions. To get started, we map each simplex, Q , to a *representative point*, $\text{center}(Q) \in \mathbb{R}^n$, and we study the resulting point process with tools from integral geometry. In the *conventional approach*, $\text{center}(Q)$ is chosen in an isometry-equivariant manner, for example as the center of mass. In contrast, in this paper we map Q to the center of the smallest sphere that passes through the vertices of the simplex and does not enclose any points of X . This definition is not isometry-equivariant as $\text{center}(Q)$ depends on more than just the simplex, so we call this the *context-sensitive approach*. Note that this mapping is generally not injective — not even if we assume that the points of X are in general position — but there is topological meaning in the incidences. Indeed, all simplices that map to the same point form an interval in the face poset, and if this interval contains two or more simplices, then adding them to the last Alpha complex does not change the homotopy type. In addition to studying the resulting point process, we need to understand the intervals of simplices that share the same representative point. In other words, we study the discrete Morse function of the Delaunay mosaic and get topologically refined stochastic information on its simplices. The idea of the context-sensitive approach thus originates in discrete Morse theory, which was introduced by Forman in [27], later

generalized to intervals by Freij in [28], and recently applied to Delaunay complexes in [5].

Outline. In this survey, we summarize the results obtained with the context-sensitive approach, particularly focusing on the remaining open questions and on numerical simulations. In Section 2, we study unweighted Delaunay mosaics, and in addition to counting all simplices, we count those with circumscribed spheres of radius at most r , which amounts to studying Alpha complexes. In Section 3, we extend the study to Wrap complexes, which exploit the flow defined by discrete Morse theory of the Delaunay mosaic to reconstruct shapes from data. In Section 4, we turn to the weighted Poisson–Delaunay mosaics generated from slices of higher-dimensional unweighted Voronoi tessellations. In Section 5, we address questions about the sphericity regions in weighted and unweighted Voronoi tessellations. In Section 6, we conclude the main part of this paper. We collect background material that may be helpful to the non-specialist in three appendices, discussing tessellations, mosaics, and complexes in Appendix A, introducing discrete Morse theory and homology groups in Appendix B, and explaining probabilistic concepts in Appendix C.

2 Poisson–Delaunay Mosaics

In this section, we present fundamental probabilistic results on random Delaunay mosaics and the Alpha complexes, which are their subcomplexes. As it turns out, it can be easier to study the Alpha complexes first and to combine the obtained insights to gain an understanding of random Delaunay mosaics. We begin with a brief introduction of the main geometric and topological concepts and follow up with probabilistic results from the literature and with numerical data collected from extensive computational experiments. More detailed introductions to the geometric, topological, and probabilistic background can be found in the appendices.

Geometric and topological concepts. We write \mathbb{R}^n for the n -dimensional Euclidean space and let $X \subseteq \mathbb{R}^n$ be a locally finite set. The *Voronoi tessellation* of X covers \mathbb{R}^n with closed convex domains, one for each point in X . Assuming general position, the *Delaunay mosaic* of X , denoted $\text{Del}(X)$, is the nerve of the Voronoi tessellation, geometrically realized by mapping each domain to the corresponding point in X . By construction, $\text{Del}(X)$ is a simplicial complex with simplices of dimension 0 to n . Each p -simplex is the convex hull of $p + 1$ points, and we find it convenient to identify the simplex with its set of vertices, $Q \subseteq X$. It is not difficult to see that for each simplex $Q \in \text{Del}(X)$, there is an $(n - 1)$ -dimensional sphere such that all points of Q lie on the sphere, and no point of X lies inside the sphere. There is a unique smallest such sphere, which we refer to as the *Delaunay sphere* of Q . Letting $\mathcal{R}(Q)$ be the radius of the Delaunay sphere, we get a function $\mathcal{R}: \text{Del}(X) \rightarrow \mathbb{R}$, which we call the (*Delaunay*) *radius function* of X . For each $r \in \mathbb{R}$, the *Alpha complex* of X for r is the corresponding sublevel set of the radius function: $\text{Alpha}_r(X) = \mathcal{R}^{-1}[0, r]$.

The Delaunay mosaic can be partitioned into *intervals*, which consist of simplices with common Delaunay sphere. The simplices inside an interval share the radius, hence every Alpha complex either contains all simplices in an interval or none of them. Writing ℓ for the minimum dimension and m for the maximum dimension of any simplex in an interval, we say the interval has *type* (ℓ, m) . If $\ell = m$, we call the interval *singular* and its only simplex a *critical simplex* of \mathcal{R} . The critical simplices determine the topology of the Alpha complexes as measured by their homology groups; see Appendix B for details. Indeed, if we construct the Delaunay mosaic incrementally, in the order of non-decreasing radius, we preserve the Betti numbers whenever we add a non-singular interval, and we change exactly one Betti

number whenever we add a critical p -simplex, namely we increase the p -th Betti number by 1 if the simplex *gives birth* of a p -cycle, and we decrease the $(p - 1)$ -st Betti number by 1 if the simplex *gives death* of a $(p - 1)$ -cycle.

Probabilistic background. As proved in [26], the radius of the typical interval of a given type is Gamma-distributed. To state this more formally, consider a stationary Poisson point process X with intensity $\rho > 0$ in \mathbb{R}^n , and write $c_{\ell,m}^n(r)$ for the expected number of intervals of type (ℓ, m) in the Alpha complex $\text{Alpha}_r(X)$ that lie inside a Borel region $\Omega \subseteq \mathbb{R}^n$. Write $\Gamma(m)$ for the Gamma function, $\gamma(m; t)$ for its lower incomplete version, and call $\tilde{\gamma}(m; t) = \gamma(m; t)/\Gamma(m)$ the *regularized lower incomplete Gamma function*. Then there exist constants $C_{\ell,m}^n$ such that

$$c_{\ell,m}^n(r) = \tilde{\gamma}(m; \rho \nu_n r^n) \cdot C_{\ell,m}^n \cdot \rho \|\Omega\| \quad (1)$$


for every $\ell > 0$ and every $r \geq 0$, in which $\nu_n = \pi^{\frac{n}{2}}/\Gamma(\frac{n}{2} + 1)$ is the volume of the unit ball in \mathbb{R}^n . Using (7) in Appendix B, we can transform this into a statement about the number of simplices. Writing $d_j^n(r)$ for the expected number of j -simplices in $\text{Alpha}_r(X)$ in Ω , we have

$$d_j^n(r) = \sum_{\ell=0}^j \sum_{m=j}^n \binom{m-\ell}{m-j} \tilde{\gamma}(m; \rho \nu_n r^n) \cdot C_{\ell,m}^n \cdot \rho \|\Omega\| \quad (2)$$

for $j > 0$ and every $r \geq 0$. For convenience, we also define the expected number of simplices per unit volume, D_j^n , using the formula $d_j^n(\infty) = D_j^n \rho \|\Omega\|$. The proof of (2) can be found in [26], where the constants $C_{\ell,m}^n$ are computed explicitly for $n \leq 4$. By saying “in Ω ” we mean that the center of the Delaunay sphere lies in Ω . An intuitively clear but technical fact [26, Appendix A] is that this condition can be replaced by “lying inside Ω ” or “intersecting Ω ”, at the cost of weakening (1) and (2) by adding $o(\|\Omega\|)$ on their right-hand sides. We use this to estimate the distribution numerically.

Computational experiments. We present experimental results in 2 and 3 dimensions. Figure 1 shows a 2-dimensional Poisson–Delaunay mosaic restricted to a square window. For the computation, we chose the square window of size 300×300 in \mathbb{R}^2 or the cube window of size $60 \times 60 \times 60$ in \mathbb{R}^3 . To avoid boundary effects, we impose periodic boundary conditions in our simulations. An instance of the Poisson point process with intensity $\rho = 1$ is sampled and the geometric software library CGAL [54] is used to compute Voronoi tessellations and Delaunay mosaics. We count the intervals and simplices and compare the experimentally observed constants $C_{\ell,m}^n$ and D_j^n with their mathematically derived values. In \mathbb{R}^2 , the Euler characteristic implies $D_0^2 = 1$, $D_1^2 = 3$, $D_2^2 = 2$. For a Poisson point process, we expect that half the triangles are acute and the other half are obtuse [39], which implies $C_{1,1}^2 = 2$, $C_{1,2}^2 = 1$, $C_{2,2}^2 = 1$. Averaging the experimentally observed numbers over 1000 runs, we match these predictions with an accuracy of at least two positions after the decimal point. In \mathbb{R}^3 , we average the observed numbers over 100 runs; see Table 1 for a comparison of the results.

Intervals and simplices. Besides the total densities of intervals and simplices, we are interested in their dependence on the radius. Analytic formulas for these dependencies can be found in [26], but no such formulas are known for the variances. As shown in [39], the standard deviations are small if compared to the expected values, but exact values are not available. Figure 2 gives the experimentally observed densities in \mathbb{R}^2 and in \mathbb{R}^3 . Looking at the graphs in the left panels, we get the number of intervals of any specific type with radius between r_1 and r_2 per unit volume as the area below the corresponding curve and

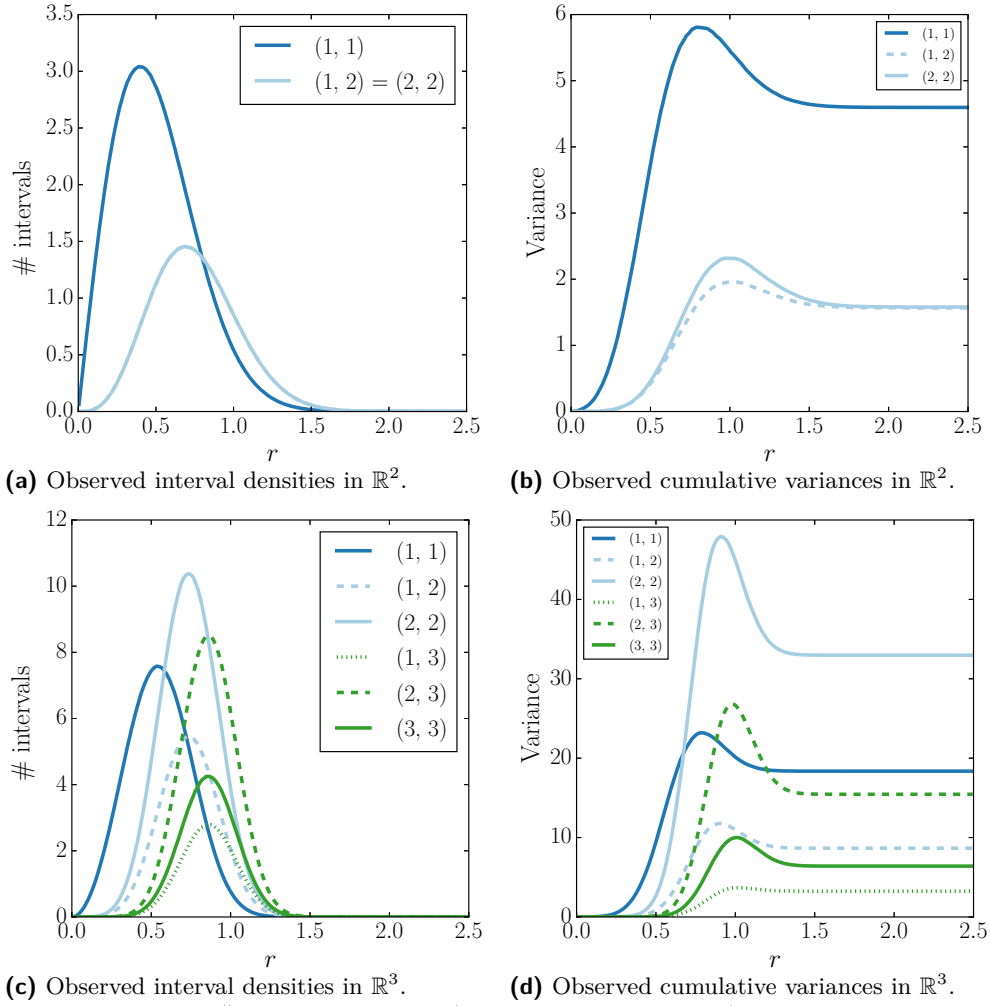


■ **Figure 1** Poisson–Delaunay mosaic. The color distinguishes between an Alpha subcomplex and the simplices whose Delaunay spheres have radii exceeding the threshold.

$n = 3$	theoretical	exp.
$C_{0,0}^3$	$1 = 1.00$	1.00
$C_{1,1}^3$	$4 = 4.00$	4.00
$C_{1,2}^3$	$\frac{9}{16}\pi^2 - 3 \approx 2.55$	2.55
$C_{2,2}^3$	$\frac{3}{16}\pi^2 + 3 \approx 4.85$	4.85
$C_{1,3}^3$	$\frac{69}{560}\pi^2 \approx 1.22$	1.22
$C_{2,3}^3$	$\frac{3}{8}\pi^2 \approx 3.70$	3.70
$C_{3,3}^3$	$\frac{3}{16}\pi^2 \approx 1.85$	1.85
D_0^3	$1 = 1.00$	1.00
D_1^3	$\frac{24}{35}\pi^2 + 1 \approx 7.77$	7.77
D_2^3	$\frac{48}{35}\pi^2 \approx 13.54$	13.53
D_3^3	$\frac{24}{35}\pi^2 \approx 6.77$	6.77

■ **Table 1** Mathematically derived [26] and experimentally estimated values for $C_{\ell,m}^n$ and D_j^n in $n = 3$ dimensions. The numbers in the right column are averaged over 100 mosaics in \mathbb{R}^3 with 6 277 766 simplices on average.

above the segment $[r_1, r_2]$. These graphs are computed by normalizing the corresponding histograms that bin the intervals with radii between $i/100$ and $(i + 1)/100$ for i ranging from 0 to 250, averaging over the same number of runs as before. We fit regularized lower incomplete Gamma functions to these graphs, using the `curve_fit`-function of Python 2.7’s `scipy`-module, which is based on least squares optimization. As stated in (2), the densities of the simplices are linear combinations of the densities of the intervals, so we fit linear combinations of Gamma functions; see Section 3 for their discussion. Since these distributions are known theoretically, it is not surprising that we get an excellent fit, but the precise quantification of the error is still useful as it calibrates the error we get for some of the theoretically unknown densities we discuss later. We have made no attempts to fit analytic



interval	distribution	experimental fit	fitting error
(1, 1)	$C_{1,1}^3 \cdot \tilde{\gamma}(1; \frac{4}{3}\pi r^3)$	$4.00 \cdot \tilde{\gamma}(1.00; 4.19r^3)$	6.96×10^{-5}
(1, 2)	$C_{1,2}^3 \cdot \tilde{\gamma}(2; \frac{4}{3}\pi r^3)$	$2.55 \cdot \tilde{\gamma}(2.00; 4.19r^3)$	4.01×10^{-5}
(2, 2)	$C_{2,2}^3 \cdot \tilde{\gamma}(2; \frac{4}{3}\pi r^3)$	$4.85 \cdot \tilde{\gamma}(2.00; 4.19r^3)$	6.68×10^{-5}
(1, 3)	$C_{1,3}^3 \cdot \tilde{\gamma}(3; \frac{4}{3}\pi r^3)$	$1.22 \cdot \tilde{\gamma}(3.00; 4.19r^3)$	3.23×10^{-5}
(2, 3)	$C_{2,3}^3 \cdot \tilde{\gamma}(3; \frac{4}{3}\pi r^3)$	$3.70 \cdot \tilde{\gamma}(3.00; 4.19r^3)$	5.50×10^{-5}
(3, 3)	$C_{3,3}^3 \cdot \tilde{\gamma}(3; \frac{4}{3}\pi r^3)$	$1.85 \cdot \tilde{\gamma}(3.00; 4.19r^3)$	3.13×10^{-5}

■ **Figure 2** The observed densities of the intervals as functions of the radius on the *left*, and the corresponding variances normalized by the expected number of points on the *right*, showing the experimental results in \mathbb{R}^2 on the *top* and in \mathbb{R}^3 in the *middle*. The table *below* the panels compares the experimentally fit distributions in \mathbb{R}^3 with the mathematically derived formulas. The constants in the table are rounded to two digits after the decimal point. We quantify the error by taking the square root of the average squared deviation of the sampled points from the fitted curve, which estimates the L^2 -difference between the fitted and sampled density functions.

curves to the graphs for the variances, shown in the right panels of Figure 2, which are computed with a window of size 100×100 in \mathbb{R}^2 and of size $15 \times 15 \times 15$ in \mathbb{R}^3 , averaged over 10 000 runs and normalized by the window size, which is also the expected number of points. Trying different window sizes (result not shown), we get almost the same graphs, which implies that the variances are proportional to the number of points.

Observe that the expected number of obtuse triangles (or intervals of type $(1, 2)$) in \mathbb{R}^2 is the same as the expected number of acute triangles (or intervals of type $(2, 2)$), for every $r \geq 0$. In contrast, the variance for obtuse triangles is consistently smaller than that for acute triangles.

Critical simplices and Betti numbers. Turning our attention to the topology of the Alpha complexes, we consider the critical simplices that give birth and that give death, the Betti numbers, and the variances of the Betti numbers as functions of the radius threshold; see Appendix B for a detailed introduction of these concepts. No analytic expressions are known for these densities, so Figure 3 just shows the experimental results. The graphs of birth- and death-giving critical simplices are visually similar to those we get for the intervals, but our attempt to fit Gamma-functions gives fitting errors of order 10^{-3} or worse, which suggests that unlike the critical simplices, their two types do not follow Gamma-distributions. We get the p -th Betti number as the difference between the number of birth-giving p -simplices and the number of death-giving $(p+1)$ -simplices. Since the latter are not Gamma-distributed, neither are the former, with the exception of the n -th Betti number, which equals the number of critical n -simplices.

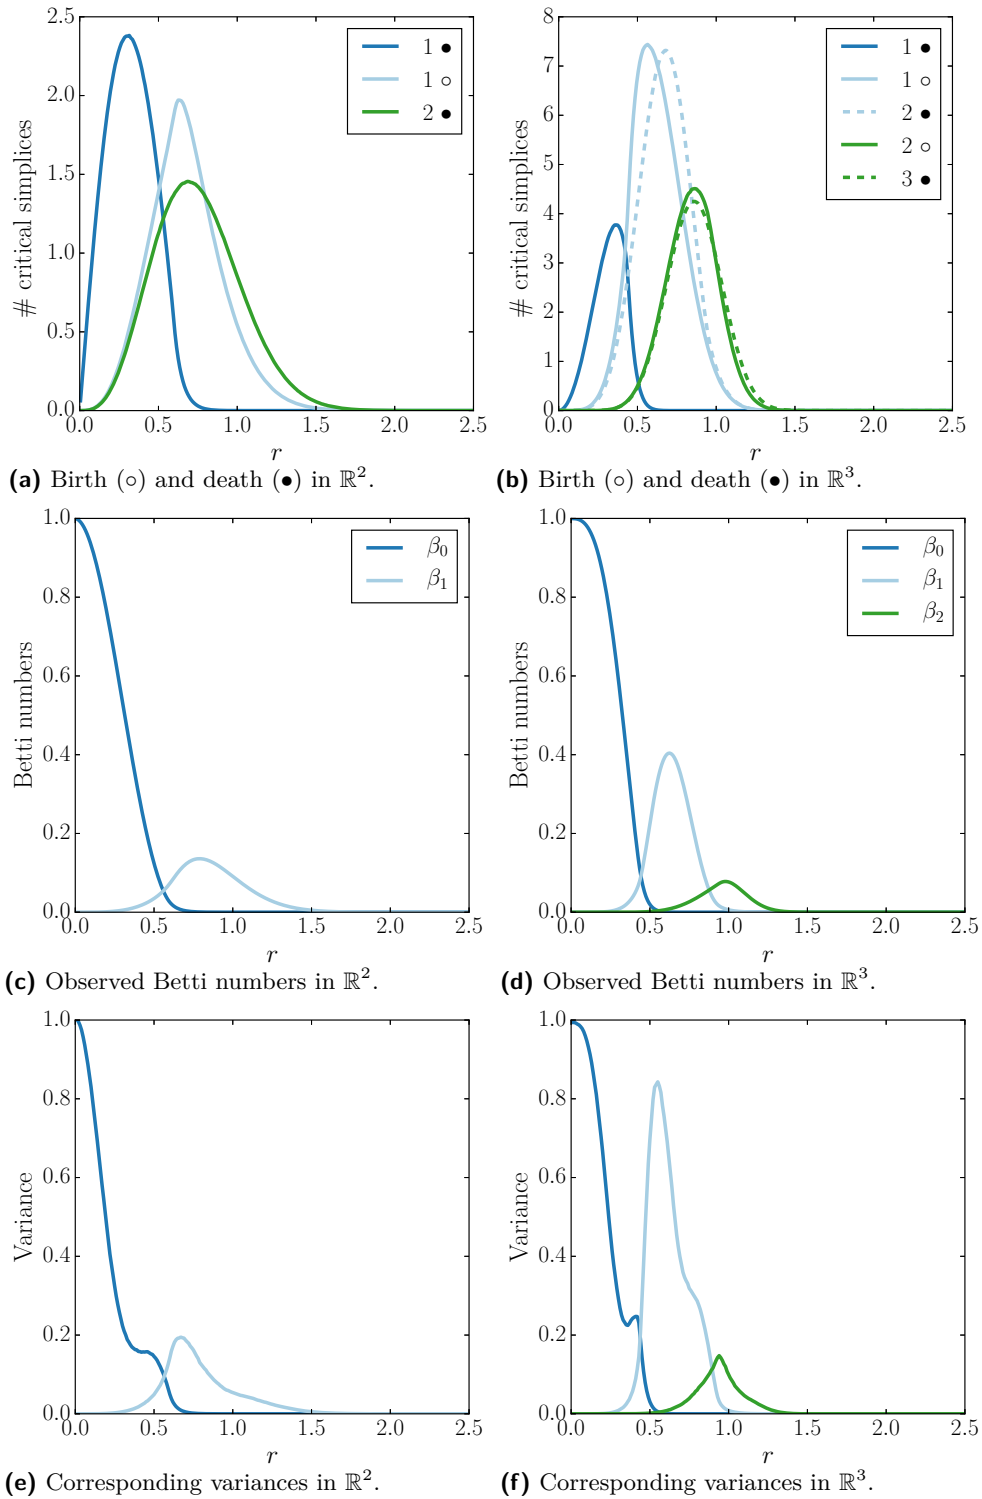
3 Poisson–Wrap Complexes

In this section, we study the Wrap complexes of Poisson point processes. As described in detail in Appendix B, each Wrap complex is contained in and homotopy equivalent to the Alpha complex for the same radius.

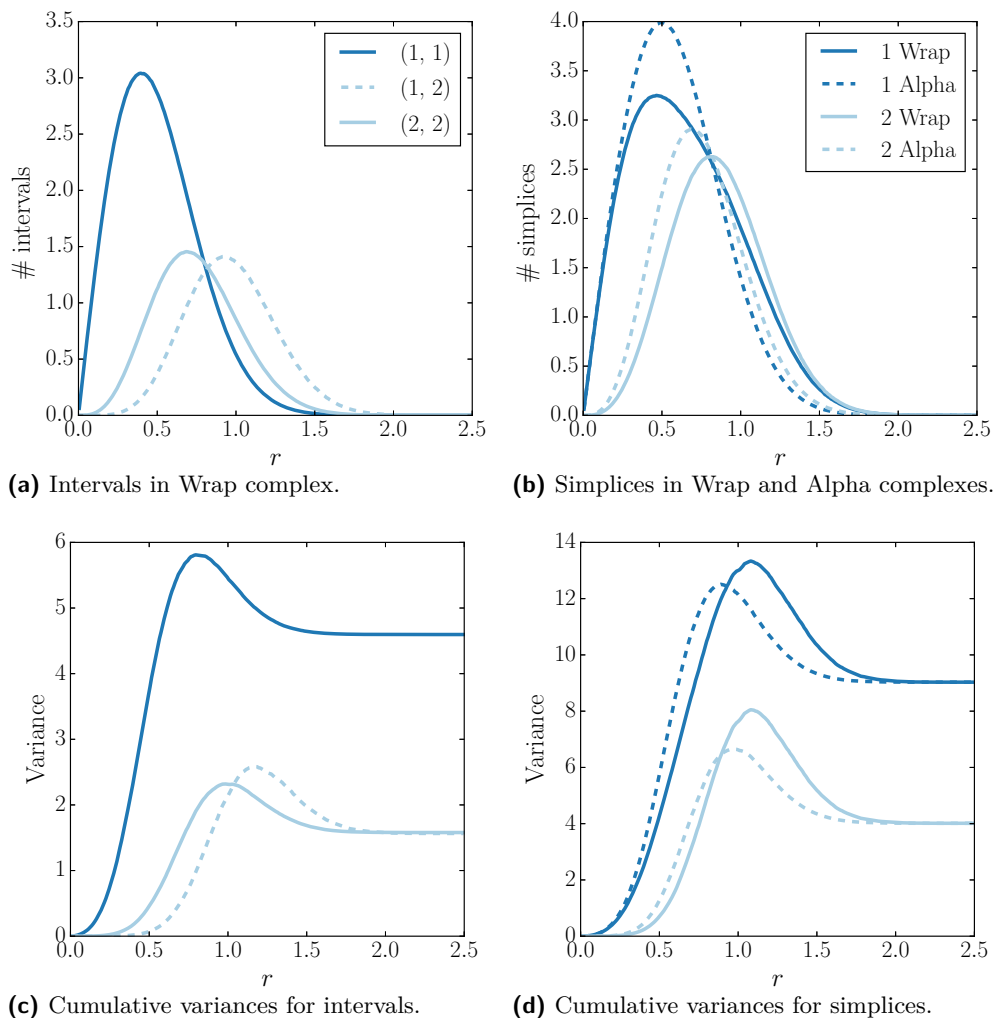
Definitions of Wrap complex. Let $X \subseteq \mathbb{R}^n$ be locally finite and in general position so that $\text{Del}(X)$ is a simplicial complex and $\mathcal{R}: \text{Del}(X) \rightarrow \mathbb{R}$ induces a well-defined partition of the Delaunay mosaic into intervals. For every $r \geq 0$, the *Wrap complex* of X for r , denoted $\text{Wrap}_r(X)$, is the smallest subcomplex of $\text{Del}(X)$ that contains all critical simplices with radius $\mathcal{R}(Q) \leq r$ and has only complete intervals. Clearly $\text{Wrap}_r(X) \subseteq \text{Alpha}_r(X)$, and since both complexes contain the same critical simplices, the two have the same homotopy type and therefore isomorphic homology groups and equal Betti numbers. It follows that the experimental results shown in Figure 3 apply without change to Wrap complexes as well.

It will be useful to unpack this definition, which we do by introducing a partial order on the simplices in the Delaunay mosaic. Inside an interval, the order climbs down the Hasse diagram, and between intervals, it climbs up the Hasse diagram. For each critical simplex Q of \mathcal{R} , we define the *lower set*, denoted $\downarrow Q$, as the simplices that precede Q in this order. The Wrap complex of X for r is then the union of the lower sets of all critical simplices in $\text{Alpha}_r(X)$.

This definition suggests we introduce the *Wrap radius function*, $\mathcal{R}_W: \text{Del}(X) \rightarrow \mathbb{R}$, which maps each simplex, P , to the minimum radius of a critical simplex, Q , that satisfies $P \in \downarrow Q$. By construction, \mathcal{R}_W agrees with \mathcal{R} on the critical simplices, and we have $\mathcal{R}(P) \leq \mathcal{R}_W(P)$ for all $P \in \text{Del}(X)$. With this definition, we have $\text{Wrap}_r(X) = \mathcal{R}_W^{-1}[0, r]$. Since $\text{Wrap}_r(X) \subseteq \text{Alpha}_r(X)$, the number of intervals of any type contained in the Wrap complex is less than or equal to those contained in the Alpha complex for the same radius.

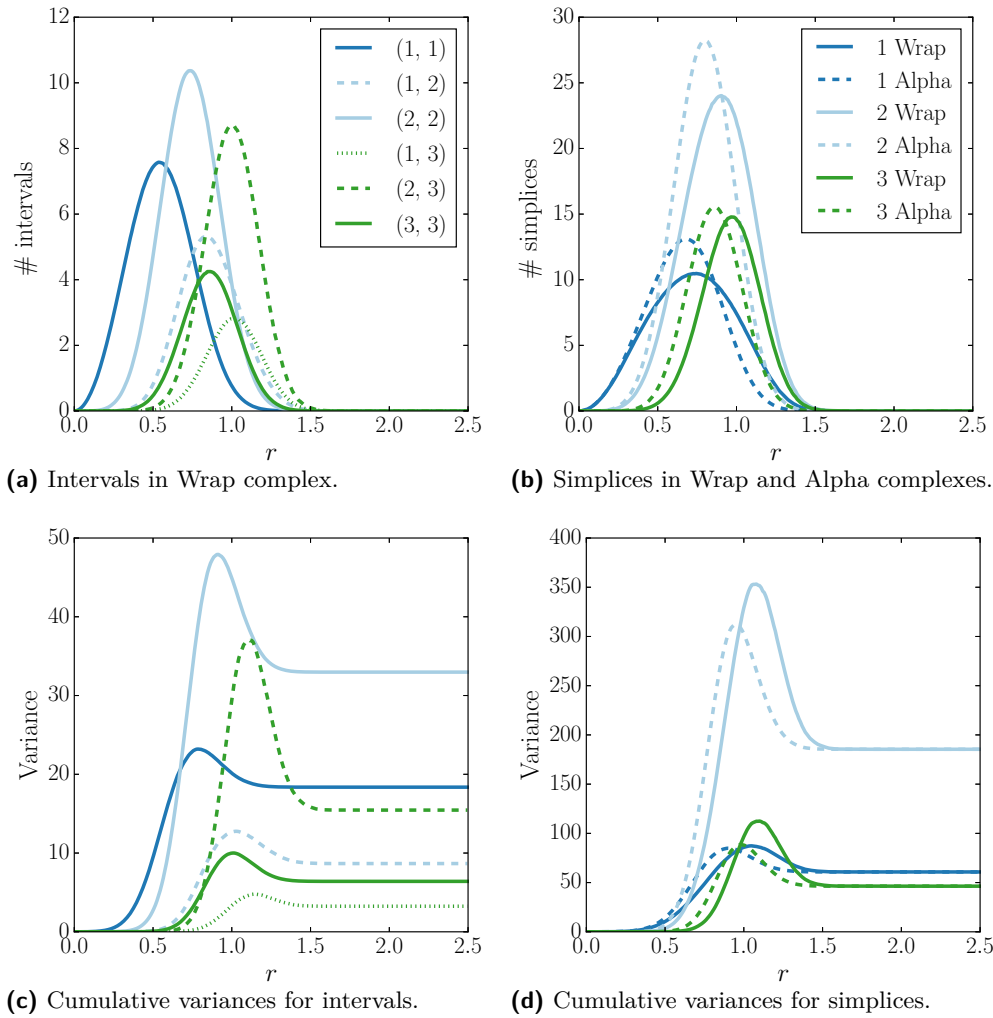


■ **Figure 3** *Top row:* experimentally observed densities of birth- and death-giving critical simplices in \mathbb{R}^2 on the *left* and in \mathbb{R}^3 on the *right*. Except for the top-dimensional simplices, which all give death, they do not seem to follow Gamma-distributions. The numbers and types of critical simplices determine the Betti numbers. *Middle and bottom rows:* the observed Betti numbers of the sublevel sets and the corresponding variances, again in \mathbb{R}^2 on the *left* and in \mathbb{R}^3 on the *right*.



interval	experimental fit	fitting error
(1, 1)	$2.00 \cdot \tilde{\gamma}(1.00; 3.14r^2)$	1.37×10^{-5}
(1, 2)	$0.62 \cdot \tilde{\gamma}(3.01; 3.41r^2) + 0.38 \cdot \tilde{\gamma}(4.07; 3.29r^2)$	2.41×10^{-5}
(2, 2)	$1.00 \cdot \tilde{\gamma}(2.00; 3.14r^2)$	1.51×10^{-5}

■ **Figure 4** Experimentally observed densities of intervals and simplices in the Wrap complexes of a Poisson point process in \mathbb{R}^2 and the corresponding variances, which are normalized by the expected number of points. The table below the panels shows that for the interval densities there are linear combinations of Gamma functions with surprisingly tight fit.



interval	experimental fit	fitting error
(1, 1)	$4.00 \cdot \tilde{\gamma}(1.00; 4.19r^3)$	6.96×10^{-5}
(1, 2)	$1.78 \cdot \tilde{\gamma}(2.99; 5.28r^3) + 0.77 \cdot \tilde{\gamma}(4.48; 3.95r^2)$	7.52×10^{-5}
(2, 2)	$4.85 \cdot \tilde{\gamma}(2.00; 4.19r^3)$	6.68×10^{-5}
(1, 3)	$0.90 \cdot \tilde{\gamma}(4.10; 3.96r^3) + 0.31 \cdot \tilde{\gamma}(6.15; 4.40r^2)$	6.89×10^{-5}
(2, 3)	$0.56 \cdot \tilde{\gamma}(4.11; 6.19r^3) + 2.94 \cdot \tilde{\gamma}(4.98; 4.23r^2) + 0.20 \cdot \tilde{\gamma}(6.34; 7.31r^2)$	8.28×10^{-5}
(3, 3)	$1.85 \cdot \tilde{\gamma}(3.00; 4.19r^3)$	3.13×10^{-5}

■ **Figure 5** Experimentally observed densities of intervals and simplices in the Wrap complexes of a Poisson point process in \mathbb{R}^3 and the corresponding variances, which are normalized by the expected number of points. The table on the *bottom* shows that for the interval densities there are linear combinations of Gamma functions with surprisingly tight fit.

It follows that the corresponding densities for the Wrap complexes lean to the right when compared to the densities for the Alpha complexes. This can indeed be seen by comparing the graphs in the upper left panels of Figures 4 and 5 with the graphs in the left two panels of Figure 2. Similarly, the densities of the simplices in the Wrap complexes lean to the right when compared to the densities of the simplices in the Alpha complexes. To facilitate this comparison, we draw all graphs in the upper right panels of Figures 4 and 5: the solid graphs for the Wrap complexes together with the dashed graphs for the Alpha complexes.

Lower sets. We measure the lower set of a critical simplex in two ways: with its cardinality and with its diameter. The motivation for the two measures is their relevance in computing the Wrap complexes. When we increase the radius threshold and thus add a critical simplex to the Wrap complex, the change can be as large as the cardinality of its lower set. When we study the Wrap complex for data within a window, we need to worry about boundary effects if the lower set of a critical simplex lies only partially within the window. For example, if the diameter of a typical lower set scales linearly with the window side, then we can expect noticeable disturbances of the results due to boundary effects. Fortunately, the diameters are typically small, which we confirmed in several computational experiments. We are, however, lacking any theoretical justification of the findings, which are quantified in Table 2. We chose *p-values* to represent the results. To explain them, consider the second row of the table, for *p-value* 0.05. The first two entries in the row say that the lower sets of 95% = 1 − 0.05 of the critical edges in \mathbb{R}^2 have cardinality at most 1 and the lower sets of a possibly different 95% of the critical edges have diameter at most 1.95. The next two entries give the maximum cardinality and maximum diameter for the 95-percentile of critical triangles in \mathbb{R}^2 , and the rest of the row gives the numbers for simplices in \mathbb{R}^3 . Accordingly, the first row, for *p-value* 0.00, gives the maximum cardinalities and the maximum diameters we observe in the experiment.

<i>p-value</i>	<i>n</i> = 2				<i>n</i> = 3					
	<i>j</i> = 1		<i>j</i> = 2		<i>j</i> = 1		<i>j</i> = 2		<i>j</i> = 3	
0.00	1	5.13	93	5.17	1	3.37	479	3.37	835	3.56
0.05	1	1.95	16	2.39	1	1.79	19	2.03	94	2.20
0.10	1	1.71	13	2.17	1	1.64	13	1.89	69	2.08
0.25	1	1.33	7	1.80	1	1.38	7	1.67	42	1.88
0.50	1	0.94	4	1.42	1	1.10	4	1.43	25	1.65
0.75	1	0.61	4	1.07	1	0.82	4	1.18	16	1.42

■ **Table 2** Each double-column shows the maximum cardinality (*left*) and the maximum diameter (*right*) of the lower sets within the given percentile. The average total number of simplices in the simulation is 540 034 in \mathbb{R}^2 and 6 277 766 in \mathbb{R}^3 . As before, we average over 1 000 experiments in a window of size 300×300 in \mathbb{R}^2 , and over 100 experiments in a window of size $60 \times 60 \times 60$ in \mathbb{R}^3 .

We close this section with an experimental observation that is not reflected in the figures and the table. Recall that for $1 \leq p \leq n - 1$, there are two kinds of *p*-simplices in the Delaunay mosaic: those that give birth and the others that give death. The data in \mathbb{R}^2 and \mathbb{R}^3 suggest that the lower sets of the death-giving *p*-simplices tend to be smaller than those of the birth-giving *p*-simplices. This makes intuitive sense because the death-giving simplices tend to precede the birth-giving simplices and therefore have smaller radii. There seems to be a positive correlation between radius and size of a lower set, which supports the conjecture.

Delaunay mosaic. Indeed, the further a point is from \mathbb{R}^k , the higher its chance to be redundant. Since sampling in \mathbb{R}^n is costly, we restrict our attention to a tube neighborhood, $\mathbb{R}^k \times r \cdot \mathbb{B}^{n-k} \subseteq \mathbb{R}^n$. Sampling the points with density $\rho = 1$ inside the tube, we aim at choosing the radius to maximize the fraction of non-redundant points while minimizing the risk of missing a non-redundant point. To study this question, we sample points in \mathbb{R}^n and record the distances of the non-redundant points from \mathbb{R}^k . Table 3 presents the results for $k = 1, 2$ and $k + 1 \leq n \leq 10$. The respective top rows give the maximum distance of a

p -value	$n = 2$	3	4	5	6	7	8	9	10
$k = 1$									
0.00	2.30	1.60	1.40	1.24	1.22	1.20	1.19	1.20	1.20
0.05	0.88	0.88	0.88	0.90	0.92	0.95	0.97	1.00	1.03
0.25	0.54	0.66	0.72	0.77	0.81	0.85	0.89	0.92	0.95
0.50	0.33	0.50	0.60	0.66	0.72	0.77	0.81	0.85	0.89
$k = 2$									
0.00		1.55	1.34	1.25	1.22	1.20	1.20	1.21	1.22
0.05		0.82	0.87	0.90	0.93	0.95	0.98	1.01	1.03
0.25		0.57	0.69	0.76	0.81	0.85	0.89	0.93	0.96
0.50		0.37	0.55	0.64	0.71	0.77	0.81	0.86	0.89

■ **Table 3** The percentile of points sampled in \mathbb{R}^n that are further from \mathbb{R}^k than the given distance but are nevertheless non-redundant. For every dimension n , the points in \mathbb{R}^n , both redundant and non-redundant, that are closer to \mathbb{R}^k than the furthest non-redundant point, comprise less than 10% of all generated points.

non-redundant point from \mathbb{R}^1 and from \mathbb{R}^2 . Reading these rows from left to right, we see that they first decrease and then increase. This effect is caused by the interaction of the average distance, which increases with n , and the variance, which decreases with n . Indeed, for increasing ambient dimension, the distances of non-redundant points get progressively more concentrated around the radius of the ball with unit volume.

Filtering points. In order to minimize the risk of missing a non-redundant point in the experiments, the vast majority of the generated points have to be redundant. However, the high number of redundant points makes the computation prohibitively expensive, despite these points having no influence on the weighted Delaunay mosaic. It is therefore necessary to eliminate the majority of redundant points for which we can guarantee their redundancy prior to computing the weighted Delaunay mosaic.

We use the following observation to eliminate some of the redundant points before generating the Voronoi diagram. Let $k = 1$, p be a point of the process in \mathbb{R}^n , and \bar{p} its projection to \mathbb{R}^1 . Note that p determines a unique hyperplane, P , in \mathbb{R}^n that is orthogonal to \mathbb{R}^1 and contains p . It is straightforward to see, that if there are two points s and t of the process in different half-spaces determined by P , such that both s and t are closer to \bar{p} than p , then every point in \mathbb{R}^1 is closer to either s or t than to p , meaning the Voronoi cell of p does not intersect \mathbb{R}^1 . We can therefore eliminate p from further computations. Similarly, for $k = 2$, there exist two orthogonal hyperplanes, P_1 and P_2 , in \mathbb{R}^n that contain p and are both orthogonal to \mathbb{R}^2 (we can pick any such pair of orthogonal hyperplanes). Note that P_1 and P_2 divide \mathbb{R}^n into four quadrants. If there are points of the process in all four quadrants that are closer to $\bar{p} \in \mathbb{R}^2$ than p to \bar{p} , then we can eliminate p from further computations as well.

In practice, we keep a test subset of already generated points against which we test the redundancy of all other points as described. The points that are marked for elimination are guaranteed to be redundant in the weighted Delaunay mosaic, and we therefore disregard

them for the purposes of computation. However, typically not all redundant points are eliminated in this way, the exact proportion depending on the amount and spatial distribution of points in the test set. We were typically able to eliminate more than 95% of the generated points, which allows for a feasible computation time.

Probabilistic background. Recall formulas (1) and (2), which give the expected numbers of intervals and simplices in an unweighted Poisson–Delaunay mosaic with radius at most r inside a Borel region Ω . The generalizations of these formulas to the weighted case can be found in [25]. They assert the existence of constants $C_{\ell,m}^{k,n}$ such that the expected numbers of intervals and simplices satisfy

$$c_{\ell,m}^{k,n}(r) = \tilde{\gamma}\left(m + 1 - \frac{k}{n}; \rho\nu_n r^n\right) \cdot C_{\ell,m}^{k,n} \cdot \rho^{\frac{k}{n}} \|\Omega\|, \quad (3)$$

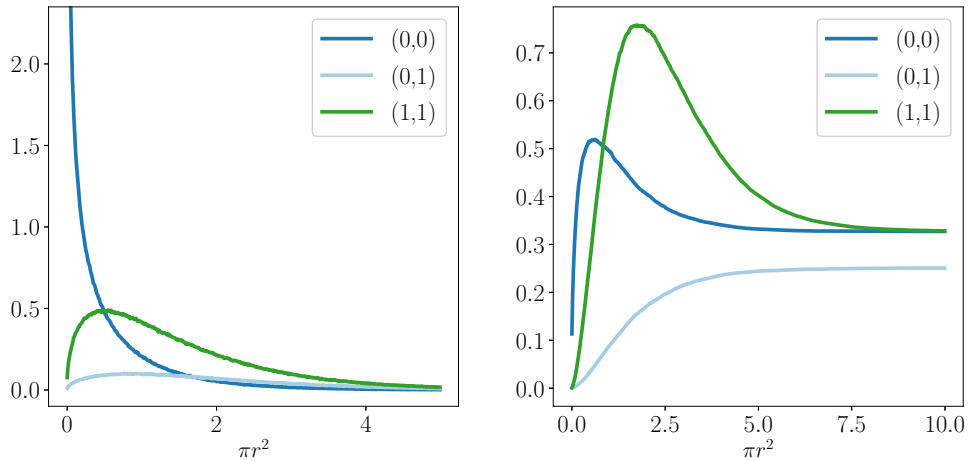
$$d_\ell^{k,n}(r) = \sum_{m=j}^n \tilde{\gamma}\left(m + 1 - \frac{k}{n}; \rho\nu_n r^n\right) \sum_{\ell=1}^j \binom{m-\ell}{m-j} C_{\ell,m}^{k,n} \cdot \rho^{\frac{k}{n}} \|\Omega\|, \quad (4)$$

$$d_j^{k,n}(\infty) = D_j^{k,n} \cdot \rho^{\frac{k}{n}} \|\Omega\|, \quad (5)$$


with $D_j^{k,n} = \sum_{m=j}^n \sum_{\ell=1}^j C_{\ell,m}^{k,n}$, in which $\tilde{\gamma}$ is the regularized lower incomplete Gamma function, and Ω is a Borel region in \mathbb{R}^k . We set $\rho = 1$ in all our experiments. Table 4 gives the constants for $k = 1, 2$ and for a few values of $n \geq k + 1$, each rounded to the nearest two digits after the decimal point. It compares these values with the experimentally estimated values, which are averaged over 10 000 runs, using a segment of length 300 in \mathbb{R}^1 and 1 000 runs, using a square of size 60×60 in \mathbb{R}^2 . We see only small discrepancies between the rounded and the experimentally estimated values, which gives us confidence that the formulas in (3) to (5) are correct.

	$n = 2$		$n = 3$		$n = 5$		$n = 8$		$n = 10$	
	th	exp	th	exp	th	exp	th	exp	th	exp
$k = 1$										
$C_{0,0}^{1,n}$	1.00	1.00	1.09	1.09	1.22	1.22	1.32	1.32	1.37	1.37
$C_{0,1}^{1,n}$	0.27	0.27	0.36	0.36	0.45	0.45	0.51	0.51	0.54	0.54
$C_{1,1}^{1,n}$	1.00	1.00	1.09	1.09	1.22	1.22	1.32	1.32	1.37	1.37
$D_0^{1,n}$	1.27	1.27	1.46	1.45	1.67	1.67	1.84	1.84	1.91	1.91
$D_1^{1,n}$	1.27	1.27	1.46	1.45	1.67	1.67	1.84	1.84	1.91	1.91
$k = 2$										
$C_{0,0}^{2,n}$			1.11	1.11	1.38	1.38	1.66	1.66	1.79	1.79
$C_{0,1}^{2,n}$			0.26	0.26	0.54	0.54	0.77	0.77	0.86	0.87
$C_{1,1}^{2,n}$			2.47	2.48	3.30	3.30	4.09	4.09	4.44	4.44
$C_{0,2}^{2,n}$			0.09	0.09	0.21	0.21	0.31	0.31	0.35	0.35
$C_{1,2}^{2,n}$			1.46	1.46	2.13	2.13	2.74	2.74	3.01	3.01
$C_{2,2}^{2,n}$			1.37	1.37	1.92	1.92	2.43	2.43	2.66	2.66
$D_0^{2,n}$			1.46	1.46	2.13	2.13	2.74	2.74	3.01	3.01
$D_1^{2,n}$			4.37	4.37	6.38	6.38	8.22	8.22	9.03	9.03
$D_2^{2,n}$			2.92	2.92	4.25	4.25	5.48	5.48	6.02	6.02

■ **Table 4** Each double-column shows the theoretically derived values for $C_{\ell,m}^{k,n}$ and $D_j^{k,n}$ on the *left* and the corresponding experimentally computed values on the *right*. In $k = 1$ dimension, we use an interval of length 300 and compute the constants by averaging over 10 000 mosaics. In $k = 2$ dimensions, we use a square of size 60×60 and compute the constants by averaging over 1 000 mosaics.

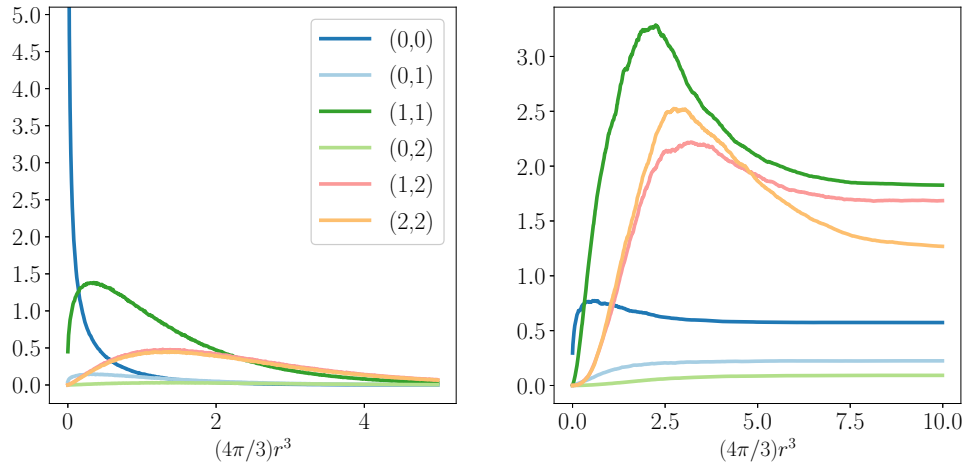


(a) Densities of intervals for $k = 1$ and $n = 2$. (b) Corresponding cumulative variances.




(c) Densities of intervals for $k = 1$ and $n = 10$. (d) Corresponding cumulative variances.

■ **Figure 7** Experimentally observed densities of intervals and corresponding variances of weighted Delaunay mosaics in \mathbb{R}^1 . The mosaic is constructed from an n -dimensional unweighted Delaunay mosaic, with $n = 2$ on the *top* and $n = 10$ on the *bottom*. The graphs are scaled by the expected total number of points in the window. Note that we use the volume of the ball of radius r as the parameter, which also stands for the expected number of points in a ball of this size. With this notation, it is easier to compare the graphs for different ambient dimensions.



(a) Densities of intervals for $k = 2$ and $n = 3$. (b) Corresponding cumulative variances.



(c) Densities of intervals for $k = 2$ and $n = 10$. (d) Corresponding cumulative variances.

■ **Figure 8** Experimentally observed densities of intervals and corresponding variances of weighted Delaunay mosaics in \mathbb{R}^2 . The mosaic is constructed from an n -dimensional unweighted Delaunay mosaic, with $n = 3$ on the *top* and $n = 10$ on the *bottom*. We use the same scaling as in Figure 7.

Intervals, simplices, and Betti numbers. Following Section 2, we count the intervals and simplices in weighted Alpha complexes of Poisson point processes and compare the obtained graphs with the theoretical predictions stated in (3) and (4); see Figures 7 and 8.

Recall that a Boolean model has the same homotopy type as the corresponding weighted Alpha complex. In $k = 1$ dimension, this implies that in the finite case the two have equally many components. In our setting, we consider a Poisson point process in \mathbb{R}^n , with $n \geq 2$, we center a ball of radius r at each point, we intersect the union of balls with a straight line, and we ask for the expected density of components (segments) in this intersection. This is also the density of the 0-th Betti number of the 1-dimensional weighted Alpha complex defined by the same process, the same line, and the same radius. In contrast to the unweighted case, not every vertex of the weighted Delaunay mosaic is critical. To count, we observe that every critical vertex gives birth to a component, every critical edge gives death to a component, and every vertex-edge pair extends a component thus preserving their number. The 0-th Betti number for radius r is therefore the number of critical vertices of radius at most r minus the number of critical edges of radius at most r ; compare the left two panels of Figure 9 for the cases $k = 1$ and $n = 2, 10$. Since the densities of critical vertices and of critical edges are known analytically [25], an expression for the 0-th Betti number follows. The situation in $k \geq 2$ dimensions is more complicated and no analytic results are available; see the right two panels of Figure 9 for the cases $k = 2$ and $n = 3, 10$.

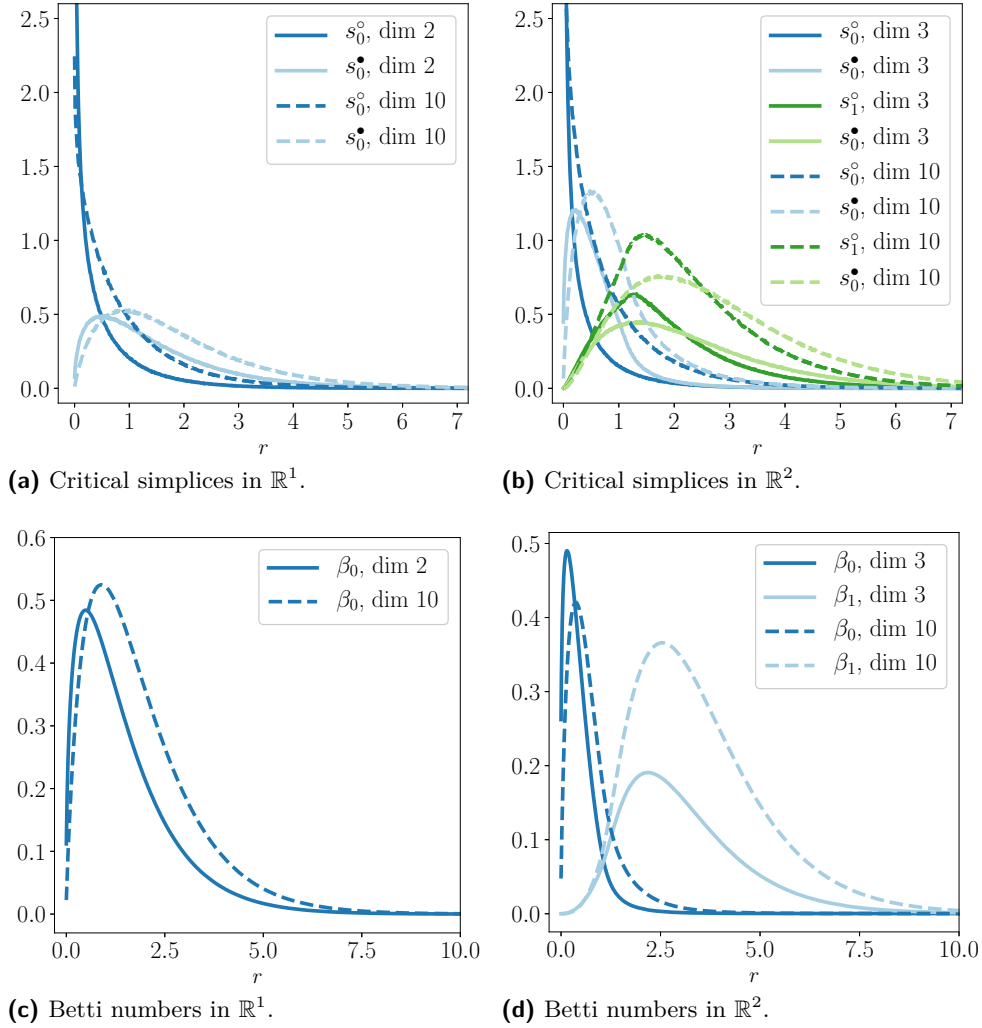
5 Weighted Poisson–Voronoi Tessellations

The motivation for the material collected in this section is the observation that the vertices of a weighted Delaunay mosaic are not distributed as a Poisson point process. Indeed, there is a repulsive force implied by the slice construction that tends to distribute the points in a more regular fashion, and progressively so with growing dimension n ; see Figure 10. We provide statistics to quantify this observation.


Distance to neighbors. As before, we write $Y \subseteq \mathbb{R}^k \times \mathbb{R}$ for a locally finite set of weighted points, which we assume is obtained from a stationary Poisson point process $X \subseteq \mathbb{R}^n$. We recall that the density of edges in the weighted Delaunay mosaic is given in (5). To start, we focus on the distance between adjacent vertices. In $k = 1$ dimension, we measure the distance between contiguous vertices along the real line, drawing the observed distribution for $n = 2, 3, 5, 8, 13, 21, 34$ in the left panel of Figure 11. Extending this to $k = 2$ dimensions, we plot the lengths of the edges in the weighted Delaunay mosaic in the right panel. Both for $k = 1$ and for $k = 2$, the distributions get progressively sharper with increasing n . We quantify this phenomenon by listing the corresponding variances in Table 5.

variance	$n = 2$	3	5	8	13	21	34
distance in \mathbb{R}^1	0.23	0.14	0.09	0.07	0.07	0.06	0.06
distance in \mathbb{R}^2		0.17	0.10	0.08	0.07	0.06	0.06
area in \mathbb{R}^2		0.23	0.18	0.15	0.13	0.11	0.10
length in \mathbb{R}^2		0.13	0.09	0.07	0.06	0.06	0.05

■ **Table 5** The variances of the distances between adjacent vertices in 1- and 2-dimensional weighted Voronoi tessellations in the two *top rows*, and the variances of the area and the edge length in 2-dimensional weighted Voronoi tessellations in the two *bottom rows*. The density of the generating Poisson point process is set to 1 for all n .



■ **Figure 9** Experimentally observed densities of birth- and death-giving critical simplices (*top row*) and Betti numbers (*bottom row*) in the k -dimensional weighted Alpha complexes constructed from a Poisson point process in \mathbb{R}^n . Note that the radii are normalized, so that for ambient dimension n , a simplex with radius r is represented as the volume of a ball with radius r in dimension n . Specifically, given a simplex with radius r , the values on the horizontal axis are πr^2 for $n = 2$, $\frac{4}{3}\pi r^3$ for $n = 3$, and $\frac{1}{120}\pi^5 r^{10}$ for $n = 10$.




■ **Figure 10** Two weighted Poisson–Voronoi tessellations in the plane, with ambient dimension $n = 3$ on the *left* and $n = 10$ on the *right*. In contrast to the graphs in Figures 11 and 12, the processes in \mathbb{R}^n are scaled to achieve an equal density of points in \mathbb{R}^2 . Note that many vertices are located outside the domain they generate.

Geometric size. Rather than measuring the length of edges, we may quantify the vertex distribution by measuring the volume of Voronoi domains. More sensitive to the shape of these domains is a direct comparison of the volume and surface area with that of the round ball. Traditionally, this is formulated as a dimensionless quantity between 0 and 1: the *sphericity* of a body, A , of dimension $m \geq 2$, is

$$\text{Sph}(A) = m\nu_m^{1/m} \frac{\text{Vol}_m(A)^{(m-1)/m}}{\text{Vol}_{m-1}(\text{bd } A)}; \quad (6)$$

see [51]. Note that for the m -dimensional unit ball, we have $\text{Vol}_m(A) = \nu_m$, $\text{Vol}_{m-1}(\text{bd } A) = \sigma_m$, and $\sigma_m/\nu_m = m$, so $\text{Sph}(A) = 1$. Observe also that one over the sphericity is the surface area of A over the surface area of the ball that has the same volume as A , which is the reason that $\text{Sph}(A)$ is sometimes referred to as the *isoperimetric quotient* [37]. Extending the geometric results on the Poisson–Voronoi mosaics from [11, 12, 39] with experimental means, we obtain the graphs in Figure 12. The two top panels show the distributions of the area of a typical polygon and the length of a typical edge in a 2-dimensional weighted Voronoi tessellation; see Table 5 for the corresponding variances, which decrease with increasing n . The two bottom panels show the distributions of the sphericity and the area-weighted sphericity of a typical weighted Voronoi polygon. For the latter measure, the value of the graph at $0 \leq s \leq 1$ is the fraction of the total area covered by polygons with sphericity s . There is no clear trend for the distributions of the sphericity, which, if anything, decrease with increasing n . In contrast, the area-weighted sphericity clearly increases with increasing n . To quantify these trends, we show the average sphericities and the weighted average sphericities in Table 6.

Combinatorial size. Geometric shape information can also be gleaned from the combinatorics of the boundary of a Voronoi domain. In \mathbb{R}^2 , the average number of edges per domain or, equivalently, the average degree of a vertex in the dual mosaic, is 6. It is interesting to observe how the distribution evolves as the ambient dimension, n , increases; see Figure 13. We observe that for $n \geq 5$, the most common degree is 5, which is only possible if sufficiently



■ **Figure 11** The densities of distances between adjacent vertices in weighted Delaunay mosaics: the length of intervals along the line on the *left* and the length of edges of the mosaic in the plane on the *right*. The weighted tessellations are constructed as slices of unweighted Voronoi tessellations in $n = 2, 3, 5, 8, 13, 21, 34$ dimensions.

average	$n = 3$	5	8	13	21	34
sphericity	0.835	0.823	0.817	0.814	0.812	0.811
weighted sphericity	0.879	0.889	0.893	0.896	0.898	0.899

■ **Table 6** The average and area-weighted average sphericity of a typical polygon in a 2-dimensional weighted Voronoi tessellation.

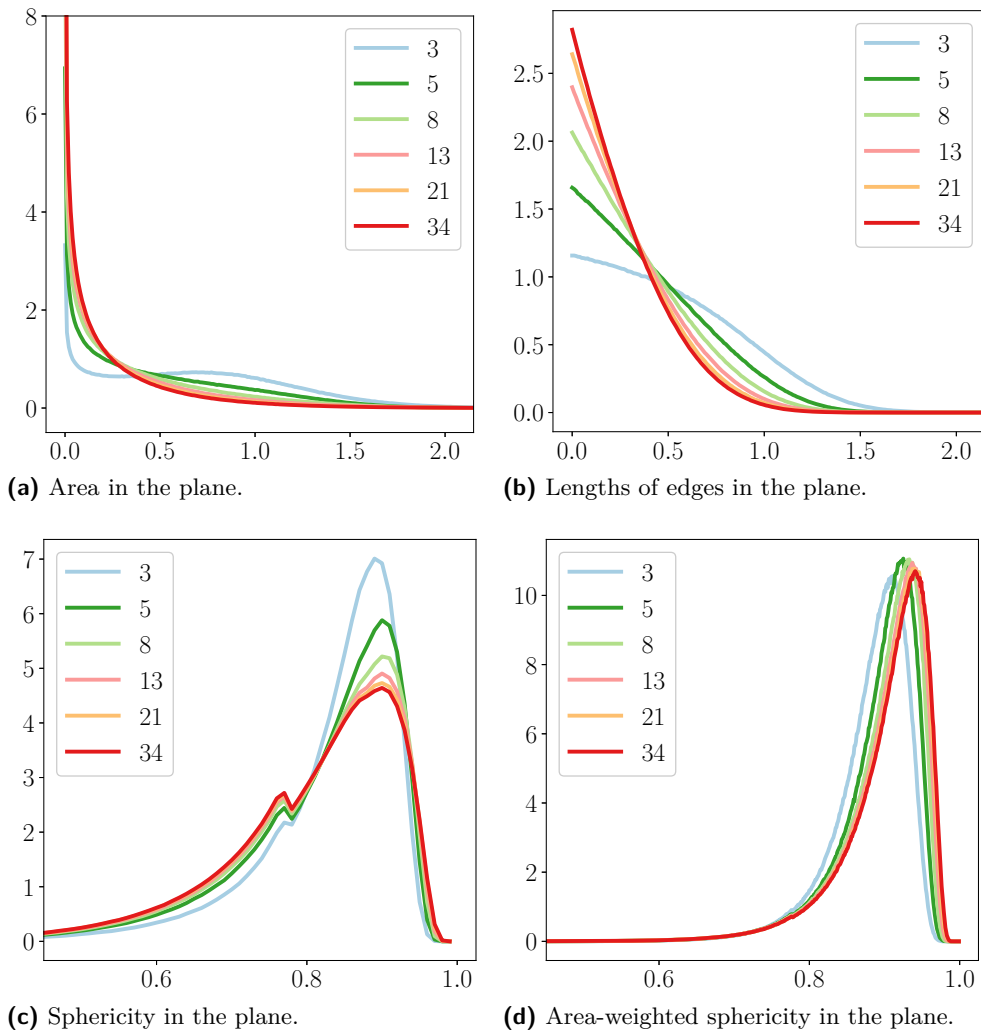
many high-degree vertices bring the average up to almost 6; see Figure 10.

6 Discussion


This work adds an experimental flavor to the mathematical analysis of Poisson–Delaunay mosaics started in [26]. The extension of the initial work on Poisson–Delaunay mosaics to weighted mosaics [25] and to order- k mosaics [23] illustrates the power of a simple idea: to count simplices through intrinsically defined representative points and to group them as dictated by discrete Morse theory. While this idea has been applied to geometric structures beyond those discussed in this paper, it has not yet reached its full potential, which includes the formation of bridges between different areas of stochastic study. For example, the work on random inscribed polytopes in [24] connects Euclidean with Fisher geometry, and the duality between Alpha complexes and Boolean models suggests unexplored connections to percolation theory. We conclude this paper with a list of questions aimed at expanding the scope of the idea and shed light on its limitations.

Variance. Theoretically it is clear that in many settings considered in this paper the variance is negligible compared to the expectation. Nevertheless, the variance distinguishes between the uniform distribution and the Poisson point process on the sphere [48]. Can our experimental results on the variance be complemented with analytic formulas?

Random Wrap complexes. Figure 4 gives experimental evidence that the intervals and simplices in a random Wrap complex follow simple linear combinations of Gamma



■ **Figure 12** *Left and right in the top row:* the densities of the area of a typical polygon and of the length of a typical edge in the weighted Voronoi tessellation. *Left and right in the bottom row:* the densities of the sphericity and the area-weighted sphericity of a typical polygon in the weighted Voronoi tessellation. The tessellations are constructed as slices of unweighted Voronoi tessellations in $n = 2, 3, 5, 8, 13, 21, 34$ dimensions.



■ **Figure 13** The normalized distribution of face degrees in weighted Voronoi tessellations in $k = 2$ dimensions. For example in $n = 3$ dimensions, the proportion of triangular faces is 6.23%, the proportion of quadrangular faces is 13.60% and so on. The tessellations are constructed as slices of unweighted Voronoi tessellations in $n = 3, 5, 8, 13, 21, 34$ dimensions.

distributions. Can this be proved or disproved analytically?

Lower sets. When we sort the critical simplices in the order of non-decreasing radius, the death-giving simplices tend to precede the birth-giving simplices of the same dimension. This suggests that the lower sets of the former type of simplices tend to be smaller than the lower sets of the latter type. Can this be proven analytically?

Constants in weighted Delaunay mosaics. We have the constants $C_{\ell,m}^{k,n}$ for $k = 1, 2$ and all corresponding combinations of ℓ , m , and n [25] and also for $k = n = 3$ and all corresponding combinations of ℓ and m [26]. Can the methods in these two papers be extended to compute all constants for $k = 3$?

Probabilistic relations. Our experiments suggest that for a stationary Poisson point process in \mathbb{R}^n , we have $C_{2,2}^{2,n} = C_{0,0}^{2,n} + C_{0,1}^{2,n}$ for all $n \geq 2$; see Table 4. Equivalently, we have $C_{1,2}^{2,n} = C_{0,0}^{2,n} + C_{0,1}^{2,n} + C_{0,2}^{2,n}$. Prove these relations and find additional ones that hold for Poisson–Delaunay mosaics but not for more general Delaunay mosaics.

Neighborhood size. Given a typical vertex of degree j in a Poisson–Delaunay mosaic in \mathbb{R}^3 , Aboav [1] and Weaire [52] study the expected degree of a neighbor. It is plausible that this function decreases with increasing j , but [35] provide experimental evidence that it increases for small j , reaches its maximum for $j = 12$, and then decreases. How does the function behave for weighted Delaunay mosaics and for dimensions k beyond 3?

Slice construction. Fixing k and letting n go to infinity, does the weighted Voronoi tessellation obtained by taking a k -dimensional slice of the unweighted Voronoi tessellation of a Poisson point process in \mathbb{R}^n approach a limiting distribution of domains? If yes, in what sense?

A common framework for many of the above questions is provided by the persistence diagram of the radius function [19, 21]. Each point in this diagram is determined by two critical simplices, so the density of the latter is twice the density of the points in the diagram. Experiments suggest that for Poisson point processes the persistence diagrams have characteristic shapes. An interesting quantity is the difference in radii between two

matched critical simplices, whose maximum was studied in [8]. The points in this diagram carry topological and therefore global meaning about the data. This makes questions about the distribution of the points difficult to approach with tools from probability theory, and it is not surprising that most of the interesting questions are yet untouched.

References

- 1 Aboav, D. (1970). The arrangement of grains in a polycrystal. *Metallography*, **3**:383–390.
- 2 Aurenhammer, F. (1987). Power diagrams: properties, algorithms, and applications. *SIAM J. Comput.*, **16**:93–105.
- 3 Aurenhammer, F. (1991). Voronoi diagrams — a survey of a fundamental geometric data structure. *ACM Comput. Surveys*, **23**:345–405.
- 4 Aurenhammer, F. and Imai, H. (1988). Geometric relations among Voronoi diagrams. *Geom. Dedicata*, **27**:65–75.
- 5 Bauer, U. and Edelsbrunner, H. (2017). The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.*, **369**:3741–3762.
- 6 Baumstark, V. and Last, G. (2009). Gamma distributions for stationary Poisson flat processes. *Adv. Appl. Prob.*, **41**:911–939.
- 7 Bobrowski, O. and Adler, R. (2014). Distance functions, critical points, and topology for some random complexes. *Homol. Homotopy Appl.*, **16**:311–344.
- 8 Bobrowski, O., Kahle, M., and Skraba, P. (2017). Maximally persistent cycles in random geometric complexes. *Ann. Appl. Probab.*, **4**:2032–2060.
- 9 Bobrowski, O. and Weinberger, S. (2017). On the vanishing of homology in random Čech complexes. *Random Structures Algorithms*, **51**:14–51.
- 10 Burtseva, L., Werner, F., Valdez, B., Pestriakov, A., Romero, R., and Petranovskii, V. (2015). Tessellation methods for modeling material structure. *Appl. Math. Mech.*, **756**:426–435.
- 11 Calka, P. (2003a). An explicit expression for the distribution of the number of sides of the typical Poisson–Voronoi cell. *Adv. Appl. Prob. (SGSA)*, **35**:863–970.
- 12 Calka, P. (2003b). Precise formulae for the distributions of the principal geometric characteristics of the typical cells of a two-dimensional Poisson–Voronoi tessellation and a Poisson line process. *Adv. Appl. Prob. (SGSA)*, **35**:551–562.
- 13 Carlsson, G. (2009). Topology and data. *Bull. Amer. Math.*, **46**:255–308.
- 14 Crofton, M. W. (1868). On the theory of local probability, applied to straight lines drawn at random in the plane: the methods used being also extended to the proof of certain new theorems in the integral calculus. *Phil. Trans. Royal Soc. London*, **158**:181–199.
- 15 Decreusefond, L., Ferraz, E., Randriam, H., and Vergne, A. (2014). Simplicial homology of random configurations. *Adv. Appl. Prob.*, **46**:1–23.
- 16 Delaunay, B. (1934). Sur la sphère vide. A la mémoire de Georges Voronoï. *Bulletin de l'Académie des Sciences de l'URSS. Classe des sciences mathématiques et na*, 793–800.
- 17 Dey, T. K. (2011). *Curve and Surface Reconstruction. Algorithms with Mathematical Analysis*. Cambridge Univ. Press, England.
- 18 Edelsbrunner, H. (2003). Surface reconstruction by wrapping finite sets in space. In *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, ed.: Aronov, B., Basu, S., Pach, J., and Sharir, M., 379–404. Springer, Heidelberg.
- 19 Edelsbrunner, H. and Harer, J. L. (2010). *Computational Topology. An Introduction*. Amer. Math. Soc., Providence, Rhode Island.
- 20 Edelsbrunner, H., Kirkpatrick, D. G., and Seidel, R. (1983). On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, **IT-29**:551–559.
- 21 Edelsbrunner, H., Letscher, D., and Zomorodian, A. J. (2002). Topological persistence and simplification. *Discrete Comput. Geom.*, **28**:511–533.
- 22 Edelsbrunner, H. and Mücke, E. P. (1994). Three-dimensional alpha shapes. *ACM Trans. Graphics*, **13**:43–72.


- 23 Edelsbrunner, H. and Nikitenko, A. (2018a). Poisson–Delaunay mosaics of order k . *Discrete Comput. Geom.*, under review.
- 24 Edelsbrunner, H. and Nikitenko, A. (2018b). Random inscribed polytopes have similar radius functions as Poisson–Delaunay mosaics. *Ann. Appl. Probab.*, **28** (2018): 3215–3238.
- 25 Edelsbrunner, H. and Nikitenko, A. (2018c). Weighted Poisson–Delaunay mosaics. In preparation.
- 26 Edelsbrunner, H., Nikitenko, A., and Reitzner, M. (2017). Expected sizes of Poisson–Delaunay mosaics and their discrete Morse functions. *Adv. Appl. Prob.*, **49**:745–767.
- 27 Forman, R. (1998). Morse theory for cell complexes. *Adv. Math.*, **134**:90–145.
- 28 Freij, R. (2009). Equivariant discrete Morse theory. *Discrete Math.*, **309**:3821–3829.
- 29 Hatcher, A. (2002). *Algebraic Topology*. Cambridge Univ. Press, England.
- 30 Kahle, M. (2011). Random geometric complexes. *Discrete Comput. Geom.*, **45**:553–573.
- 31 Kahle, M. (2014). Topology of random simplicial complexes: a survey. *AMS Contemp. Math.*, **620**:201–222.
- 32 Kingman, J. F. C. (1993). *Poisson Processes*. Oxford Univ. Press, Oxford, England.
- 33 Lautensack, C. (2007). *Random Laguerre Tessellations*. PhD thesis, Math. Dept., Univ. Karlsruhe, Germany.
- 34 Lautensack, C. and Zuyev, S. (2008). Random Laguerre tessellations. *Adv. Appl. Prob. (SGSA)*, **40**:630–650.
- 35 Lazar, E. A., Mason, J. K., MacPherson, R. D., and Srolovitz, D. J. (2013). Statistical topology of three-dimensional Poisson–Voronoi cells and cell boundary networks. *Phys. Rev. E*, **88**:063309.
- 36 Leray, J. (1945). Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl.*, **24**:95–167.
- 37 Lutz, F. H., Mason, J. K., Lazar, E. A., and MacPherson, R. D. (2017). Roundness of grains in cellular microstructures. *Phys. Rev. E*, **96**:023001.
- 38 Meijering, J. (1953). Interface area, edge length, and number of vertices in crystal aggregates with random nucleation. *Philips Res. Rep.*, **8**:270–290.
- 39 Miles, R. E. (1970). On the homogeneous planar Poisson point process. *Math. Biosci.*, **6**:85–127.
- 40 Miles, R. E. (1971). Isotropic random simplices. *Adv. Appl. Prob.*, **3**:353–382.
- 41 Milnor, J. (1963). *Morse Theory*. Annals of Math. Studies, Princeton Univ. Press, New Jersey.
- 42 Møller, J. (1989). Random tessellations in \mathbb{R}^d . *Adv. Appl. Prob.*, **21**:37–73.
- 43 Nikitenko, A. (2017). *Discrete Morse Theory for Random Complexes*. PhD thesis, IST Austria, Klosterneuburg, Austria.
- 44 Okabe, A., Boots, B., Sugihara, K., Chiu, S. N., and Kendall, D. G. (2008). *Spatial Tessellations*. John Wiley & Sons, Inc., second edition.
- 45 Prunet, N. and Meyerowitz, E. M. (2016). Genetics and plant development. *C.R. Biologies*, **339**:240–246.
- 46 Schneider, R. and Weil, W. (2008). *Stochastic and Integral Geometry*. Springer, Berlin, Germany.
- 47 Sibson, R. (1980). A vector identity for the Dirichlet tessellation. *Math. Proc. Camb. Phil. Soc.*, **87**:151–155.
- 48 Stemeseder, J. (2014). *Random Polytopes with Vertices on the Sphere*. PhD thesis, Math. Dept., Univ. Salzburg, Austria.

- 49** Voronoi, G. (1907). Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier Mémoire: Sur quelques propriétés des formes quadratiques positive parfaites. *J. Reine Angew. Math.*, **133**:97–178.
- 50** Voronoi, G. (1908). Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième Mémoire: Recherches sur les paralléloèdres primitifs. *J. Reine Angew. Math.*, **134**:198–287.
- 51** Wadell, H. (1935). Volume, shape and roundness of quartz particles. *The Journal of Geology*, **43**:250–280.
- 52** Weaire, D. (1974). Some remarks on the arrangement of grains in polycrystals. *Metallography*, **7**:157–160.
- 53** Wendel, J. G. (1962). A problem in geometric probability. *Math. Scand.*, **11**:109–111.
- 54** The CGAL Project (2017). *CGAL User and Reference Manual*. CGAL Editorial Board, 4.10 edition.

A Voronoi Tessellations and Delaunay Mosaics

In this appendix, we present the definitions and basic properties of Voronoi tessellations, Delaunay mosaics, Alpha complexes, and their weighted variants.

The unweighted case. Let X be a locally finite set in \mathbb{R}^n . The *Voronoi domain* of $x \in X$ consists of all points for which x minimizes the Euclidean distance: $\text{dom}(x) = \{a \in \mathbb{R}^n \mid \|a - x\| \leq \|a - y\| \text{ for all } y \in X\}$. Every Voronoi domain is a closed, possibly unbounded, full-dimensional, convex polyhedron. The *Voronoi tessellation* of X is the collection of Voronoi domains of points in X [49, 50]. As illustrated in Figure 14, the domains in the tessellation have disjoint interiors, they possibly overlap along shared faces, and they cover the entire \mathbb{R}^n . We refer to [44] for a discussion of this concept and of its many variants. The *Delaunay mosaic* is the dual of the Voronoi tessellation [16], and it may be viewed geometrically, as a collection of cells, or combinatorially, as a collection of subsets of X . We prefer the latter view and define $\text{Del}(X)$ as the collection of maximal subsets $Q \subseteq X$ such that the Voronoi domain of the points in Q have a given non-empty intersection: $\text{dom}(Q) = \bigcap_{x \in Q} \text{dom}(x) \neq \emptyset$ and $\text{dom}(Q) \neq \text{dom}(R)$ for every proper superset R of Q . For example, if X is the set of 4 vertices of a square in \mathbb{R}^2 , then $\text{Del}(X)$ consists of four singletons, four pairs, and one quadruplet. On the other hand, if X is in general position — which in the plane includes that no 4 points lie on a common circle — then $\text{Del}(X)$ is a simplicial complex, which in the combinatorial setting means that $Q \in \text{Del}(X)$ implies that every subset of Q belongs to $\text{Del}(X)$. More specifically, in this case $\text{Del}(X)$ is isomorphic to the nerve of the Voronoi tessellation and obtained by mapping every Voronoi domain to its generating point. The Delaunay mosaic has a natural geometric realization in \mathbb{R}^n , which we obtain by mapping every $Q \in \text{Del}(X)$ to the convex hull of Q . This is the geometric view of the Delaunay mosaic.




■ **Figure 14** The Voronoi tessellation decomposes the plane as well as the union of disks into convex regions. The edges and vertices of the Delaunay mosaic are superimposed, and the Alpha subcomplex is shown by thickening its edges and shading its triangles.

Now fix $r > 0$, write $B_r(x)$ for the closed ball with center x and radius r , and let $X_r = \bigcup_{x \in X} B_r(x)$ be the union of these balls. Clipping each ball to within the corresponding Voronoi domain gives us a convex decomposition of the union: $X_r = \bigcup_{x \in X} [B_r(x) \cap \text{dom}(x)]$; see Figure 14. The *Alpha complex* of X and r , denoted $\text{Alpha}_r(X)$, consists of all cells in

the Delaunay mosaic whose clipped Voronoi domains have a non-empty common intersection [20, 22]; see again Figure 14. By the Nerve Theorem of algebraic topology [36], X_r and $\text{Alpha}_r(X)$ have the same homotopy type; see also [19, 29]. Similarly, X_r has the same homotopy type as the Čech complex of X and r , which by definition is the nerve of the balls $B_r(x)$. This is the justification why in topological data analysis the two complexes are often used as proxies for the union of balls [13, 22].

The Delaunay mosaic can also be defined directly, without the introduction of Voronoi tessellations. Call an $(n-1)$ -sphere S in \mathbb{R}^n *empty* if all points of X lie either on or outside S . Then $Q \subseteq X$ belongs to $\text{Del}(X)$ iff there is an empty sphere S with $Q = X \cap S$. Indeed, every such sphere has its center in the interior of $\text{dom}(Q)$, and every point in the interior of $\text{dom}(Q)$ is the center of such a sphere. Furthermore, every point in $\text{dom}(Q)$ is the center of an empty sphere that passes through all points of Q and possibly through some additional points of X . This slightly larger set contains a unique smallest sphere, which we call the *Delaunay sphere* of Q . The *Delaunay radius function*, $\mathcal{R}: \text{Del}(X) \rightarrow \mathbb{R}$, maps every Q to the radius of its Delaunay sphere. Its sublevel sets are the Alpha complexes: $\text{Alpha}_r(X) = \mathcal{R}^{-1}[0, r]$. Observe that \mathcal{R} is *non-decreasing*, by which we mean that $P, Q \in \text{Del}(X)$ and $P \subseteq Q$ implies $\mathcal{R}(P) \leq \mathcal{R}(Q)$. Hence, $Q \in \text{Alpha}_r(X)$ implies that all faces of Q also belong to $\text{Alpha}_r(X)$. This shows that $\text{Alpha}_r(X)$ is indeed a subcomplex of $\text{Del}(X)$ and not just a subset. We refer to Appendix B for the topological significance of \mathcal{R} .

The weighted case. In the weighted setting, we use real weights to control the influence a point has on its surrounding. The extra degree of freedom permits better approximations of observed space decompositions, such as cell cultures in plants [45] and microstructures of materials [10], to name just two. For ease of distinction from the unweighted case, we write k for the dimension of the space, we write $Y \subseteq \mathbb{R}^k \times \mathbb{R}$ for a locally finite set, and for each $(y, w_y) \in Y$ we call y the *point* and w_y the *weight*. We will however frequently abuse notation and write $y \in Y$ instead. Let $\pi_y: \mathbb{R}^k \rightarrow \mathbb{R}$ be defined by $\pi_y(a) = \|a - y\|^2 - w_y$ and call $\pi_y(a)$ the *power distance* of a from y . For $w_y = 0$, the power distance equals the squared Euclidean distance from y . For positive weight, the weighted point is conveniently visualized by drawing the zero-set of π_y , which is the sphere with center y and radius $\sqrt{w_y}$.



■ **Figure 15** The weighted Voronoi tessellation decomposes the plane as well as the union of disks into convex regions. The edges and vertices of the weighted Delaunay mosaic are superimposed, and the weighted Alpha subcomplex is shown by thickening its edges and shading its triangles.

The *weighted Voronoi domain* of y consists of all points for which y minimizes the power distance: $\text{dom}(y) = \{a \in \mathbb{R}^k \mid \pi_y(a) \leq \pi_z(a) \text{ for all } z \in Y\}$. Similar to the unweighted case, $\text{dom}(y)$ is a closed, possibly unbounded, convex polyhedron, but it is not necessarily full-dimensional. Indeed, if y is the only point with non-zero weight in Y , then its domain is larger than in the unweighted case if $w_y > 0$ and smaller if $w_y < 0$. The latter case includes the possibility that $\text{dom}(y)$ is empty. The *weighted Voronoi tessellation* of Y is the collection of weighted Voronoi domains; see Figure 15. The *weighted Delaunay mosaic* of Y , denoted $\text{Del}(Y)$, is the dual of the weighted Voronoi tessellation, which we define as the maximal subsets $Q \subseteq Y$ such that the weighted Voronoi domains of the points in Q have a given non-empty intersection: $\text{dom}(Q) = \bigcap_{y \in Q} \text{dom}(y) \neq \emptyset$ and $\text{dom}(Q) \neq \text{dom}(R)$ for every proper superset R of Q . Since weighted Voronoi domains can be empty, the vertex set of $\text{Del}(Y)$ is a subset and not necessarily the entire set Y . Assuming the weighted points are in general position — for an appropriate definition of this notion — the weighted Delaunay mosaic is again isomorphic to the nerve of the weighted Voronoi tessellation, and it can be geometrically realized by mapping each domain to the generating point; see Figure 15.

Now fix $r \in \mathbb{R}$, write $B_r(y)$ for the set of points that satisfy $\pi_y(a) \leq r^2$, and let $Y_r = \bigcup_{y \in Y} B_r(y)$ be the union of these balls. Clipping the balls, we get again a convex decomposition: $Y_r = \bigcup_{y \in Y} [B_r(y) \cap \text{dom}(y)]$; see Figure 15. The *weighted Alpha complex* of Y and r , denoted $\text{Alpha}_r(Y)$, is again the dual of these clipped domains. The Nerve Theorem still applies, so Y_r , $\text{Alpha}_r(Y)$, and the nerve of the $B_r(y)$ all have the same homotopy type. Observe that the weighted Voronoi tessellation and the weighted Delaunay mosaic do not change if we add the same constant value to the weight of every point in Y . In contrast, this operation generally changes the weighted Alpha complex as well as the nerve of the balls.

We now present an alternative description of the weighted concepts that has the advantage that the points and their weights can be selected in a single process. Let $k < n$ and write \mathbb{R}^k for the space spanned by the first k coordinate vectors of \mathbb{R}^n . For a point $y \in \mathbb{R}^k$ with weight $w_y \leq 0$, let $x = x(y)$ be a point in \mathbb{R}^n whose orthogonal projection to \mathbb{R}^k is y and whose distance from \mathbb{R}^k is $\|x - y\| = \sqrt{-w_y}$. Then $\pi_y(a) = \|a - x\|^2$ for every point $a \in \mathbb{R}^k$. In other words, if every point y in $Y \subseteq \mathbb{R}^k \times \mathbb{R}$ has non-positive weight, we can find a set of unweighted points, $X \subseteq \mathbb{R}^n$, such that the weighted Voronoi tessellation of Y is the intersection of \mathbb{R}^k with the (unweighted) Voronoi tessellation of X ; see [4, 47]. Similarly, we can construct $\text{Del}(Y)$ from $\text{Del}(X)$, which we explain by calling an $(n - 1)$ -sphere in \mathbb{R}^n *anchored* if its center lies in \mathbb{R}^k . Recall that $Q \subseteq X$ belongs to $\text{Del}(X)$ iff there is an empty sphere S with $Q = X \cap S$. By adding the requirement that S be anchored, we get exactly the cells $Q \in \text{Del}(X)$ whose projections to \mathbb{R}^k belong to $\text{Del}(Y)$. Similarly, we call the smallest empty anchored sphere that passes through all points of Q the *anchored Delaunay sphere* of Q . Accordingly, the *weighted Delaunay radius function*, $\mathcal{R}: \text{Del}(Y) \rightarrow \mathbb{R}$, maps every cell that is the projection of $Q \in \text{Del}(X)$ to the radius of the anchored Delaunay sphere of Q . Finally, the *weighted Alpha complex* of Y and r is the sublevel set of this function: $\text{Alpha}_r(Y) = \mathcal{R}^{-1}[0, r]$; see again Figure 15. We note that the above construction is predicated on the assumption that all weights are non-positive, but this is not a limitation of generality since we can add a constant value to all weights without changing the tessellation and the mosaic.

The Crofton connection. There is a direct connection between the density of top-dimensional simplices in a k -dimensional weighted Delaunay mosaic and the density of the $(n - k)$ -dimensional skeleton of the n -dimensional Voronoi tessellation from which the mosaic is obtained. This connection is the classic Crofton Formula of integral geometry [14, 46], as we now explain. Observe that for a stationary Poisson point process in \mathbb{R}^n , the statistics

of the k -dimensional weighted Delaunay mosaic does not depend on the specific k -plane we use. We can therefore integrate over all k -planes. Counting the top-dimensional simplices in the k -dimensional weighted Delaunay mosaics is the same as counting the vertices in the k -dimensional weighted Voronoi tessellations. By Crofton's Formula, this number integrates to a constant times the volume of the $(n - k)$ -skeleton of the unweighted n -dimensional tessellation.

As it turns out, it is relatively easy to compute the expected density of the $(n - k)$ -skeleton, so we use the Crofton Formula backward, deriving the expected density of top-dimensional simplices in a weighted Poisson–Delaunay mosaic. This density implies the expected density of $(k - 1)$ -simplices, and for $k = 2$, the expected density of vertices in the weighted mosaic. To get similar results for the intervals, we need traditional integral geometric methods, as described in [25].

B Discrete Morse Theory and Homology

In the generic case, the radius function on a weighted or unweighted Delaunay mosaic satisfies the axioms of a generalized discrete Morse function. This motivates us to define these functions and discuss their basic properties in this second appendix. In addition, we introduce homology groups and Betti numbers, which we use to state the discrete counterparts of the classic Morse inequalities.

Generalized discrete Morse theory. The inspiration stems from classical Morse theory [41], which studies manifolds through the behavior of generic smooth functions on them. There are several ways to transport the smooth theory to the piecewise linear setting, and we follow the more radical proposal by Forman [27] that formulates the discrete theory in combinatorial terms entirely. More precisely, we present a slightly generalized version of the original theory [28].

Let K be a simplicial complex, and recall that the *face lattice* of K is the partial order on its simplices, in which $P \leq Q$ if P is a face of Q . The *Hasse diagram* is the reduced graph representation of the face lattice, in which we draw an edge from P to Q if $P \leq Q$ and $\dim P = \dim Q - 1$. An *interval* is a maximal collection of simplices that have a common *lower bound* and a common *upper bound*: $[P, R] = \{Q \in K \mid P \leq Q \leq R\}$. We call a function $f: K \rightarrow \mathbb{R}$ *non-decreasing* if $P \leq Q$ implies $f(P) \leq f(Q)$. A *level set*, $f^{-1}(r)$, is a subset of K , and a *step* of f is a maximal subset of $f^{-1}(r)$ that induces a connected subgraph of the Hasse diagram. Finally, a non-decreasing f is a *generalized discrete Morse function* if every step is an interval. Its *generalized discrete gradient* is the corresponding partition of K into intervals. We remark that it is often possible to assume that each level set is an interval, but sometimes such an assumption seems unnatural, which is our motivation to introduce the notion of a step.

The intervals are combinatorially simple. Indeed, if $\ell = \dim P$ and $m = \dim R$, then $[P, R]$ contains $2^{m-\ell}$ simplices, namely $\binom{m-\ell}{m-j}$ simplices of dimension j , for $\ell \leq j \leq m$. We call (ℓ, m) the *type* of the interval. Knowing the number of intervals, of each type, it is therefore easy to compute the number of simplices in the complex. Writing $c_{\ell, m}$ for the number of intervals of type (ℓ, m) , the number of j -simplices in K is

$$s_j = \sum_{\ell=0}^j \sum_{m=j}^{\infty} \binom{m-\ell}{m-j} c_{\ell, m}. \quad (7)$$

The difference between the original discrete Morse theory of Forman [27] and the generalization proposed in [28] is that the former limits its steps to intervals of size 1 and 2, while the

latter permits intervals of all possible sizes. The intervals of size 1 play a special role, so we call such an interval *singular*, the simplex it contains *critical*, and the function value of this simplex a *critical value* of f . Their special role is best appreciated by considering the construction of K by adding the simplices in increasing order of function values. When we add the simplices of a non-singular interval, then the inverse operation can be realized as a deformation retraction, which implies that that complex has the same homotopy type before the operation as after the operation; see [19, 29] for background on these concepts. Writing K_i and $K_{i+1} = K_i \cup [P, R]$ for the two complexes, we say K_{i+1} *collapses* to K_i , denoted $K_{i+1} \searrow K_i$, and we call the addition of $[P, R]$ an *anticollapse*. The following lemma is analogous to the classical theorem about the retractability of sublevel sets in smooth Morse theory. Write $K_r = f^{-1}(-\infty, r]$ for the subcomplex that consists of all simplices which value at most r .

► **Lemma 1** (Collapsibility [27, 28]). *Let $f: K \rightarrow \mathbb{R}$ be a generalized discrete Morse function on a simplicial complex. If the half-open interval $(r, s] \subseteq \mathbb{R}$ contains no critical value of f , then $K_s \searrow K_r$.*

As mentioned earlier, $K_s \searrow K_r$ implies that the two complexes have the same homotopy type. Indeed, a stronger statement is implied: there is a CW-complex whose cells are in bijection with the critical simplices of f such that the subcomplex of cells whose critical simplices have function value at most r has the same homotopy type as K_r , for every $r \in \mathbb{R}$. This CW-complex is called the *Morse complex* of f .

Homology and Morse inequalities. We use the language of homology to talk about how a space or a complex is connected. The comprehensive introduction of this formalism is beyond the scope of this paper, and we refer to [19, 29] for further background. However, if we limit ourselves to $\mathbb{Z}/2\mathbb{Z}$ coefficients — which amounts to using modulo-2 arithmetic — then homology groups can be explained in purely combinatorial terms, as we now do. A p -*chain* is a collection of p -simplices in K , and the *sum* of two p -chains is the symmetric difference of the two collections. The *boundary* of a p -simplex is the collection of its $(p-1)$ -faces, and the *boundary* of a p -chain is the sum of the boundaries of its simplices. Writing C_p for the group of p -chains and $\partial_p: C_p \rightarrow C_{p-1}$ for the p -th boundary operator, we get the *chain complex*, $\dots \rightarrow C_p \rightarrow C_{p-1} \rightarrow \dots$. A p -*cycle* is a p -chain with empty boundary, $z \in \ker \partial_p$, and a p -*boundary* is the boundary of a $(p+1)$ -chain, $b \in \text{img } \partial_{p+1}$. The boundaries and cycles form subgroups of each other and of the chain group: $\text{img } \partial_{p+1} \subseteq \ker \partial_p \subseteq C_p$. We finally get the p -*th homology group* by calling cycles that differ by a boundary equivalent: $H_p = \ker \partial_p / \text{img } \partial_{p+1}$. Its *rank* is the p -*th Betti number*, which for $\mathbb{Z}/2\mathbb{Z}$ coefficients is the binary logarithm of the order: $\beta_p = \log_2 |H_p|$.

While we define Betti numbers in terms of the simplices in K , they are in fact independent of how we triangulate a space. More generally even, Betti numbers can be defined without triangulation, and two spaces or complexes that have the same homotopy type also have the same Betti numbers. For example, if K is a complex in \mathbb{R}^2 , then $\beta_0(K)$ is the number of components and $\beta_1(K)$ is the number of holes. Similarly, if K is a complex in \mathbb{R}^3 , then $\beta_0(K)$, $\beta_1(K)$, and $\beta_2(K)$ are its numbers of components, tunnels, and voids. Indeed, these are the only possibly non-zero Betti numbers of complexes in \mathbb{R}^2 and \mathbb{R}^3 .


As mentioned earlier, non-singular intervals of a generalized discrete Morse function preserve the homotopy type and therefore the Betti numbers. Therefore, there must be a connection between the Betti numbers and the critical simplices, as we now explain. To get started, let s_p be the number of p -simplices in K , and define the *Euler characteristic* as the alternating sum of simplex numbers: $\chi = \sum_p (-1)^p s_p$. If we construct K one interval at a time,

every non-singular interval preserves χ simply because $\sum_{p=\ell}^m (-1)^{p-\ell} \binom{m-\ell}{m-p} = (1-1)^{m-\ell} = 0$. On the other hand, if we add a critical p -simplex, the Euler characteristic changes by $(-1)^p$. Writing c_p for the number of critical p -simplices, we thus get $\chi = \sum_p (-1)^p c_p$. We extend this relation to Betti numbers by observing that $s_p = s_p^\circ + s_p^\bullet$, in which s_p° counts the critical simplices that give *birth* to p -cycles, and s_p^\bullet counts the critical simplices that give *death* to $(p-1)$ -cycles. In other words, when we add a critical simplex to the complex, either β_p increases by 1, or β_{p-1} decreases by 1. Either way, χ changes by $(-1)^p$, which implies the *discrete Euler–Poincaré formula*: $\chi = \sum_p (-1)^p \beta_p$. The implied relation between the alternating sum of critical simplex numbers and of Betti numbers is traditionally stated as a strengthening of the last in a sequence of (*strong*) *discrete Morse inequalities*:

$$\sum_{p=0}^q (-1)^{q-p} \beta_p \leq \sum_{p=0}^q (-1)^{q-p} c_p, \quad (8)$$

for all $q \geq 0$. For example, the first inequality asserts $\beta_0 \leq c_0$, while the second inequality asserts $\beta_1 - \beta_0 \leq c_1 - c_0$, which implies $\beta_1 \leq c_1$. Indeed, we get $\beta_q \leq c_q$, for all $q \geq 0$, which are sometimes referred to as the *weak discrete Morse inequalities*.

The Delaunay setting. The prime example of generalized discrete Morse functions in this paper are the radius functions of weighted and unweighted Delaunay mosaics [5]. We recall that in the unweighted case, $\mathcal{R}(Q)$ is the radius of the smallest empty sphere that passes through all points of Q , and in the weighted case, it is the radius of the smallest empty anchored sphere that passes through the preimages of all weighted points in Q . We consider first the unweighted case and let $X \subseteq \mathbb{R}^n$ be locally finite and in general position. Let $R \in \text{Del}(X)$ and recall that $\text{dom}(R)$ is the intersection of the Voronoi domains of the points in R . The center of the smallest empty circumsphere of R is the point $a \in \text{dom}(R)$ that is closest to the affine hull of R . If $a \in \text{aff } R$, then R is the upper bound of an interval, and otherwise it is not. Assuming R is an upper bound, we now describe how to find the corresponding lower bound and thus the entire interval. We call $Q \subseteq R$ a *facet* if $\dim Q = \dim R - 1$, and we say Q is *visible* from a if $\text{aff } Q$ separates a from the unique point in $R \setminus Q$ within $\text{aff } R$. By assumption of general position, $\text{aff } Q$ contains neither a nor the point in $R \setminus Q$, so there is no ambiguity. The lower bound of the interval with upper bound R is the subset $P \subseteq R$ of points that belong to all visible facets. This completely describes the generalized discrete gradient of the Delaunay radius function in the unweighted case. Figure 16 illustrates the two possible configurations for a triangle in \mathbb{R}^2 : either $P = R$, in which case R is critical, or P is an edge of the triangle. It is not possible that R has two visible edges, because this contradicts that the three points lie on a circle centered at a . More generally, an m -simplex R cannot have more than $m-1$ visible facets.

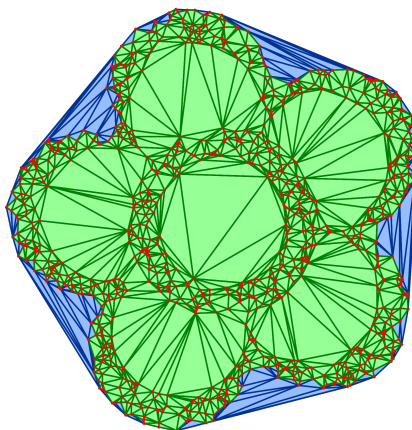


■ **Figure 16** From *left to right*: the triangle has 0, 1, 2 visible edges. Only the first two cases can occur in an unweighted Delaunay mosaic in \mathbb{R}^2 .

We consider second the weighted case, letting $X \subseteq \mathbb{R}^n$ be locally finite and in general position, letting \mathbb{R}^k be spanned by the first $k < n$ coordinate axes of \mathbb{R}^n , and writing $Y \subseteq \mathbb{R}^k \times \mathbb{R}$ for the weighted points obtained by orthogonal projection. The situation is similar to the unweighted case, except that we restrict ourselves to anchored spheres, which we recall are spheres in \mathbb{R}^n whose centers lie in \mathbb{R}^k . For a simplex $R \subseteq Y$, the relevant sphere is the smallest empty anchored sphere that passes through the preimages of the weighted points in R . If such a sphere does not exist, then $R \notin \text{Del}(Y)$. Otherwise, we let $a \in \text{dom}(R)$ be the point closest to the affine hull of R , which is an affine subspace of \mathbb{R}^k . If $a \in \text{aff } R$, then R is the upper bound of an interval of $\mathcal{R}: \text{Del}(Y) \rightarrow \mathbb{R}$. As in the unweighted case, the corresponding lower bound is the subset $P \subseteq R$ of points that belong to all facets of R visible from a .

The main difference to the unweighted case is that vertices of $\text{Del}(Y)$ are no longer necessarily critical. Indeed, an m -simplex $R \in \text{Del}(Y)$ may have as many as m visible facets. If it has m visible facets, then their intersection is a single point, which implies that the corresponding lower bound is a vertex of the weighted Delaunay mosaic.

The Wrap complex. An application of discrete Morse theory is the *Wrap complex* studied in Section 3; see [18] for the original paper on the subject and [17] for a discussion of alternative methods. To construct the Wrap complex of $X \subseteq \mathbb{R}^n$ and $r \in \mathbb{R}$, denoted $\text{Wrap}_r(X)$, we start with $\text{Alpha}_r(X) \subseteq \text{Del}(X)$ and then collapse all intervals of the radius function that can be collapsed; see Figure 17. Specifically, if $[P, R]$ is a non-singular interval whose simplices all belong to the current complex, and P is not face of any simplex outside of $[P, R]$, then we remove all simplices in this interval from the complex and repeat. The final result is the Wrap complex, and it does not depend on the sequence in which we collapse the intervals. Equivalently, we can grow $\text{Wrap}_r(X)$ from the critical simplices of radius at most r . To do this, we add all faces to get a complex, we complete the partially added intervals, and we repeat. Again, it does not matter in which sequence the simplices are added and the intervals are completed. A more formal definition of the Wrap complex can be found in Section 3.



■ **Figure 17** The colors distinguish the Wrap complex for $r = \infty$ (green) from the Alpha complex for the same radius (green and blue), which for $r = \infty$ is the entire Delaunay mosaic.

C Randomness and Expectation

In this appendix, we present the probabilistic background used throughout this paper. Our preferred model is a Poisson point process in Euclidean space, which we introduce first. Besides such point processes, we also consider points uniformly sampled on spheres.

Poisson point process. This is a natural extension of uniformly sampled points from compact to possibly unbounded domains. A *homogeneous* or *stationary Poisson point process*, X , with density $\rho > 0$ in \mathbb{R}^n is characterized by the following two properties:

1. the number of points sampled within a Borel set of measure $\|\Omega\|$ is Poisson distributed; that is: $\mathbb{P}[|X \cap \Omega| = k] = \rho^k \|\Omega\|^k e^{-\rho \|\Omega\|} / k!$,
2. the numbers of points in any finite collection of pairwise disjoint Borel sets are independent;

see [32] for a good introduction to the topic. We construct Voronoi tessellations and Delaunay mosaics, so it is important that X be locally finite. This is indeed the case with probability 1. Furthermore, X is in general position with probability 1. We use the following four notions of *general position*:

- A. No $(n - 1)$ -sphere in \mathbb{R}^n passes through more than $n + 1$ points of X .
- B. For any $P \subseteq Q \subseteq X$ with $p = \dim P$, $q = \dim Q$, and $p < q \leq n$, the center of the $(p - 1)$ -sphere defined by P is different from the center of the $(q - 1)$ -sphere defined by Q .
- C. Fixing a linear subspace $\mathbb{R}^k \subseteq \mathbb{R}^n$, no $(n - 1)$ -plane orthogonal to \mathbb{R}^k passes through more than k points of X .
- D. For any \mathbb{R}^k as in C and $P \subseteq Q$ as in B, the center of the anchored p -sphere orthogonal to \mathbb{R}^k defined by P is different from the center of the anchored q -sphere orthogonal to \mathbb{R}^k defined by Q .


We need A to get Delaunay mosaics that are simplicial, and we need A and B to get Delaunay radius functions that are generalized discrete Morse. We need A and C to get weighted Delaunay mosaics that are simplicial, and we need A, C, and D to get weighted Delaunay radius functions that are generalized discrete Morse.

Random mosaics. Our main focus is on counting the simplices and on determining the radius distribution of a typical simplex of a random mosaic. As explained in Section 1, we adjust the conventional notion of representative point and Palm distribution to carry topological meaning. For example, for the Poisson–Delaunay mosaic, we define the representative point of a simplex $Q \in \text{Del}(X)$ as the center of the smallest empty sphere that passes through all points of Q . We call this the *Delaunay sphere* of Q , we refer to its center as the *center* of Q , and we call its radius the *radius* of Q . The *typical Delaunay simplex* of dimension j is therefore defined as a random j -simplex uniformly chosen from all j -simplices of the Poisson–Delaunay mosaic with center in some open bounded Borel set $\Omega \subseteq \mathbb{R}^n$, conditioned on the existence of such simplices. With this definition, the properties of the typical j -simplex can depend on the choice of Ω , but this is not the case for the properties we are interested in; compare the results surveyed in this paper with [42, 46].

The centers of the simplices form yet another point process in \mathbb{R}^n , and we count simplices by studying its intensity. Notice however that this point process is not necessarily *simple*, even for X in general position; that is: even in the generic case representative points of different simplices can be the same. We make it generically simple by mapping the intervals of the radius function rather than the individual simplices to \mathbb{R}^n . This motivates the probabilistic analysis of intervals, which we will address shortly. The new point process is however homogeneous — both for the intervals but also for the simplices — which can either be seen

directly from the homogeneity of the underlying Poisson point process, or from the results described in [25, 26].

Random inscribed simplices. The difference between the point process for intervals and for simplices motivates the following probabilistic question: Letting Q be a set of $m + 1$ points chosen uniformly and independently on the unit sphere in \mathbb{R}^m , what is the probability that ℓ facets of $\text{conv } Q$ are visible from the origin? Following Wendel [53], we consider the 2^{m+1} m -simplices obtained by either retaining or centrally reflecting each point in Q ; see Figure 18. Assuming $0 \in \text{conv } Q$, we can decompose the m -simplex into the $m + 1$ cones of



■ **Figure 18** *Left:* the inscribed triangle, UVW , and the three triangles obtained by reflecting one of the vertices through the origin. *Right:* the three triangles obtained from UVW by reflecting two of the vertices, and the triangle obtained by reflecting all three vertices.

0 over the facets. The m -dimensional volume of $\text{conv } Q$ is thus the sum of the volumes of these cones. Importantly, the volume of each of the 2^{m+1} m -simplices is a sum of the same cone-volumes, except that each cone either appears with the coefficient $+1$ or -1 in this sum. More specifically, the sign pattern is either the same as for the vertices, or it is the opposite. For example, in \mathbb{R}^2 , this implies that 2 of the 8 triangles contain the origin and 6 triangles have one visible edge each. This implies that the probability of the triangle to be acute is $\frac{1}{4}$ and the probability of the triangle to be obtuse is $\frac{3}{4}$. This should be compared with the fact that an expected half of the triangles in a Poisson–Delaunay mosaic in \mathbb{R}^2 are acute and the other half are obtuse. The difference is explained by the re-parametrization necessary to transform one setting to the other. Things get more complicated in higher dimensions, and we refer to [26] for a complete analysis in dimension $m \leq 4$. The approach of Wendel extends to the weighted case, in which the points of Q are no longer required to lie on the unit sphere, and a complete analysis in dimension $m \leq 2$ can be found in [25].