CHROMATIC ALPHA COMPLEXES

SEBASTIANO CULTRERA DI MONTESANO[1], ONDŘEJ DRAGANOV[2], HERBERT EDELSBRUNNER[3], AND MORTEZA SAGHAFIAN[4]

Abstract. Motivated by applications in the medical sciences, we study finite chromatic sets in Euclidean space from a topological perspective. Based on the persistent homology for images, kernels and cokernels, we design provably stable homological quantifiers that describe the geometric micro- and macro-structure of how the color classes mingle. These can be efficiently computed using chromatic variants of Delaunay and alpha complexes, and code that does these computations is provided.

1. Introduction

This paper takes a topological approach to quantifying spatial interactions between several point sets, which we distinguish by color. The aim is the development of a mathematical language to answer questions like: "how, how often, and at what scale do blue points surround groups of red points?", or "are there cycles made out of blue, red, and green points that make essential use of all three colors?". We tackle these questions from a multi-scale homological perspective, with the goal of disentangling patterns such as the ones shown in Figure 1.

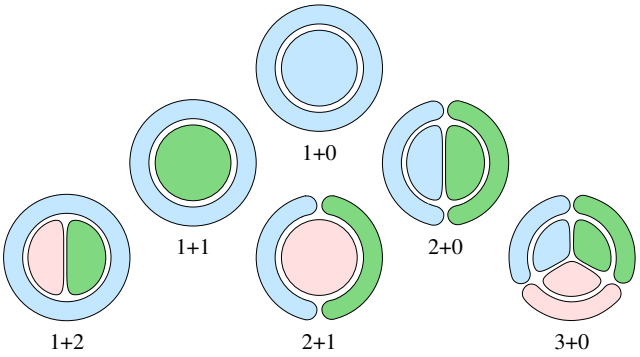

Figure 1. Mingling patterns distinguished by the number of colors needed to form a cycle and the number of additional colors needed to fill this cycle. The drawings are caricatures of similar patterns for cycles different from circles and fillings different from disks. The patterns are but a first attempt to differentiate types of interactions, and they are by no means precise or exhaustive. For example, two additional colors can fill a cycle in at least two different ways (see the pattern of type 1+2): in a collaboration as suggested in the drawing, or each individually, like two different patterns of type 1+1.

One of the motivations for this work is the recent growth of interest in *spatial biology*, which combines the biological properties of cells with their locations. An example is the *tumor immune microenvironment* [3] in cancer research, which focuses on the interplay between tumor and immune cells. Can we identify as well as quantify patterns in the interaction between cell types that correlate with clinical outcomes? Another biological process that raises similar mathematical questions is *cell sorting* (the natural segregation of cells by types), which in early development is studied for instance in [17] and in the context of somitogenesis is mentioned in [18]. This motivates the study of *chromatic point sets*, in which the points represent cells and colors represent their types.

Our specific approach is based on the formulation of persistent homology in terms of growing balls [11], which is often used in topological data analysis to describe and quantify spatial arrangements of (monochromatic) point sets. The idea behind the construction is to transform a discrete set of points into a

[1,2,3,4]ISTA (Institute of Science and Technology Austria), Klosterneuburg, Austria

E-mail addresses: [1]sebastiano.cultrera@ist.ac.at, [2]ondrej.draganov@ist.ac.at, [3]herbert.edelsbrunner@ist.ac.at, [4]morteza.saghafian@ist.ac.at.

Key words and phrases. Topological data analysis, Delaunay mosaic, alpha complex, chromatic sets, persistent homology, kernel/image/cokernel persistent homology, radius function, discrete Morse theory, exact sequences.

nested sequence of topological spaces. This is achieved by growing balls centered at the input points (from zero to infinite radius), which yields a sequence of progressively larger shapes. Such a nested sequence of spaces is called a *filtration*, and it becomes a sequence of vector spaces connected by linear maps when we apply homology with field coefficients. The latter is also known as a *persistence module*. Studying the induced maps rather than just individual vector spaces, we can not only identify radii at which topological features appear and disappear, but also pair these events to quantify for how long each feature persists in the filtration.

When the data is bi-chromatic, say *red* and *blue*, we have two sets of growing balls at our disposal. A natural way to relate them is to consider the inclusion map between the union of balls of one color, say the blue ones, into the union of balls of both colors as their radii grow. Like in the mono-chromatic case, we apply homology with field coefficients and get two persistence modules together with connecting maps, which are induced by inclusions relating the two filtrations. The connecting maps carry important information about the mingling of the two point sets. For example, if an essential cycle is present in the blue filtration at a certain radius, it may or may not also be present in the red and blue filtration. If it is, such a cycle will be in the image of the connecting map, while if it is not, it will be in the kernel. More generally, given a pair of filtrations related by inclusions, we can look at the persistent homology of the subspace, the full space, the relative space, as well as the kernels, the images, and the cokernels of the connecting maps [8]. We call the resulting collection the *6-pack of persistence diagrams*, which we use to capture different aspects of the mingling between geometric sets. One contribution of this paper is the study of relations between the six persistence diagrams composing the 6-pack, such as linear relations between their 1-*norms*, which are the sums of persistences of the points in the diagrams (Theorem 5.3).

Just like alpha complexes are a possible discrete model for the union of balls in the mono-chromatic setting [12, 13], we seek a chromatic variant that enables the computation of the 6-pack of persistent diagrams. At first sight, this seems problematic as the red alpha complex does not include into the alpha complex of the union of red and blue points. Similarly, taking the red Delaunay subcomplex of the full Delaunay mosaic does not work either, as it does not capture the homotopy type of the union of red balls. We circumvent these limitations by using a third type of complex, the *chromatic Delaunay mosaic*, which was introduced for two colors by Reani and Bobrowski [19] and which we extended beyond two colors in [4]. This mosaic uses an extra dimension for each color beyond the first to capture the interaction between colors. Counter-balancing the increase in dimension, [4] showed that the complexity of the mosaic is moderate for a small number of colors. For example, this paper gives linear bounds on the expected size for points in two dimensions randomly colored by a constant number of colors. Building on the results in [4], we show that the chromatic Delaunay mosaic can be equipped with a radius function whose sublevel sets capture the alpha complexes of different color classes as well as their interactions. Within this setting, we show that the radius function on the chromatic Delaunay mosaic can be computed in linear time assuming the dimension and the number of colors is constant (Theorem 3.9), and that it has the combinatorial structure of a generalized discrete Morse function (Theorem 4.6). Code that implements these algorithms is available at [10].

The entire development could have been based on chromatic variants of the Čech complex, with almost no differences, except that the complexes would be significantly larger, making computational experiments of the kind presented in this paper infeasible. Similarly, we could have used chromatic variants of the Vietoris–Rips complex, but the complexes would again be significantly larger, and we would have to cope with topological artifacts, which at this time are not understood. More recently, Dowker complexes and witness complexes have been suggested as possible candidates for encoding spatial relations in the tumor microenvironment [20]. The main limitation of these complexes is their lack of stability: perturbing a point set can produce a very different filtration. Moreover, the Dowker complex is limited to the study of two interacting point sets, while the witness complex requires a choice of "landmark points" and it is not clear how to choose those in practice.

Outline. Section 2 reviews the alpha complex in the mono-chromatic case. Section 3 extends this construction to the chromatic case. Section 4 proves that the radius function on the chromatic Delaunay mosaic is generalized discrete Morse. Section 5 studies the persistent homology of the chromatic alpha complexes, with an emphasis on the two and three colors settings. Section 6 concludes the paper.

2. Mono-chromatic Point Sets

In this section, we recall several standard definitions and relevant results used in topological data analysis of point sets with no extra color labels.

2.1. **Voronoi Tessellation and Delaunay Complex.** Letting $A \subseteq \mathbb{R}^d$ be a finite set of points, the *Voronoi domain* of $a \in A$, denoted $\mathrm{dom}(a, A)$, is the set of points $x \in \mathbb{R}^d$ that satisfy $\|x - a\| \leq \|x - b\|$ for all $b \in A$. Observe that $\mathrm{dom}(a, A)$ is the intersection of finitely many closed half-spaces and therefore a closed convex polyhedron. The *Voronoi tessellation* of A, denoted $\mathrm{Vor}(A)$, is the collection of Voronoi domains defined by the points in A. These domains cover $\mathbb{R}^d$ while their interiors are pairwise disjoint. Nevertheless, a collection of these polyhedra may overlap in a shared face, which we refer to as a *Voronoi cell*. For a generic set, A, the dimension of a Voronoi cell is determined by the number of Voronoi domains that contain it.

Definition 2.1 (Conventional Genericity)**.** We call a point set, $A \subseteq \mathbb{R}^d$, *generic* if for every subset, B, of at most $d + 1$ points, the smallest circumsphere of B does not pass through any point of $A \setminus B$.

Then, indeed, the common intersection of any $p + 1$ Voronoi domains is either empty or a convex polyhedron of dimension $d - p$. Furthermore, the condition prohibits right-angled triangles and similarly special higher-dimensional simplices, as required for the proof of Proposition 2.4. Note that our notion of genericity allows for more than $p + 1$ points on a p-dimensional affine subspace.

The *Delaunay complex*, denoted $\mathrm{Del}(A)$, is the simplicial complex with vertex set A that contains a simplex for each collection of Voronoi domains with non-empty common intersection. It is isomorphic to the nerve of the Voronoi domains,

$$\mathrm{Nerve}(\mathrm{Vor}(A)) = \{\nu \subseteq \mathrm{Vor}(A) \mid \bigcap \nu \neq \emptyset\}. \tag{2.1}$$

Notice that the Delaunay complex generally differs from the *dual* of the Voronoi tessellation, which contains a p-dimensional cell for each $(d - p)$-dimensional intersection of Voronoi domains. We call this dual the *Delaunay mosaic*: it contains the convex hull of a subset of points as a cell whenever the corresponding collection of Voronoi domains is maximal with this necessarily non-empty common intersection. However, for a generic set of points, the Delaunay mosaic is the Delaunay complex of the points. Throughout this paper, we will work with the Delaunay complex rather than the mosaic, and we will appeal to genericity in cases this is necessary. Both, the Delaunay complex and the Delaunay mosaic can also be characterized with empty spheres passing through points. A $(d - 1)$-sphere, S, is *empty* if all points of A lie on or outside the sphere, so there are no points inside the sphere—that is, no points in the interior of the ball bounded by the sphere. Furthermore, we say the sphere *passes through* the points that lie on the sphere; see Figure 4 for an empty 1-sphere passing through three points.

Lemma 2.2. Let $A \subseteq \mathbb{R}^d$ be a finite set of points, and $\nu \subseteq A$. Then ν is a simplex in the Delaunay complex iff there exists an empty sphere that passes through all points in ν. The convex hull of ν is a cell in the Delaunay mosaic iff there is an empty sphere that passes through the points in ν and through no other points of A.

Proof. We prove the first claim and omit the argument for the Delaunay mosaic. Let $\nu \in \mathrm{Del}(A)$. By definition of Delaunay complex, $\bigcap_{a \in \nu} \mathrm{dom}(a, A) \neq \emptyset$, and we let x be a point in this common intersection of Voronoi domains. Then x has the same distance to all points in ν and the same or a larger distance to all other points in A. Hence, x is the center of an empty sphere that passes through all points in ν, and possibly also through other points of A. Each of the above implications can be reversed, which implies that x is in the intersection of Voronoi domains of points in ν iff x is the center of an empty ball that passes through ν. □

2.2. **Alpha Complex.** There can be more than one empty sphere passing through the vertices of a simplex, $\nu \in \mathrm{Del}(A)$, but there is a unique *smallest empty sphere* that passes through the points in ν. This yields a *radius function* on the Delaunay complex, $\mathrm{Rad}\colon \mathrm{Del}(A) \to \mathbb{R}$, which maps each simplex to the radius of the smallest empty sphere that passes through its vertices. The *alpha complex*, $\mathrm{Alf}_r(A) \subseteq \mathrm{Del}(A)$, is the sublevel set consisting of all simplices with radius at most r. Note that for $r \leq R$ we have $\mathrm{Alf}_r(A) \subseteq \mathrm{Alf}_R(A)$.

Let $B_r(a)$ be the d-ball in $\mathbb{R}^d$ centered at a with radius r. The *Voronoi ball* of $a \in A$ with radius r is this ball clipped to within the Voronoi domain: $B_r^V(a, A) = B_r(a) \cap \mathrm{dom}(a, A)$. Using the correspondence in Lemma 2.2, it is straightforward to observe that the alpha complex, $\mathrm{Alf}_r(A)$, is isomorphic to the nerve of the Voronoi balls of A with radius r. The Nerve Theorem [5] then implies that

$$\mathrm{Alf}_r(A) \simeq \bigcup\nolimits_{a \in A} B_r^V(a, A) = \bigcup\nolimits_{a \in A} B_r(a). \tag{2.2}$$

Furthermore, a generalization of the theorem guarantees that the homotopy equivalences for different radii, $r \leq R$, commute with the inclusions on both sides. From the perspective of topology, we can equivalently study the union of growing balls or its discrete counterpart, the growing alpha complex.

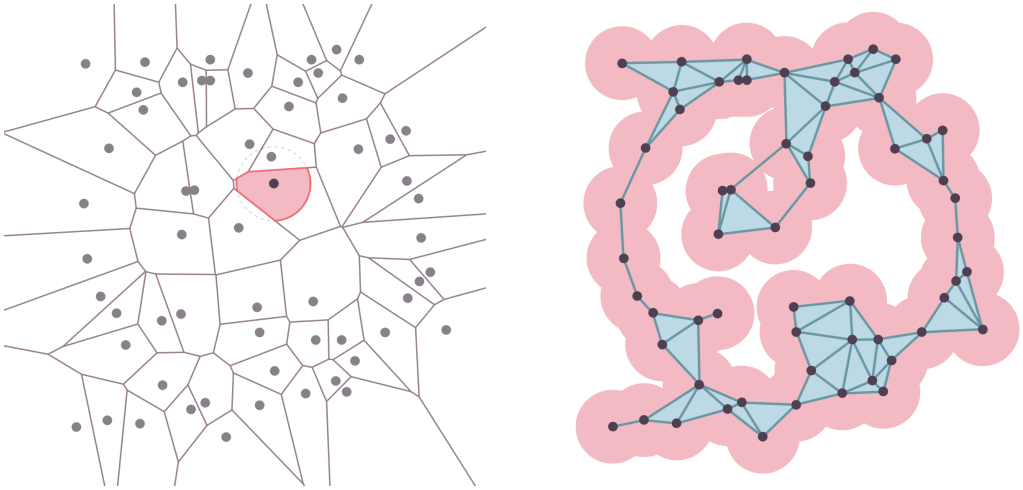

FIGURE 2. On the *left*: a set, $A \subseteq \mathbb{R}^2$, together with its Voronoi tessellation, $\mathrm{Vor}(A)$, and one Voronoi ball highlighted. On the *right*: the union of disks, $\bigcup_{a \in A} B_r(a)$, with the alpha complex, $\mathrm{Alf}_r(A)$, superimposed.

For a finite point set, the Delaunay complex is finite, and hence the filtration of different alpha complexes is finite. To better understand the structure of this filtration, it is helpful to understand the level sets of the radius function. Given simplices $\alpha \subseteq \gamma$ in a simplicial complex K, write $[\alpha, \gamma]$ for the simplices β that satisfy $\alpha \subseteq \beta \subseteq \gamma$; that is: $[\alpha, \gamma]$ is an *interval* in the face poset of K. Given a monotonic function $f\colon K \to \mathbb{R}$, an *interval of* f is an interval on which f is constant, and it is *maximal* if it is not strictly contained in a larger interval of f.

Definition 2.3. A monotonic function on a simplicial complex, $f\colon K \to \mathbb{R}$, is *generalized discrete Morse* if the maximal intervals of f partition K.

Equivalently, f is generalized discrete Morse if every level set, $K_t = f^{-1}(t)$, is a disjoint union of maximal intervals.

Proposition 2.4 ([1, Corollary 4.6])**.** Let $A \subseteq \mathbb{R}^d$ be finite and generic, as defined in Definition 2.1. Then $\mathrm{Rad}\colon \mathrm{Del}(A) \to \mathbb{R}$ is a generalized discrete Morse function.

The significance of this result is that we can construct each alpha complex by adding one interval at a time. If this interval consists of two or more simplices, then the addition does not affect the homotopy type of the complex. Indeed, the complex before is a deformation retract of the complex after the addition of the interval. On the other hand, if the interval consists of a single simplex, then the addition of this simplex changes the homology of the complex in a controlled manner.

3. Chromatic Point Sets

The main concept in this section is the chromatic alpha complex, which generalizes the bi-chromatic construction in [19] to three and more colors. A crucial ingredient is the radius function on the chromatic Delaunay complex, whose sublevel sets are the chromatic alpha complexes. The section follows the logical structure of Section 2 to clearly showcase the analogies between the mono-chromatic and more general, chromatic settings. Indeed, we will see that the chromatic definitions match the definitions in the previous section when we only consider one color.

A *chromatic point set* is a mapping $\chi\colon A \to \sigma$, in which $A \subseteq \mathbb{R}^d$ is a finite point set, and σ is a set of colors. We usually write $s = \#\sigma - 1$ and $\sigma = \{0, 1, \ldots, s\}$. Furthermore, we write $A_j = \chi^{-1}(j)$ for the subset of points of a given color $j \in \sigma$. We fix this notation throughout the section.

3.1. **Chromatic Voronoi Tessellation and Chromatic Delaunay Complex.** The *chromatic Voronoi tessellation*, $\mathrm{Vor}(\chi)$, is the collection of Voronoi domains, $\mathrm{dom}(a, A_{\chi(a)})$, for all points $a \in A$. In other words, $\mathrm{Vor}(\chi)$ is the union of the $\mathrm{Vor}(A_j)$, over all $j \in \sigma$. Differently colored Voronoi domains can have overlapping interiors. Indeed, every point in $\mathbb{R}^d$ is covered by at least $s+1$ different domains from $\mathrm{Vor}(\chi)$, namely at least one domain for each color.

The *chromatic Delaunay complex*, denoted $\mathrm{Del}(\chi)$, contains a simplex $\nu \subseteq A$ if the common intersection of the corresponding domains is non-empty. It is isomorphic to the nerve of the chromatic Voronoi tessellation. A direct analogy of the characterization of the Delaunay complex with empty spheres is the characterization of the chromatic Delaunay complex with what we call empty stacks. A *σ-stack* in $\mathbb{R}^d$ is a collection of $s+1$ concentric $(d-1)$-spheres, one for each color in σ; see Figure 3. We drop σ from the notation if it is clear from the context. The *radius* of the stack is the maximum radius of its spheres, and its *center* is the common center of the spheres. We label the spheres S_j, $j \in \sigma$, and say the stack is *empty* if S_j is empty of points in $A_j = \chi^{-1}(j)$, for each $j \in \sigma$. We say the stack *passes through* $\nu \subseteq A$ if S_j passes through all the points of $\nu \cap A_j$, for each color $j \in \sigma$.

Lemma 3.1. Let $\chi\colon A \to \sigma$ be a chromatic point set in $\mathbb{R}^d$, write $A_j = \chi^{-1}(j)$, and let $\nu \subseteq A$ be a collection of points. Then $\nu \in \mathrm{Del}(\chi)$ iff there exists an empty stack of spheres that passes through ν.

Proof. Let $\nu_j = \nu \cap A_j$ be the j-colored points in ν, for each $j \in \sigma$. By Lemma 2.2, the existence of an empty sphere, S_j, with center x that passes through ν_j is equivalent to x being in the intersection of the corresponding Voronoi domains: $x \in \bigcap_{a \in \nu_j} \mathrm{dom}(a, A_j)$. Therefore, there exists an empty stack passing through ν centered at x iff $x \in \bigcap_{a \in \nu_j} \mathrm{dom}(a, A_j)$ for each $j \in \sigma$. This is the defining property of ν being in $\mathrm{Del}(\chi)$, namely that $\bigcap_{a \in \nu} \mathrm{dom}(a, A_{\chi(a)})$ is non-empty. □

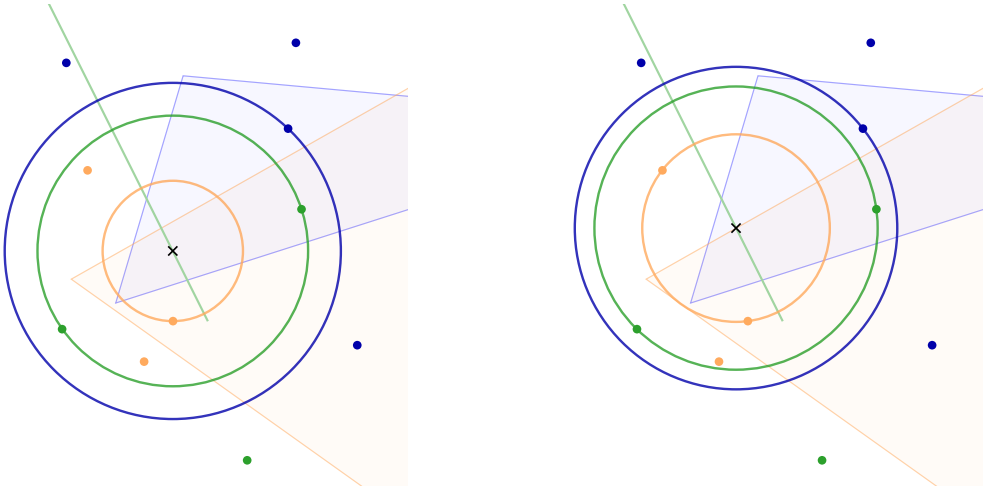

FIGURE 3. Two empty stacks in $\mathbb{R}^2$ that pass through one blue point, two green points, and one orange point forming a simplex $\nu \in \mathrm{Del}(\chi)$. (In fact, the stack on the *right* passes through *two* orange points, so it also passes through the *one* orange point that lies on the *left* orange circle.) The set of centers of all empty stacks that pass through these four points is the intersection of three Voronoi cells: a blue 2-cell, a green 1-cell, and an orange 2-cell. The *right panel* shows the smallest empty stack in this collection: its center lies on the boundary of the intersection of Voronoi cells, which is the reason why one of its circles passes through an extra point.

3.2. **Chromatic Alpha Complex.** Like in the mono-chromatic setting, we define chromatic alpha complexes as sublevel sets of the radius function defined on the chromatic Delaunay complex. We recall that the radius of a stack is the radius of its largest sphere.

Definition 3.2. Let $\chi\colon A \to \sigma$ be a chromatic point set, and $\mathrm{Rad}\colon \mathrm{Del}(\chi) \to \mathbb{R}$ the *radius function* defined by mapping $\nu \in \mathrm{Del}(\chi)$ to the radius of the smallest empty stack that passes through ν. The *chromatic alpha complex* of χ with radius $r \in \mathbb{R}$ is $\mathrm{Alf}_r(\chi) = \mathrm{Rad}^{-1}[0, r]$.

Using the empty spheres and empty stacks characterizations (Lemma 2.2 and Lemma 3.1), we can see a clear relation between alpha complexes and chromatic alpha complexes. If there exists an empty $(d-1)$-sphere, S, of radius r passing through $\nu \subseteq A$, then there also exists an empty stack of radius r passing

through ν. Indeed, we can take $S_j = S$ for each $j \in \sigma$. Similarly, an empty sphere, S, that passes through points $\nu \subseteq A_j$ of the same color is itself an empty stack when we set S_i to be a sphere with zero radius for $i \neq j$. However, the same simplex can have a different radius in $\mathrm{Del}(A)$ and in $\mathrm{Del}(\chi)$: the smallest empty sphere can have strictly larger radius than the smallest empty stack passing through the same points; see Figure 4. We formulate the above observation in a slightly more general form.

Lemma 3.3. Let $\chi\colon A \to \sigma$ be a chromatic point set in $\mathbb{R}^d$, and $\tau \subseteq \sigma$ a subset of the colors.

(i) Let $\eta\colon \sigma \to \tau$ be a merging of colors. Then $\mathrm{Alf}_r(\eta \circ \chi) \subseteq \mathrm{Alf}_r(\chi)$ for all r.

(ii) Let $\chi|\tau\colon \chi^{-1}(\tau) \to \tau$ be the restriction of χ to colors in τ. Then $\mathrm{Alf}_r(\chi|\tau) \subseteq \mathrm{Alf}_r(\chi)$ for all r.

In particular, $\mathrm{Alf}_r(A) \subseteq \mathrm{Alf}_r(\chi)$, and $\mathrm{Alf}_r(A_j) \subseteq \mathrm{Alf}_r(\chi)$ for every $j \in \sigma$, in which $A_j = \chi^{-1}(j)$.

Proof. To see (i), let $(S_i)_{i\in\tau}$ be an empty stack of radius r that passes through the points in $\nu \subseteq A$; it witnesses $\nu \in \mathrm{Alf}_r(\eta \circ \chi)$. Then $(S_{\eta(j)})_{j\in\sigma}$ is an empty stack of radius r that passes though the points in ν; it witnesses $\nu \in \mathrm{Alf}_r(\chi)$. To see (ii), let $(S_i)_{i\in\tau}$ be an empty stack of radius r that witnesses $\nu \in \mathrm{Alf}_r(\chi|\tau)$. Adding zero-radius spheres for the colors $j \in \sigma \setminus \tau$, we get an empty stack that witnesses $\nu \in \mathrm{Alf}_r(\chi)$. □

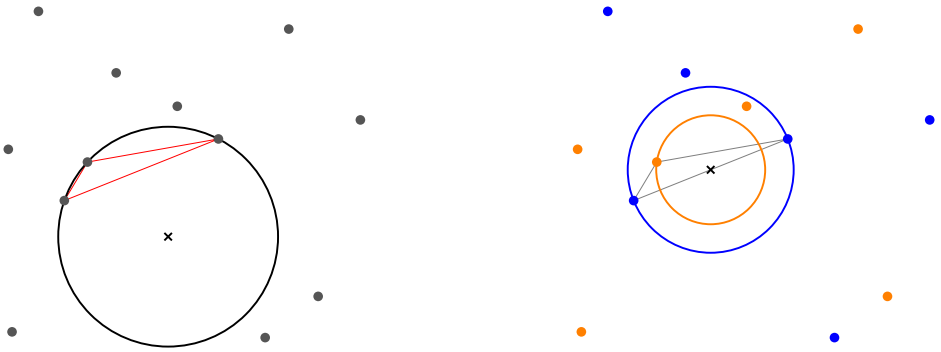

FIGURE 4. An obtuse triangle with two blue points and an orange point at the obtuse angle. On the *left:* the smallest empty sphere that passes through the three points. It has strictly larger radius than the smallest empty stack that passes through the three points, which is shown on the *right.* Therefore, the triangle belongs to both, the Delaunay complex and the chromatic Delaunay complex, but it has a different value in the two radius functions.

An important reason why the alpha complex is useful in the mono-chromatic setting is its correspondence to the union of balls growing from the input points. From the topological point of view, studying the growing union of balls is equivalent to studying the growing alpha complex. In the following, we draw an analogous connection for chromatic alpha complexes. One important distinction is that there is more structure to be preserved in the chromatic setting: not only the topological spaces themselves, but also how they are related to each other. For example, for a bi-chromatic point set as in Figure 5, we study the inclusion of the union of the blue disks into the union of all disks. We prove that we can equivalently study the inclusions of the blue alpha complex into the chromatic alpha complex.

For a point $a \in A$ in a chromatic set $\chi\colon A \to \sigma$, we define its *(chromatic) Voronoi ball* of radius r as the intersection of the ball of radius r with the Voronoi domain within its color class:

$$B_r^V(a,\chi) = B_r^V(a, A_{\chi(a)}) = B_r(a) \cap \mathrm{dom}(a, A_{\chi(a)}). \tag{3.1}$$

Let $\nu \subseteq A$ be a set of points. Like in the proof of Lemma 3.1, we observe that x is the center of an empty stack of radius r passing through ν iff x is contained in the intersection of the Voronoi balls of radius r centered at the points in ν; that is: $x \in \bigcap_{a\in\nu} B_r^V(a,\chi)$. This implies that $\mathrm{Alf}_r(\chi)$ is isomorphic to the nerve of all Voronoi balls $B_r^V(a,\chi)$, $a \in A$. Since the union of the Voronoi balls is the same as the union of the balls, the Nerve Theorem yields the following:

Lemma 3.4. $\mathrm{Alf}_r(\chi) \simeq \mathrm{Alf}_r(A)$, and both are homotopy equivalent to the union of balls, $\bigcup_{a\in A} B_r(a)$.

Unlike the alpha complex, $\mathrm{Alf}_r(A)$, the chromatic alpha complex, $\mathrm{Alf}_r(\chi)$, contains $\mathrm{Alf}_r(A_j)$ as a subcomplex for each color $j \in \sigma$. We claim that this reflects the inclusion of the union of j-colored balls into the union of all balls, which allows us to study that inclusion on the discrete side. Since the complexes involved are defined as nerves, we can use a version of Nerve Theorem to show that the inclusions commute with the homotopy equivalences.

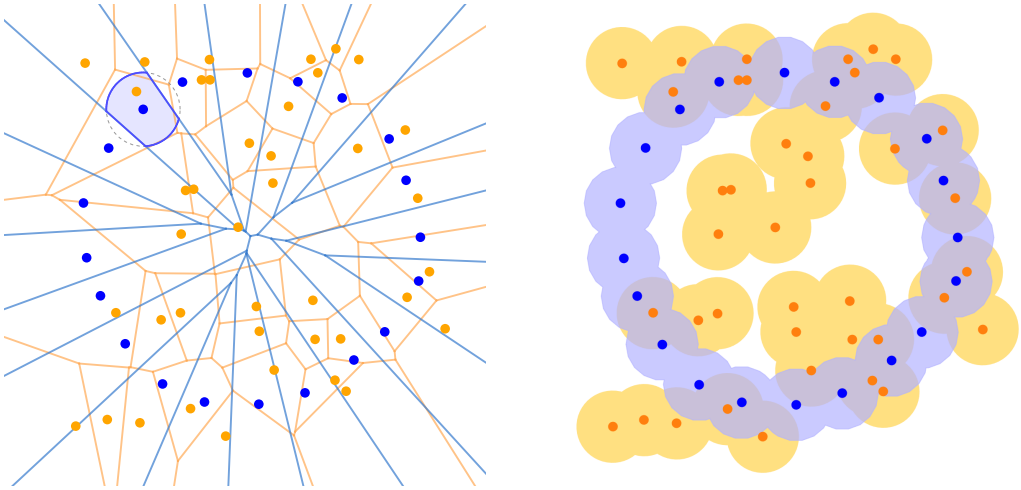

FIGURE 5. On the *left*: a chromatic set together with Voronoi tessellations of the blue and orange points overlaid, and one chromatic Voronoi ball highlighted. On the *right*: the union of blue disks and the union of orange disks; we study, e.g., the inclusion of the blue area into the union of all the disks.

Lemma 3.5. For every color $j \in \sigma$ and radius r, the following diagram commutes:

$$\begin{array}{ccc} \mathrm{Alf}_r(\chi) & \xleftarrow{\simeq} & \bigcup_{a\in A} B_r(a) \\ \uparrow & & \uparrow \\ \mathrm{Alf}_r(A_j) & \xleftarrow{\simeq} & \bigcup_{a\in A_j} B_r(a) \end{array}$$

See Theorem 3.8 in Section 3.4 for a generalization of this statement and a proof.

3.3. **Lifting Construction.** Next we recall a lifting construction that sheds light on the structure of the chromatic Delaunay complex and can be used for its computation [4]. We generalize the construction for two colors in [19] to any number of colors.

Let $\chi\colon A \to \sigma$ be a chromatic set of points, with $A \subseteq \mathbb{R}^d$ finite and colors $\sigma = \{0, 1, \ldots, s\}$. To separate the colors, we use points $u_0, u_1, \ldots, u_s$ in $\mathbb{R}^s$. For convenience, we assume these points are the vertices of the standard s-simplex embedded in $\mathbb{R}^s$, but the construction would work for any affinely independent collection of $s+1$ points. Let $\mathbb{R}^d$ and $\mathbb{R}^s$ be spanned by the first d and last s coordinate vectors of $\mathbb{R}^{d+s}$, respectively; that is: we treat $\mathbb{R}^d$ and $\mathbb{R}^s$ as orthogonal subspaces of $\mathbb{R}^{d+s}$. Write $A_j = \chi^{-1}(j)$, and set $A_j^\triangle = A_j + u_j$ for $0 \le j \le s$. Then $A^\triangle = A_0^\triangle \cup A_1^\triangle \cup \ldots \cup A_s^\triangle$ is a finite set in $\mathbb{R}^{d+s}$, and we call it the *chromatic lifting of* χ. We claim that the chromatic Delaunay complex, $\mathrm{Del}(\chi)$, is the standard Delaunay complex, $\mathrm{Del}(A^\triangle)$, after identifying the lifted vertices with their original counterparts. On one hand, this gives us a straightforward way to compute $\mathrm{Del}(\chi)$—we lift the points and use a standard algorithm to compute the Delaunay complex—and on the other hand, it gives a more intuitive view on the structure of $\mathrm{Del}(\chi)$; see Figures 6 and 7. The claimed equality is easy to prove when we use the characterizations via empty stacks and empty spheres.

Lemma 3.6. Let $\chi\colon A \to \sigma$ be a chromatic point set in $\mathbb{R}^d$ and $A^\triangle \subseteq \mathbb{R}^{d+s}$ be its chromatic lifting. There exists an empty stack of $s+1$ $(d-1)$-spheres that pass through the points in $\nu \subseteq A$ iff there exists an empty $(d+s-1)$-sphere that passes through the lifted points in $\nu^\triangle = \{a + u_{\chi(a)} \mid a \in \nu\}$.

Proof. There is a 1-to-1 correspondence between the stacks of concentric $(d-1)$-spheres in $\mathbb{R}^d$ and the $(s+d-1)$-spheres in $\mathbb{R}^{d+s}$ that have a non-empty intersection with $\mathbb{R}^d + u_j$ for each $j \in \sigma$. Indeed, if S is such an $(s+d-1)$-sphere, then we get the stack by setting S_j to be the intersection of S with $\mathbb{R}^d + u_j$ projected back to $\mathbb{R}^d$, for each $j \in \sigma$, and if $(S_j)_{j\in\sigma}$ is a stack, its spheres share a common center, so we can find a sphere S whose intersection with $\mathbb{R}^d + u_j$ is $S_j + u_j$, for each $j \in \sigma$.

This correspondence implies that S is empty of points in $A^\triangle$ and passes through the points of $\nu^\triangle$ iff $(S_j)_{j\in\sigma}$ is empty of points in χ and passes through the points in ν. □

Lemmas 2.2, 3.1, and 3.6 imply the following result.

Corollary 3.7. Let $\chi\colon A \to \sigma$ be a chromatic point set and A' its chromatic lifting. Then $\mathrm{Del}(\chi)$ and $\mathrm{Del}(A^\triangle)$ are isomorphic, with the isomorphism defined by mapping $a \in A$ to $a^\triangle = a + u_{\chi(a)}$.

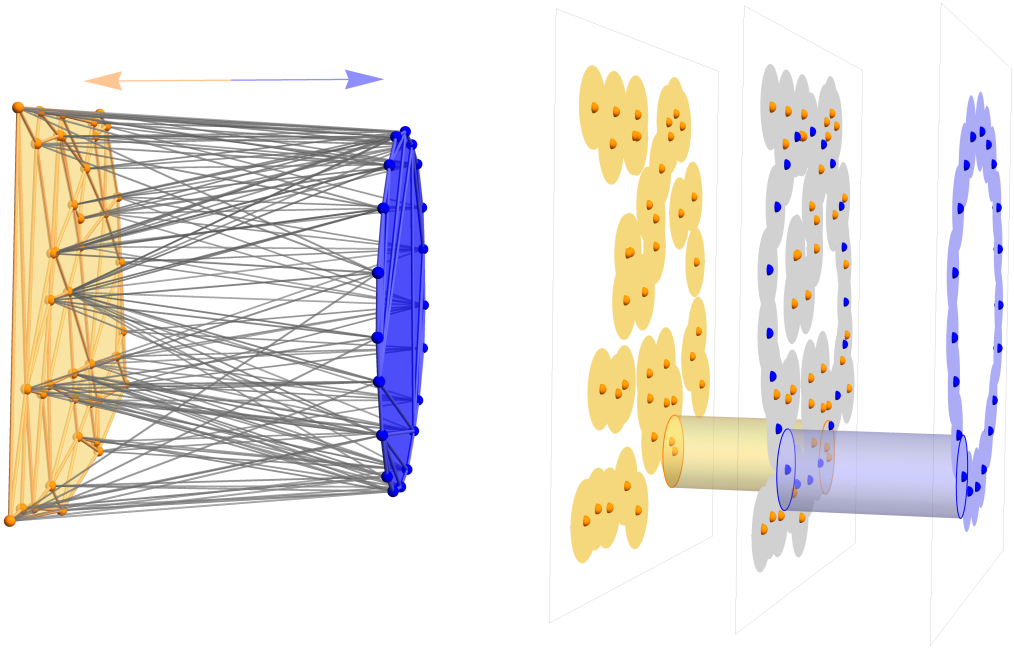

FIGURE 6. The lifting for points in $\mathbb{R}^2$ with two colors. *Left:* the chromatic Delaunay complex embedded in three dimensions using the lifted points as vertices—it is isomorphic to the Delaunay complex of the lifted points. The triangles and tetrahedra with more than one color are left unfilled for clarity. *Right:* the union of disks of radius r, for each color separately on the two sides, and for both colors together in the middle. As defined shortly in Section 3.4, each plank connects a disk in the middle with the same disk on one of the sides. We show one plank for each color.

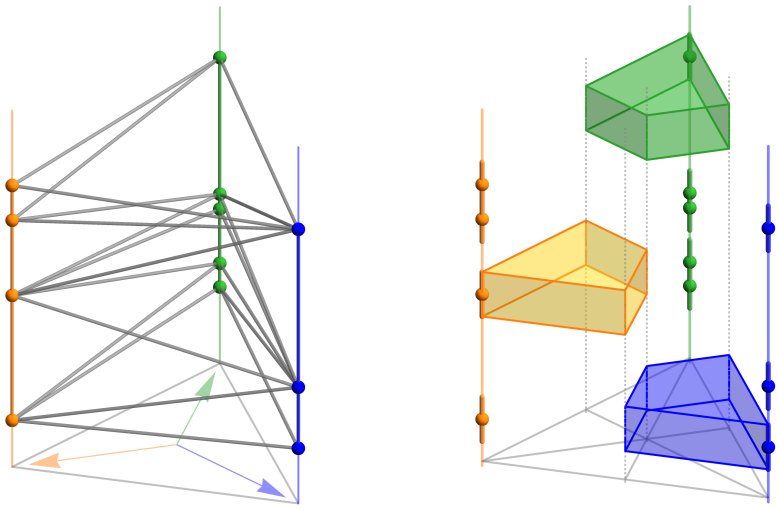

FIGURE 7. The lifting for points in $\mathbb{R}^1$ with three colors. *Left:* the chromatic Delaunay complex embedded in three dimensions using the lifted points as vertices—it is isomorphic to the Delaunay complex of the lifted points. The triangles and tetrahedra are left unfilled for clarity. *Right:* the union of segments of length $2r$, for each color separately along the three lines. Note the parallel lines emanating from the midpoints of the edges and the barycenter of the standard triangle. As defined shortly in Section 3.4, each plank is a quadrangular prism that connects one of the intervals with its projections to three of these parallel lines. We show one plank for each color.

The lifting construction provides an interesting insight independent of the focus of the paper. With the minor modification of using the vertices of the standard s-simplex in $\mathbb{R}^{s+1}$, the lifting can be interpreted as the *one-hot encoding* used in machine learning to vectorize categorized data. Corollary 3.7 states that concatenating the spatial coordinates with the one-hot encoding of the category (color) for each point has a concrete geometric meaning using just the regular Euclidean metric. This is surprising especially because the pairwise distances between the lifted points do not seem particularly meaningful.

3.4. **Chromatic Subcomplexes.** The inclusion of one color into all others, $\mathrm{Alf}_r(A_j) \hookrightarrow \mathrm{Alf}_r(\chi)$, has two major drawbacks: it effectively uses only two colors (j versus the rest), which is too little information to detect any of the tri-chromatic patterns in Figure 1, and it is asymmetric by construction. To overcome these issues, we call σ the *color simplex*, write Σ for the complex of faces of σ, and for any subcomplex, $\Gamma \subseteq \Sigma$, define the Γ-*subcomplex* of the chromatic alpha complex:

$$\mathrm{Alf}_r(\chi, \Gamma) = \{\nu \in \mathrm{Alf}_r(\chi) \mid \chi(\nu) \in \Gamma\}. \tag{3.2}$$

For example, if Γ consists of all subsets of size $t+1$ or less, then $\mathrm{Alf}_r(\chi, \Gamma)$ is the collection of all simplices whose vertices have at most $t+1$ different colors, and we call this the $(t+1)$-*chromatic subcomplex* of the alpha complex. This choice of Γ is symmetric, as it prefers no colors over any other colors. For $t = 0$, $\mathrm{Alf}_r(\chi, \Gamma)$ is the disjoint union of the $s+1$ mono-chromatic alpha complexes, and we will see shortly that studying the inclusion $\mathrm{Alf}_r(\chi, \Gamma) \hookrightarrow \mathrm{Alf}_r(\chi)$ is equivalent to studying the natural map

$$\Big[\bigcup\nolimits_{a\in A_0} B_r(a)\Big] \sqcup \Big[\bigcup\nolimits_{a\in A_1} B_r(a)\Big] \sqcup \ldots \sqcup \Big[\bigcup\nolimits_{a\in A_s} B_r(a)\Big] \longrightarrow \bigcup\nolimits_{a\in A} B_r(a), \tag{3.3}$$

acting as inclusion on each color. For example, this map captures the loops composed of points of *any one* of the colors that are filled by points of the *other* colors. When the homology functor is applied, such a loop becomes a non-trivial homology class that maps to zero.

We need definitions to gain intuition and give meaning to Γ-subcomplexes for Γ more general than just the vertices in the color simplex. We will make use of the chromatic lifting defined in the previous section, and instead of growing balls around the points, we grow what we call planks around the lifted points, which are balls in $\mathbb{R}^d$ extruded into the s extra color dimensions.

To begin, we fix a chromatic lifting of χ, as in Section 3.3, with points $u_0, u_1, \ldots, u_s \in \mathbb{R}^s$. With slight abuse of notation, we write Σ for the simplicial complex that consists of all simplices spanned by any subset of the $s+1$ points, and $|\Sigma|$ for its underlying space. The *barycenter* of a simplex is the average of its vertices. A *chain* in Σ is a nested sequence of its simplices, which gives a sequence of points (the barycenters), and taking their convex hull, we get a new simplex. The collection of simplices obtained this way from chains of Σ is the *barycentric subdivision* of Σ, denoted $\mathrm{Sd}\,\Sigma$; see Figure 7 where we see the barycentric subdivision of a triangle at the bottom of the right drawing. The *star* of u_j in $\mathrm{Sd}\,\Sigma$, denoted $\mathrm{st}(u_j, \mathrm{Sd}\,\Sigma)$, is the underlying space of all simplices in $\mathrm{Sd}\,\Sigma$ whose corresponding chains in Σ contain only simplices that share u_j. For example, the star of a vertex of a barycentrically subdivided triangle is the convex quadrangle that is the union of the two triangles spanned by the vertex, the barycenters of the two incident edges, and the barycenter of the triangle in Σ; see Figure 7. Note that a collection of vertex stars restricted to $|\Gamma|$ have a non-empty intersection iff their colors form a simplex in Γ. With

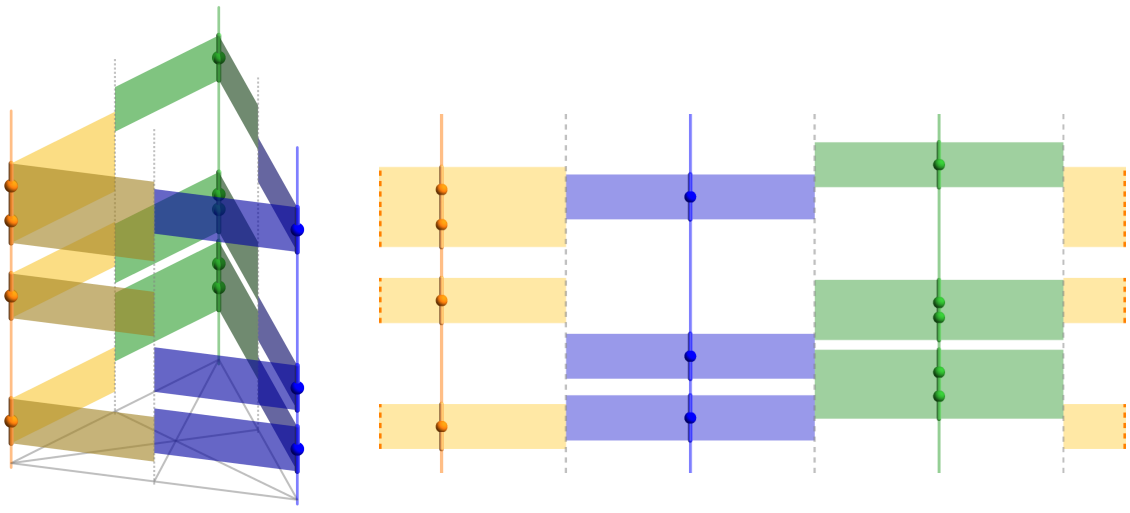

FIGURE 8. The union of bi-chromatic planks for a tri-chromatic point set in $\mathbb{R}^1$. It is homotopy equivalent to the bi-chromatic subcomplex of the corresponding chromatic alpha complex. On the *left*, we show the planks in the sides of the triangular prism erected on top of the barycentrically subdivided color triangle. On the *right*, the three sides of the triangular prism are unfolded into the plane, and the planks are glued along the *orange dashed* lines. A similar unfolding one dimension higher helps us to understand the situation for 2-dimensional data in Figure 9.

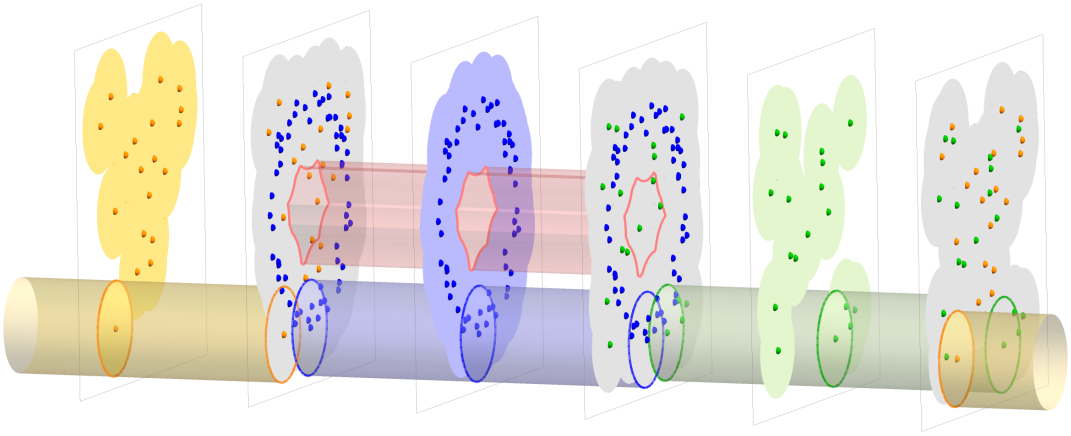

FIGURE 9. The union of bi-chromatic planks for a tri-chromatic point set in $\mathbb{R}^2$. For clarity, only one plank per color is shown—every blue, green and yellow disk is connected to its gray counterparts via cylinders. By construction, the planks are subsets of the boundary faces of a 4-dimensional triangular prism. Similar to Figure 8, we unfold the 3-dimensional boundary so we can illustrate the planks in $\mathbb{R}^3$, as shown. Observe the highlighted 2-hole in the middle: this is a topological feature that captures a loop created by one color (*blue*) and filled by each of the other two colors. This is one variant of the pattern of type 1+2 from Figure 1.

these notions, we are ready to define the planks, which are instrumental to relate the union of balls with the subcomplexes of the chromatic alpha complex. We have three progressively smaller variants: the first for the entire Σ, the second restricted to a subcomplex $\Gamma \subseteq \Sigma$, and the third further restricted to within the mono-chromatic Voronoi domains:

$$\text{Plank}_r(a,\chi) = B_r(a) \times \text{st}(\chi(a), \text{Sd}\,\Sigma), \tag{3.4}$$

$$\text{Plank}_r(a,\chi,\Gamma) = B_r(a) \times \big(\text{st}(\chi(a), \text{Sd}\,\Sigma) \cap |\Gamma|\big), \tag{3.5}$$

$$\text{Plank}^V_r(a,\chi,\Gamma) = B^V_r(a,\chi) \times \big(\text{st}(\chi(a), \text{Sd}\,\Sigma) \cap |\Gamma|\big). \tag{3.6}$$

For technical reasons, we will need convex planks, but depending on Γ, the latter two variants are not necessarily convex. This can be fixed by taking the convex hulls of the planks in $\mathbb{R}^{d+s}$. Importantly a collection of planks of the type defined in (3.5) or (3.6) has a non-empty common intersection iff the convex hulls of these planks have a non-empty common intersection.

The theorem we are about to prove follows from the Nerve Theorem, which has a long history and many versions, which vary in assumptions and generality. The historically first instances appeared in the papers by Leray [16], Borsuk [5], and Weil [21]. We use a more recent version [2, Thm B/3.11], which, in particular, also talks about diagrams induced by inclusions. It assumes closed and convex sets and requires that every non-empty intersection of sets contains a point that is preserved in all inclusions.

Theorem 3.8. Let $\chi\colon A \to \sigma$ be a chromatic point set, and $\Gamma \subseteq \Sigma$ a subcomplex of the color simplex.

(1) For every radius r, the union of Γ-planks is homotopy equivalent to the Γ-subcomplex of the chromatic alpha complex: $\bigcup_{a\in A} \text{Plank}_r(a,\chi,\Gamma) \simeq \text{Alf}_r(\chi,\Gamma)$.

(2) The homotopy equivalences commute with inclusions. Specifically, if $r \le r'$ and $\Gamma \subseteq \Gamma'$, then the following two diagrams commute:

$$\begin{array}{ccc} \bigcup_{a\in A} \text{Plank}_{r'}(a,\chi,\Gamma) & \xleftrightarrow{\simeq} & \text{Alf}_{r'}(\chi,\Gamma) \\ \uparrow & & \uparrow \\ \bigcup_{a\in A} \text{Plank}_{r}(a,\chi,\Gamma) & \xleftrightarrow{\simeq} & \text{Alf}_{r}(\chi,\Gamma) \end{array} \qquad \begin{array}{ccc} \bigcup_{a\in A} \text{Plank}_{r}(a,\chi,\Gamma') & \xleftrightarrow{\simeq} & \text{Alf}_{r}(\chi,\Gamma') \\ \uparrow & & \uparrow \\ \bigcup_{a\in A} \text{Plank}_{r}(a,\chi,\Gamma) & \xleftrightarrow{\simeq} & \text{Alf}_{r}(\chi,\Gamma) \end{array}$$

Proof. Each union of planks of the type defined in (3.5) remains unchanged if we replace them by planks of the type defined in (3.6). We claim that for a fixed $r > 0$ and $\Gamma \subseteq \Sigma$, the nerve of the latter types of planks is homotopy equivalent to the Γ-subcomplex of the chromatic alpha complex:

$$\text{Nerve}(\{\text{Plank}^V_r(a,\chi,\Gamma) \mid a \in A\}) \cong \text{Alf}_r(\chi,\Gamma). \tag{3.7}$$

Indeed, $\nu \subseteq A$ has intersecting such planks iff their Voronoi balls intersect and the stars of the colors clipped to $|\Gamma|$ intersect. As argued before Lemma 3.4, the former happens iff $\nu \in \text{Alf}_r(\chi)$, and as

mentioned earlier in this subsection, the latter happens iff $\chi(\nu) \in \Gamma$. Together, these are the two defining conditions for $\nu \in \mathrm{Alf}_r(\chi, \Gamma)$. Assuming the two assumptions for the Nerve Theorem as stated in [2, Thm B/3.11] as satisfied, we thus have

$$\bigcup_{a \in A} \mathrm{Plank}_r(a, \chi, \Gamma) = \bigcup_{a \in A} \mathrm{Plank}_r^V(a, \chi, \Gamma) \simeq \mathrm{Nerve}(\{\mathrm{Plank}_r^V(a, \chi, \Gamma) \mid a \in A\}) \cong \mathrm{Alf}_r(\chi, \Gamma). \quad (3.8)$$

It thus remains to show that the two assumptions for the Nerve Theorem are indeed satisfied. The planks in (3.6) are not necessarily convex, but we can replace each by its convex hull without changing any common intersection of two or more of them. To satisfy the second assumption, we need a point in every non-empty common intersection of planks that is preserved by all relevant inclusions in the claimed diagrams. For $\mathrm{Plank}_r^V(a, \chi, \Gamma)$ we take the point $p_a^\triangle = a + u_{\chi(a)}$. For a collection, $\nu \subseteq A$, we find the radius, r, and the point, p_ν, such that $\bigcap_{a \in \nu} B_r^V(a, \chi) = p_\nu$, and we set $p_\nu^\triangle$ equal to p_ν plus the barycenter of the simplex spanned by the a_j with $j \in \nu$. Since $p_\nu^\triangle$ is the first point that appears in the common intersection of the growing planks, $p_\nu^\triangle$ is also contained in the common intersection if we substitute $r' \geq r$ for r or $\Gamma' \supseteq \Gamma$ for Γ. □

3.5. **Algorithm for Chromatic Alpha Complex.** This section discusses how to compute the chromatic alpha complexes. The procedure has two parts: first the construction of the chromatic Delaunay complex, and second the computation of the radius function on this complex. Following Corollary 3.7, the chromatic Delaunay complex is computed as a standard Delaunay complex of the chromatic lifting of the input chromatic point set; see [23] for fast and widely available code and [4] for bounds on the complexity for chromatic point sets.[1] Below we describe the computation of the radius function, and argue that for generic data (see Definition 2.1) in fixed dimension and with a constant number of colors, the algorithm takes time linear in the size of the chromatic Delaunay complex.

For the algorithm, it is convenient to work with the squared radius, which we use for the reminder of this section. We fix a chromatic point set, $\chi\colon A \to \sigma$, with $A \subseteq \mathbb{R}^d$ and $\#\sigma = s + 1$. Let $\nu \in \mathrm{Del}(\chi)$ be a simplex, $\tau = \chi(\nu)$ its colors, and $P = \bigcap_{v \in \nu} \mathrm{dom}(v, A_{\chi(v)})$ the intersection of the chromatic Voronoi domains of its vertices. Then the squared radius of the smallest empty stack that passes through the points in ν—the value of ν under the squared radius function—is the result of convex optimization:

$$\mathrm{Rad}^2(\nu) = \min_{x \in P} \max_{v \in \nu} \|x - v\|^2. \quad (3.9)$$

We write the convex optimization as an algorithm that avoids the explicit construction of P and can be implemented using exact arithmetic. For each $j \in \tau$, consider the affine subspace $E_j \subseteq \mathbb{R}^d$ consisting of the points x equidistant to all points in ν with color j, and the function $e_j\colon E_j \to \mathbb{R}$ that maps $x \in E_j$ to the squared distance to any point of ν with color j. Observe that the intersection of these affine subspaces, $E = \bigcap_{j \in \tau} E_j$, is the smallest affine subspace that contains P. The pointwise maximum function, $e\colon E \to \mathbb{R}$ defined by $e(x) = \max_{j \in \tau} e_j(x)$, is a strictly convex function with unique minimum, $y \in E$. If $y \in P$, then $e(y) = \min_{x \in P} e(x)$, so $\mathrm{Rad}^2(\nu) = e(y)$. Otherwise, $\min_{x \in P} e(x)$ is attained on the boundary of P, which implies that $\mathrm{Rad}^2(\nu)$ is the smallest $\mathrm{Rad}^2(\mu)$ over all cofaces μ of ν in $\mathrm{Del}(\chi)$. Note that to query whether $y \in P$, we only need to check whether the stack centered in y going through the points ν is empty, which is easy.

To formalize the algorithm, we write $S(x, r)$ for the $(d-1)$-sphere with center x and radius r in $\mathbb{R}^d$. The algorithm visits the simplices of $\mathrm{Del}(\chi)$ in the order of decreasing dimension:

```
for p = s + d downto 1 do
  for each p-simplex ν ∈ Del(χ) do
    STEP 1: construct the affine spaces E_j, j ∈ χ(ν), and E = ⋂_{j∈χ(ν)} E_j;
    STEP 2: construct e_j : E_j → ℝ, for each j ∈ χ(ν), and e : E → ℝ;
    STEP 3: find the unique minimum of e, the point y ∈ E;
    STEP 4: if S(y, √e_j(y)) is empty of A_j, for each j ∈ χ(ν)
              then Rad(ν) = e(y)
              else Rad(ν) = min{Rad(μ) | ν ⊆ μ, μ ∈ Del(χ), dim μ = p + 1}
            endif
  endfor
endfor
```

[1] Some implementations do not allow many points lying on the same affine subspace, and for others it slows down the computations. An alternative is to slightly perturb the lifted points, compute the Delaunay complex, and then only keep the down-set of those maximal simplices that span all colors.

Assume now that the chromatic lifting of the points is generic as in Definition 2.1—note that this is implied by *chromatic genericity* used in the next section; Definition 4.1. Assuming constant d and s, every step takes only constant time, except Step 4, which loops over cofaces both for checking emptiness and for determining the coface with smallest squared radius. We will see shortly that Step 4 takes constant time in the amortized sense.

Theorem 3.9. Let $A \subseteq \mathbb{R}^d$, $\chi\colon A \to \sigma$ be finite chromatic point set in general position, $s = \dim \sigma$, and m the number of simplices in $\mathrm{Del}(\chi)$. Assuming d and s are constants, $\mathrm{Rad}\colon \mathrm{Del}(\chi) \to \mathbb{R}$ can be computed in $O(m)$ time.

Proof. The body of the algorithm is executed once for each simplex $\nu \in \mathrm{Del}(\chi)$. It is easy to see that Steps 1 and 2 take only constant time each. To see the same for Step 3, we observe that y is the center of the smallest sphere that encloses all vertices of ν and whose center lies on E. The latter condition can be enforced by reflecting ν through E and adding its vertices to the points to be enclosed. The number of points to be considered is at most $2(d + s + 1) = O(1)$, so we can compute the smallest enclosing sphere in constant time with the miniball algorithm [22] or indeed a brute-force algorithm that checks all possibilities. In Step 4 we loop through cofaces of ν, both for checking emptiness of stacks, and in the to find the minimum radius in the `else`-clause. There can be many such cofaces for any individual ν, but any $(p+1)$-simplex $\mu \in \mathrm{Del}(\chi)$ is a coface of only $p + 2$ simplices. Since $p + 2 \leq d + s + 1$, this implies that in total we run at most $O\big((d + s + 1) \cdot m\big) = O(m)$ tests. □

4. Chromatic Radius Functions are Generalized Discrete Morse

As before, we assume that $\chi\colon A \to \sigma$ is a chromatic point set with finite $A \subseteq \mathbb{R}^d$ and $\sigma = \{0, 1, \ldots, s\}$, and we write $\mathrm{Rad}\colon \mathrm{Del}(\chi) \to \mathbb{R}$ for the chromatic radius function. By Definition 2.3, we call Rad a generalized discrete Morse function if every level set is a union of disjoint maximal intervals. The purpose of this section is to prove that Rad is indeed generalized discrete Morse, provided A satisfies a genericity condition, which we introduce first.

4.1. Chromatic Genericity.

A common genericity condition for a mono-chromatic set in $\mathbb{R}^d$ requires that any p-sphere passes through at most $p + 2$ points of the set. We extend this notion so it can be applied to the chromatic case.

Definition 4.1 (Chromatic Genericity)**.** Call a finite set $A \subseteq \mathbb{R}^d$ *chromatically generic* if every $k + 1$ concentric $(d-1)$-spheres contain at most $d + k + 1$ points, and the intersection with any affine p-plane contains at most $p + k + 1$ of these points.

Letting $k + 1$ be the number of spheres in a stack, and m the number of points on these spheres, we sometimes call $m - k - 1$ the *surplus* of the configuration. Definition 4.1 limits the surplus to d, or to p, respectively. We will use two equivalent formulations of chromatic genericity. To formulate them, write $E(B)$ for the maximal affine subspace whose points are equidistant to all points in B. It is the common intersection of all bisecting hyperplanes of any two points in B.

Lemma 4.2. Assume $A \subseteq \mathbb{R}^d$ is finite. Then the following two conditions are equivalent to A being chromatically generic:

(a) if $B_0, B_1, \ldots, B_k$ are non-empty disjoint subsets of A, then either $E = E(B_0) \cap E(B_1) \cap \ldots \cap E(B_k)$ is empty or the codimension of E is equal to the surplus, namely $\sum_{j=0}^{k} \#B_j - (k + 1)$;
(b) if $S_0, S_1, \ldots, S_k$ are concentric spheres, $B_j = S_j \cap A$, and $c_j \in B_j$ is an arbitrary but fixed choice for $0 \leq j \leq k$, then the vectors $\{b - c_j \mid b \in B_j \setminus \{c_j\}, 0 \leq j \leq k\}$ are linearly independent.

Proof. We establish the equivalences by showing that the chromatic genericity of A implies (a), that (a) implies (b), and that (b) implies the chromatic genericity of A.

Chromatic genericity ⇒ (a). We show the contrapositive. Let $B_0, B_1, \ldots, B_k$ be non-empty disjoint subsets that violate (a), and suppose that they minimize the surplus among all such violating collections. To violate (a), E must be non-empty and at least one of the sets must contain more than one point. Suppose $\#B_0 \geq 2$, let $x \in B_0$, write $B_0' = B_0 \setminus \{x\}$, and note that (a) holds for $B_0', B_1, \ldots, B_k$, by extremal assumption. The surplus of the latter collection is $\sum_{j=0}^{k} \#B_j - k$, which is therefore the codimension of $E' = E(B_0') \cap E(B_1) \cap \ldots \cap E(B_k)$. It is also the codimension of E, since E is contained in E' and the codimensions differ by at most one. It follows that the two spaces coincide. Write $B' = B_0' \sqcup B_1 \sqcup \ldots \sqcup B_k$, let H be the affine hull of B, and set $p = \dim H$. Let t be the smallest number for which there exist $t + 1$ concentric spheres, $S_0, S_1, \ldots, S_t$ in H, such that each of $B_0', B_1, \ldots, B_k$ lies on one of them. We

have $t \leq k$ since we may choose the common center of the spheres in E'. We claim $\#B' \geq p+t+1$. Assuming $\#B' < p+t+1$, the codimension of E' is less than p, so $\dim E' + \dim H > d$, which implies the existence of a line, L, common to H and E'. We can therefore find indices $0 \leq i < j \leq t$ and points $y \in S_i$ and $z \in S_j$ such that the bisector of y, z intersects L in a point, o. Choosing the common center of the spheres at o, we thus get only t spheres, which contradicts the choice of t. So $\#B' \geq p+t+1$, as claimed. Since the constructed stack of spheres is centered at a point in $E' = E \subseteq E(B_0)$, the sphere on which B'_0 lies also contains x. We thus have a stack of $t+1$ spheres in H that passes through more than $p+t+2$ points, which shows that A is not chromatically generic.

(a) $\Rightarrow$ (b). Assume S_j, B_j, c_j are as in (b). For each j, we write $U_j = \{b - c_j \mid b \in B_j \setminus \{c_j\}\}$ and note that $E(B_j)$ is a translate of the orthogonal complement of $\operatorname{span} U_j$. Writing $U = U_0 \cup U_1 \cup \ldots \cup U_k$, we thus get E as a translate of the orthogonal complement of $\operatorname{span} U$. By (a), the codimension of E is $\sum_{j=0}^{k} \#B_j - (k+1)$, which is therefore the dimension of $\operatorname{span} U$. But this is also the number of vectors in U, which implies that they are linearly independent, as required to get (b).

(b) $\Rightarrow$ chromatic genericity. Let $m = \sum_{j=0}^{k} \#B_j$. By (b), we get $m - (k+1)$ linearly independent vectors in $\mathbb{R}^d$. Therefore $m \leq d+k+1$. If all m points lie in a p-dimensional affine subspace, then the dimension of the span of the vectors is at most p, which implies $m \leq p+k+1$, as required for the chromatic genericity. □

We call the condition in Definition 4.1 generic, tacitly implying that it is satisfied by almost all finite sets. We prove that this is indeed the case.

Lemma 4.3. For each positive integer, n, the family of sets of n points in $\mathbb{R}^d$ that violate chromatic genericity, as a subset of $\mathbb{R}^{nd}$, is a finite union of sets with dimension at most $nd-1$.

Proof. It is enough to consider the d-dimensional condition in Definition 4.1. Indeed, a configuration that violates the p-dimensional condition implies at least $p+2$ points in an affine p-plane, and the sets that contain such $p+2$ points belong to a subset of dimension at most $nd-1$ of $\mathbb{R}^{nd}$. Consider $k+1$ concentric spheres and $d+k+2$ points on these spheres in $\mathbb{R}^d$. The surplus of this configuration is $d+1$, which violates chromatic genericity. Assigning the sum of squares of the d coordinates as a $(d+1)$-st coordinate to each point, we get $d+k+2$ points on $k+1$ parallel hyperplanes in $\mathbb{R}^{d+1}$. For each hyperplane, pick one of its points and take the difference vectors to the other points on this hyperplane. This gives a total of $(d+k+2)-(k+1) = d+1$ linearly dependent vectors in $\mathbb{R}^{d+1}$. Writing $(x_{i,1}, x_{i,2}, \ldots, x_{i,d+1})$ for the i-th vector, the $d+1$ vectors satisfy

$$\det \begin{bmatrix} x_{1,1} & x_{1,2} & \ldots & x_{1,d+1} \\ x_{2,1} & x_{2,2} & \ldots & x_{2,d+1} \\ \vdots & \vdots & \ddots & \vdots \\ x_{d+1,1} & x_{d+1,2} & \ldots & x_{d+1,d+1} \end{bmatrix} = 0.$$

This is a polynomial in the coordinates of $\mathbb{R}^{nd}$ that is not everywhere zero. Hence, its zero-set is a subspace of dimension at most $nd-1$.

We have such a polynomial for any $d+k+2$ points and their partition into $k+1$ sets. This is a finite collection as k is bounded by n. For a set of n points in $\mathbb{R}^d$ to be chromatically generic, it suffices to avoid the resulting finite number of zero-sets, each of dimension at most $nd-1$. □

4.2. **Convex Optimization.** The approach mimics the proof in [1, Section 4] that the radius function on the Delaunay mosaic of a mono-chromatic point set is generalized discrete Morse. As before, $\chi\colon A \to \sigma$ is a d-dimensional point set with $s+1$ colors. Given a collection of points $\nu \in \operatorname{Del}(\chi)$, we define the *smallest empty circumstack* as the solution to an optimization problem with variables $z \in \mathbb{R}^d$ for the center of the stack, and $r = (r_0, r_1, \ldots, r_s)$ for the radii of the spheres in the stack:

$$\begin{aligned} \underset{z,r}{\text{minimize}} \quad & \max\{r_0, r_1, \ldots, r_s\}, \\ \text{subject to} \quad & \|x - z\| = r_{\chi(x)} && \text{for } x \in \nu, \\ & \|x - z\| \geq r_{\chi(x)} && \text{for } x \in A \setminus \nu. \end{aligned}$$

We want to turn this into a differentiable convex optimization problem. Since the maximum function is not differentiable, we introduce a new variable, b. The constraints need to be either inequalities, $g \leq 0$, for convex differentiable g, or equalities or inequalities, for affine g. We switch to squared distances and radii so that $g_x(z) = \|x-z\|^2 - r^2_{\chi(x)}$ is differentiable and strictly convex, but the inequalities are in

the wrong direction. We substitute new variables, $a = (a_0, a_1, \ldots, a_s)$, for the radii, constrain b to be smaller than or equal to all a_i, and get an equivalent optimization problem, whose constraints are affine with respect to the new variables:

$$\begin{aligned} a_j &= \|z\|^2 - r_j^2 \quad \text{for all } j = 0, 1, \ldots, s, \\ g_x(z, a) \;&=\; \|x - z\|^2 - r_{\chi(x)}^2 \;=\; \|x\|^2 - 2\langle x, z\rangle + \|z\|^2 - r_{\chi(x)}^2 \;=\; \|x\|^2 - 2\langle x, z\rangle + a_{\chi(x)}. \end{aligned}$$

For any $\nu \in \mathrm{Del}(\chi)$, we then get the following differentiable convex optimization problem, (P_ν), in which we write $h_j(a, b) = b - a_j$:

$$\begin{array}{lll} \underset{z,a,b}{\text{minimize}} & f(z, b) = \|z\|^2 - b, & \\ \text{subject to} & h_j(a, b) \leq 0 & \text{for } j = 0, 1, \ldots, s, \\ & g_x(z, a) = 0 & \text{for } x \in \nu, \\ & g_x(z, a) \geq 0 & \text{for } x \in A \setminus \nu. \end{array}$$

We formulate the task this way in order to make use of duality, which is a powerful tool in convex optimization [6, Chapter 5.2]. The *Lagrange dual problem* of (P_ν), denoted (D_ν) and with variables $\lambda = (\lambda_x)_{x \in A}$ and $\mu = (\mu_j)_{j \in \sigma}$, is the following:

$$\begin{array}{lll} \underset{\lambda,\mu}{\text{maximize}} & G(\lambda, \mu) = \inf\limits_{z,a,b} f(z, b) + \sum\limits_{x \in A} \lambda_x g_x(z, a) + \sum\limits_{j \in \sigma} \mu_j h_j(a, b), & \\ \text{subject to} & \lambda_x \leq 0 & \text{for } x \in A \setminus \nu, \\ & \mu_j \geq 0 & \text{for } j \in \sigma. \end{array}$$

Because the sums in (D_ν) are both non-positive, the value of the dual problem for a feasible solution is always smaller than or equal to the value of the primal problem for any of its feasible solutions. This is, in particular, true for the optimal values, and the difference between those is referred to as the *optimal duality gap*. Under the chromatic genericity conditions of Definition 4.1, the gap for problem (P_ν) is guaranteed to be zero and attained by some z, a, b, λ, μ. This is because (P_ν) is convex and satisfies the Slater's condition [6, Section 5.2.3], which states that there exists a feasible solution with all inequalities strict. For (P_ν), this translates to the claim that if there exists an empty circumstack of ν, then there also exists an empty circumstack of ν that contains no points from $A \setminus \nu$, which is implied by Lemma 4.2 (a)—indeed, if spaces of equidistant points intersect generically, then so do Voronoi cells, and the desired stack can be centered at any point in the interior of $\bigcap_{v \in \nu} \mathrm{dom}(v, A_{\chi(v)})$.

Since (P_ν) is, in addition to convexity, also differentiable, points z, a, b, λ, μ are primal and dual optima iff they satisfy the *Karush-Kuhn-Tucker* (KKT) conditions [6, Section 5.5.3]:

$$z, a, b \text{ is primal feasible, and } \lambda, \mu \text{ is dual feasible,} \tag{4.1}$$

$$\lambda_x \cdot g_x(z, a) = 0 \text{ for all } x \in A, \text{ and } \mu_j \cdot h_j(a, b) = 0 \text{ for all } j \in \sigma, \tag{4.2}$$

$$\nabla f(z, b) + \sum\nolimits_{x \in A} \lambda_x \nabla g_x(z, a) + \sum\nolimits_{j \in \sigma} \mu_j \nabla h_j(a, b) = 0. \tag{4.3}$$

Note that the second condition implies that $\lambda_x \neq 0$ only if $g_x(z, a) = 0$, i.e., if x lies on the stack. Similarly, $\mu_j \neq 0$ only if the j-colored sphere has maximum radius among the spheres on the stack. Below we give the gradients at point (z, a, b) needed in the last condition:

$$\begin{array}{ll} \nabla f = (2z; 0, \ldots, 0; -1), & \\ \nabla g_x = (-2x; 0, \ldots, 0, 1, 0, \ldots, 0; 0) & \text{with 1 at the position corresponding to } a_{\chi(x)}, \\ \nabla h_j = (0; 0, \ldots, 0, -1, 0, \ldots, 0; 1) & \text{with } -1 \text{ at the position corresponding to } a_j. \end{array}$$

Putting all the above observations together—the existence of zero gap solutions, the KKT conditions, and the gradients—we get the following:

Lemma 4.4. Let $\chi\colon A \to \sigma$, with $A \subseteq \mathbb{R}^d$ chromatically generic, and $\nu \in \mathrm{Del}(\chi)$. Let z, a, b describe an empty circumstack of ν, and let $\eta \supseteq \nu$ contain all points of A that lie on this stack. Then this is the

smallest empty circumstack of ν iff there exist λ_x for $x \in A$ and μ_j for $j \in \sigma$ such that:

$$\sum\nolimits_{x\in\eta} \lambda_x x = z, \tag{4.4}$$

$$\sum\nolimits_{x\in\eta_j} \lambda_x = \mu_j, \text{ in which } \eta_j = \eta \cap \chi^{-1}(j), \tag{4.5}$$

$$\sum\nolimits_{j\in\sigma} \mu_j = 1, \tag{4.6}$$

$$\lambda_x \leq 0 \text{ for } x \in A \setminus \nu, \text{ and } \lambda_x = 0 \text{ if } x \in A \setminus \eta, \tag{4.7}$$

$$\mu_j \geq 0 \text{ for all } j \in \sigma, \text{ and } \mu_j = 0 \text{ if the } j\text{-th sphere does not have maximum radius.} \tag{4.8}$$

Proof. We argue the five conditions in reverse order, from (4.8) to (4.4). The inequality in (4.8) just rewrites the second condition in (D_ν), and the strengthening equality is implied by the second slackness condition in (4.2). Similarly, the inequality in (4.7) just rewrites the first condition in (D_ν), and the strengthening equality is implied by the first slackness condition in (4.2). For the remaining three conditions, we plug the gradients into (4.3):

$$(2z; 0, \ldots, 0; -1) + \sum\nolimits_{x\in A} \lambda_x(-2x; 0, \ldots, 1, \ldots, 0; 0) + \sum\nolimits_{j\in\sigma} \mu_j(0; 0, \ldots, -1, \ldots, 0; 1) = 0. \tag{4.9}$$

Comparing the last coordinates, we get (4.6), and comparing the coordinates that correspond to color j, we get (4.5). Combining (4.5) and (4.6), we get $\sum_{x\in A} \lambda_x = \sum_{x\in\eta} \lambda_x = 1$ because $\lambda_x = 0$ if $x \in A \setminus \eta$. Now comparing the first coordinates of the gradients, we get (4.4). □

Since (4.5) and (4.6) imply $\sum_{x\in\eta} \lambda_x = 1$, the dual solution expresses the center of the stack as an affine combination of the points on the spheres in this stack; see (4.4). It pays to unpack this interpretation by distinguishing the spheres with maximum radius from the others. Let $\sigma' \subseteq \sigma$ be the colors whose spheres have maximum radius and write $\sigma'' = \sigma \setminus \sigma'$. For each $i \in \sigma''$, we have $\mu_i = 0$, so the sum of the corresponding λ_x vanishes, which suggests we interpret the corresponding combination as a vector: $\sum_{x\in\eta_i} \lambda_x x = \sum_{x\in\eta_i} \lambda_x(x - y)$, in which $y = y(i)$ is an arbitrary but fixed point in η_i. For each $j \in \sigma'$, the corresponding combination of the points is $\sum_{x\in\eta_j} \lambda_x x$. With this, we can rewrite (4.4) as

$$z = \sum_{j\in\sigma'} \sum_{x\in\eta_j} \lambda_x x + \sum_{i\in\sigma''} \sum_{x\in\eta_i} \lambda_x(x - y(i)).$$

The λ_x in the first sum add up to 1, so we can interpret this first sum as an affine combination, while we think of the second sum as a vector that moves us from this affine combination to the center of the smallest empty circumstack.

4.3. **Proof of Generalized Discrete Morse Property.** The crucial insight that turns the dual solution into a proof that the chromatic radius function is generalized Morse is the following: when we remove a point x from ν, this only affects Condition (4.7) in Lemma 4.4. Therefore, if $\lambda_x \leq 0$, then the smallest empty circumstack of ν is still the smallest empty circumstack of $\nu \setminus \{x\}$. The idea is that we remove all points from ν with non-positive coefficient and thus obtain the minimum of the interval that contains ν. For this, it is important that we identify the points uniquely, but this is guaranteed by the chromatic genericity of the points, which ascertains that the optimal dual solution is unique.

Lemma 4.5. Let $\chi\colon A \to \sigma$ be a chromatic point set in $\mathbb{R}^d$, and let z, a, b describe an empty stack that passes through the points in $\eta \subseteq A$. If A is chromatically generic, then there exists at most one set of parameters λ_x and μ_j, with $x \in A$ and $j \in \sigma$, that satisfies the conditions of Lemma 4.4.

Proof. By Conditions (4.7) and (4.8), we have $\lambda_x = 0$ if $x \in A \setminus \eta$ and $\mu_j = 0$ if $j \in \sigma''$, in which we recall that $\sigma = \sigma' \sqcup \sigma''$ and σ' are the colors whose spheres have maximum radius. It thus suffices to show that the linear relation (4.3) restricted to points $x \in \eta$ and colors $j \in \sigma'$,

$$\nabla f(z, b) + \sum\nolimits_{x\in\eta} \lambda_x \nabla g_x(z, a) + \sum\nolimits_{j\in\sigma'} \mu_j \nabla h_j(a, b) = 0,$$

has at most one solution. We do this by showing that the ∇g_x and ∇h_j are linearly independent. Writing these vectors as the columns of a matrix, we perform elementary column operations to make it obvious that the columns are linearly independent. First simplify the notation by assuming $\sigma' = \{0, 1, \ldots, k\}$, and replace ∇h_j by $\nabla h_j - \nabla h_0$, for $1 \leq j \leq k$. The resulting k vectors are the respective first columns of blocks 1 to k of the matrix in Table 1. Furthermore, we replace ∇g_x by $\nabla g_x + \nabla h_j - \nabla h_0$ for every $x \in \eta_j$ and $1 \leq j \leq k$, which effectively moves the 1 in row $d + j + 1$ to row $d + 1$. Recall that the first d coordinates of ∇g_x are those of $-2x$. We may replace them by the coordinates of x without affecting

	block 0				block 1				block k				block $k+1$			
$1, 2, \ldots, d$	0	c_0	$x - c_0$	$\ldots$	0	$x - c_0$	$\ldots$	$\ldots$	0	$x - c_0$	$\ldots$		c_{k+1}	$x - c_{k+1}$	$\ldots$	$\ldots$
$d+1$	-1	1			1			$\ldots$	1							$\ldots$
$d+2$					-1			$\ldots$								$\ldots$
$\vdots$								$\vdots$								$\vdots$
$d+k$								$\ldots$	-1							$\ldots$
$d+k+1$								$\ldots$					1			$\ldots$
$\vdots$								$\vdots$								$\vdots$
$d+s+2$	1							$\ldots$								$\ldots$

TABLE 1. The columns in this matrix are the gradient vectors after combining them as explained in the proof of Lemma 4.5. Zero entries are left blank. There are $s+1$ blocks of columns, one for each color. The respective first columns of the first $k+1$ blocks contain ∇h_0 and $\nabla h_j - \nabla h_0$, for $1 \leq j \leq k$. The points of colors 0 to k all lie on the same sphere, and we get vectors by subtracting the same point, c_0, from each such point. For each color $j \geq k+1$, we get vectors from the points in block j by subtracting an arbitrary but fixed point $c_j \in B_j$.

the linear independence of the vectors. Finally, choose an arbitrary but fixed $c_0 \in \eta_0$, and replace ∇g_x by $\nabla g_x - \nabla g_{c_0}$, for all $x \neq c_0$ in η_j with $0 \leq j \leq k$. Similarly, for each other color, $k+1 \leq j \leq s$, choose an arbitrary but fixed point, $c_j \in \eta_j$, and replace ∇g_x by $\nabla g_x - \nabla g_{c_j}$, for each $x \in \eta_j \setminus \{c_j\}$; see the first row of the matrix in Table 1.

It is now easy to see that the columns in the matrix are linearly independent. To begin, collect all columns that start with $x - c_j$. Their topmost d positions contain the vectors considered and found linearly independent in Lemma 4.2 (b). All the remaining columns have their unique pivots in the $s+2$ rows below the top d rows, so adding them preserves the linear independence. □

Guaranteeing uniqueness of the dual solution is, indeed, necessary to have the simplices organized in intervals. See Figure 10 for an example of points not chromatically generic whose radius function is not generalized discrete Morse. The common center of the two circles is the point z, which can be expressed as an affine combination satisfying conditions in Lemma 4.4 in more than one way. We need $\lambda_a + \lambda_b = 1$ and $\lambda_c + \lambda_d = 0$. We can express z as a combination of a, b, and not use c, d at all, we can start at a combination of a, b to the left of a and then move to z along a positive multiple of the vector $d - c$, or we can start at a combination of a, b to the right of b and move to z along a negative multiple of $d - c$. This leads to $(\lambda_a, \lambda_b, \lambda_c, \lambda_d)$ having signatures $(+, +, 0, 0)$, $(+, -, -, +)$, and $(-, +, +, -)$, respectively. As explained in the caption of Figure 10, this ambiguity prevents the formation of intervals.

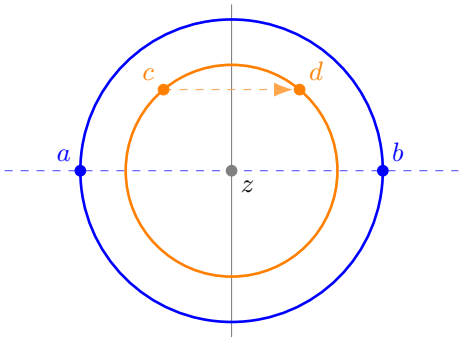

FIGURE 10. An example showing that chromatic genericity is needed for the radius function to be generalized discrete Morse. Two *blue* points, a, b, and two *orange* points, c, d, share a common bisector and, therefore, violate the chromatic genericity condition in Definition 4.1. Indeed, the four points also lie on a common circle. The two shown circles belong to the smallest empty circumstack of the simplex $abcd$. This stack is still the smallest empty circumstack for the edges ab, ad, and bc, but not of any single vertex. Hence, the set of simplices that share the radius with $abcd$ does not have a unique minimum and is therefore not an interval.

We are now ready to argue that the radius function on the chromatic Delaunay complex is generalized discrete Morse, provided the points are chromatically generic. To this end, we construct the interval

that contains a given simplex, ν, in the Delaunay complex of $\chi\colon A \to \sigma$. Starting with the smallest empty circumstack of ν, we add a sphere for every color that is not yet represented and contains a point at distance at most the radius of the stack from its center. Specifically, we add the unique sphere that shares the center with the other spheres and passes through the closest point of that color. Call this the *augmented smallest empty circumstack* of ν, and write $\nu_{\max}$ for the points its spheres pass through. This augmented stack is the smallest empty circumsphere of $\nu_{\max}$. By Lemma 4.5, there is a unique dual solution, λ, μ. This dual solution assigns a coefficient λ_x to each point $x \in A$, and we write $\nu_{\min}$ for the points with $\lambda_x > 0$. By construction, $\nu \subseteq \nu_{\max}$, and by Condition (4.7) in Lemma 4.4, $\nu_{\min} \subseteq \nu$.

Theorem 4.6. Let $A \subseteq \mathbb{R}^d$ be chromatically generic and $\chi\colon A \to \sigma$ a chromatic point set. Then the chromatic radius function, $\mathrm{Rad}\colon \mathrm{Del}(\chi) \to \mathbb{R}$, is generalized discrete Morse.

Proof. Given $\nu \in \mathrm{Del}(\chi)$, we show that the simplices ν' that satisfy $\nu_{\min} \subseteq \nu' \subseteq \nu_{\max}$ are the unique interval of Rad that contains ν. Let λ, μ be the parameters of the dual solution to the smallest empty circumstack of $\nu_{\max}$. By Lemma 4.5, λ and μ are unique, and by Lemma 4.4, we get the same dual solution of every ν' that satisfies $\nu_{\min} \subseteq \nu' \subseteq \nu_{\max}$. This implies $\mathrm{Rad}(\nu') = \mathrm{Rad}(\nu)$ for every such ν'.

Next we show that $\mathrm{Rad}(\nu'') < \mathrm{Rad}(\nu)$ for every $\nu'' \subseteq \nu_{\max}$ that does not contain $\nu_{\min}$. Consider the smallest empty circumstack of ν'', and let λ'', μ'' be the dual solution. By Lemma 4.4, we get $\lambda''_x \leq 0$ for every $x \in \nu_{\min} \setminus \nu''$, but we have $\lambda_x > 0$ by construction of $\nu_{\min}$. By Lemma 4.5, this implies that the smallest empty circumstack of ν'' is different from that of ν. Since the two smallest empty circumstacks are different, they have different centers and, by strict convexity, different radii, so $\mathrm{Rad}(\nu'') < \mathrm{Rad}(\nu)$.

It remains to show that $\mathrm{Rad}(\nu''') > \mathrm{Rad}(\nu)$ for every $\nu''' \supseteq \nu_{\min}$ that is not contained in $\nu_{\max}$. The points $y \in \nu''' \setminus \nu_{\max}$ do not lie on the smallest empty circumstack of $\nu_{\max}$, which implies that the smallest empty circumstack of ν''' is different from that of ν. Since ν''' contains $\nu_{\min}$, the stack of ν''' must therefore be larger than that of ν. As before, the smallest empty circumstacks are different, which implies they have different centers and radii, so $\mathrm{Rad}(\nu''') > \mathrm{Rad}(\nu)$. □

We close this section with two remarks. The first notes that the smallest empty circumstack of a maximum, $\nu_{\max}$, is also the smallest circumstack of $\nu_{\max}$ (without enforcing emptiness). This is evident from the formulation as an optimization problem: all constaints $g_x \geq 0$ are inactive so all corresponding λ_x vanish. If we remove these constraints, the same parameters z, a, b, λ, μ still satisfy the KKT conditions and therefore describe the same optimal solution.

Given a chromatically generic $A \subseteq \mathbb{R}^d$, we may compare the Delaunay complexes and their radius functions in the chromatic and the mono-chromatic cases. As stated in Lemma 3.4, the sublevel sets at matching thresholds have the same homotopy type. Since the type changes whenever we add a critical simplex—which is characterized by $\nu_{\min} = \nu_{\max}$—this implies that the two radius functions have the same critical values. The optimization perspective reveals that the critical values belong to the same critical simplices. Indeed, the smallest empty circumstack of a chromatic critical simplex, ν, is in fact a single circumsphere: if $\nu_{\max} = \nu_{\min}$, then its dual solution has $\lambda_x > 0$ for all $x \in \nu$, so $\mu_j > 0$ for all $j \in \chi(\nu)$, and all spheres share the same, maximum radius. Since the chromatic radius is bounded from above by the mono-chromatic radius, both agree on ν. To see that ν is also critical in the mono-chromatic case, note that problem (P_ν), but with inequalities $h_j \leq 0$ changed to equalities, has the mono-chromatic radius as the optimum. The only change in the dual problem is that we lose inequalities on μ_j. Since Lemma 4.5 guarantees uniqueness regardless of the inequalities, this implies that $\lambda_x > 0$ also for the modified dual problem, so ν is critical also in the mono-chromatic case.

5. Persistent Homology of Chromatic Alpha Complexes

In this section, we review the background needed to turn the chromatic alpha complexes into persistence diagrams, and we advocate the use of six such diagrams, which we refer to as a 6-pack. In addition, we discuss the relations between the diagrams in a 6-pack, as well as relations between different 6-packs arising from different choices of the subcomplex.

5.1. **Background: Persistent Homology.** The goal of this subsection is to introduce the framework of persistent homology [11], together with its kernel, image, and cokernel generalizations [8]. We keep the formalism to a minimum by limiting ourselves to simplicial complexes and $\mathbb{Z}/2\mathbb{Z}$ coefficients.

Homology with $\mathbb{Z}/2\mathbb{Z}$ Coefficients. Loosely speaking, homology is an algebraic framework that defines and counts holes in a shape. Given a simplicial complex, K, a *p-chain* is a subset of p-simplices. The *sum* of two p-chains is the symmetric difference of the two sets: if a p-simplex belongs to both chains, the two copies erase each other, as $1 + 1 = 0$ in modulo-2 arithmetic. The *boundary* of a p-simplex is the set of $(p-1)$-dimensional faces, which is a $(p-1)$-chain. The p-chains with the sum operation form a group, $\mathsf{C}_p(K)$, and the *boundary operator*, $\partial_p \colon \mathsf{C}_p(K) \to \mathsf{C}_{p-1}(K)$, maps a p-chain to the sum of its simplices' boundaries. A *p-cycle* is a p-chain with empty boundary, a *filling* of this p-cycle is a $(p+1)$-chain whose boundary is the p-cycle, and a *p-boundary* is a p-cycle for which there exits a filling. The p-boundaries and the p-cycles form groups by themselves, and since every p-boundary is a p-cycle, and every p-cycle is a p-chain, we get three nested groups: $\mathsf{B}_p(K) \subseteq \mathsf{Z}_p(K) \subseteq \mathsf{C}_p(K)$. Two p-cycles are *homologous* if their sum has a filling or, equivalently, adding a p-boundary to one p-cycle gives the other p-cycle. Being homologous is an equivalence relation, whose equivalence classes are the elements of the *p-th homology group*: $\mathsf{H}_p(K) = \mathsf{Z}_p(K)/\mathsf{B}_p(K)$. All mentioned groups are vector spaces, so the *ranks* are their dimensions. Of particular importance is the *p-th Betti number* of K, which is the rank of the p-th homology group: $\operatorname{rank} \mathsf{H}_p(K) = \operatorname{rank} \mathsf{Z}_p(K) - \operatorname{rank} \mathsf{B}_p(K)$.

Let L be a subcomplex of K. Relative homology describes the connectivity of the topological pair (K, L), which geometrically represents K with the subspace L identified as a single point. The chain groups are the quotients $C_p(K, L) = C_p(K)/C_p(L)$. Cycle and boundary subgroups are defined as before, and their quotients are the *relative homology groups of the pair*, denoted $H_p(K, L)$. The homology groups and their relative cousins are related by the following long exact sequence:

$$\ldots \to \mathsf{H}_p(L) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K, L) \to \mathsf{H}_{p-1}(L) \to \ldots \tag{5.1}$$

A well known basic property of long exact sequences is that the alternating sum of dimensions of the vector spaces vanishes.

Lemma 5.1. Let $L_i \subseteq K_i$ be simplicial complexes. Then

$$\sum\nolimits_{p \in \mathbb{Z}} (-1)^p [\operatorname{rank} \mathsf{H}_p(L_i) - \operatorname{rank} \mathsf{H}_p(K_i) + \operatorname{rank} \mathsf{H}_p(K_i, L_i)] = 0.$$

Proof. By definition of exactness, the rank of each homology group in (5.1) can be written as the sum of two non-negative integers such that it shares one with the preceding group and the other with the succeeding group along the sequence. Since only finitely many groups have non-zero ranks, this implies that the alternating sum of ranks vanishes. □

Persistent Homology. In the following, let $f \colon K \to \mathbb{R}$ be monotonic, with values $r_1 < r_2 < \ldots < r_n$, and let $K_i = f^{-1}(-\infty, r_i]$ be its i-th sublevel set. Applying the p-th homology functor to $\emptyset = K_0 \subseteq K_1 \subseteq \ldots \subseteq K_n$, we get a sequence of vector spaces:

$$\mathsf{H}_p(K_0) \to \ldots \to \mathsf{H}_p(K_{i-1}) \to \mathsf{H}_p(K_i) \to \ldots \to \mathsf{H}_p(K_{j-1}) \to \mathsf{H}_p(K_j) \to \ldots \to \mathsf{H}_p(K_n).$$

There is one such sequence for each dimension, p. The inclusions $K_i \subseteq K_j$ induce maps $\mathsf{f}_i^j \colon \mathsf{H}_p(K_i) \to \mathsf{H}_p(K_j)$ for all $0 \leq i \leq j \leq n$. This sequence is called a *persistence module*. It can be written as a direct sum of indecomposable modules of the form $\ldots \to 0 \to \mathsf{k} \to \ldots \to \mathsf{k} \to 0 \to \ldots$, where $\mathsf{k} = \mathbb{Z}/2\mathbb{Z}$, all maps between these 1-dimensional vector spaces are identities, and all others are zero maps. Each indecomposable module has a concrete interpretation, namely a *birth* followed by a *death* of a homology class. Specifically, we have such an indecomposable module from position i to position $j-1$ if

- there is a class, $\gamma \in \mathsf{H}_p(K_i)$ that does not belong to the image of f_{i-1}^i, and
- $\mathsf{f}_i^{j-1}(\gamma)$ does not belong to the image of f_{i-1}^{j-1}, but $\mathsf{f}_i^j(\gamma)$ belongs to the image of f_{i-1}^j.

We say γ is *born* at K_i and *dies entering* K_j. We record this information with the point $(f(r_i), f(r_j))$, noting that the second coordinate is ∞ if the class is born but never dies. The resulting multi-set of points in the extended plane is the *p-th persistence diagram* of f, denoted $\operatorname{Dgm}_p(f)$. Sometimes, we drop the index and write $\operatorname{Dgm}(f)$ for the disjoint union of the $\operatorname{Dgm}_p(f)$ over all dimensions, p.

If L is a subcomplex of K, we get a filtration, L_i, by restricting f to L. The inclusions of the pairs $(K_0, L_0) \subseteq (K_1, L_1) \subseteq \ldots \subseteq (K_n, L_n)$ give rise to a sequence of relative homology groups,

$$\mathsf{H}_p(K_0, L_0) \to \ldots \to \mathsf{H}_p(K_{i-1}, L_{i-1}) \to \mathsf{H}_p(K_i, L_i) \to \ldots \to \mathsf{H}_p(K_n, L_n).$$

Applying the above definitions to this sequence yields the *p-th relative persistent diagram.*

An important property of the persistence diagram is its stability. Specifically, the bottleneck distance between the diagrams of $f, g\colon K \to \mathbb{R}$ is bounded from above by the L_∞-distance between the two maps:

$$W_\infty(\mathrm{Dgm}_p(f), \mathrm{Dgm}_p(g)) \leq \|f - g\|_\infty; \tag{5.2}$$

see [7]. The *persistence* of a point in the persistence diagram is the vertical distance to the diagonal, $|f(r_j) - f(r_i)|$, and the 1-*norm* is the sum of persistences of the points in the diagram, denoted $\|\mathrm{Dgm}(f)\|_1$. To cope with points at infinity, we use a *cut-off*, C, which we effectively substitute for ∞ (and for birth- and death-values larger than the threshold). This gives finite 1-norms and preserves relationships implied by exact sequences, as expressed in Theorem 5.3.

Kernels, Images, and Cokernels. Let $L \subseteq K$ be simplicial complexes, $f_K\colon K \to \mathbb{R}$ monotonic, and $f_L\colon L \to \mathbb{R}$ the restriction of f_K to L. Taking sublevel sets, we get two parallel persistence modules and maps from one module to the other:

$$\begin{array}{ccccccccccc}
\mathsf{H}_p(K_0) & \to & \ldots & \to & \mathsf{H}_p(K_i) & \to & \mathsf{H}_p(K_{i+1}) & \to & \ldots & \to & \mathsf{H}_p(K_n) \\
\uparrow & & \ldots & & \uparrow & & \uparrow & & \ldots & & \uparrow \\
\mathsf{H}_p(L_0) & \to & \ldots & \to & \mathsf{H}_p(L_i) & \to & \mathsf{H}_p(L_{i+1}) & \to & \ldots & \to & \mathsf{H}_p(L_n).
\end{array}$$

Write $\kappa_i\colon \mathsf{H}_p(L_i) \to \mathsf{H}_p(K_i)$ for the vertical maps, which are induced by the inclusions $L_i \subseteq K_i$, for $0 \leq i \leq n$. These maps have kernels, images, and cokernels, which form persistence modules of their own:

$$\begin{array}{ccccccccccc}
\ker_p \kappa_0 & \to & \ldots & \to & \ker_p \kappa_i & \to & \ker_p \kappa_{i+1} & \to & \ldots & \to & \ker_p \kappa_n, \\
\mathrm{im}_p \kappa_0 & \to & \ldots & \to & \mathrm{im}_p \kappa_i & \to & \mathrm{im}_p \kappa_{i+1} & \to & \ldots & \to & \mathrm{im}_p \kappa_n, \\
\mathrm{cok}_p \kappa_0 & \to & \ldots & \to & \mathrm{cok}_p \kappa_i & \to & \mathrm{cok}_p \kappa_{i+1} & \to & \ldots & \to & \mathrm{cok}_p \kappa_n.
\end{array}$$

These persistence modules were introduced and studied in [8]. Following the notation in that paper, we write $\mathrm{Dgm}(\ker f_L \to f_K)$, $\mathrm{Dgm}(\mathrm{im}\, f_L \to f_K)$, and $\mathrm{Dgm}(\mathrm{cok}\, f_L \to f_K)$ for the corresponding persistence diagrams. These diagrams are also stable under perturbations of the monotonic function, and they can be computed efficiently. We omit details and refer to [8], in particular for the matrix reduction algorithms, which we have implemented [10] to study the meaning of these derived persistence diagrams for chromatic point sets.

5.2. **6-pack of Persistent Diagrams.** The main new concept in this section is a collection of six related persistence diagrams, which we use to quantify the way different point sets mingle. We call this collection a 6-*pack*. A 6-pack can be defined for any pair of topological spaces $L \subseteq K$ with a filtration on K. We explain the construction on a concrete example: the blue circle on an orange background in Figure 11. Let $K = \mathrm{Del}(\chi)$ be the chromatic Delaunay complex of the portrayed chromatic set, and let $L \subseteq K$ be the blue subcomplex, consisting of those simplices in K that only have blue vertices. Let $f_K\colon K \to \mathbb{R}$ be the chromatic radius function, and write f_L and $f_{K,L}$ for the restrictions of f_K to L and $K \setminus L$. The radius function and its restrictions give rise to three persistence modules, and we get

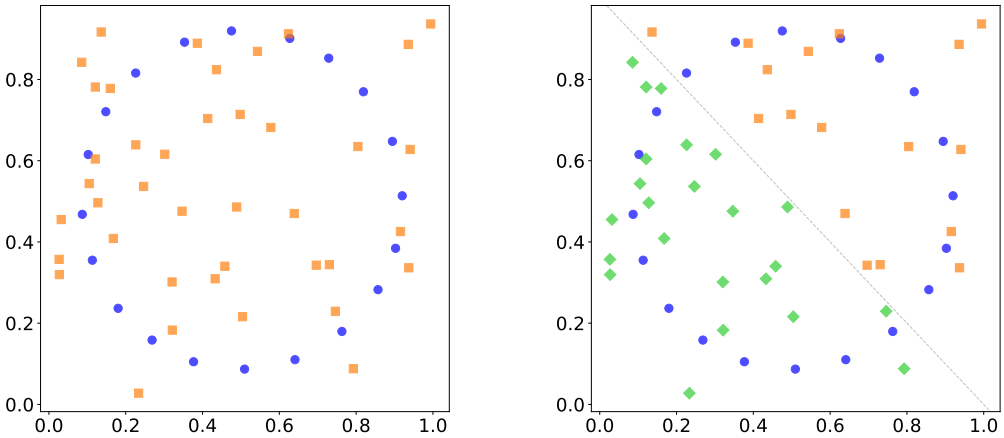

FIGURE 11. A bi-chromatic point set on the *left*, and a tri-chromatic point set on the *right*. The *dotted* line indicates the separation of *green* from *orange* points that form the background for the *blue* circle.

three additional persistence modules for the kernel, the image, and the cokernel of the map on homology induced by the inclusion $L \subseteq K$; see Section 5.1. The persistence diagrams in the 6-pack are arranged as in Table 2, in a manner that lends itself to comparing the information between them.

kernel: $\mathrm{Dgm}(\ker f_L \to f_K)$	relative: $\mathrm{Dgm}(f_{K,L})$	cokernel: $\mathrm{Dgm}(\mathrm{cok}\, f_L \to f_K)$
domain: $\mathrm{Dgm}(f_L)$	image: $\mathrm{Dgm}(\mathrm{im}\, f_L \to f_K)$	codomain: $\mathrm{Dgm}(f_K)$

TABLE 2. The arrangement of the persistence diagrams in the 6-pack for the pair $L \subseteq K$ in two rows and three columns. Read the six positions in a circle so that the diagram of the domain lies between those of the kernel and the image, the diagram of the image lies between those of the domain and the codomain, etc.

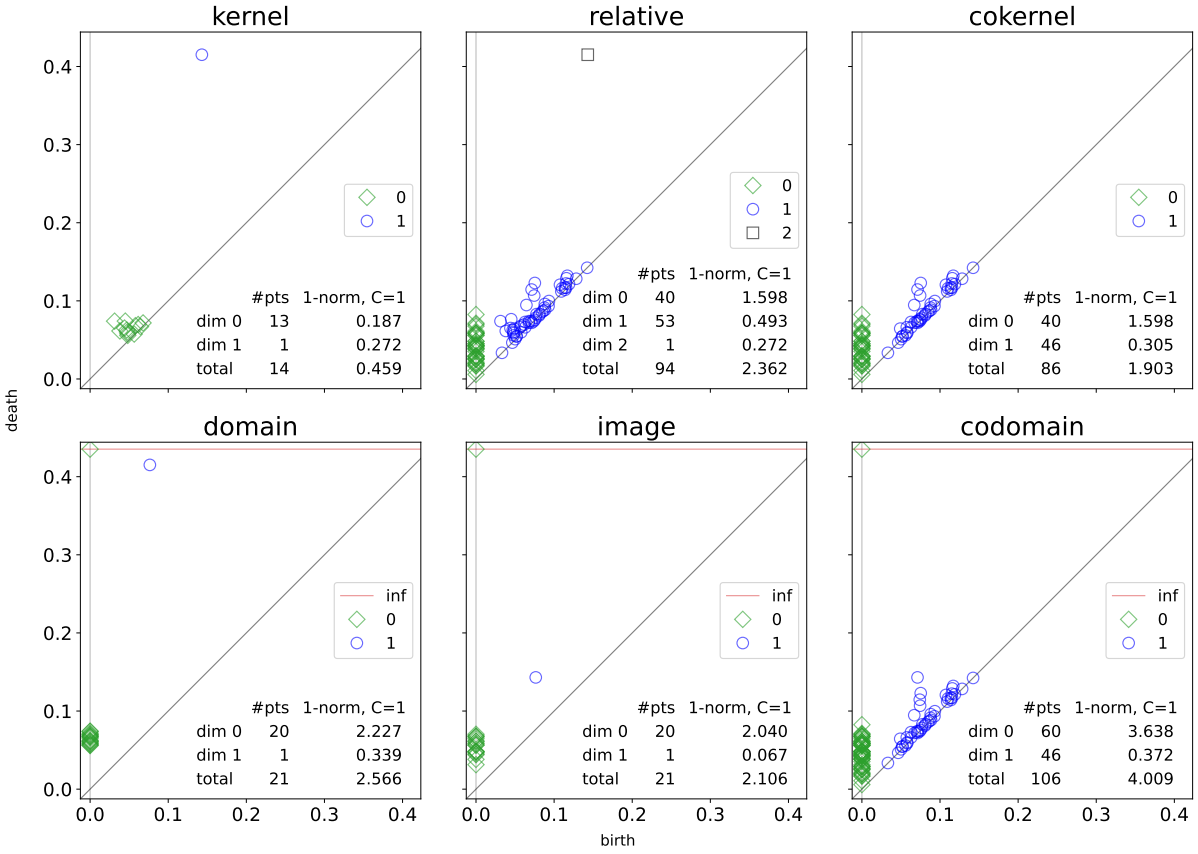

FIGURE 12. The 6-pack for the bi-chromatic point set in the *left panel* of Figure 11. The domain, L, is the blue subcomplex of the codomain, K, which is the 3-dimensional chromatic Delaunay mosaic of the *blue* and *orange* points.

Figure 12 displays the 6-pack for the point set in the left panel of Figure 11 with the blue subcomplex chosen as L. Not surprisingly, the circle of blue points gives rise to a persistent 1-cycle in L captured in the diagram of the domain. At the time of its birth, this 1-cycle includes into a non-trivial 1-cycle in K, so we get a point with the same birth-coordinate in the diagram of the image. When the circle is filled by orange disks, it becomes a trivial 1-cycle in K, which is marked by its death in the image and the simultaneous birth in the kernel. Eventually, the blue circle is filled by blue disks, so it dies in the domain and simultaneously in the kernel. To summarize, the point (a, c) in the diagram of the domain splits into two points, (a, b) in the diagram of the image, and (b, c) in the diagram of the kernel. While the split into two like this is a common phenomenon, not all points split in this manner; see the relations in the next subsection. The point (b, c) can also be seen one dimension higher in the relative persistence diagram of the pair. Indeed, there is a non-bounding 2-cycle in the quotient space once the blue circle is filled by orange disks. Similarly, the point (a, b) can also be found in the diagram of the codomain. Both

occurrences of (a, b) correspond to the 1-cycle representing the blue circle in homology, which explains why the point is missing in the diagram of the cokernel.

Note that other natural choices of L are the orange subcomplex or the union of the blue and orange subcomplexes, which is a choice that is symmetric with respect to the colors. For two colors, these are the three possible Γ-subcomplexes defined in Section 3.4. We now revisit some of these observations in a more general setting, where the pair of topological spaces, $L \subseteq K$, is not necessarily formed by chromatic complexes.

5.3. **Relations Between Diagrams in a 6-pack.** The inclusion of sublevel sets $L_i \subseteq K_i$ induces a map on homology $\kappa_i \colon \mathsf{H}(L_i) \to \mathsf{H}(K_i)$. This map has a component in each dimension, p, and we write $\ker_p \kappa_i$, $\operatorname{im}_p \kappa_i$, $\operatorname{cok}_p \kappa_i$ for the kernel, image, cokernel of κ_i in dimension p.

Lemma 5.2. Let $L_i \subseteq K_i$ be simplicial complexes and $\kappa_i \colon \mathsf{H}(L_i) \to \mathsf{H}(K_i)$ the induced map on homology. For each dimension, p, there are short exact sequences

$$0 \to \ker_p \kappa_i \to \mathsf{H}_p(L_i) \to \operatorname{im}_p \kappa_i \to 0, \tag{5.3}$$

$$0 \to \operatorname{im}_p \kappa_i \to \mathsf{H}_p(K_i) \to \operatorname{cok}_p \kappa_i \to 0, \tag{5.4}$$

$$0 \to \operatorname{cok}_p \kappa_i \to \mathsf{H}_p(K_i, L_i) \to \ker_{p-1} \kappa_i \to 0. \tag{5.5}$$

Proof. The first two exact sequences are obvious from the definitions and the isomorphism theorem. To see the third exact sequence, we recall the long exact sequence of the pair, see Equation (5.1). Working with field coefficients, all homology groups are vector spaces and thus split. In particular, $\mathsf{H}_p(L_i) \simeq \ker_p \kappa_i \oplus \operatorname{im}_p \kappa_i$, in which $\ker_p \kappa_i$ and $\operatorname{im}_p \kappa_i$ are the images of the incoming and outgoing maps. We can therefore substitute $\ker_p \kappa_i \to 0 \to \operatorname{im}_p \kappa_i$ for $\mathsf{H}_p(L_i)$. By the same token, we substitute $\operatorname{im}_p \kappa_i \to 0 \to \operatorname{cok}_p \kappa_i$ for $\mathsf{H}_p(K_i)$, and we remove $0 \to \operatorname{im}_p \kappa_i \to \operatorname{im}_p \kappa_i$ to get

$$\ldots \to \ker_p \kappa_i \to 0 \to \operatorname{cok}_p \kappa_i \to \mathsf{H}_p(K_i, L_i) \to \ker_{p-1} \kappa_i \to 0 \to \operatorname{cok}_{p-1} \kappa_i \to \ldots,$$

which contains the required third short exact sequence. □

It follows that the ranks of $\ker_p \kappa_i$ and $\operatorname{im}_p \kappa_i$ add up to the rank of $\mathsf{H}_p(L_i)$, etc. This implies relations between the 1-norms of corresponding persistence diagrams.

Theorem 5.3. Let $L \subseteq K$ be simplicial complexes, $f_K \colon K \to \mathbb{R}$ monotonic, and f_L, $f_{K,L}$ the restrictions of f_K to L and $K \setminus L$. For each dimension, p, and any fixed cut-off for the 1-norms, $C > 0$,

$$\|\mathrm{Dgm}_p(f_L)\|_1 = \|\mathrm{Dgm}_p(\ker f_L \to f_K)\|_1 + \|\mathrm{Dgm}_p(\operatorname{im} f_L \to f_K)\|_1, \tag{5.6}$$

$$\|\mathrm{Dgm}_p(f_K)\|_1 = \|\mathrm{Dgm}_p(\operatorname{im} f_L \to f_K)\|_1 + \|\mathrm{Dgm}_p(\operatorname{cok} f_L \to f_K)\|_1, \tag{5.7}$$

$$\|\mathrm{Dgm}_p(f_{K,L})\|_1 = \|\mathrm{Dgm}_p(\operatorname{cok} f_L \to f_K)\|_1 + \|\mathrm{Dgm}_{p-1}(\ker f_L \to f_K)\|_1. \tag{5.8}$$

Proof. We prove (5.6). We write $0 \leq r_1, r_2, \ldots, r_n$ for the values of f_K smaller than C. In addition, set $r_0 = -\infty$ and use the cut-off $r_{n+1} = C$ for the 1-norms. Letting $L_i = {f_L}^{-1}[0, r_i]$, note that $L_i = {f_L}^{-1}[0, r]$ for all $r_i \leq r < r_{i+1}$, and hence the ranks of the various groups are constant between two contiguous values. We can therefore write the 1-norm of $\mathrm{Dgm}_p(f_L)$ as a sum of n contributions, and similar for the 1-norms of the kernel and image diagrams:

$$\|\mathrm{Dgm}_p(f_L)\|_1 = \sum\nolimits_{i=0}^{n} (r_{i+1} - r_i) \operatorname{rank} \mathsf{H}_p(L_i), \tag{5.9}$$

$$\|\mathrm{Dgm}_p(\ker f_L \to f_K)\|_1 = \sum\nolimits_{i=0}^{n} (r_{i+1} - r_i) \operatorname{rank} \ker_p \kappa_i, \tag{5.10}$$

$$\|\mathrm{Dgm}_p(\operatorname{im} f_L \to f_K)\|_1 = \sum\nolimits_{i=0}^{n} (r_{i+1} - r_i) \operatorname{rank} \operatorname{im}_p \kappa_i. \tag{5.11}$$

We thus get (5.6) from (5.3). With the same argument applied to K_i, image, and cokernel, we get (5.7) from (5.4), and applied to (K_i, L_i), cokernel, and kernel, we get (5.8) from (5.5). □

We note that similar equations do not hold for the 0-norm, which counts the points in the diagrams. Putting the equations in Theorem 5.3 together yields a vanishing alternating sum:

$$\sum\nolimits_{p \in \mathbb{Z}} (-1)^p \left[\|\mathrm{Dgm}_p(f_L)\|_1 - \|\mathrm{Dgm}_p(f_K)\|_1 + \|\mathrm{Dgm}_p(f_{K,L})\|_1 \right] = 0.$$

While there are relations between the diagrams in a 6-pack, no single diagram is necessarily determined by the others. Figure 13 shows one such example.

Further relations among the diagrams in a 6-pack are suggested by the case-by-case analysis for simultaneous occurrence of births and deaths in various groups provided in [8]. For example, consider

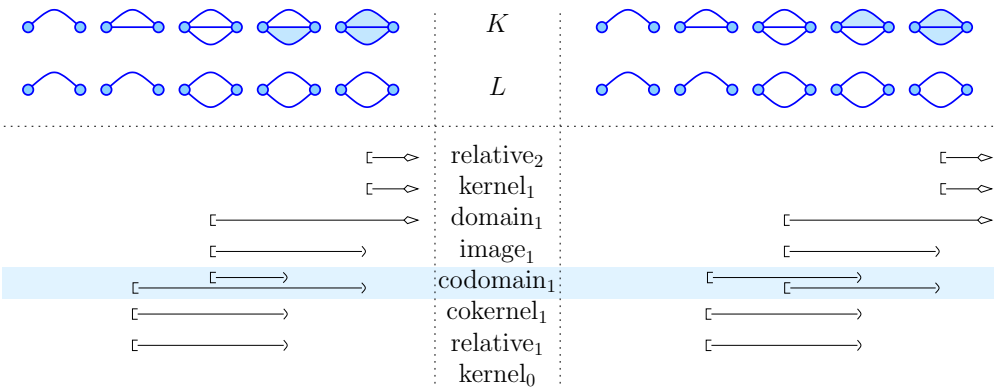

FIGURE 13. Example showing that five diagrams do not imply the sixth. The two filtrations differ by a single 2-dimensional cell added in the respective fourth steps of the filtrations. Correspondingly, five of the 1-dimensional persistence diagrams (shown as barcodes) are the same, while the highlighted diagrams of the codomain differ on the two sides.

the triple $\ker \kappa_i$, $\mathsf{H}(L_i)$, $\operatorname{im} \kappa_i$. At a given radius, the rank of each group can change by at most one. The short exact sequence (5.3) reduces the twenty-six non-trivial combinations of changes down to only six. Out of those, [8] gives examples for five of them and shows that the sixth, death-nothing-birth, cannot occur because a death in $\ker \kappa_i$ always implies a death in $\mathsf{H}(L_i)$. This is an additional relation, which is not implied directly by (5.3). The same case is excluded for the triple in (5.4). Analogously, one can show that the case death-nothing-birth is excluded for the triple $\operatorname{cok}_p \kappa_i$, $\mathsf{H}_p(K_i, L_i)$, $\ker_{p-1} \kappa_i$.

5.4. **Relations Between 6-packs of a Triplet.** The framework described so far is amenable to a pair of complexes $L \subseteq K$, filtered by a monotonic function. This section addresses the next simplest case: when we have a sequence of three nested complexes, $M \subseteq L \subseteq K$, which gives rise to four long exact sequences:

$$\ldots \to \mathsf{H}_p(L) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K, L) \to \mathsf{H}_{p-1}(L) \to \ldots, \tag{5.12}$$

$$\ldots \to \mathsf{H}_p(M) \to \mathsf{H}_p(K) \to \mathsf{H}_p(K, M) \to \mathsf{H}_{p-1}(M) \to \ldots, \tag{5.13}$$

$$\ldots \to \mathsf{H}_p(M) \to \mathsf{H}_p(L) \to \mathsf{H}_p(L, M) \to \mathsf{H}_{p-1}(M) \to \ldots, \tag{5.14}$$

$$\ldots \to \mathsf{H}_p(L, M) \to \mathsf{H}_p(K, M) \to \mathsf{H}_p(K, L) \to \mathsf{H}_{p-1}(L, M) \to \ldots. \tag{5.15}$$

To shed light on how they relate to each other, we draw them as sine-like curves, each directed from left

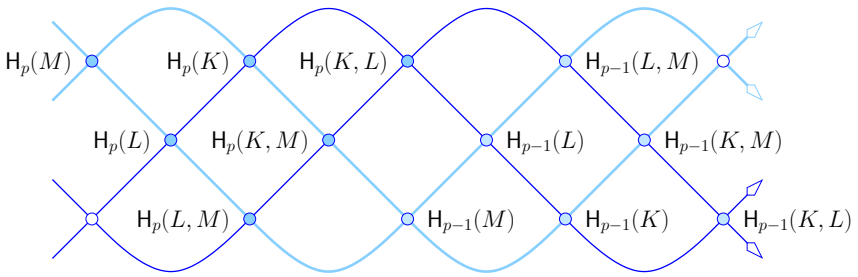

FIGURE 14. The four exact sequences for three complexes drawn along sine-like curves in the plane. After each half-period, the dimension of the homology group drops by one.

to right, with the homology groups sitting at the crossings between the curves; see Figure 14. Observe that the upper left triangular diagram commutes, which implies

$$\ker [\mathsf{H}_p(M) \to \mathsf{H}_p(L)] \subseteq \ker [\mathsf{H}_p(M) \to \mathsf{H}_p(K)], \tag{5.16}$$

$$\operatorname{im} [\mathsf{H}_p(M) \to \mathsf{H}_p(K)] \subseteq \operatorname{im} [\mathsf{H}_p(L) \to \mathsf{H}_p(K)] \tag{5.17}$$

for all dimensions p. Similar inclusions follow from the commutativity of the other regions in the arrangement of curves. The four inclusions that give rise to the sequences (5.12) to (5.15) yield four 6-packs, among which six diagrams appear twice, namely, $\operatorname{Dgm}(f_K)$, $\operatorname{Dgm}(f_L)$, $\operatorname{Dgm}(f_M)$, $\operatorname{Dgm}(f_{K,L})$,

$\mathrm{Dgm}(f_{K,M})$, $\mathrm{Dgm}(f_{L,M})$. Therefore, we have eighteen unique diagrams, some of which are closely related.

5.5. **A Tri-chromatic Case Study.** While two colors give rise to interesting patterns, more colors do more so. With the increase in the number of colors, there is an explosive increase of configurations to study. We suggest looking at the relations between *k-chromatic subcomplexes* of $\mathrm{Del}(\chi)$, which are the subcomplexes composed of all simplices with at most k colors, as defined in Section 3.4. In this section, we focus on the tri-chromatic case, with colors $\sigma = \{0, 1, 2\}$. Let M be the 1-chromatic subcomplex, L the 2-chromatic subcomplex, and K the full 3-chromatic Delaunay complex. As before, $f_K \colon K \to \mathbb{R}$ is the chromatic squared radius function, and f_L, f_M, $f_{K,L}$, $f_{K,M}$, $f_{L,M}$ are its restrictions. A cycle can be formed by points of 1, 2, or 3 colors, and it can be filled by points of 0, 1, or 2 additional colors. Requiring that the sum of two numbers is at most 3, we get the six mingling types sketched in Figure 1. Note that these six patterns are not independent. For example, the pattern 1+2 also gives rise to pattern 1+0, because the cycle gets filled by its own color eventually. However, different patterns corresponding to the same cycle will generally have different persistence, which quantifies which patterns is a better fit. Without a claim on completeness, we list where in the 6-packs one can find prominent cases of each of these six patterns.

- Case **1+0:** $\mathrm{Dgm}(f_M)$. The complex M is the disjoint union of the three mono-chromatic Delaunay mosaics. The diagram records the mono-chromatic cycles.
- Case **2+0:** $\mathrm{Dgm}(\mathrm{cok}\, f_M \to f_L)$. The complex L contains all mono- and bi-chromatic cycles, and it shares the former with M. Therefore, we look at the cokernel to keep only the cycles that need two colors to be formed. A cycle dies either when it is filled by its own two colors, or when one of the two colors suffices to form a homologous cycle.
- Case **3+0:** $\mathrm{Dgm}(\mathrm{cok}\, f_L \to f_K)$. As in the previous case, we look at the cokernel to capture cycles that are formed by all three colors, but not yet by any two.
- Case **1+1:** $\mathrm{Dgm}(\ker f_M \to f_L)$. When a cycle formed by one color is filled by an additional one, a birth in this diagram occurs, and the feature persists until it is filled by its own color.
- Case **2+1:** $\mathrm{Dgm}(\ker f_{L,M} \to f_{K,M})$. The idea is similar to Case 1+1: we look at cycles formed by two colors that are filled when also using the third. Unlike the previous case, we consider the quotient spaces to filter out the mono-chromatic cycles.
- Case **1+2:** $\mathrm{Dgm}(\mathrm{cok}\, f_{L,M} \to f_{K,M})$. Mono-chromatic p-cycles filled by the other colors appear in the pair (K, M) as $(p+1)$-cycles. Those that are filled by exactly one other color also appear in (L, M). We use the cokernel to filter them out.

We now look more closely at the concrete example displayed in the right panel of Figure 11: a circle of blue points with split background of green and orange points; compare with the mingling pattern 1+2. Focusing on this pattern, we search the 6-pack of the inclusion of the pairs $(L, M) \subseteq (K, M)$ in Figure 15. As suggested in Case 1+2 above, we expect a clear signal in the cokernel diagram, and indeed we see a single prominant point representing a 2-dimensional relative class. By construction, this class is born when the mono-chromatic 1-cycle is filled with two extra colors, and its persistence indicates how much longer it takes to fill the 1-cycle with just one extra color. Compare this with the even more prominent point in the diagram of the codomain, $\mathrm{Dgm}(f_{K,M})$. This point represents the same 1-cycle, but it expresses different information because it is not sensitive to whether the 1-cycle is filled by one or two additional colors.

It is interesting to interpret the two high persistence points in the diagram of the domain, which records classes in $\mathsf{H}(L_i, M_i)$. Since the background consists of two colors, it fills the blue 1-cycle with only one additional color twice, once with green and another time with orange. Both classes die at the same moment, namely when the blue cycle is filled by its own color.

As described in the caption of Figure 1, Case 1+2 could also be interpreted as two overlapping instances of Case 1+1. For example, consider mixing the orange and green points in the right panel of Figure 11. This pattern is captured in $\mathrm{Dgm}(f_L)$ as explained in Figure 9.

6. Discussion

The main contribution of this paper is the extension of the theory of alpha complexes to the setting where points are assigned a label. We prove structural results about the radius function on the chromatic Delaunay mosaic and provide an implementation that facilitates its use in applications. The

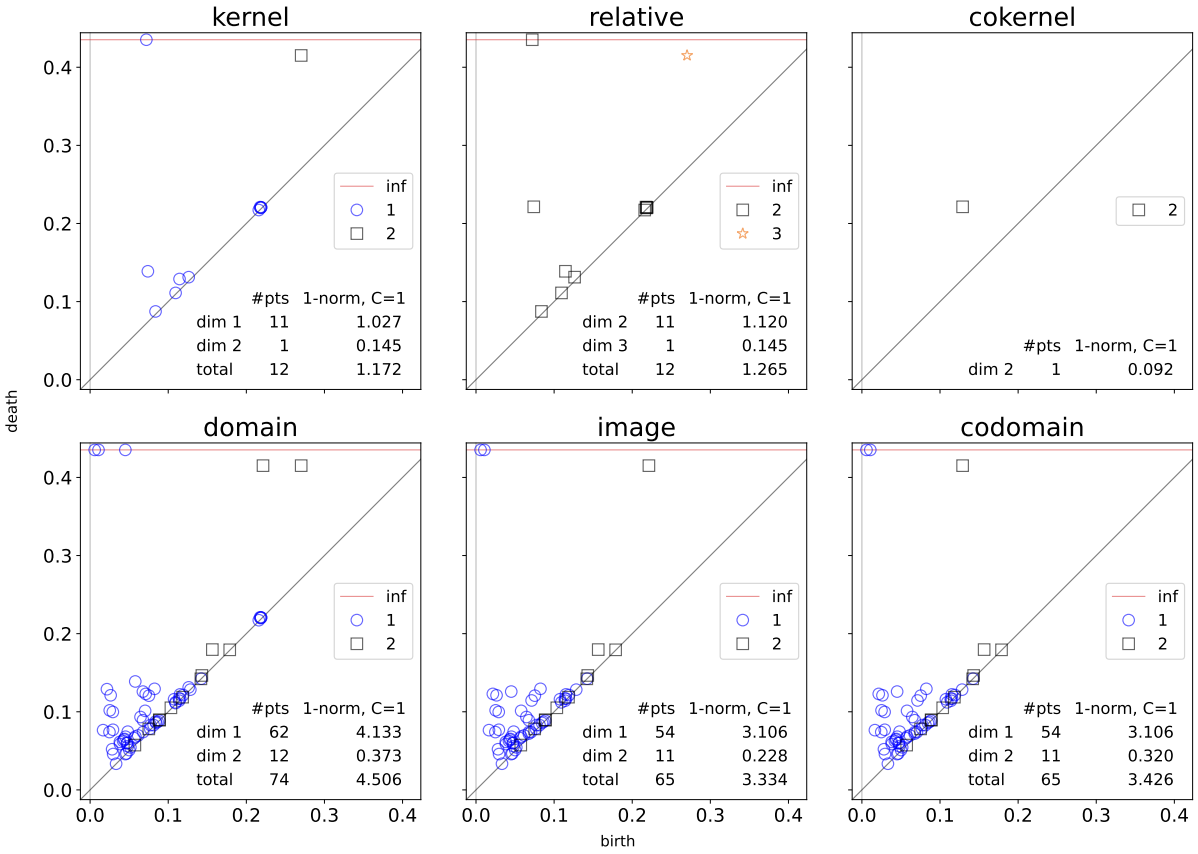

FIGURE 15. The 6-pack of $(L, M) \subseteq (K, M)$ for the data in the *right panel* of Figure 11. M, L, and K are the 1-, 2- and 3-chromatic subcomplexes of the chromatic Delaunay complex.

work reported in this paper suggests new directions of mathematical research aimed at solidifying our understanding of the chromatic setting. We list two possible directions.

- Develop a **chromatic variant of Forman's discrete Morse theory** [14]. Two concrete questions are the extension of the collapsibility of the Čech complex to the alpha complex proved in the mono-chromatic case [1] and the further collapse of $\mathrm{Alf}_r(\chi)$ to $\mathrm{Alf}_r(A)$.
- In many biological questions, the mingling between different populations of cells changes over time; see e.g. the study of cell segregation in early development [17] and an early topological approach in [15]. It would therefore be useful to **extend the vineyard algorithm** [9] to the chromatic setting introduced in this paper.

REFERENCES

[1] U. BAUER AND H. EDELSBRUNNER. The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.* **369** (2017), 3741–3762.

[2] U. BAUER, M. KERBER, F. ROLL AND A. ROLLE. A unified view on the functorial nerve theorem and its variations. *Expositiones Mathematicae* **41** (2023), 125503.

[3] M. BINNEWIES ET AL. Understanding the tumor immune microenvironment (TIME) for effective therapy. *Nat. Med.* **24** (2018), 541–550.

[4] R. BISWAS, S. CULTRERA DI MONTESANO, O. DRAGANOV, H. EDELSBRUNNER AND M. SAGHAFIAN. On the size of chromatic Delaunay mosaics. `arXiv:2212.03121 [math.CO]`, 2022.

[5] K. BORSUK. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.* **35** (1948), 217–234.

[6] S. BOYD AND L. VANDENBERGHE. *Convex Optimization.* Cambridge Univ. Press, Cambridge, England, 2004.

[7] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.

[8] D. COHEN-STEINER, H. EDELSBRUNNER, J. HARER AND D. MOROZOV. Persistent homology for kernels, images, and cokernels. In "Proc. 20th Ann. ACM-SIAM Sympos. Discrete Alg., 2009", 1011–1020.

[9] D. COHEN-STEINER, H. EDELSBRUNNER AND D. MOROZOV. Vines and vineyards by updating persistence in linear time. In "Proc. 22nd Ann. Sympos. Comput. Geom., 2006", 119–126.

[10] O. DRAGANOV AND M. MAHINI. Chromatic-tda. `github.com/OnDraganov/chromatic-tda`, 2023.

[11] H. Edelsbrunner and J.L. Harer. *Computational Topology. An Introduction.* Amer. Math. Soc., Providence, Rhode Island, 2010.
[12] H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551-–559.
[13] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72
[14] R. Forman. Morse theory for cell complexes. *Adv. Math.* **134** (1998), 90–145.
[15] M. Kerber and H. Edelsbrunner. 3d kinetic alpha complexes and their implementation. In "Proc. Mtg. Algorithm Engin. Experiments. 2013", 70–77.
[16] J. Leray. Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl.* **24** (1945), 95–167.
[17] J.-L. Maître et al. Adhesion functions in cell sorting by mechanically coupling the cortices of adhering cells. *Science* **338** (2012), 253–256.
[18] Y. Miao et al. Reconstruction and deconstruction of human somitogenesis in vitro. *Nature* **614** (2023), 500–508.
[19] Y. Reani and O. Bobrowski. A coupled alpha complex. *J. Comput. Geom.* **14** (2023), 221–256.
[20] B.J. Stolz, J. Dhesi, J.A. Bull, H.A. Harrington, H.M. Byrne and I.H.R. Yoon. Relational persistent homology for multispecies data with application to the tumor microenvironment. `arXiv:2308.06205 [math.AT]`, 2023.
[21] A. Weil. Sur les théorèmes de de Rham. *Comment. Math. Helv.* **26** (1952), 119–145.
[22] E. Welzl. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science*, ed.: H. Maurer, Springer LNCS **555**, 1991, 359–370.
[23] The CGAL Project. *CGAL User and Reference Manual.* CGAL Editorial Board, 4.10 edition, 2017.