

The Adaptive Topology of a Digital Image

Herbert Edelsbrunner^{*†‡}
herbert.edelsbrunner@ist.ac.at

Olga Symonova^{*}
olga.symonova@ist.ac.at

^{*}*Institute of Science and Technology Austria
Klosterneuburg, Austria*

[†]*Departments of Computer Science and of Mathematics
Duke University, Durham, North Carolina, USA*

[‡]*Geomagic
Research Triangle Park, North Carolina, USA*

Abstract—In order to enjoy a digital version of the Jordan Curve Theorem, it is common to use the closed topology for the foreground and the open topology for the background of a 2-dimensional binary image. In this paper, we introduce a single topology that enjoys this theorem for all thresholds decomposing a real-valued image into foreground and background. This topology is easy to construct and it generalizes to n -dimensional images.

Keywords—image processing; digital topology; homology, persistence; cubical subdivisions; Voronoi tessellations; perturbations.

I. INTRODUCTION

The main result of this paper is in digital topology, a subfield within image processing that concerns itself with the micro-structure of digital images; that is: the rules with which local neighborhoods are formed. The topic is important because the micro-structure has macroscopic consequences [12]. A general desire is that the macro-structure of an image satisfies basic properties of classical topology. This motivates the now widely accepted convention that defines the foreground of a binary image as a closed set and the background as an open set. In two dimensions, this is equivalent to using the 8-neighborhood rule for the foreground and the 4-neighborhood rule for the background. With this convention, the foreground forms a loop if and only if it separates the background into two components — a digital version of the Jordan Curve Theorem [11]. This convention is required because the pixels form a non-simple complex. Indeed, in a simple complex the maximum number of cells with non-empty common intersection is 3, while 4 pixels (unit squares) can share a common vertex. In three and higher dimensions, the gap widens between the at most $n + 1$ cells having non-empty intersection in a simple complex and the 2^n cubical cells sharing a vertex in an n -dimensional digital image. To address the problem, we follow Kovalevsky [13] in treating the image as a

complex and not just as a collection of cubes; see also [1], [17]. However, this does not suffice to compute persistence diagrams that satisfy fundamental symmetries implied by Lefschetz and Alexander duality.

Our concrete motivation for addressing the digital topology problem outlined above is the study of root systems of agricultural plants [14]. The goal is to extract phenotypical parameters and to relate them to the genotypical make-up of the plant [10]. Reconstructing the 3-dimensional shape from a series of 2-dimensional photographs, we represent the root system as a set of voxels (unit cubes) in \mathbb{R}^3 [20]. Beyond static root shapes, we also look at time series and address the growth behavior of root systems [9]. In this context, we consider functions defined on voxel representations of root systems, which are examples of 3-dimensional images. Among other things, we use the persistence diagrams of these functions to recognize tips and forks and to assess their robustness. Similarly, we use the diagram to recognize loops and voids, which should not be present but often are. Noting that the persistence diagram of a real-valued function encapsulates information about all binary images obtained for different thresholds, we recognize the need for a more sophisticated treatment of the micro-structure of the image, one that is consistent with fundamental topological symmetries for *all* thresholds.

Results. Our main result is the specification of the micro-structure of an n -dimensional image in terms of a simplicial complex, and a proof that the implied PL function gives persistence diagrams that satisfy fundamental topological symmetries. We state these results now, giving the proof in Section III. First we need definitions.

Letting \mathcal{U}'_n be the (infinite) set of unit cubes in the regular grid subdivision of \mathbb{R}^n , we define a *digital image* as a real-valued function, $\varphi : \mathcal{U}_n \rightarrow \mathbb{R}$, where \mathcal{U}_n is a finite subset of \mathcal{U}'_n . Given a threshold, $r \in \mathbb{R}$, the corresponding *binary image* is $\psi_r : \mathcal{U}_n \rightarrow \{0, 1\}$, defined by $\psi_r(U) = 1$ if $\varphi(U) \geq r$ and $\psi_r(U) = 0$ if $\varphi(U) < r$. Writing \cup for the union of the cubes in \mathcal{U}_n , the binary image decomposes \cup into

This research is partially supported by the National Science Foundation (NSF) under grant DBI-0820624, and the Research Network Programme of the European Science Foundation (ESF).

the *foreground*, $\mathbb{F}_r = \bigcup \psi_r^{-1}(1)$, and the *background*, $\mathbb{B}_r = \mathbb{U} - \mathbb{F}_r$. By construction, the foreground is closed and the background is open. Both change only at a finite number of thresholds, giving two *filtrations*: $\mathbb{F}_r \subseteq \mathbb{F}_s$ and $\mathbb{B}_s \subseteq \mathbb{B}_r$, for $r \leq s$. Applying the homology functor, we get a *persistence module*,

$$0 \rightarrow \dots \rightarrow H(\mathbb{F}_r) \rightarrow \dots \rightarrow H(\mathbb{U}, \mathbb{B}_r) \rightarrow \dots \rightarrow 0.$$

whose information is summarized in the *persistence diagram* of the image, denoted as $\text{Dgm}(\varphi)$. The diagram of this particular filtration is easier to define than to compute because the sets \mathbb{B}_r are not closed. For the next step, it is convenient to extend φ to the entire grid, $\varphi' : \mathcal{U}'_n \rightarrow \mathbb{R}$, assuming that φ and φ' are injective and agree on the cubes in \mathcal{U}_n , and that $\varphi'(U') < \varphi'(U)$ whenever U belongs to \mathcal{U}_n and U' does not. The vertices of the *dual complex* of φ' are the cubes in \mathcal{U}'_n . Working inductively, in the direction of increasing dimension, we add the p -simplices by considering all sets of 2^p n -cubes that share a common $(n-p)$ -face. They form a dual p -dimensional cube, and we let U_0 be the vertex with maximal function value. By induction, all faces of the p -cube have already been triangulated, and we triangulate the p -cube by adding the cones of U_0 over the $(p-1)$ -simplices in the triangulation of the opposite faces. After $n+1$ steps, we arrive at the dual complex of φ' , and we define the *dual complex* of φ , denoted as \mathcal{D} , as the full subcomplex whose vertices are the n -cubes in \mathcal{U}_n . Writing \mathbb{D} for the union of its simplices, we construct the *dual PL function*, $\delta : \mathbb{D} \rightarrow \mathbb{R}$, which agrees with φ on the vertices. We thus arrive at a PL function on a simplicial complex for which the persistence diagram is defined by the standard filtrations of excursion and incursion sets. We also have:

MAIN THEOREM. $\text{Dgm}(\delta) = \text{Dgm}(\varphi)$.

The computational significance of the result is the existence of standard algorithms for computing the diagram of δ [7, Chapter VII]. In contrast, the direct computation of the diagram of φ is difficult because the neighborhood of an n -cube changes as it migrates from the background to the foreground. The mathematical significance of the result is that the diagram satisfies the fundamental topological symmetries, in particular the reflection properties implied by Lefschetz and Alexander duality; see [5] and [8]. Contrarily, the diagram computed from the complex formed by the cubes and their faces does not enjoy these properties.

Outline. Section II presents the background from image processing and algebraic topology. Section III proves the Main Theorem. Section IV discusses an application to the root system architecture of agricultural plants. Finally, Section V concludes the paper.

II. BACKGROUND

In this section, we describe the background from image processing and from algebraic topology that provides the context for the work in this paper; see [19] and [7], [16] for further information. Since the relevant differences between 2-, 3-, and higher-dimensional images are minor, we cast all discussions in \mathbb{R}^n , where $n > 0$ is an arbitrary but fixed integer.

Images. We think of an image as a discrete representation of a continuous function. Specifically, we let $U_0 = [-\frac{1}{2}, \frac{1}{2}]^n$ be the n -dimensional unit cube centered at the origin, \mathbb{Z}^n the set of points with integer coordinates, and $\mathcal{U}'_n = U_0 + \mathbb{Z}^n$ the subdivision of \mathbb{R}^n into integer translates of the unit n -cube. In \mathbb{R}^2 and \mathbb{R}^3 , the elements of this subdivision are sometimes referred to as *pixels* (picture elements) and as *voxels* (volume elements). We call \mathcal{U}'_n the *cubical subdivision* of \mathbb{R}^n . To construct an image, we let $\mathcal{U}_n \subseteq \mathcal{U}'_n$ be a finite subset.

DEFINITION. An n -dimensional image is a real-valued function on a finite subset of the cubical subdivision of \mathbb{R}^n , $\varphi : \mathcal{U}_n \rightarrow \mathbb{R}$.

Collecting all faces of the n -cubes in \mathcal{U}_n , which are k -cubes for $0 \leq k \leq n$, we get a *cubical complex*, which we denote as \mathcal{U} . Finally, we write $\mathbb{U} = |\mathcal{U}|$ for the *underlying space* of the complex, which we define as the union of interiors of the cells of different dimensions. For now, \mathbb{U} is just the union of the n -cubes in \mathcal{U}_n , but we will see more interesting underlying spaces shortly. We think of φ as an approximation of a continuous function on \mathbb{R}^n , one in which the size of the cubes dictates the resolution. For example, we may construct a piecewise constant function $\bar{\varphi} : \mathbb{U} \rightarrow \mathbb{R}$ by defining

$$\bar{\varphi}(x) = \max_{x \in U \in \mathcal{U}_n} \varphi(U), \quad (1)$$

where the maximum is over all n -cubes that contain the point. For later reference, we call $\bar{\varphi}$ the *PC function* of the image φ . A *binary image* is a digital image that takes only two values, 0 and 1. Given a threshold, r , we get a binary image, $\psi : \mathcal{U}_n \rightarrow \{0, 1\}$, by setting

$$\psi(U) = \begin{cases} 1 & \text{if } \varphi(U) \geq r, \\ 0 & \text{if } \varphi(U) < r, \end{cases} \quad (2)$$

for all n -cubes $U \in \mathcal{U}_n$. We note that the PC function of ψ can be defined by setting $\bar{\psi}(x)$ to 1, if $\bar{\varphi}(x) \geq r$, and to 0, if $\bar{\varphi}(x) < r$.

Digital topology. The primary topological concern in image processing is *whether* things are connected, and to a lesser extent *how* they are connected. For this reason, the topology is usually expressed by specifying which n -cubes are neighbors. For example, the 26-neighborhood rule in \mathbb{R}^3 says that two cubes are neighbors if they share a square

face, an edge, or just a vertex. We take a different approach and specify topologies as subsets of the cubical complex. For example, the *closed topology* of a set of n -cubes consists of all their faces, which are cubes of dimension $0 \leq k \leq n$. Assuming a binary image $\psi : \mathcal{U}_n \rightarrow \{0, 1\}$, we let $\mathcal{F} \subseteq \mathcal{U}$ be the closed topology of $\bar{\psi}^{-1}(1)$ and define the *foreground* as its underlying space, $\mathbb{F} = |\mathcal{F}|$. In contrast, the *open topology* contains a k -cube only if all n -cubes that share the k -cube are in the set. Letting $\mathcal{B} \subseteq \mathcal{U}$ be the open topology of $\bar{\psi}^{-1}(0)$, we have $\mathcal{B} = \mathcal{U} - \mathcal{F}$, and we define the *background* as its underlying space, $\mathbb{B} = |\mathcal{B}|$. Note that \mathcal{F} is a subcomplex of \mathcal{U} , while \mathcal{B} is generally only a subset of \mathcal{U} .

The closed and the open topologies are complementary; that is: if we partition the n -cubes into two collections and we use the closed topology for one and the open topology for the other, then we get two complementary spaces. This is convenient as it lends itself to topological properties that we intuitively expect. For example, the Jordan Curve Theorem says that a loop in \mathbb{R}^2 decomposes the plane, separating points inside from points outside the loop. If the foreground is defined by pixels that outline a loop, we may expect the same property. Indeed, using the closed topology for the foreground and the open topology for the background, we get a separation between the pixels inside and outside the loop formed by the foreground.

Homology. We quantify the connectivity of a space using its homology groups, whose elements are classes of homologous cycles; see e.g. [16]. While there are different formalisms that make this definition precise, they all lead to the same algebraic framework. Using the formalism of *singular homology*, a p -dimensional chain is the sum of finitely many continuous maps of the standard p -simplex into the space. The *boundary* of this chain is the sum of the restrictions of these maps to the boundary of the p -simplex. Adding maps and allowing for cancellations, it is possible that the boundary is empty, in which case the p -chain is called a p -cycle. If there is a $(p + 1)$ -chain whose boundary is that p -cycle, then the p -cycle is called a p -boundary. Finally, two p -cycles are *homologous* if their sum is a p -boundary. We thus get equivalence classes of homologous p -cycles, and the main technical insight is that these classes form a group that can be used to quantify the p -dimensional connectivity of the space. To remove the last bit of ambiguity, we specify the coefficient group used to add chains. While there are other choices, we prefer to work with a field, for example with $\mathbb{Z}/2\mathbb{Z}$, which is integer addition modulo 2. For a field, the homology groups are vector spaces, and the ranks of the groups are the dimensions of the vector spaces. It is customary to write H_p for the p -th homology group, and to refer to its rank as the p -th Betti number, $\beta_p = \text{rank } H_p$. We write $H = \bigoplus_p H_p$ for the direct sum of the homology groups of all dimensions.

Filtrations and persistence. We return to the real-valued image, $\varphi : \mathcal{U}_n \rightarrow \mathbb{R}$, and we recall its PC function, $\bar{\varphi}$, as defined in (1). For each threshold, $r \in \mathbb{R}$, set $\mathbb{F}_r = \bar{\varphi}^{-1}[r, \infty)$ and $\mathbb{B}_r = \bar{\varphi}^{-1}(-\infty, r)$, as in Section I. Let s_1 to s_m be the values in $\varphi(\mathcal{U}_n)$, and consider an interleaved sequence, $r_0 > s_1 > r_1 > \dots > s_m > r_m$. Simplifying the notation, we write $\mathbb{U}^i = \mathbb{F}_{r_i}$ and $\mathbb{U}_i = \mathbb{B}_{r_i}$. Since the foreground increases and the background decreases, we have $\mathbb{U}^0 \subseteq \mathbb{U}^1 \subseteq \dots \subseteq \mathbb{U}^m$ and $\mathbb{U}_m \subseteq \mathbb{U}_{m-1} \subseteq \dots \subseteq \mathbb{U}_0$. Replacing each foreground by the direct sum of its homology groups, $H(\mathbb{U}^i) = \bigoplus_p H_p(\mathbb{U}^i)$, and each background by the direct sum of its relative homology groups, $H(\mathbb{U}, \mathbb{U}_i) = \bigoplus_p H_p(\mathbb{U}, \mathbb{U}_i)$, we get a linear sequence of homology groups connected by maps induced by inclusion:

$$0 = H(\mathbb{U}^0) \rightarrow \dots \rightarrow H(\mathbb{U}^m) \quad (3)$$

$$= H(\mathbb{U}, \mathbb{U}_m) \rightarrow \dots \rightarrow H(\mathbb{U}, \mathbb{U}_0) = 0. \quad (4)$$

Using coefficients in a field, the homology groups are vector spaces, and the maps between them are linear. As described in [5], there is a canonical definition when a homology class is born and when it dies in the filtration. Its *persistence* is the absolute difference between the function values at which the class is born and dies. The totality of this information is represented by the *persistence diagram* of φ , denoted as $\text{Dgm}(\varphi)$, which is a multiset of points (called *dots*) in a double covering of the plane. Each




Figure 1. Left: the image maps every cube in the stylized root system to its vertical distance from a horizontal line. Right: the persistence diagram with three dots for the branches and one dot for the loop. The shaded pair of rectangles corresponds to the level set defined for $\frac{d+e}{2}$.

dot has two coordinates, the value at the birth and the value at the death of the classes it represents. Since we use the entire range of values twice, first coming down in (3) and second going up in (4), we need the double covering to distinguish between classes that are born and die coming down, classes that are born coming down and die going up, and classes that are born and die going up.

Symmetries. The Jordan Curve Theorem has various generalizations, one being Lefschetz duality, which relates the homology of a manifold with boundary with the relative homology of the pair consisting of the manifold and its boundary. As proved in [5], Lefschetz duality implies $\text{Dgm}(f) = \text{Dgm}^T(f)$ whenever f is defined on a manifold

without boundary. Here, T is the transformation that reflects each dot across the vertical axis and replaces its dimension by the complementary dimension. While the domain of a digital image cannot be a manifold without boundary, its boundary can. The symmetry implied by Lefschetz duality is therefore indirectly relevant as follows. Suppose the digital image can be extended to a map f on \mathbb{R}^n that has no homological critical values, such as for example every Cartesian coordinate function. Then the diagram of the function restricted to the boundary is the disjoint union of the diagram of the image and of its reflection: $\text{Dgm}(f|_{\partial U}) = \text{Dgm}(f|_U) \sqcup \text{Dgm}^T(f|_U)$; see [8]. Another useful property of the diagram is the possibility to read the homology of level sets. As illustrated in Figure 1, the homology of $\varphi^{-1}(\frac{d+e}{2})$ can be determined by counting the dots in the symmetric difference of the rectangles that determine the homology of the corresponding excursion set and the relative homology of the pair consisting of the domain and the incursion set. For this purpose, the dimension of the dots in the right half of the diagram are lowered by one [3]. In the particular example, we count three dots on the left and one dot on the right. After adjusting the dimension, all four dots are 0-dimensional, which is consistent with the level set consisting of four components and otherwise having no non-trivial homology. We close this section by showing that the mentioned




Figure 2. Middle: a torus decomposed into 16 pixels. Left: the diagram computed with the cubical complex. Note that we have two identical dots near the top, corresponding to two 1-dimensional classes. Right: the diagram computed with the dual complex.

properties are easily lost if we do not pay attention to the micro-structure of the image. Consider the complex that decomposes the 2-dimensional torus into 16 pixels, as depicted in Figure 2. The function maps the 14 white pixels to a common high value and the two dark pixels to a common low value. Computing persistence with the cubical complex and the dual complex, we get the diagrams on the left and on the right. We see the violation of the symmetry implied by Lefschetz duality in the left diagram. The algorithm requires closed incursion sets and therefore does not see the 1-dimensional relative class represented by curves connecting the two dark pixels. We mention that [18] compute persistence with the cubical complex, but they consider only a quarter of the diagram, namely the left triangle, so Lefschetz duality does not apply.

III. THE PROOF

We take two steps to prove the Main Theorem stated in Section I. The first step is geometric, relating the dual complex of an image with a weighted Voronoi tessellation of the space. The second step is topological, relating the foregrounds and backgrounds of the digital image with the excursion and incursion sets of the dual PL function.

Voronoi tessellation. We can think of the cells in the cubical subdivision as the Voronoi cells of their centers. Making essential use of this interpretation, we assign weights to the centers to produce an approximating simple complex. We begin with definitions. Given a discrete collection of points in \mathbb{R}^n , the (unweighted) Voronoi cell of a point v_i in this collection is the set of points $x \in \mathbb{R}^n$ for which the Euclidean distance to v_i is no larger than to any other point in the collection: $\|x - v_i\| \leq \|x - v_j\|$ for all v_j . We generalize this concept by assigning a real weight, w_i , to each point v_i . The *weighted Voronoi cell* of v_i is then the set V_i of points $x \in \mathbb{R}^n$ that satisfy

$$\|x - v_i\|^2 - w_i \leq \|x - v_j\|^2 - w_j \quad (5)$$

for all j ; see e.g. [6]. The set of such cells is the *weighted Voronoi tessellation* of the collection of points. Its cells are n -dimensional, closed but possibly unbounded, convex polyhedra. Together with their common intersections, these cells form a complex. If the points and their weights satisfy a genericity condition, then the intersection of any $i+1$ cells is either empty or has dimension $n-i$. In particular, this means that the common intersection of $n+2$ or more cells is empty. Complexes with this property are said to be *simple*.

We use the fact that for any discrete collection of points, we can find weights such that the weighted Voronoi tessellation forms a simple complex. Indeed, the weight assignments that do not lead to simple complexes form a measure zero subset of all assignments. To see this, consider the geometric primitives used to construct a Voronoi tessellation: the *orientation test* and the *in-sphere test*. The first takes $n+1$ points as input and decides whether the last point lies on the positive or negative side of the hyperplane spanned by the first n points. Writing $x_{i,j}$ for the j -th coordinate of the i -th point, this is decided by the sign of the determinant of

$$\Delta = \begin{bmatrix} 1 & x_{1,1} & \dots & x_{1,n} \\ 1 & x_{2,1} & \dots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n+1,1} & \dots & x_{n+1,n} \end{bmatrix}. \quad (6)$$

We have a degeneracy if $\det \Delta = 0$, which happens with probability zero. Similarly, the in-sphere test takes a sequence of $n+2$ points as input and decides whether the last point lies inside or outside the $(n-1)$ -dimensional sphere passing through the other $n+1$ points. Here, we need a second matrix,

$$\Lambda = \begin{bmatrix} 1 & x_{1,1} & \cdots & x_{1,n} & \sum x_{1,j}^2 \\ 1 & x_{2,1} & \cdots & x_{2,n} & \sum x_{2,j}^2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & x_{n+2,1} & \cdots & x_{n+2,n} & \sum x_{n+2,j}^2 \end{bmatrix}, \quad (7)$$

and the decision is based on the sign of the product of the two determinants. We have a degeneracy if $\det \Delta = 0$ or $\det \Lambda = 0$, which again happens only with probability zero. To construct the *weighted* Voronoi tessellation, we use the same orientation test, and we modify the in-sphere test by subtracting the weights from the last column of Λ . A geometric interpretation can be found in [6, Chapter 5].

Structural stability. Given the cubical subdivision of \mathbb{R}^n , we are interested in assigning weights to the centers of the n -cubes such that the weighted Voronoi tessellation perturbs the subdivision. We construct the perturbation depending on an image, $\varphi : \mathcal{U}_n \rightarrow \mathbb{R}$. Assuming φ is injective, we extend it to an injective function $\varphi' : \mathcal{U}'_n \rightarrow \mathbb{R}$, where \mathcal{U}'_n is the cubical subdivision, φ and φ' agree on the cubes in \mathcal{U}_n , and $\varphi'(U') < \varphi'(U)$ whenever $U \in \mathcal{U}_n$ and $U' \notin \mathcal{U}_n$. We can therefore index the infinitely many n -cubes such that $\varphi'(U_1) > \varphi'(U_2) > \dots$. As mentioned earlier, the cubical subdivision is the Voronoi tessellation of the centers u_i of the U_i . To define the perturbation, we let v_i be the point u_i with the only difference that v_i has a weight, $w_i = \varepsilon^i$. Here, ε is a positive real number, which we constrain to guarantee that our construction is independent of the particular choice of ε , and i is the index of the n -cube in the above defined ordering. In other words, we require that the combinatorial structure of the weighted Voronoi tessellation does not depend on the particular choice of ε . To write this more explicitly, we let I be a sequence of $n+1$ distinct positive integers, we let ℓ be a positive integer not in I , and we write $\Lambda_{I,\ell}(\varepsilon)$ for the matrix in (7), assuming of course that the first $n+1$ input points are the ones with indices in I , that the $(n+2)$ -nd input point has index ℓ , and that the weights are assigned using ε . Similarly, we write Δ_I , noting that this matrix does not depend on ε . We say I *labels* a vertex in the weighted Voronoi tessellation of ε if $\det \Delta_I \neq 0$ and $\det \Delta_I \cdot \det \Lambda_{I,\ell}(\varepsilon) < 0$ for all $\ell \notin I$.

DEFINITION. An interval $(0, a)$ has the *stable in-sphere property* if a sequence of $n+1$ distinct indices I labels a vertex in the weighted Voronoi tessellation of ε iff it labels a vertex in the tessellation of ε' , for all $0 < \varepsilon, \varepsilon' < a$.

It is not difficult to see that the stable in-sphere property implies that the weighted Voronoi tessellations defined by ε and ε' are structurally the same. To formalize this concept, we consider the weighted Voronoi tessellation of the points in \mathbb{Z}^n , with weights assigned as described and ε chosen in $(0, a)$. Let $\mathcal{V}_n(\varepsilon)$ be the subset of cells generated by

centers of cubes in \mathcal{U}_n , and let $\mathcal{D}(\varepsilon)$ be the nerve of $\mathcal{V}_n(\varepsilon)$, geometrically realized in \mathbb{R}^n by mapping each cell to its generating integer point. We have $\mathcal{D}(\varepsilon) = \mathcal{D}(\varepsilon')$ for all $0 < \varepsilon < \varepsilon' < a$ since $(0, a)$ has the stable in-sphere property. To see that $\mathcal{D}(\varepsilon)$ is equal to the dual complex, \mathcal{D} , introduced in Section I, we observe that the perturbation reduces to the lexicographical ordering of the n -cubes used in the construction of \mathcal{D} .

Stable in-sphere property. We show that there exist non-empty open intervals that have the stable in-sphere property.

STABLE IN-SPHERE LEMMA. In \mathbb{R}^n , the interval $(0, a_n)$, with $a_n = \frac{1}{(n+1)!+1}$, has the stable in-sphere property.

Proof: Consider the 2^n vertices of the unit n -cube whose center is the origin of \mathbb{R}^n . Picking $n+2$ of the points, we get a matrix whose entries in the middle n columns are $\pm \frac{1}{2}$, while all entries in the first column are 1, and all entries in the last column are $\frac{n}{4}$. Subtracting the weights from the last column, which we assume for simplicity are w_1, w_2, \dots, w_{n+2} , we get

$$\Lambda(\varepsilon) = \begin{bmatrix} 1 & \pm \frac{1}{2} & \cdots & \pm \frac{1}{2} & \frac{n}{4} - \varepsilon \\ 1 & \pm \frac{1}{2} & \cdots & \pm \frac{1}{2} & \frac{n}{4} - \varepsilon^2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \pm \frac{1}{2} & \cdots & \pm \frac{1}{2} & \frac{n}{4} - \varepsilon^{n+2} \end{bmatrix}. \quad (8)$$

After subtracting $\frac{n}{4}$ times the first column from the last, we see that $\det \Lambda(\varepsilon)$ is a polynomial in ε in which the coefficient of ε^i is plus/minus the determinant of the matrix obtained by removing the i -th row and the last column. This determinant is $1/2^n$ times an integer between $-(n+1)!$ and $(n+1)!$. Since $\det \Delta_I \neq 0$, by assumption, not all coefficients are zero. Let k be the smallest index for which we get a non-zero coefficient, and note that it has absolute value at least $1/2^n$. The absolute value of $\det \Lambda(\varepsilon)$ is therefore at least

$$\frac{\varepsilon^k}{2^n} - \frac{(n+1)!}{2^n} \sum_{j=k+1}^{n+2} \varepsilon^j = \frac{\varepsilon^k}{2^n} - \frac{(n+1)!}{2^n} \frac{\varepsilon^{k+1} - \varepsilon^{n+3}}{1 - \varepsilon}.$$

Since $0 < \varepsilon < \frac{1}{(n+1)!+1}$, this is larger than 0. In other words, the sign of the k -th coefficient decides the sign of the entire determinant. Since the former does not depend on ε , neither does the latter, as required.

Finally, we note that the special choices we made along the way do not limit the generality of the argument. First, we can restrict ourselves to the vertices of a unit cube, because no other sets of $n+2$ integer points define matrices that are relevant. Second, we can assume that this unit cube is centered at the origin, because translations in \mathbb{R}^n do not affect the computation. Third, the final calculation generalizes to any choice of $n+2$ indices. \square

Foreground. We now switch gears and use the geometric constructions to show that the digital image and its dual PL function have identical persistence diagrams. We begin

with the first half (3) of the extended filtration, which consists of the homology groups of the foregrounds, $\mathbb{U}^i = \mathbb{F}_{r_i} = \bar{\varphi}^{-1}[r_i, \infty)$; see Section II. Correspondingly, we write $\mathbb{D}^i = \delta^{-1}[r_i, \infty)$ for the *excursion sets* of the dual PL function. Our goal is to show that we can substitute δ for φ when we compute the persistence diagram. We prepare the proof by defining the *perturbed image*. To begin, we introduce $\phi : \mathcal{V}_n(\varepsilon) \rightarrow \mathbb{R}$, defined by setting $\phi(V_i) = \varphi(U_i)$ for the n -dimensional cell generated by v_i . Writing \mathbb{V} for the union of cells in $\mathcal{V}_n(\varepsilon)$, the corresponding PC function, $\bar{\phi} : \mathbb{V} \rightarrow \mathbb{R}$, maps each point $x \in \mathbb{V}$ to the maximum value of the cells in $\mathcal{V}_n(\varepsilon)$ that share x . Given a threshold, r_i , we introduce the *perturbed foreground*, $\mathbb{V}^i = \bar{\phi}^{-1}[r_i, \infty)$.

EXCURSION SET LEMMA. The homology groups of the foreground and excursion set defined by the same threshold are isomorphic: $H(\mathbb{U}^i) \simeq H(\mathbb{D}^i)$, for $0 \leq i \leq m$.

Proof: To begin, let \mathcal{D}^i be the nerve of the cells in $\mathcal{V}_n(\varepsilon)$ with value at least r_i . Clearly, \mathcal{D}^i is a subcomplex of \mathcal{D} , namely the full subcomplex defined by the vertices with value at least r_i . The underlying space of \mathcal{D}^i is a subset of \mathbb{D}^i , and there is a deformation retraction from the latter to the former; see [7, p. 136]. It follows that both have the same homotopy type. Similarly, the Nerve Theorem implies that $|\mathcal{D}^i|$ and \mathbb{V}^i have the same homotopy type and, by transitivity, so do \mathbb{D}^i and \mathbb{V}^i .

It remains to show that \mathbb{U}^i and \mathbb{V}^i have isomorphic homology groups. Here, we exploit that the centers of the n -cubes in the foreground receive larger weights than those in the background. It follows that the foreground encroaches upon the background; that is: $\mathbb{U}^i \subseteq \mathbb{V}^i$. Hence, every cycle in \mathbb{U}^i is also a cycle in \mathbb{V}^i . To prove the converse, we note that \mathbb{V}^i depends on ε , but we dropped the parameter from the notation because we get identical nerves for all choices of ε in $(0, a_n)$. We can therefore choose ε as small as we like, which pushes the cycles in \mathbb{V}^i arbitrarily close to \mathbb{U}^i . Using a compactness argument, for the representatives of a homology class, we conclude that the cycles also exist in \mathbb{U}^i . Finally, \mathbb{U}^i and \mathbb{V}^i could fail to have isomorphic homology groups if a non-bounding cycle in one space bounds in the other. This cannot happen for a bounding cycle in \mathbb{U}^i since the chain it bounds also belongs to \mathbb{V}^i . In the other direction, we use a compactness argument, as before. \square

Background. Here, we are concerned with the second half (4) of the extended filtration, which consists of the homology groups of the backgrounds, $\mathbb{U}_i = \mathbb{B}_{r_i} = \bar{\varphi}^{-1}(-\infty, r_i]$. Note that \mathbb{U}_i is a union of n -cubes minus some of their faces and is therefore not necessarily closed. Thus \mathbb{U}_i cannot be represented by a subcomplex of the cubical subdivision which would induce closed topology. Taking singular homology defined by finite chains, we get groups that are isomorphic to the regular homology groups of a

suitable closed subset of \mathbb{U}_i . Correspondingly, we write $\mathbb{D}_i = \delta^{-1}(-\infty, r_i]$ for the *incursion set* of the dual PL function defined by r_i , and we introduce the *perturbed background*, $\mathbb{V}_i = \bar{\phi}^{-1}(-\infty, r_i]$.

INCURSION SET LEMMA. The relative homology groups of the background and incursion set defined by the same threshold are isomorphic: $H(\mathbb{U}, \mathbb{U}_i) \simeq H(\mathbb{D}, \mathbb{D}_i)$, for $0 \leq i \leq m$.

Proof: The first half of the proof concerns the absolute homology groups of the background and the incursion set. The analysis is similar to the proof of the Excursion Set Lemma. Let \mathcal{D}_i be the nerve of the cells in $\mathcal{V}_n(\varepsilon)$ with value at most r_i , and note that it is a full subcomplex of \mathcal{D} . Its underlying space, $|\mathcal{D}_i|$, is a subset of \mathbb{D}_i , and there is a deformation retraction implying that both have the same homotopy type. By the Nerve Theorem, $|\mathcal{D}_i|$ and \mathbb{V}_i have the same homotopy type, and by transitivity, so do \mathbb{D}_i and \mathbb{V}_i . Next, we show that \mathbb{U}_i and \mathbb{V}_i have isomorphic homology groups. In contrast to the situation for excursion sets, we do not have the inclusion of \mathbb{U}_i in \mathbb{V}_i , nor do we have the inclusion the other way round. Indeed, \mathbb{V}_i is locally contained where \mathbb{U}_i is open, and it locally contains where \mathbb{U}_i is closed. In spite of this difference, we can exploit the possibility to choose $\varepsilon > 0$ as small as we like, and use a compactness argument to show that $H(\mathbb{U}_i)$ and $H(\mathbb{V}_i)$ are isomorphic.

The second half of the proof considers five consecutive terms in the exact sequence of the pair $(\mathbb{U}, \mathbb{U}_i)$, with the relative homology group in the middle position. In parallel, we consider the sequence for the pair $(\mathbb{D}, \mathbb{D}_i)$:

$$\begin{array}{ccccccccc} H_p(\mathbb{U}_i) & \rightarrow & H_p(\mathbb{U}) & \rightarrow & H_p(\mathbb{U}, \mathbb{U}_i) & \rightarrow & H_{p-1}(\mathbb{U}_i) & \rightarrow & H_{p-1}(\mathbb{U}) \\ \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\ H_p(\mathbb{D}_i) & \rightarrow & H_p(\mathbb{D}) & \rightarrow & H_p(\mathbb{D}, \mathbb{D}_i) & \rightarrow & H_{p-1}(\mathbb{D}_i) & \rightarrow & H_{p-1}(\mathbb{D}). \end{array}$$

The above argument implies vertical isomorphisms between the first two and the last two groups. To prove that each square in the diagram commutes, it would be convenient if the vertical isomorphisms were induced by inclusions. They are not, but we can change the drawing of \mathcal{D} so that they are. Specifically, we take the barycentric subdivision of \mathcal{D} , and we map each vertex to the center of the corresponding cube, which is the common intersection of n -cubes and therefore of dimension $0 \leq k \leq n$. Because of the degenerate intersections, this can result in identical realizations of different abstract simplices. More important than not having a global geometric realization in \mathbb{R}^n is that now all simplices of \mathcal{D}_i are subsets of \mathbb{U}_i . Using this drawing, all vertical maps are induced by inclusions. All horizontal maps in the diagram are induced by inclusions, which implies that all squares commute. Since the horizontal sequences are exact, by construction, we now have everything ready to apply the Steenrod Five Lemma and to conclude that the middle vertical map is also an isomorphism, as claimed; see [16, p. 140]. \square

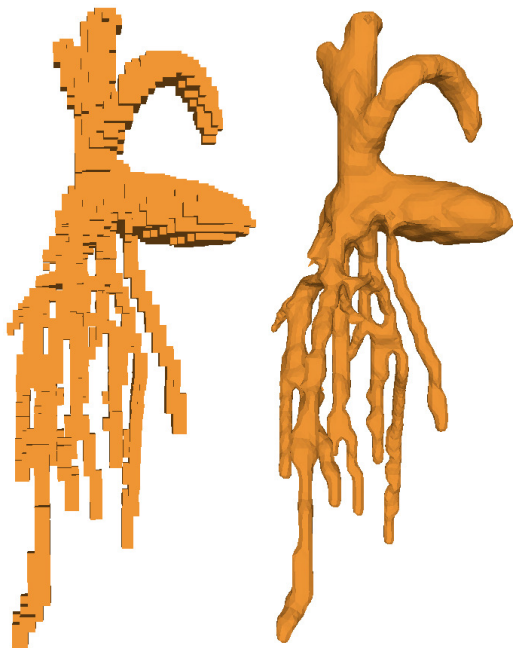


Figure 3. The 3-dimensional reconstruction of a root system. Left: the set of voxels that forms the model of the root system. Right: a more intuitive, smooth surface representation whose only purpose is the visualization of the model.

Final step. We are now ready to finish the proof of the Main Theorem stated in Section I. Consider the filtration defined by the foregrounds and backgrounds of the first function, and the filtration defined by the excursion sets and incursion sets of the second function. Connecting them with the vertical isomorphisms implied by the above two lemmas, we get

$$\begin{array}{ccccccc}
 0 & \rightarrow & \dots & \rightarrow & H(\mathcal{U}^i) & \rightarrow & \dots & \rightarrow & H(\mathcal{U}, \mathcal{U}_i) & \rightarrow & \dots & \rightarrow & 0 \\
 \uparrow & & & & \uparrow & & & & \uparrow & & & & \uparrow \\
 0 & \rightarrow & \dots & \rightarrow & H(\mathcal{D}^i) & \rightarrow & \dots & \rightarrow & H(\mathcal{D}, \mathcal{D}_i) & \rightarrow & \dots & \rightarrow & 0
 \end{array}$$

The horizontal maps are induced by inclusions. Furthermore, using the drawing of \mathcal{D} described in the proof of the Incursion Set Lemma, the vertical isomorphisms are induced by inclusions. It follows that all squares in the diagram commute. The claim follows by the Persistence Equivalence Theorem in [7, p. 159].

IV. APPLICATION

To illustrate the results of this paper, we construct the dual complex and the dual PL map of the 3-dimensional image in Figure 3. The set of voxels models a root system reconstructed from a series of 2-dimensional images [20]. Choosing a voxel near the top as the *origin*, we compute a real-valued function that approximates the geodesic distance (inside the root system) from the origin. The dual complex and PL map are then constructed as described in Section I.

Using the PL map, we compute the persistence diagram, which contains a wealth of useful information, as we now

explain. We distinguish between white, gray, and black dots representing 0-, 1-, and 2-dimensional homology classes, respectively; see Figure 4. Each white dot in the left triangle




Figure 4. The persistence diagram of the geodesic distance function defined for the 3-dimensional model of a root system in Figure 3.

represents a branch in the root system (a 0-dimensional homology class in the sequence of excursion sets), with vertical distance from the baseline equal to the length of this branch. Each gray dot in the same triangle represents a ceiling depression, perhaps right below a fork. Similarly, each black dot in the right triangle represents a floor depression. Since depressions are mere artefacts of the model, it is not surprising that they are all very close to the baseline. The dots in the middle diamond represent homology classes of the model, as opposed to of excursion or incursion sets of the geodesic distance function. There is only one white dot, which shows that the model consists of a single connected component, as it should. All other dots in the square mark mistakes that happened in the reconstruction process. In particular, the gray dots represent loops, which are typically caused by a failure to separate almost crossing branches due to a lack of resolution. The black dots represent voids, which are caused by fluctuations of the estimated density function used in the reconstruction. The horizontal distance from the axis of the diagram represents the size of the loop or void. In the case at hand, all dots are close to the axis and therefore represent only minor flaws in the model. To repair them, we can simply fill the voids with voxels, but to remove loops in a manner that reflects the original root system is more difficult and will be discussed elsewhere.

V. DISCUSSION

The main contribution of this paper is the definition of the adaptive topology of a real-valued image, represented by the dual simplicial complex, and a proof that the persistence diagram of its implied PL function agrees with the one defined by closed foregrounds and open backgrounds. We are therefore justified to call it *the* persistence diagram of the image. Algorithms that construct the persistence diagram from the filtration of excursion and incursion sets can be found in [7, Chapter VII]. Assuming the dimension is a constant, the number of simplices in the dual complex is at

most proportional to the number of n -cubes. The mentioned algorithms take worst case time at most proportional to the cubic power of the number of n -cubes, but are observed to be much faster in practice. Furthermore, there are sub-cubic algorithms that are either difficult to implement [15], whose running time depends on the output [4], or that are restricted to particular kinds of data [7, Chapter VII.2].

REFERENCES

- [1] M. ALLILI, K. MISCHAIKOW AND A. TANNENBAUM. Cubical homology and the topological classification of 2D and 3D imagery. *In* “Proc. IEEE Internat. Conf. Image Process., 2001”, 173–176.
- [2] P. BENDICH, H. EDELSBRUNNER AND M. KERBER. Computing robustness and persistence for images. *IEEE Trans. Visual. Comput. Graphics* **16** (2010), 1251–1260.
- [3] P. BENDICH, H. EDELSBRUNNER, D. MOROZOV AND A. PATEL. Homology and robustness of level and interlevel sets. *Homology, Homotopy, and Applications*, to appear.
- [4] C. CHEN AND M. KERBER. An output-sensitive algorithm for persistent homology. *In* “Proc. 27th Ann. Sympos. Comput. Geometry, 2011”, 207–215.
- [5] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Extending persistence with Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79–103.
- [6] H. EDELSBRUNNER. *Geometry and Topology for Mesh Generation*. Cambridge Univ. Press, England, 2001.
- [7] H. EDELSBRUNNER AND J. L. HARER. *Computational Topology. An Introduction*. Amer. Math. Soc., Providence, Rhode Island, 2010.
- [8] H. EDELSBRUNNER AND M. KERBER. Alexander duality for functions: the persistent behavior of land and water and shore. *In* “Proc. 28th Ann. Sympos. Comput. Geometry, 2012”, to appear.
- [9] H. EDELSBRUNNER, O. SYMONOVA AND C. TOPP. Reconstruction and analysis of growing plant roots. Manuscript under construction, IST Austria, Klosterneuburg, Austria, 2011.
- [10] A. IYER-PASCUZZI, O. SYMONOVA, Y. MILEYKO, Y. HAO, H. BELCHER, J. HARER, J. S. WEITZ AND P. BENFEY. Imaging and analysis platforms for automatic phenotyping and classification of plant root systems. *Plant Physiol.* **152** (2010), 1148–1157.
- [11] E. KHALIMSKY, R. KOPPERMAN AND P. R. MEYER. Computer graphics and connected topologies on finite ordered sets. *Topology Appl.* **36** (1990), 1–17.
- [12] R. KLETTE AND A. ROSENFELD. *Digital Geometry*. Morgan Kaufmann, San Francisco, California, 2004.
- [13] V. A. KOVALEVSKY. Finite topology as applied to image analysis. *Comput. Vision, Graph., Im. Process.* **46** (1989), 141–161.
- [14] J. LYNCH. Root architecture and plant productivity. *Plant Physiol.* **109** (1995), 7–13.
- [15] N. MILOSAVLJEVIĆ, D. MOROZOV AND P. ŠKRABA. Zigzag persistent homology in matrix multiplication time. *In* “Proc. 27th Ann. Sympos. Comput. Geometry, 2011”, 216–225.
- [16] J. R. MUNKRES. *Elements of Algebraic Topology*. Perseus, Cambridge, Massachusetts, 1984.
- [17] N. PASSAT, M. COUPRIE, L. MAZO AND G. BERTRAND. Minimal simple sets: a new concept for topology-preserving transformations. *In* “Proc. 1st Workshop Combin. Top. Image Context, 2008”.
- [18] V. ROBINS, P. J. WOOD AND A. P. SHEPPARD. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *IEEE Trans. Patt. Anal. Mach. Intell.* **33** (2011), 1646–1658.
- [19] M. SONKA, V. HLAVAC AND R. BOYLE. *Image Processing, Analysis and Machine Vision*. Second edition, PWS Publishing, Pacific Grove, California, 1999.
- [20] Y. ZHENG, S. GU, H. EDELSBRUNNER, C. TOMASI AND P. BENFEY. Detailed reconstruction of 3D plant root shape. *In* “Proc. 13th Internat. Conf. Comput. Vision, 2011”, 2026–2033.