

Persistent Homology: Theory and Practice

Herbert Edelsbrunner* and Dmitriy Morozov†

Abstract. Persistent homology is a recent grandchild of homology that has found use in science and engineering as well as in mathematics. This paper surveys the method as well as the applications, neglecting completeness in favor of highlighting ideas and directions.

2010 Mathematics Subject Classification. Primary 55N99; Secondary 68W30.

Keywords. Algebraic topology, homology groups, distance, stability, algorithms; scale, shape analysis, topology repair, high-dimensional data.

1. Introduction

Built on a sequence of spaces and the corresponding homology groups with homomorphism between them, persistence assesses the interval within which a homology class contributes. Among other situations, this ability is useful when a space is not fixed but depends on the scale of the observation, which is a common scenario in the sciences. After a brief review of the historical development, we sketch characteristics of the method.

History. Like many other concepts in mathematics, persistent homology has a beginning but also a historical root system that comes into sight when we increase the resolution of the inquiry. This is precisely what persistent homology does for a much more general class of spaces: it synthesizes the different views as aspects into a single consistent reality that spans a range of scales. We mention three main historical tracks in the root system of persistent homology. In 1990, Patrizio Frosini and collaborators introduced size functions, a formalism that is equivalent to 0-dimensional persistent homology [24]. The main direction of the pursuant work is on shape analysis and its applications in computer vision and medical imaging; see a recent survey [4]. In 1999, Vanessa Robins studied the homology of sampled spaces and described the images of homomorphisms induced by inclusions as persistent homology groups [39]. In 2000, Edelsbrunner, Letscher and Zomorodian independently introduced persistent homology together with a fast algorithm and the diagram [22], as we will discuss later. Both of these works were inspired by the computational notion of alpha shapes [21, 23] and the related Betti number algorithm [16]. Within mathematics, there is a distinct relationship with

*This research is partially supported by NSF under grant DBI-0820624, by ESF under the Research Networking Programme, and by the Russian Government Project 11.G34.31.0053.

†This work is partially supported by the DOE Office of Science ASCR under award number KJ0402-KRD047 and contract number DE-AC02-05CH11231.

spectral sequences, originally introduced in 1946 by Jean Leray [34]. Motivated by its significant applications, persistent homology has found repeated exposure in the popular mathematics literature [5, 27, 45] and features prominently in a recent text on computational topology [20].

Perspectives. In a nutshell, persistent homology expands the relationship between a topological space and its homology groups to that between a function and its persistence diagram. The latter relationship gives rise to a rich theory which invites different perspectives if studied with different mind-sets.

Mathematics: We see an extension of the algebraic theory of homology forming a bridge to measure theory. The extension is inspired by Morse theoretic reasoning taken to the algebraic level of homology groups connected by maps.

Computation: The persistent homology groups are computed by reducing the boundary matrices of complexes. Indeed, all algebraic relationships have parallels in the matrix representation.

Applications: The matrices give fast algorithms and the algebra leads to scale-dependent measurements of spaces. Importantly, these measurements are stable, they can be used to compare and analyze shapes, and they can be exploited to repair faulty topology.

Depending on the interest, we focus on different aspects of the method. We present the material in two main sections, first explaining the theoretical framework of persistent homology, and second sketching four applications selected to highlight different aspects of the theory.

2. The Theory

In this section, we discuss the mathematical and computational representation of topological spaces, the algebra obtained by applying the homology functor, the implications of this construction to measuring topology, and algorithms that compute the groups in this algebra.

2.1. Spaces and Functions. Persistent homology is applied to a filtered space or, equivalently, to the sequence of sublevel sets of a real-valued function on this space. We discuss how to extract both ingredients, a space and a function, from different kinds of data.

Data. Sometimes the data already comes as a real-valued function, such as *digital images*. They are usually laid out on a regular integer grid, in which every cell records the locally averaged intensity value of the measured light field. Most common are 2D images in which the cells are squares, but also time-series of 2D images and 3D images are widely used. Indeed, digital images form one of the

most important classes of data as they are inexpensive to acquire and they probe nature in exquisite detail.

Another prevailing form of input data are *point clouds*, finite subsets of some ambient space, most often Euclidean. Each point represents a sequence of measurements of an individual in a population. We typically want to understand the overall shape of the cloud, for example, by measuring the topology of the space we get by thickening each point to a ball and taking the union. Equivalently, we may introduce the *distance function* that maps each point of the ambient space to its distance from the nearest data point. Letting α be the radius of the balls, we get the union as the sublevel set, defined as the set of points with function value at most α . A crucial property of this construction is its stability: if the input data follows an underlying law that appears as a shape in the ambient space, then the function we construct is close to the distance function defined by that shape and thus facilitates the study of the latter.

A third class of input data are *shapes*, subsets of ambient space that satisfy regularity conditions of one kind or another. A common subclass consists of surfaces in \mathbb{R}^3 , e.g., obtained by collecting points on the boundary of a solid object with a 3D scanner and connecting the points to a surface by interpolation. As in the point cloud case, the function is typically constructed in a second step, perhaps to highlight or define features of the shape, such as protrusions or cavities. In the case of a surface, popular such functions are the *mean* and the *Gaussian curvature*, as well as the *eccentricity* [30]. There are plenty of other possibilities — with special constructions for special purposes — such as the *elevation function* defined in terms of the persistent homology of the 2-parameter family of height functions in \mathbb{R}^3 [1].

Complexes. Following a long-standing tradition in topology, we work with complexes to represent continuous spaces. Common examples are *CW-*, *cubical* and *simplicial complexes*, to name a few. Cubical complexes have already been mentioned as the basis of digital images. They consist of cubes of various dimensions, with the requirement that with every p -cube, the complex also contains the $2p$ $(p - 1)$ -cubes that are its faces. Significant improvements in the efficiency of computations are gained if we store cubical subdivisions hierarchically, such as in quad- and oct-trees [40]. The CW- and simplicial complexes are extreme examples on opposite ends of the spectrum. The CW-complexes allow for complicated cells glued to each other in complicated ways and thus facilitate representations of spaces with only a few cells. In contrast, all cells in a simplicial complex are simplices, and any two are glued along a single shared face or not at all. In spite of the frequently required large number of simplices, the local simplicity of these complexes lends itself to efficient computations.

Every simplicial complex has an abstract and a geometric side, and it is useful to fully exploit both. Take, for example, the *nerve* of a finite collection of convex sets; that is: the system of subcollections with non-empty common intersection. This is an abstract simplicial complex since every collection U in the nerve implies the membership of the subsets of U . A particularly useful collection of convex sets are

the *Voronoi cells* of a finite set of points in \mathbb{R}^n [43]. Assuming general position, the maximum number of Voronoi cells with non-empty common intersection is $n + 1$. In this case, the nerve has a natural geometric realization, known as the *Delaunay triangulation* of the points in \mathbb{R}^n [15]. Specifically, for each U in the nerve of the Voronoi cells, the Delaunay triangulation contains the convex hull of the points whose Voronoi cells are in U . This simplicial complex supports computations of the Euclidean distance function defined by the points. Indeed, the sublevel set of a threshold $\alpha > 0$ is a union of balls of radius α , one around each point. Intersecting each ball with the corresponding Voronoi cell gives another collection of convex sets, and its nerve is isomorphic to a subsystem of the nerve of the Voronoi cells. Its geometric realization is known as the α -*complex* [21, 23], which is, of course, a subcomplex of the Delaunay triangulation.

There are many situations in which the Delaunay triangulation is not defined, or we cannot afford to compute it. A popular alternative is the *Vietoris-Rips complex*, which exists whenever we have the distances between pairs of points. Given a threshold $a > 0$, this complex contains a simplex spanned by $p + 1$ points iff every two of these points are at distance at most a from each other. Equivalently, the Vietoris-Rips complex for parameter a is the flag complex built on the set of edges with length at most a .

2.2. Algebra. The classic theory of homology maps a topological space to an abelian group which, in the case of coefficients in a field, is a vector space. Having a filtered space, we get a sequence of vector spaces, together with linear maps induced by inclusion. This is the basic set-up for persistent homology, which we now describe.

Homology. The theory of homology is a classic subject within algebraic topology, which is described in most of the standard texts, including Munkres [38] and Hatcher [28]. The construction begins with a *chain group*, C_p , whose elements are the p -chains, which for a given complex are formal sums of the p -dimensional cells. The *boundary homomorphism*, $\partial_p : C_p \rightarrow C_{p-1}$, maps each p -chain to the sum of the $(p - 1)$ -dimensional faces of its p -cells, which is a $(p - 1)$ -chain. Writing the groups and maps in sequence, we get the *chain complex*:

$$\dots \xrightarrow{\partial_{p+2}} C_{p+1} \xrightarrow{\partial_{p+1}} C_p \xrightarrow{\partial_p} C_{p-1} \xrightarrow{\partial_{p-1}} \dots \quad (1)$$

The kernels and the images of the boundary homomorphisms are the *cycle* and the *boundary groups*. A fundamental property of the boundary homomorphism is that its square is zero, $\partial_p \circ \partial_{p+1} = 0$. Therefore, for every p , the boundaries form a subgroup of the cycles, and we can take the quotient, which gives a group whose elements are classes of homologous cycles. This is the p -*th homology group*, denoted as H_p , where p is again the dimension. We assume coefficients in a field, F , so that $H_p = F \oplus F \oplus \dots \oplus F = F^{\beta_p}$ is a vector space over F , with $\beta_p = \text{rank } H_p$ known as the p -*th Betti number*. For a topological space, \mathbb{X} , we write $H_p(\mathbb{X})$ and $\beta_p(\mathbb{X})$ for its p -th homology group and Betti number. They are defined for every integer, p , but if the dimension of \mathbb{X} is n , then the only possibly non-trivial homology groups

are for $0 \leq p \leq n$. Accordingly, we have $\beta_p = 0$ unless $0 \leq p \leq n$. To simplify the notation, we will often suppress the dimension and write $H(\mathbb{X}) = \bigoplus_p H_p(\mathbb{X})$ for the direct sum.

Let $\mathbb{X}_0 \subseteq \mathbb{X}$ be a topological subspace. Every cycle in \mathbb{X}_0 is also a cycle in \mathbb{X} , although it may be trivial in the latter without being trivial in the former. The inclusion of \mathbb{X}_0 in \mathbb{X} induces a linear map on the homology groups, $\varphi : H(\mathbb{X}_0) \rightarrow H(\mathbb{X})$. We will also consider the pair of spaces, $(\mathbb{X}, \mathbb{X}_0)$, whose (relative) homology is obtained by identifying cycles that differ only inside \mathbb{X}_0 . We have again a linear map induced by inclusion, $\psi : H(\mathbb{X}) \rightarrow H(\mathbb{X}, \mathbb{X}_0)$. Furthermore, there is a third linear map, $D : H(\mathbb{X}, \mathbb{X}_0) \rightarrow H(\mathbb{X}_0)$, such that

$$\dots \xrightarrow{D_{p+1}} H_p(\mathbb{X}_0) \xrightarrow{\varphi_p} H_p(\mathbb{X}) \xrightarrow{\psi_p} H_p(\mathbb{X}, \mathbb{X}_0) \xrightarrow{D_p} H_{p-1}(\mathbb{X}_0) \xrightarrow{\varphi_{p-1}} \dots \quad (2)$$

is *exact*, by which we mean that the image of every map is the kernel of the next map. This particular sequence is the *exact sequence of the pair* $(\mathbb{X}, \mathbb{X}_0)$. It is a compact expression of how the relative homology of the pair is related to the (absolute) homology of the two spaces.

Filtrations. The basic set-up for persistent homology consists of a *filtered space*, a nested sequence of subspaces that begins with the empty and ends with the complete space [22, 48]. Writing $\emptyset = \mathbb{X}_0 \subseteq \mathbb{X}_1 \subseteq \dots \subseteq \mathbb{X}_m = \mathbb{X}$, we apply the homology functor, which for each space gives a vector space and for each inclusion gives a linear map:

$$0 = H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow H(\mathbb{X}_m) = H(\mathbb{X}), \quad (3)$$

referring to this sequence as a *persistence module*. It is instructive to split the module into *indecomposable summands* of the form $0 \rightarrow F \rightarrow \dots \rightarrow F \rightarrow 0$, where every nonzero map is the identity. There is a unique such decomposition whose direct sum gives the original module. Each summand can be interpreted as the *birth* of a homology class at its first non-zero term and the *death* of the same class right after its last non-zero term. More precisely, the summand represents an entire coset of classes that are born and die together, but we prefer to simplify language by talking about generators. It should be clear that the module above is not necessarily exact. In fact, it is exact iff each summand is of the form $0 \rightarrow F \rightarrow F \rightarrow 0$, consisting of precisely two non-zero terms. Of particular significance is the length of a summand, which measures the duration of the corresponding class. We refer to it as the *persistence* of the homology class. When a filtration results from a function, we often define persistence not as the number of non-zero terms but rather as the absolute difference between the function values at the birth and the death. A related concept are the *persistent homology groups*, which are the images under the composition of the linear maps. For example, the image of $H(\mathbb{X}_i)$ in $H(\mathbb{X}_j)$ is such a group, and its rank is the number of indecomposable summands whose births happen at or before $H(\mathbb{X}_i)$ and whose deaths happen after $H(\mathbb{X}_j)$.

Since (3) ends with a possibly non-trivial group, some homology classes may never die. We set the value at the death to ∞ , but doing so deprives us of a

meaningful measure of the duration of such a class. Alternatively, we may add relative homology classes constituting a second pass:

$$0 = H(\mathbb{X}_0) \rightarrow \dots \rightarrow H(\mathbb{X}_m) \rightarrow H(\mathbb{X}, \mathbb{X}^m) \rightarrow \dots \rightarrow H(\mathbb{X}, \mathbb{X}^0) = 0, \quad (4)$$

where \mathbb{X}^i is the closure of $\mathbb{X} - \mathbb{X}_i$. This is the *extended persistence module* as introduced in [12]. Decompositions into summands, births, and deaths are defined as before. Now every class that is born also dies. We distinguish between three kinds: classes that are born and die during the first pass, classes that are born during the first pass and die during the second pass, and classes that are born and die during the second pass. The second kind comprises all classes of the entire space, \mathbb{X} , which are precisely the ones that were born but did not die in (3).

Beyond homology groups, the above decomposition holds for any linear sequence of vector spaces. In this context, we note the connection between persistence and quiver representations observed in [6]. A fundamental result for quivers states that the orientation of maps between vector spaces does not affect the structure of the indecomposable summands [17]. This implies that a module can be replaced by a sequence in which any two contiguous vector spaces are connected by a map — either from left to right, or from right to left. Such generalized sequences, referred to as *zigzag modules*, elucidate the relationship between the extended persistence and the homology of interlevel sets of scalar functions [7].

2.3. Measuring. The persistence of a homology class is the length of the interval that supports it. The connection to applications is that the persistence measurement carries useful information about spaces, functions, and data. A particularly useful property of this measurement is its stability under perturbations of the function, as we explain in this section.




Figure 1. Left: the height function with six critical points on a topological sphere. We also show five interleaving level sets and highlight one sublevel set. Right: the persistence diagram with two finite and two infinite points. The wedge anchored at $(x, 0)$ contains two points labeled 0 and one point labeled 1, implying that the highlighted sublevel set has $\beta_0 = 2$ and $\beta_1 = 1$. Below: the barcode representation of the same information.

Persistence diagrams. The splitting into indecomposable summands suggests a combinatorial representation as a multi-set of points in the extended 2-dimensional plane. Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be continuous. With a minor modification of the original construction, we build this multi-set by adding a copy of the point $\frac{1}{2}(y+x, y-x)$ for each summand with birth at x and death at y ; see Figure 1. We refer to this multi-set as the *persistence diagram* of the function f , denoting it as $\text{Dgm}(f)$, or as $\text{Dgm}_p(f)$ if it is restricted to classes of dimension p . Instead of the points in the plane, we sometimes draw the intervals defined such that a point $u \in \text{Dgm}(f)$ is contained in a wedge anchored at $(x, 0)$ iff x is contained in the interval of u ; see Figure 1. The multi-set of such intervals is the *barcode* of the function f . Using this wedge, we can determine the Betti numbers of the corresponding sublevel set simply by counting the points of the diagram it contains. More generally, the points of the diagram contained in the wedge anchored at $\frac{1}{2}(y+x, y-x)$ determine the rank of the persistent homology group defined by the inclusion of the sublevel set for x in the sublevel set for y .

Some modifications are in order if we substitute the extended module (4) for the ordinary module (3). Since we get each value twice, we get a multi-set in a double covering of the plane. As a benefit of the complication, we can read the ranks of the persistent homology groups of all sublevel and superlevel sets as well as of all level and interlevel sets of the function; see [7].

Stability. To compare the diagrams of two functions, $f, g : \mathbb{X} \rightarrow \mathbb{R}$, we may use the *Wasserstein distance* between them, which is defined as the q -th root of the infimum, over all matchings between the points, of the sum of q -th powers of the edge lengths:

$$W_q(\text{Dgm}(f), \text{Dgm}(g)) = \inf_{\gamma} \left(\sum_{u \in \text{Dgm}(f)} \|u - \gamma(u)\|_{\infty}^q \right)^{\frac{1}{q}}, \quad (5)$$

where q is a positive real number; see e.g. [42]. In the limit, for q going to infinity, we get the *bottleneck distance*, which is the length of the longest edge in the best matching. For these definitions to make sense, we add infinitely many copies of every point on the horizontal axis to the diagrams; they guarantee that there are bijections between the multi-sets. An important property of the bottleneck distance is its stability with respect to perturbations. Specifically, we have

$$W_{\infty}(\text{Dgm}(f), \text{Dgm}(g)) \leq \|f - g\|_{\infty}, \quad (6)$$

whenever f and g are both *tame*, by which we mean that they have only finitely many critical values, and all sublevel sets have finite rank homology groups. This is the Bottleneck Stability Theorem first proved in [11]. A word of caution is in order: (6) implies that the critical value pairs that define the points in the diagram are stable, but it *does not* imply that the critical values or the critical points are stable. In fact, they are *not*.

While the bottleneck distance leads to a very general stability result, it has drawbacks in practice because it is sensitive to only the worst edge in the best

matching. The other Wasserstein distances do not imply stability for quite as general a class of functions, but they do so for interesting classes, such as for Lipschitz functions [13]. A different extension of the stability result — from tame functions to parametrized families of vector spaces — appears in [10].

2.4. Computation. An alternative to the algebraic description of homology based on chain complexes is the computational description based on boundary matrices. The algorithms form the bridge that connects the rich field of algebraic topology with applications, as discussed in Section 3.

Matrices and ranks. The p -th boundary matrix, denoted as D_p , is a computationally convenient representation of the p -th boundary homomorphism, ∂_p . Its columns are indexed by the p -dimensional cells, its rows by the $(p-1)$ -dimensional cells, and $D_p[i, j]$ stores the coefficient of the i -th $(p-1)$ -cell in the boundary of the j -th p -cell. Recall that a p -chain is a formal sum of p -cells. Writing it as a column vector, c , we can multiply with the matrix to get its boundary, $D_p c$, again written as a column vector. By construction, the column space of D_p is isomorphic to the group of $(p-1)$ -boundaries. Similarly, the null space of D_p is isomorphic to the group of p -cycles. Since the p -th homology group is the quotient of the p -th cycle group over the p -th boundary group, we get its rank as the dimension of the null space of D_p minus the dimension of the column space of D_{p+1} . To compute these dimensions, we put the boundary matrices into *normal form* in which an initial segment of the diagonal contains 1's while the rest of the matrix is zero. To do this, we use elementary row and column operations:

1. exchange two rows or two columns;
2. add a row to another row or a column to another column;
3. multiply a row or a column with a coefficient from the field.

Similar to Gauss-Jordan elimination, we apply these operations to move a 1 to the upper-left corner and to zero out its row and its column. The normal form is then completed by recursing on the smaller matrix obtained by removing the lead row and the lead column. Of course, the recursion halts when the remaining matrix is empty or zero. The row operations can be summarized by multiplying the boundary matrix from the left, and the column operations by multiplying from the right. This gives $N_p = U_{p-1} D_p V_p$, where N_p is the matrix in normal form, which provides all the information we need:

$$\begin{aligned} z_p &= \#\text{zero columns in } N_p &= \text{rank of null space;} \\ b_p &= \#\text{non-zero columns in } N_{p+1} &= \text{rank of column space;} \\ \beta_p &= z_p - b_p &= \text{rank of homology group.} \end{aligned}$$

The auxiliary matrices, U_{p-1} and V_p , provide additional information, which is sometimes useful. In particular, the last z_p columns of V_p give a basis of the p -th cycle group, and the first b_{p-1} columns of the inverse, U_{p-1}^{-1} , give a basis of the $(p-1)$ -st boundary group.

Preserving order. We can do more with less: we can compute homology as well as persistence while stopping short of reducing the boundary matrix to normal form. To describe how this works, we put all boundary information into a single matrix, D . We assume that the topological space is constructed one cell at a time, making sure that each cell is preceded by its faces. Denoting the corresponding ordering of the cells by $\sigma_1, \sigma_2, \dots, \sigma_m$, $D[i, j]$ is the coefficient of σ_i in the boundary of σ_j . We reduce D with a subset of the column operations, refraining from exchanging columns and adding columns from left to right. The algorithm pays special attention to the lowest non-zero entry in each column, which we may assume is 1. If all lowest 1s appear in distinct rows, then the matrix is *reduced*. To get D into this form, we iterate through the columns from left to right, reducing each column by subtracting multiples of conflicting preceding columns. The greedy nature of the process ensures that the resulting matrix is reduced. As before, we can express the operations as a multiplication with another matrix: $R = DV$, where R is reduced and V is invertible and upper-triangular. While this decomposition is not unique, the lowest 1s in R are unique [14]. We get all information from their number and their locations within the reduced matrix. Similar to before, the number of zero columns that belong to p -cells is the rank of the p -th cycle group, and the number of lowest 1s in columns of p -cells is the rank of the $(p - 1)$ -st boundary group. But we can extract more:

- adding σ_j gives birth to a homology class iff column j of R is zero;
- in contrast, adding σ_j kills a homology class iff column j is non-zero; letting $R[i, j]$ be its lowest 1, σ_j kills the class born with the addition of σ_i .

If non-zero, column j of R contains a cycle representative of the dying class; it is the boundary of the chain in column j of V . To prove these relationships, we assume again that $R[i, j]$ is the lowest 1 in column j . The representative cycle of the dying class thus contains σ_i , which implies that it did not exist before the addition of σ_i . All classes born before the addition of σ_i cannot die with the addition of σ_j , else we could prove inductively that R is not yet reduced.

The above greedy algorithm is due to [22]. Its running time is the total squared persistence of the filtration. In the worst case, it is proportional to m^3 but shows significantly better performance in practice. The worst-case time can be improved to m^ω , where $\omega = 2.372\dots$ [37], which is the currently best upper bound on the complexity of matrix multiplication [46]; see also Strassen [41] for a milestone paper in the sequence of improvements. These algorithms work for arbitrary boundary matrices, while we can sometimes exploit special structure to get faster algorithms. For example, if our space is a 2-manifold, we can use Poincaré duality and limit the computation to 0-dimensional homology. In this case, a combinatorial algorithm maintaining disjoint sets computes persistence in time proportional to $m \log m$ [26]. A common special case are regular cubical grids used in image processing. The algorithm in [44] takes full advantage of the possibility to compute boundaries implicitly, through subscript computations in an array. Additional savings are possible if we use hierarchical cubical complexes [3], such as quad- and oct-trees.

3. The Practice

In this section, we discuss four applications of persistent homology: the first to atomic structures highlighting the role of scale, the second to human jaws illustrating derived metrics, the third to root systems controlling topological connectivity, and the fourth to natural images mapping data to high dimensions.

3.1. The Atomic Structure of Material. Nature is full of structures that possess features on multiple scales. Persistent homology quantifies scale and can be used to measure the relative abundance of one scale to another. In this section, we approach simulated organic material from this angle, following the work of MacPherson and Schweinhart [36], who take steps toward characterizing the statistical distribution of scale.

Pockets and cages. Let \mathbb{X} be a union of finitely many closed balls in \mathbb{R}^3 . We may think of \mathbb{X} as the geometric model of a protein, as commonly used in structural molecular biology [33]. More interesting for biological questions than the model itself is, in many ways, its complement. The cavities of the model are prime candidate areas for interactions with small ligands and other proteins. Here, ‘cavity’ is an informal term for a depression or a partially protected area of the surface that is still accessible from the outside. In an effort to make this intuition concrete, [19] introduces the notion of a *pocket*, which is a subset of $\mathbb{R}^3 - \mathbb{X}$ that turns into a void under uniform thickening of \mathbb{X} ; see [35] for a biologically motivated study of their volume and shape. Similar to the evolution under thickening, which can be complicated, pockets exhibit hierarchical structure. Without going into detail, we note that each pocket of \mathbb{X} corresponds to a point in the 2-nd persistence diagram of the distance function, $d_{\mathbb{X}} : \mathbb{R}^3 \rightarrow \mathbb{R}$, defined by the geometric model. A point $u = \frac{1}{2}(y + x, y - x)$ corresponds to a void that forms at the thickening radius x and disappears at the radius y . The existence of this point does not contradict the possibility of the void splitting up into two at a radius $x < r < y$, with one of the two voids disappearing at $r < s < y$. In this particular case, we have a side pocket that corresponds to another point, $\frac{1}{2}(s + r, s - r)$, in $\text{Dgm}_2(d_{\mathbb{X}})$.

This interpretation of points in $\text{Dgm}_2(d_{\mathbb{X}})$ raises the question about the meaning of points in the 0-th and the 1-st diagrams. We find a common metaphor by interpreting their geometric realizations as *cages* with dimension and scale. For example, the pocket corresponding to u cages a ball of radius between x and y ; it does not have enough space for a ball of radius larger than y , and it cannot prevent a ball of radius smaller than x from escaping. Similarly, a point $v = \frac{1}{2}(y + x, y - x)$ in $\text{Dgm}_1(d_{\mathbb{X}})$ cages an endless tube of cross-section radius between x and y . Indeed, we can move the tube through the partial loop, but we cannot remove it unless we find a place where its cross-section has radius less than x . Finally, a point $w = \frac{1}{2}(y + x, y - x)$ in $\text{Dgm}_0(d_{\mathbb{X}})$ cages a closed surface uniformly thickened to radius between x and y .

Random polymers. We construct idealized geometric models of polymers iteratively, at each step randomly adding a unit ball to the growing structure. We call the result a *branched polymer* if the new ball is glued at a single point, and this point is chosen uniformly at random. As illustrated in Figure 2, the set of points on the boundary that are available for gluing can be constructed by first doubling the radius and second shrinking the boundary back to the original model. This




Figure 2. A collection of 12 touching unit disks. The set of points on the boundary where a 13-th disk can be glued without creating any additional intersection is constructed using the boundary of the union of the disks with twice the radius.

set is a union of open patches on the spheres bounding the balls, and its area can be computed using software based on alpha shapes [31]. To have a comparison, we introduce *Brownian trees*, which are constructed the same way except that the uniform distribution over the mentioned set is replaced by another distribution that takes into account the difficulty of reaching a point with a unit ball approaching the union by Brownian motion from infinity. The branched peptides and the Brownian trees are easy to distinguish. Indeed, a branched peptide cannot have a large void; a ball that could comfortably fit inside would be added sooner or later. In contrast, a Brownian tree can protect a large void with narrow entrances. Letting $d_P, d_T : \mathbb{R}^3 \rightarrow \mathbb{R}$ be the distance functions defined by a branched peptide and by a Brownian tree, we therefore expect points with large persistence in $\text{Dgm}(d_T)$ but not in $\text{Dgm}(d_P)$.

There is a less obvious difference in the horizontal or *scale* direction. Let $P_p(x)$ be the number of points in $\text{Dgm}_p(d_P)$ with death value plus birth value at least $2x$, and let T_p be the similarly defined function for the Brownian tree. MacPherson and Schweinhart find experimental evidence that P_1 and P_2 are both roughly a constant times $\frac{1}{x^2}$, which is the motivation to say that the branched peptides have *persistence dimension 2*, both for 1-cages and for pockets. No such exponent seems to exist for Brownian trees.

3.2. The Shape of a Human Jaw. Its stability suggests the Wasserstein distance between persistence diagrams as a similarity measure for shapes. Indeed, it is difficult to compare shapes directly, but it is easy to compute and compare persistence diagrams for suitably chosen functions. We discuss this approach by, first, presenting a relation between persistence and the Gromov-Hausdorff distance and, second, reviewing an application to human jaws.

Comparing shapes and metrics. The comparison or fitting of shapes arises in many walks of life — too many to warrant an example. For solid shapes, we may focus on prominent protrusions and cavities, but this is less effective when we deal with flexible shapes. An important subclass of the latter has a boundary that folds but does not stretch or shrink. More formally, this boundary is a space with a constant metric. Following this line of thought, Chazal et al. [9] consider finite metric spaces, \mathbb{X} and \mathbb{Y} , and use persistence diagrams as their stable signatures. We need some definitions to state their results. A *correspondence* between \mathbb{X} and \mathbb{Y} is a subset of $\mathbb{X} \times \mathbb{Y}$ whose projections back to \mathbb{X} and to \mathbb{Y} are the entire spaces. The *Gromov-Hausdorff distance* between the two spaces is

$$\text{GH}(\mathbb{X}, \mathbb{Y}) = \frac{1}{2} \inf \sup \left| \|x - x'\|_{\mathbb{X}} - \|y - y'\|_{\mathbb{Y}} \right|, \quad (7)$$

where we take the infimum over all correspondences, γ , and the supremum over all pairs (x, y) and (x', y') in γ . With this definition in place, we consider the Vietoris-Rips complex for a distance threshold, $a \geq 0$. Specifically, we draw an edge between any two points at distance at most a , and we let $\text{Rips}_a(\mathbb{X})$ be the flag complex defined by these edges; see Section 2.1. Varying a , we let $R_{\mathbb{X}}$ be the resulting sequence of spaces, and, applying the homology functor, we get a persistence module characterized by the persistence diagram, which we denote as $\text{Dgm}(R_{\mathbb{X}})$. A consequence of the main result in [9] is a relation between the bottleneck distance and the Gromov-Hausdorff distance:

$$W_{\infty}(\text{Dgm}(R_{\mathbb{X}}), \text{Dgm}(R_{\mathbb{Y}})) \leq \text{GH}(\mathbb{X}, \mathbb{Y}). \quad (8)$$

We thus get a lower bound on a quantity that is generally difficult to compute and to approximate. If we accept the Gromov-Hausdorff distance as a reasonable comparison of metric spaces, we can use the bottleneck distance between persistence diagrams to disprove that two spaces are similar. But because the inequality is one-sided, we cannot prove their similarity.

Average and individual jaws. An important shape in orthodontics is the human jaw. Comparisons between them have several practical applications: one being the recognition of medical conditions, such as the Habsburger chin; another is the monitoring of ongoing treatments. Traditionally, this comparison is done with a standardized version of the landmark method in statistical shape analysis [18]. Here, we describe an enhancement of this method using persistent homology, as employed by Gamble and Heo [25]. In this particular study, they consider a collection of $N = 240$ jaw bones, each represented by $k = 22$ landmark points chosen

by an expert for their clinical relevance. The points are labeled and denoted as u_i^j , for $1 \leq i \leq k$ and $1 \leq j \leq N$. After aligning the jaw bones in \mathbb{R}^3 , we average the landmarks to get k points $u_i = (\sum_{j=1}^N u_i^j)/N$, which we call the *mean shape* of the N data sets. The k points define a Delaunay triangulation, \mathcal{D} , which is a 3-dimensional simplicial complex with probability 1. We view it as an abstract (as opposed to geometric) simplicial complex, and filter it differently for each data set. To describe this filtration, let \mathcal{D}_j be the j -th copy of the Delaunay triangulation, and define the *weight* of the edge connecting the points u_i and $u_{i'}$ as

$$\text{weight}_j(i, i') = \frac{\|u_i^j - u_{i'}^j\|}{\sum_{\ell=1}^N \|u_i^\ell - u_{i'}^\ell\|}.$$

Inspired by the construction of the Vietoris-Rips complex, we use a real threshold $a \geq 0$ and filter \mathcal{D}_j by taking the maximal subcomplex whose edges have weight at most a . This gives N filtrations of the Delaunay triangulation and, correspondingly, N persistence diagrams, one for each data set.

The final analysis is done in the space of persistence diagrams, which is an important point in this story. Fixing $q = 2$, we get the Wasserstein distance, W_q , between every pair of diagrams. Doing this for individual dimensions but also for the cumulative diagrams, the most interesting results appear in dimension 1. Switching to traditional methods, the pairwise Wasserstein distances are used to embed the data as points in \mathbb{R}^2 using multidimensional scaling [32]. Closely examining the results, Gamble and Heo find that one of the coordinates correlates with the expansion of the jaw — the treatment used on the patients. In particular, it distinguishes between the control group and the two treatment groups as they evolve over time. Interestingly, the inter-landmark distances that have the highest positive correlation with that coordinate are those that cross the mouth and measure the width of the jaw.

3.3. The Connectivity of Root Systems. A common theme in the reconstruction of shapes is topologically correct connectivity. One example are brain surfaces, which, at the commonly adopted scale, all have the topology of the 2-dimensional sphere. Another example are root systems of agricultural plants, which are thickened 1-dimensional trees. In this section, we focus on the contribution of persistent homology to the control of the topological connectivity.

Reconstruction by ordered selection. A common paradigm in the reconstruction of shapes is the selection of cells from an underlying collection, \mathcal{U} . This is often facilitated by estimating a fitness value for each cell; that is: a function $f : \mathcal{U} \rightarrow \mathbb{R}$. Given a threshold, α , we select all cells with fitness at least α . In other words, we reconstruct the shape $f^{-1}[\alpha, \infty)$. This is also the strategy in the reconstruction of root systems as described in [47]. Inevitably, there are cells with fitness value close to the threshold for which the decision depends on chance. To avoid such cases, we may put effort into improving the accuracy of the fitness function. Here, we follow an alternative approach that uses global information to influence the selection.

To make the setting more concrete, assume \mathcal{U} is a decomposition of a compact subset of \mathbb{R}^3 into unit cubes called *voxels*. The information about the root system is obtained from a collection of 2-dimensional photographs taken from different directions, each segmented into *foreground* and *background*, the former being the projection of the root system onto the plane of the camera. We construct the shape as the collection of voxels all of whose projections belong to foregrounds. To make this more realistic, we allow for ambiguity entering the setting through uncertainty about the position and the angle of the camera, imperfect lighting conditions, optical distortion, shape details that challenge the resolution of our observations, etc. Instead of a binary we get a real-valued fitness function, as discussed earlier. To shed light on the dependence of the reconstruction on the threshold, we sort the voxels in the order of non-increasing fitness, adding their square sides, edges, and vertices, making sure that every cell (of any dimension) succeeds its faces in the ordering. We call the result a *filter*, listing its cells in order as $\sigma_1, \sigma_2, \dots, \sigma_m$. Letting K_i be the complex consisting of σ_1 to σ_i , we get a filtered complex: $\emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_m$. Each complex K_i is our best choice for $f(\sigma_i) \geq \alpha > f(\sigma_{i+1})$.

How do we know that it would not be better to use a slightly different threshold or to permute some of the cells with same or similar fitness values? We use persistence to elucidate this question. When the target connectivity is clear, this perspective leads to improved local choices. For a root system, we expect $\beta_0 = 1$ and $\beta_1 = \beta_2 = 0$; that is: a connected shape without tunnels and voids.

Local reordering. To get started, we apply the homology functor to get a persistence module $0 = \mathbf{H}(K_0) \rightarrow \dots \rightarrow \mathbf{H}(K_m)$. As explained in Section 2.2, homology classes are born and die. In our case, they correspond to components, loops, and closed walls. Since the complex is built up one cell at a time, we can associate these events with individual cells.

dim = 0: a vertex gives birth to a new component; there is no other case.

dim = 1: an edge gives either birth to a loop, or it kills a component by bridging the gap to another component.

dim = 2: a square gives either birth to a wall, or it kills a loop by filling in the last opening of the tunnel.

dim = 3: a voxel kills a wall by filling in the last piece of the void it surrounds; there is no other case.

Importantly, we can associate each birth with a death, or with infinity if it marks a homology class of the last complex. We visualize these pairs as intervals in the barcode, paying special attention to the ones that contain the threshold, α . Suppose there are $\beta_0 + \beta_1 + \beta_2$ such intervals, and note that they correspond to the components, loops, and walls in K_i , where $f(\sigma_i) \geq \alpha > f(\sigma_{i+1})$. In the lucky case, we have $\beta_0 = 1$ and $\beta_1 = \beta_2 = 0$, and, therefore, a connected reconstruction of the root system, without loops or voids, as desired. Otherwise, we aim at removing

all surplus intervals, which we do by modifying the fitness values of the cells. This may lead to changes in the ordering, which we decompose into transpositions of contiguous cells, an operation we discuss next.

Suppose we increase the fitness of a cell $\sigma = \sigma_\ell$, let $\tau = \sigma_{\ell-1}$ be the cell to its left, and assume that the increase improves the fitness of σ beyond that of τ . If there is no dependence between σ and τ , then we can just transpose them. If the transposition affects the pairing we call it a *switch* and refer to [14] for a complete analysis and a fast update algorithm. The most interesting case is the switch in which σ and τ change their status: from giving birth to giving death and vice versa. Finally, if τ is a face of σ , then the transposition is prohibited, and we have to increase the fitness value of τ along with that of σ . Moving τ may have an adverse effect on the connectivity at α since τ may be the endpoint of another interval. Indeed, obstacles to repair cannot always be avoided as the general problem of optimal reconstruction is NP-hard [2]. Notwithstanding these shortcomings, the filter is an efficient mechanism for the control of the topology of the reconstructed shape. Considering the widespread need of topology repair in the applications, this presents a significant potential for the improvement of reconstruction algorithms.

3.4. The Statistics of Natural Images. After discussing low-dimensional applications, we are ready to extrapolate what we learned to dimensions beyond the visible. The need for such extensions is substantial because scientists collect progressively more and larger datasets whose meaning is hidden in the invisible dimensions. An example are cancer profiles which promise to shed new light on individual differences. In this section, we focus on high-dimensional data derived from photographs, following the work of Carlsson et al. [8].

Image statistics. To understand the variation of receptive properties of simple cells in the mammalian visual cortex, van Hateren and van der Schaaf [29] used a carefully calibrated digital camera to gather a collection of 4,212 images of natural environments — woods, open landscapes, and urban areas. Following earlier work, their research relates the statistics of such images to the cell properties and supports the proposition that the cells have evolved to process natural images. Moving toward a mathematically accessible setting, Carlsson et al. [8] consider the topology of 3-by-3 high-contrast patches extracted from these photographs. Studies show that humans look more in the regions of high spatial contrast, which justifies their emphasis. The restriction to the small patches is especially interesting. It allows to dramatically reduce the dimensionality of the problem to nine while preserving information about the global statistics of the image. The patches are selected as follows:

- (1) after picking 5,000 patches in each image, we treat each one as a vector with nine coordinates (one per pixel) and therefore as a point in \mathbb{R}^9 ;
- (2) we subtract the average from each component, noting that this puts every point on a hyperplane, which we identify with \mathbb{R}^8 , and moves low contrast

patches close to the origin of \mathbb{R}^8 ;

- (3) defining the *contrast* of a patch as the norm of the point, we select the 1,000 patches with highest contrast from each image;
- (4) normalizing by the contrast, we obtain a set of points in the 7-dimensional sphere in \mathbb{R}^8 , which we denote as \mathbb{S}^7 .

For computational reasons, the space is down-sampled further, from about four million to 50,000 points. Even after the initial filtering, the data set contains more information than we can comprehend. Therefore, it is prudent to focus on its core subsets to expose otherwise obscured phenomena. To this end, the remaining high-contrast patches are filtered by their local densities in \mathbb{S}^7 . Specifically, we compute the distance to the k -nearest neighbor for each point, and we write $X(k, P)$ for the top P percent of the points ordered by this distance measurement.

Popular subspaces. Like the knobs on a microscope, the parameters k and P control the focus of our view. At the coarsest scale, the space $X(300, 30)$ consists of a single circle, noticeable in the 1-dimensional persistence diagram of the distance function. Inspecting it, Carlsson and collaborators find that it consists of linear gradients, rotating around the center of the patch. In Figure 3, it is depicted by the patches on the two horizontal gray lines; they connect into a single circle by identifying the matching patches at their opposite ends.




Figure 3. The Klein bottle of 3-by-3 patches. The horizontal edges are glued to each other from left to right, and the vertical edges are glued with a twist.

After sharpening the view by transitioning to $X(15, 30)$, the 1-dimensional persistence diagram detects five prominent homology classes. Inspection of the point set verifies the ‘3-circle model’, suggested in earlier work by Carlsson and de Silva. In addition to the first, two more circles appear in $X(15, 30)$ and intersect the primary circle in two points each. (The first Betti number of the resulting space

is indeed five.) In Figure 3, the matching patches at the top and the bottom are identified, turning the two dotted lines into circles. The appearance of the three circles in the high-density subsample hints at the two preferences in natural images: linear intensity functions as well as vertical and horizontal directions. Turning the knobs further down and more, we fill in lower density regions. Persistent homology of the resulting point set, taken with modulo two coefficients, acquires a 2-dimensional class and retains two independent 1-cycles. The torus and the Klein bottle are the only 2-manifolds with this homology. By examining how they fit into the point set (both experimentally and theoretically), the Klein bottle prevails. Figure 3 illustrates the corresponding arrangement of the patches.

Knowing that the bulk of the points lie near a Klein bottle, Carlsson and collaborators push on to find an explicit representation of this 2-manifold inside \mathbb{S}^7 . The motivation is image compression. A point (a 3-by-3 patch) on the Klein bottle is fully specified by only two coordinates. There are not many points far from the Klein bottle, and each such point can be specified by two coordinates plus a residual description of the difference to the projection onto the Klein bottle.

4. Discussion

Persistent homology is a new mathematical concept that has received attention from inside and outside mathematics. It is our interpretation that the reason for the interest is multi-faceted. We hope that the structure of this paper has made this point clear. In particular,

- we stress the connection to *data* in Section 2.1;
- we emphasize the *algebraic* side of persistent homology in Section 2.2;
- we explain the *stability* of its diagrams in Section 2.3;
- we sketch its fundamental *algorithms* in Section 2.4;
- we shed light on the role of *scale* in Section 3.1;
- we discuss *derived metrics* facilitating the analysis of shapes in Section 3.2;
- we exhibit the control of *topological connectivity* in Section 3.3;
- and we show that *high dimensions* aid our understanding in Section 3.4.

What are the developments we may expect to push the envelope of the method in the next few years? We see multi-parameter persistence, the statistics of persistence, and persistence for dynamical systems as major thrusts of the current research. All three are driven by applications, as was persistent homology from its very beginning.

References

- [1] P. K. Agarwal, H. Edelsbrunner, J. Harer and Y. Wang. Extreme elevation on a 2-manifold. *Discrete Comput. Geom.* **36** (2006), 553–572.
- [2] D. Attali and A. Lieutier. Optimal reconstruction might be hard. In “Proc. 26th Ann. Sympos. Comput. Geom., 2010”, 334–343.
- [3] P. Bendich, H. Edelsbrunner and M. Kerber. Computing robustness and persistence for images. *IEEE Trans. Visual. Comput. Graphics* **16** (2010), 1251–1260.
- [4] S. Biasotti, L. De Floriani, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, L. Papaleo and M. Spagnuolo. Describing shapes by geometrical-topological properties of real functions. *ACM Comput. Surveys* **40** (2008), 1–87.
- [5] G. Carlsson. Topology and data. *Bulletin Amer. Math. Soc.* **46** (2009), 255–308.
- [6] G. Carlsson and V. de Silva. Zigzag persistence. *Found. Comput. Math.* **10** (2010), 367–405.
- [7] G. Carlsson, V. de Silva and D. Morozov. Zigzag persistent homology and real-valued functions. In “Proc. 25th Ann. Sympos. Comput. Geom., 2009”, 227–236.
- [8] G. Carlsson, T. Ishkhanov, V. de Silva and A. Zomorodian. On the local behavior of spaces of natural images. *Internat. J. Comput. Vision* **76** (2008), 1–12.
- [9] F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli and S. Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Comput. Graphics Forum* **28** (2009), 1393–1403.
- [10] F. Chazal, D. Cohen-Steiner, L. Guibas and S. Oudot. Proximity of persistence modules and their diagrams. In “Proc. 25th Ann. Sympos. Comput. Geom., 2009”, 237–246.
- [11] D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.
- [12] D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79–103.
- [13] D. Cohen-Steiner, H. Edelsbrunner, J. Harer and Y. Mileyko. Lipschitz functions have L_p -stable persistence. *Found. Comput. Math.* **10** (2010), 127–139.
- [14] 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.
- [15] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk* **7** (1934), 793–800.
- [16] C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Comput. Aided Geom. Design* **12** (1995), 771–784.
- [17] H. Derksen and J. Weyman. Quiver representations. *Notices Amer. Math. Soc.* **52** (2005), 200–206.
- [18] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. Wiley, Chichester, England, 1998.
- [19] H. Edelsbrunner, M. A. Facello and J. Liang. On the definition and the construction of pockets in macromolecules. *Discrete Appl. Math.* **88** (1998), 83–102.

- [20] H. Edelsbrunner and J. Harer. *Computational Topology. An Introduction*. Amer. Math. Soc., Providence, Rhode Island, 2010.
- [21] H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **29** (1983), 551–559.
- [22] H. Edelsbrunner, D. Letscher and A. Zomorodian. Topological persistence and simplification. In “Proc. 41st IEEE Sympos. Found. Comput. Sci., 2000”, 454–463, also *Discrete Comput. Geom.* **28** (2002), 511–533.
- [23] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.
- [24] P. Frosini. A distance for similarity classes of submanifolds of a Euclidean space. *Bulletin Australian Math. Soc.* **42** (1990), 407–416.
- [25] J. Gamble and G. Heo. Exploring uses of persistent homology for statistical analysis of landmark-based shape data. *J. Multivar. Analysis* **101** (2010), 2184–2199.
- [26] L. Georgiadis, R. E. Tarjan and R. F. Werneck. Design of data structures for mergeable trees. In “Proc. 17th Ann. ACM-SIAM Sympos. Discrete Alg., 2006”, 394–403.
- [27] R. Ghrist. Barcodes: the persistent topology of data. *Bulletin Amer. Math. Soc.* **45** (2008), 61–75.
- [28] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, England, 2002.
- [29] J. H. van Hateren and A. van der Schaaf. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. Royal Soc. London B* **265** (1998), 359–366.
- [30] A. B. Hamza and H. Krim. Geodesic object representation and recognition. *Lect. Notes Comput. Sci.* **2886** (2003), 378–387.
- [31] P. Koehl. ProShape: Understanding the shape of protein structures. Software at biogeometry.duke.edu/software/proshape/.
- [32] J. B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* **29** (1964), 1–27.
- [33] B. Lee and F. M. Richards. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.
- [34] J. Leray. L’anneau d’homologie d’une représentation and Structure de l’anneau d’homologie d’une représentation. *Les Comptes rendus de l’Académie des Sciences* **222** (1946), 1366–1368 and 1419–1422.
- [35] J. Liang, H. Edelsbrunner and C. Woodward. Anatomy of protein pockets and cavities: measurement of binding site geometry and implications for ligand binding. *Protein Science* **7** (1998), 1884–1897.
- [36] R. MacPherson and B. Schweinhart. Measuring shape with topology. arXiv:1011.2258, 2010.
- [37] N. Milosavljević, D. Morozov and P. Škraba. Zigzag persistent homology in matrix multiplication time. In “Proc. 27th Ann. Sympos. Comput. Geom., 2011”, 216–225.
- [38] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
- [39] V. Robins. Toward computing homology from finite approximations. *Topology Proceedings* **24** (1999), 503–532.

- [40] H. Samet. *Foundations of Multidimensional and Metric Data Structures*. Morgan Kaufmann, 2006.
- [41] V. Strassen. Relative bilinear complexity and matrix multiplication. *J. Reine Angew. Math.* **375/376** (1987), 406–443.
- [42] C. Villani. *Topics in Optimal Transportation*. Amer. Math. Soc., Providence, Rhode Island, 2003.
- [43] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.* **133** (1907), 97–178 and **134** (1908), 198–287.
- [44] H. Wagner, C. Chen and E. Vuçini. Efficient computation of persistent homology for cubical data. In “Proc. 4th Workshop Topology-based Methods in Data Analysis and Visualization”, 2011.
- [45] S. Weinberger. What is ... persistent homology? *Notices Amer. Math. Soc.* **58** (2011), 36–39.
- [46] V. Williams. Breaking the Coppersmith-Winograd barrier. Preprint. (2011)
- [47] 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.
- [48] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.* **33** (2005), 249–274.

Herbert Edelsbrunner, IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria,
Duke University, Durham, North Carolina, and Geomagic, Research Triangle Park,
North Carolina.

E-mail: edels@ist.ac.at

Dmitriy Morozov, Lawrence Berkeley National Lab, 1 Cyclotron Road, Berkeley, CA
94720-8139, USA

E-mail: dmitriy@mrzv.org