Homology, Homotopy and Applications, vol. 15(1), 2013, pp.51–72

HOMOLOGY AND ROBUSTNESS OF LEVEL AND INTERLEVEL SETS

PAUL BENDICH, HERBERT EDELSBRUNNER, DMITRIY MOROZOV AND AMIT PATEL

(*communicated by Robert Ghrist*)

Abstract

Given a continuous function $f\colon \mathbb{X} \to \mathbb{R}$ on a topological space, we consider the preimages of intervals and their homology groups and show how to read the ranks of these groups from the extended persistence diagram of f. In addition, we quantify the robustness of the homology classes under perturbations of f using well groups, and we show how to read the ranks of these groups from the same extended persistence diagram. The special case $\mathbb{X} = \mathbb{R}^3$ has ramifications in the fields of medical imaging and scientific visualization.

1. Introduction

The work reported in this paper has two motivations, one theoretical and the other practical. The former is the recent introduction of *well groups* in the study of mappings $f\colon \mathbb{X} \to \mathbb{Y}$ between topological spaces. Assuming a metric space of perturbations, we have such a group for each subspace $\mathbb{A} \subseteq \mathbb{Y}$, each bound $r \geqslant 0$ on the magnitude of the perturbation, and each homological dimension p. These groups, and the diagrams that they generate, extend the boolean concept of transversality to a real-valued measure we refer to as *robustness*. Using this measure, we can quantify the robustness of a fixed point of a mapping [**10**] and prove the stability of the apparent contour of a mapping from an orientable 2-manifold to $\mathbb{R}^2$ [**9**]. In this paper, we contribute to the general understanding of well groups by studying the real-valued case. Along the way, we also extend the general theory of well groups to incorporate relative well groups. Specifically,

I. We give a general definition of relative well groups given a mapping $f\colon \mathbb{X} \to \mathbb{Y}$, a number $r \geqslant 0$, and a nested pair $\mathbb{A}' \subseteq \mathbb{A}$ of subspaces of $\mathbb{Y}$, and
II. We characterize the relative well groups of $f\colon \mathbb{X} \to \mathbb{R}$ whenever $\mathbb{A}$ is an interval and $\mathbb{A}'$ is a subset of the endpoints.

This research is partially supported by the National Science Foundation (NSF) under grant DBI-0820624 and the Defense Advanced Research Projects Agency (DARPA) under grants HR0011-05-1-0057 and HR0011-09-0065.
Received January 4, 2011, revised January 10, 2012; published on March 5, 2013.
2000 Mathematics Subject Classification: 55, 68.
Key words and phrases: persistence, zigzag, levelset, homology, perturbation, well group, robustness.
Article available at `http://intlpress.com/HHA/v15/n1/a3` and `doi:10.4310/HHA.2013.v15.n1.a3`

Applications of this theoretical work are anticipated in medical imaging and scientific visualization, where data in the form of real-valued functions is common. To mention one example, it is common to acquire information about internal organs through a magnetic resonance image, which results in a 3-dimensional array of intensity values, best viewed as a function from the unit cube to the real line. The predominant method for highlighting or extracting relevant substructures of this image uses preimages of real values. Generically, these are 2-manifolds, commonly referred to as *contours* or *isosurfaces* [**12**]. Sometimes, these 2-manifolds are complemented by preimages of intervals, referred to as *interval volumes* in visualization [**11**]. In this paper, we call the preimage of a value a *level set*, and the preimage of an interval an *interlevel set*, in which the interval can be closed, open, or half-open. We contribute to the state-of-the-art by

III. Explaining how the homology of level and interlevel sets can be read off the extended persistence diagram of the function, and
IV. Describing how the robustness of features in level and interlevel sets, quantified through well groups, can be read off the same diagram.

Our results add up to a 'point calculus' in algebraic topology for mining the rich homological information contained in the extended persistence diagram of a real-valued function. The compactness of the data representation and the efficiency of the mining operations make the diagram an attractive graphical interface tool for studying 3-dimensional images. We view this tool as complementary to the contour spectra described in [**1**], which plot continuously varying quantities, such as area and volume, across the family of level sets. The most novel aspect of our diagram is the robustness information, which has previously not been available. This novelty is combined with the unprecedented ease with which homological information is accessible. There is also evidence for the practicality of the interface provided by the fast oct-tree implementation of the described concepts [**2**], which has been used to study 3-dimensional images of root systems of agricultural plants.

In Section 2, we review necessary background on persistence, zigzag modules, and well groups. In Section 3, we explain the point calculus for interlevel sets. In Section 4, we extend the point calculus to include the robustness information provided by the well groups. Finally, Section 5 concludes the paper with a brief discussion of the contributions and of future research directions.

Acknowledgements

The authors thank Vin de Silva for many helpful technical discussions leading to a strengthening of this paper.

2. Background

We divide the background material into three parts, introducing persistence and extended persistence in Section 2.1, explaining the extension to zigzag modules and level set pyramids in Section 2.2, and defining absolute and relative well groups in Section 2.3.

2.1. Forward Maps

Traditional persistent homology is based on a nested sequence of spaces, which induces a linear sequence of homology groups connected by maps from left to right. We describe this concept in two steps.

2.1.1. Persistence

The persistence of homology classes along a filtration of a topological space can be defined in a quite general context [**8**]. For this paper, we need only a particular type of filtration, one defined by the sublevel sets of a tame function. Given a real-valued function f on a compact topological space $\mathbb{X}$, we consider the filtration of $\mathbb{X}$ via the *sublevel sets* $\mathbb{X}_r(f) = f^{-1}(-\infty, r]$, for all real values r. Whenever $r < s$, the inclusion $\mathbb{X}_r(f) \hookrightarrow \mathbb{X}_s(f)$ induces maps on the homology groups $\mathsf{H}_p(\mathbb{X}_r(f)) \to \mathsf{H}_p(\mathbb{X}_s(f))$, for each dimension p. Here we will use field coefficients so that the homology groups are vector spaces over the field. Often we will suppress the homological dimension from our notation, writing $\mathsf{H}(\mathbb{X}_r(f)) = \bigoplus_p \mathsf{H}_p(\mathbb{X}_r(f))$; in this case, we will always assume that all maps $\mathsf{H}(\mathbb{X}_r(f)) \to \mathsf{H}(\mathbb{X}_s(f))$ decompose into the direct sum of maps on each factor. A real value r is called a *homological regular value* of f if there exists $\epsilon > 0$ such that the inclusion $\mathbb{X}_{r-\delta}(f) \hookrightarrow \mathbb{X}_{r+\delta}(f)$ induces an isomorphism between homology groups for every $\delta < \epsilon$. If r is not a homological regular value, then it is a *homological critical value*. We say that f is *tame* if it has finitely many homological critical values and if the homology groups of each sublevel set have finite rank. Assuming that f is tame, we enumerate its homological critical values $r_1 < r_2 < \ldots < r_n$. Choosing $n+1$ homological regular values s_i such that $s_0 < r_1 < s_1 < \ldots < r_n < s_n$, we put $\mathbb{X}_i = \mathbb{X}_{s_i}(f)$. The inclusions $\mathbb{X}_i \hookrightarrow \mathbb{X}_j$ induce maps $\mathsf{f}^{i,j} \colon \mathsf{H}(\mathbb{X}_i) \to \mathsf{H}(\mathbb{X}_j)$ for $0 \leqslant i \leqslant j \leqslant n$ and give the following filtration:

$$0 = \mathsf{H}(\mathbb{X}_0) \to \mathsf{H}(\mathbb{X}_1) \to \ldots \to \mathsf{H}(\mathbb{X}_n) = \mathsf{H}(\mathbb{X}). \tag{1}$$

We say a class $\alpha \in \mathsf{H}(\mathbb{X}_i)$ is *born* at $\mathbb{X}_i$ if $\alpha \notin \operatorname{im} \mathsf{f}^{i-1,i}$. A class α born at $\mathbb{X}_i$ is said to *die entering* $\mathbb{X}_j$ if $\mathsf{f}^{i,j}(\alpha) \in \operatorname{im} \mathsf{f}^{i-1,j}$ but $\mathsf{f}^{i,j-1}(\alpha) \notin \operatorname{im} \mathsf{f}^{i-1,j-1}$. We remark that if a class α is born at $\mathbb{X}_i$, then every class in the coset $[\alpha] = \alpha + \operatorname{im} \mathsf{f}^{i-1,i}$ is born at the same time. Of course, whenever such an α dies entering $\mathbb{X}_j$, the entire coset $[\alpha]$ also dies with it. We represent $[\alpha]$ graphically as the point (r_i, r_j) in the plane. Drawing all birth-death pairs as points, we get diagrams like the ones sketched in Figures 1 and 3. Supposing that $b \in \mathbb{R}$ is different from all homological critical values, we collect all points in the upper-left quadrant defined by (b, b) to get all classes born before b and still alive; see the left diagram in Figure 1. Their number is the rank of the homology group of the sublevel set, $\operatorname{rank} \mathsf{H}(\mathbb{X}_b(f))$.

Observe that we really need the extended plane to draw the points because some classes are born but never die, so the corresponding points have ∞ as their second coordinates. There is an elegant way around this minor annoyance, which we now describe.

2.1.2. Extended persistence

Since the filtration in (1) begins with the zero group but ends with a potentially nonzero group, it is possible to have classes that are born but never die. We call these *essential* classes, as they represent the actual homology of the space $\mathbb{X}$. To measure the persistence of the essential classes, we follow [**7**] and extend the sequence in (1) using

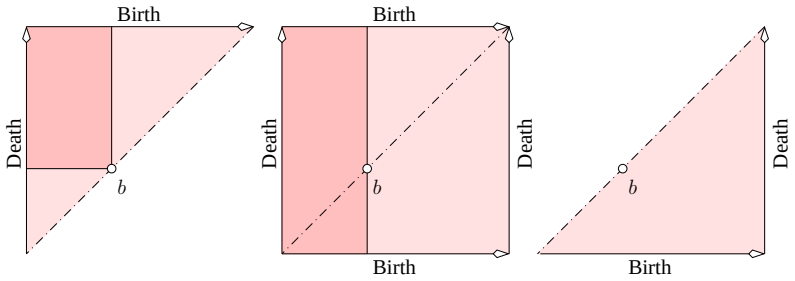

Figure 1: From left to right: the ordinary, extended, and relative subdiagrams of $\mathcal{D}(f)$. The number of points (not shown) in the dark shaded regions is equal to the rank of the homology group of the sublevel set defined by b.

relative homology groups. More precisely, we consider for each i the *superlevel set* $\mathbb{X}^i = f^{-1}[s_{n-i}, \infty)$. Note that we have $\mathbb{X}^0 = \emptyset$ and $\mathbb{X}^n = \mathbb{X}$ by compactness. For $i < j$, the inclusion $\mathbb{X}^i \hookrightarrow \mathbb{X}^j$ induces a map on relative homology $\mathsf{H}(\mathbb{X}, \mathbb{X}^i) \to \mathsf{H}(\mathbb{X}, \mathbb{X}^j)$. These maps therefore give rise to the following extended filtration:

$$0 = \mathsf{H}(\mathbb{X}_0) \to \mathsf{H}(\mathbb{X}_1) \to \ldots \to \mathsf{H}(\mathbb{X}_n) = \mathsf{H}(\mathbb{X}, \mathbb{X}^0) \to \ldots \to \mathsf{H}(\mathbb{X}, \mathbb{X}^n) = 0. \quad (2)$$

We extend the notions of birth and death in the obvious way. Since this filtration begins and ends with the zero group, all classes eventually die. We also extend the graphical representation of the information contained by forming *persistence diagrams*, which we now introduce more formally. We have such a diagram for each dimension p; see Figure 1. Each diagram is a multiset of points in the plane, containing one point (r_i, r_j) for each coset of classes that is born at $\mathbb{X}_i$ or $(\mathbb{X}, \mathbb{X}^{n-i+1})$, and dies entering $\mathbb{X}_j$ or $(\mathbb{X}, \mathbb{X}^{n-j+1})$. In some circumstances, it is convenient to add the points on the diagonal to the diagram, but in this paper, we will refrain from doing so. The persistence diagram contains three important subdiagrams, corresponding to three different combinations of birth and death location. The *ordinary subdiagram*, $\mathcal{O}_p(f)$, represents classes that are born and die during the first half of (2). The *relative subdiagram*, $\mathcal{R}_p(f)$, represents classes that are born and die during the second half. Finally, the *extended subdiagram*, $\mathcal{E}_p(f)$, represents classes that are born during the first half and die during the second half of the extended filtration. Note that points in $\mathcal{O}_p(f)$ all lie above the main diagonal while points in $\mathcal{R}_p(f)$ all lie below. On the other hand, $\mathcal{E}_p(f)$ may contain points on either side of the main diagonal. By $\mathcal{D}(f)$, we mean the points of all diagrams in all dimensions. Drawing these subdiagrams side by side can be cumbersome, and drawing them on top of each other can be confusing. In Section 3, we will introduce a new design that addresses these concerns.

2.2. Mixed Maps

We note that the homology groups in the extended filtration of (2), or in the shorter filtration of (1), are all vector spaces over a fixed field and that the maps between them are all linear maps. In [**5**], Carlsson and de Silva generalize this situation to sequences of vector spaces that are connected by maps going from left to right or from right to left. We now briefly review their work as well as the related work on level set zigzag modules in [**6**].

2.2.1. Zigzag modules

A *zigzag module* W is a finite sequence of vector spaces connected by linear maps which either go forward or backward between consecutive spaces:

$$\mathsf{W}_1 \leftrightarrow \mathsf{W}_2 \leftrightarrow \ldots \leftrightarrow \mathsf{W}_j \leftrightarrow \mathsf{W}_{j+1} \leftrightarrow \ldots \leftrightarrow \mathsf{W}_n. \tag{3}$$

If the arrow advances from W_j to W_{j+1}, then we denote the corresponding linear map as $a_j \colon \mathsf{W}_j \to \mathsf{W}_{j+1}$; otherwise, we write $b_j \colon \mathsf{W}_{j+1} \to \mathsf{W}_j$. A *submodule* U of W is a collection of linear subspaces $\mathsf{U}_j \subseteq \mathsf{W}_j$ such that $a_j(\mathsf{U}_j) \subseteq \mathsf{U}_{j+1}$ or $b_j(\mathsf{U}_{j+1}) \subseteq \mathsf{U}_j$, whichever is the case for j. A submodule U is a *summand* if there is a complementary submodule V, meaning every vector space splits as a direct sum $\mathsf{W}_j = \mathsf{U}_j \oplus \mathsf{V}_j$. The authors in [**5**] prove that every zigzag module can be split into indecomposable summands of a certain form, and, in particular, it has a basis, a concept we now describe. First, we suppose that we have, for each j, a set of elements $u_j^i \in \mathsf{W}_j$ such that the nonzero elements form a basis of W_j. In other words, we can decompose W_j into the direct sum $\mathsf{W}_j = \bigoplus_i \langle u_j^i \rangle$, noting that some of the terms on the right hand side may be zero. We use the superscripts to form correspondences between the bases. Specifically, we require $a_j(u_j^i) = u_{j+1}^i$, or $b_j(u_{j+1}^i) = u_j^i$, depending on the case. Furthermore, we assume that, for each superscript i, there exist $x \leqslant y$ such that $u_j^i \neq 0$ iff $j \in [x, y]$. In other words, for each fixed i, we have a submodule

$$\langle u_1^i \rangle \leftrightarrow \langle u_2^i \rangle \leftrightarrow \ldots \leftrightarrow \langle u_j^i \rangle \leftrightarrow \langle u_{j+1}^i \rangle \leftrightarrow \ldots \leftrightarrow \langle u_n^i \rangle \tag{4}$$

of W in which the non-zero vector spaces are 1-dimensional and form a single contiguous subsequence connected by identity maps. Calling such a submodule an *interval module*, we think of it as being in correspondence with the closed interval $[x, y]$. The collection $\{u_j^i\}$ is a *basis* for the zigzag module if W can be decomposed into the direct sum of the interval modules (4). Equivalently, the collection is a basis for W if each map a_j is the direct sum of the maps $\langle u_j^i \rangle \to \langle u_{j+1}^i \rangle$, and each map b_j is the direct sum of the maps $\langle u_{j+1}^i \rangle \to \langle u_j^i \rangle$, whichever one is defined.

Although a zigzag module W can have many different bases, the set of intervals associated to any such basis will be unique [**5**]. For example, any basis for the zigzag module given by the filtration in (1) will have one interval $[x, y]$ for each coset of classes born at $\mathbb{X}_x$ and dying entering $\mathbb{X}_y$.

2.2.2. Mayer-Vietoris diamonds

We are interested in an elementary operation that connects two minimally different zigzag modules: a *Mayer-Vietoris diamond* [**5**]. We suppose that we have two zigzag modules differing only at position j, and that at this position we have a diamond of the following form:

$$\begin{array}{ccccc} & & \mathsf{H}(\mathbb{V}, \mathbb{V}') & & \\ & \overset{a_{j-1}}{\nearrow} & & \overset{b_j}{\nwarrow} & \\ \longleftrightarrow \mathsf{H}(\mathbb{C}, \mathbb{C}') & & & & \mathsf{H}(\mathbb{D}, \mathbb{D}') \longleftrightarrow \\ & \underset{b_{j-1}}{\nwarrow} & & \underset{a_j}{\nearrow} & \\ & & \mathsf{H}(\mathbb{E}, \mathbb{E}') & & \end{array} , \tag{5}$$

where we show the more general, relative form in which the primed spaces are subspaces of the corresponding unprimed ones, and we have $\mathbb{E} = \mathbb{C} \cap \mathbb{D}$, $\mathbb{E}' = \mathbb{C}' \cap \mathbb{D}'$, $\mathbb{V} = \mathbb{C} \cup \mathbb{D}$, and $\mathbb{V}' = \mathbb{C}' \cup \mathbb{D}'$. We get the more special, absolute form by setting $\mathbb{C}' = \mathbb{D}' = \mathbb{E}' = \mathbb{V}' = \emptyset$. The name of the diamond is justified by the long exact sequence we get by reading the diamond from bottom to top and iterating through the dimensions. When the primed spaces are all empty, this gives the classic version of the Mayer-Vietoris sequence, and more generally, we get the relative version:

$$\ldots \to \mathsf{H}_p(\mathbb{E}, \mathbb{E}') \to \mathsf{H}_p(\mathbb{C}, \mathbb{C}') \oplus \mathsf{H}_p(\mathbb{D}, \mathbb{D}') \to \mathsf{H}_p(\mathbb{V}, \mathbb{V}') \to \mathsf{H}_{p-1}(\mathbb{E}, \mathbb{E}') \to \ldots;$$

see e.g. [**13**]. Importantly, this sequence is exact, which means that the image of each map equals the kernel of the next map.

Such diamonds arise in the following context. Consider again the function $f\colon \mathbb{X} \to \mathbb{R}$ and the interleaved sequence of homological regular and critical values: $s_0 < r_1 < s_1 < \ldots < r_n < s_n$. Setting $\mathsf{W}_{2j} = \mathsf{H}(f^{-1}(s_j))$ and $\mathsf{W}_{2j+1} = \mathsf{H}(f^{-1}[s_j, s_{j+1}])$, we get a zigzag module of length $2n+1$, which, following [**6**], we refer to as the *level set zigzag* of f. It starts and ends with 0 and alternates between advancing maps a_{2j} and backward maps b_{2j+1}. From this module, we can create a new one by fixing an index j, substituting $[s_j, s_{j+2}] = [s_j, s_{j+1}] \cup [s_{j+1}, s_{j+2}]$ for $s_{j+1} = [s_j, s_{j+1}] \cap [s_{j+1}, s_{j+2}]$, and leaving all other groups unchanged; of course we also reverse the two maps involving the changed space. This produces a new zigzag module which differs from the old via a Mayer-Vietoris diamond. This construction can be generalized by flipping between intersections and unions of larger intervals and pairs of intervals, thus producing a whole array of zigzag modules which differ via Mayer-Vietoris diamonds.

2.2.3. The pyramid

Starting with the level set zigzag, we get an array of zigzag modules which are best described as monotonic paths that go diagonally up and down, always from left to right. The array of such paths is connected within a pyramidal structure, which we now describe. As a graphical guide, we consider the square drawn in Figure 2. We give it a coordinate system by parameterizing the downward slope from ∞ at the upper left, to $-\infty$ in the middle, and back up to ∞ at the lower right. Similarly, we parameterize the upward slope from $-\infty$ at the lower left, to ∞ in the middle, and back to $-\infty$ at the upper right. The two slopes divide the square into four triangular regions, each containing a point with coordinates a and b for every choice of $a \leqslant b$. We interpret this point differently in each of the regions. To explain this interpretation, it is convenient to introduce a shorthand that uses open set notation for pairs of closed sets, writing $\mathbb{A} - \mathbb{A}'$ for $(\mathbb{A}, \mathbb{A}')$. Specifically,

$$\begin{aligned}
f^{-1}(x, y] &= (f^{-1}(-\infty, y], f^{-1}(-\infty, x]),\\
f^{-1}[x, y) &= (f^{-1}[x, \infty), f^{-1}[y, \infty)),\\
f^{-1}(x, y) &= (f^{-1}(-\infty, \infty), f^{-1}(-\infty, x] \cup f^{-1}[y, \infty)).
\end{aligned}$$

If a point with coordinates x and y lies in the bottom region, we think of it as the space $f^{-1}[x, y]$. However, if the point lies in the left, right, or top region, we think of it as $f^{-1}(x, y]$, $f^{-1}[x, y)$, or $f^{-1}(x, y)$, respectively. If we now take $w < x < y < z$ and consider the points (w, y), (w, z), (x, y), and (x, z), we get a Mayer-Vietoris diamond in each region; see Figure 2. This is easiest to see in the closed interval

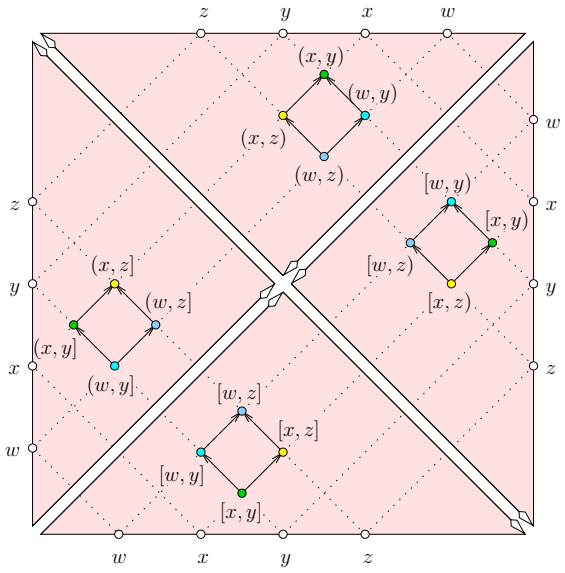

Figure 2: Points in the pyramid are absolute and relative homology groups. Monotonic paths are zigzag modules, any two of which differ by a finite number of Mayer-Vietoris diamonds.

case since $[x, y] = [w, y] \cap [x, z]$ and $[w, z] = [w, y] \cup [x, z]$. In the closed-open case, we have $[x, \infty) = [w, \infty) \cap [x, \infty)$ and $[w, \infty) = [w, \infty) \cup [x, \infty)$ as well as $[z, \infty) = [z, \infty) \cap [y, \infty)$ and $[y, \infty) = [z, \infty) \cup [y, \infty)$. Similar computations verify the diamond in the remaining two cases.

By repeated application of the diamond, we can generate any monotonic path from the one along the bottom edge of the square. Each path is thus decorated by spaces as described, and applying the homology functor gives a zigzag module of absolute and relative homology groups. The latter arise when we move the left or right end of the path, which can be done without the Mayer-Vietoris diamond because the corresponding spaces are and stay empty so that the module remains unchanged. Besides the level set zigzag along the bottom edge, we are particularly interested in the path along the upward slope, which translates into the extended filtration of (2). Its midpoint is $(-\infty, \infty)$, the center of the square, which results in $\mathsf{H}(f^{-1}(-\infty, \infty)) = \mathsf{H}(\mathbb{X})$. For this reason, we think of the center as the apex of a pyramid, as viewed from above.

Remark 1. As a partial justification for the notation with open sets, we mention that the homology group of the preimage of the interval (x, y), if computed with infinite chains, is isomorphic to the relative homology group of $(f^{-1}[x, y], f^{-1}(x) \cup f^{-1}(y))$. By excision, this is isomorphic to the relative homology group of

$$(f^{-1}(-\infty, \infty], f^{-1}(-\infty, x] \cup f^{-1}[y, \infty)).$$

2.3. Perturbations

The reader who wishes to learn how to read the homology of interlevel sets can safely skip Section 2.3 and now continue with Section 3. However, to differentiate the robust from the non-robust homological information in these readings, we need to first understand the subgroups of homology that give meaning to this concept.

2.3.1. Well groups

Suppose that we have a continuous mapping $f\colon \mathbb{X} \to \mathbb{Y}$ between topological spaces. Given a subset $\mathbb{A} \subseteq \mathbb{Y}$, we review here the definition of the well groups $\mathsf{U}_{\mathbb{A}}(r)$ for each radius $r \geqslant 0$. When $\mathbb{A}$ is clear from context, we will drop it from the notation and simply write $\mathsf{U}(r)$, by which we mean the direct sum of groups $\mathsf{U}_p(r)$, for each homological dimension p. We will also need the assumption that $f^{-1}(\mathbb{A})$ has homology groups of finite rank in each dimension. In addition to the mapping f, we assume a subspace $\mathcal{P}$ of $C(\mathbb{X}, \mathbb{Y})$, the space of continuous mappings from $\mathbb{X}$ to $\mathbb{Y}$, requiring that $\mathcal{P}$ contains f. For example, $\mathcal{P}$ might consist of all mappings homotopic to f. We assume a metric on $\mathcal{P}$ and write $\|f - h\|_{\mathcal{P}}$ for the distance between two mappings. We call h an *r-perturbation* of f if $\|f - h\|_{\mathcal{P}} \leqslant r$. Given $\mathbb{A} \subseteq \mathbb{Y}$, we introduce the *radius function*, $f_{\mathbb{A}}\colon \mathbb{X} \to \mathbb{R}$, by setting $f_{\mathbb{A}}(x)$ to the infimum value of r for which there exists an r-perturbation $h \in \mathcal{P}$ with $h(x) \in \mathbb{A}$. We filter $\mathbb{X}$ via the sublevel sets of the radius function, setting $\mathbb{X}_r(f_{\mathbb{A}}) = f_{\mathbb{A}}^{-1}[0, r]$. For $r < s$, there is a map $\mathsf{f}_{\mathbb{A}}^{r,s}\colon \mathsf{H}(\mathbb{X}_r(f_{\mathbb{A}})) \to \mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}}))$. The preimage of $\mathbb{A}$ under any r-perturbation h of f will obviously be a subset of $\mathbb{X}_r(f_{\mathbb{A}})$, and hence there is a map on homology, $\mathsf{j}_h\colon \mathsf{H}(h^{-1}(\mathbb{A})) \to \mathsf{H}(\mathbb{X}_r(f_{\mathbb{A}}))$. Given a class $\alpha \in \mathsf{H}(\mathbb{X}_r(f_{\mathbb{A}}))$ and an r-perturbation h of f, we say that α is *supported* by h if $\alpha \in \operatorname{im} \mathsf{j}_h$. The *well group* $\mathsf{U}(r) \subseteq \mathsf{H}(\mathbb{X}_r(f_{\mathbb{A}}))$ is then defined [**10**] to consist of the classes that are supported by all r-perturbations of f:

$$\mathsf{U}(r) = \bigcap_{\|h-f\|_{\mathcal{P}} \leqslant r} \operatorname{im} \mathsf{j}_h.$$

For $r < s$, the map $\mathsf{f}_{\mathbb{A}}^{r,s}$ restricts to a map $\mathsf{U}(r) \to \mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}}))$. On the other hand, $\mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}}))$ contains $\mathsf{U}(s)$ as a subgroup. It can be shown that $\mathsf{U}(s) \subseteq \mathsf{f}_{\mathbb{A}}^{r,s}(\mathsf{U}(r))$ whenever $r < s$; see [**10**]. In other words, the rank of the well group can only decrease as the threshold value increases. We call a value of r at which the rank of the well group decreases a *terminal critical value* of $f_{\mathbb{A}}$. The *well diagram* of f and $\mathbb{A}$ is the multiset of terminal critical values of $f_{\mathbb{A}}$, taking a value k times if the rank of the well group drops by k at the value. Often we will refer to this diagram as the *robustness* of the preimage $f^{-1}(\mathbb{A})$. In this paper, we focus on the case $\mathbb{Y} = \mathbb{R}$ and $\mathcal{P} = C(\mathbb{X}, \mathbb{R})$, lifting the usual metric on $\mathbb{R}$ to $\mathcal{P}$ by defining $\|f - h\|_{\mathcal{P}} = \|f - h\|_{\infty} = \sup_{x \in \mathbb{X}} |f(x) - h(x)|$. In this case, the radius function satisfies $f_{\mathbb{A}}(x) = \inf_{a \in \mathbb{A}} |f(x) - a|$. In general, the relationship between the terminal critical values and the homological critical values of $f_{\mathbb{A}}$ is not completely understood. However, if $\mathbb{Y} = \mathbb{R}$ and $\mathbb{A}$ is a point, we will see that the former is a subset of the latter. We get more complicated relationships when $\mathbb{A}$ is an interval.

2.3.2. Example

Consider the torus $\mathbb{X}$, as shown in Figure 3, along with the vertical height function $f\colon \mathbb{X} \to \mathbb{R}$ and the space $\mathbb{A} = \{a\}$. The preimage of $\mathbb{A}$, $f^{-1}(\mathbb{A}) = f_{\mathbb{A}}^{-1}(0)$, consists of

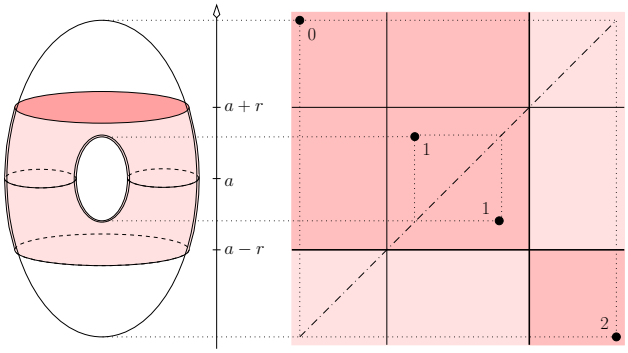

Figure 3: Left: the torus and the preimage of the interval $[a-r, a+r]$. Right: the extended persistence diagram of the vertical height function. Each point is labeled by the dimension of the corresponding homology class. The dark shaded portions of the diagram represent the homology of $f^{-1}[a-r, a+r]$.

two disjoint circles on the torus; hence there are two components and two independent 1-cycles, all belonging to the well group at radius 0. For small values of r, $\mathbb{X}_r(f_{\mathbb{A}})$ consists of two disjoint cylinders. The homology has yet to change; furthermore, although the proof will come later, all classes still belong to the well groups at these small radii.

Now consider the value of r shown in Figure 3. For this r, the sublevel set $\mathbb{X}_r = \mathbb{X}_r(f_{\mathbb{A}})$ consists of two pair-of-pants glued together along two common circles. We note that $\mathsf{H}_0(\mathbb{X}_r)$ has dropped in rank by one, while the rank of $\mathsf{H}_1(\mathbb{X}_r)$ has grown to three. In contrast, the rank of $\mathsf{U}_1(r)$ is less than or equal to one. Indeed, the function $h\colon \mathbb{X} \to \mathbb{R}$, defined by $h = f - r$, is an r-perturbation of f and the zero set of the corresponding distance function, $h_{\mathbb{A}}^{-1}(0) = f^{-1}(a+r)$, is a single closed curve. Since the rank of the first homology group of that curve is one, and since the rank of $\operatorname{im} \mathsf{j}_h$ can be no bigger than this rank, the well group $\mathsf{U}_1(r)$ can also have rank at most one. That it does in fact have rank exactly one will follow from our results in Section 4.

2.3.3. Relative well groups

Since the pyramid involves relative homology groups, it seems wise to extend the definition of well groups into the context of relative homology. While this notion is new, it follows the above ideas closely so that presenting the definition in this background section seems appropriate. Assume again that we have a continuous mapping $f\colon \mathbb{X} \to \mathbb{Y}$ between topological spaces, as well as a subspace $\mathcal{P}$ of $C(\mathbb{X}, \mathbb{Y})$ that contains f and is equipped with a metric. Given a nested pair $\mathbb{A}' \subseteq \mathbb{A}$ of subspaces of $\mathbb{Y}$, and a radius $r \geqslant 0$, we note that $\mathbb{X}'_r = \mathbb{X}_r(f_{\mathbb{A}'})$ is a subset of $\mathbb{X}_r = \mathbb{X}_r(f_{\mathbb{A}})$. For each r-perturbation h of f, there is an inclusion of pairs $(h^{-1}(\mathbb{A}), h^{-1}(\mathbb{A}')) \hookrightarrow (\mathbb{X}_r, \mathbb{X}'_r)$, which induces a map $j_h\colon \mathsf{H}(h^{-1}(\mathbb{A}), h^{-1}(\mathbb{A}')) \to \mathsf{H}(\mathbb{X}_r, \mathbb{X}'_r)$ between relative homology groups. The *relative well group* $\mathsf{U}_{(\mathbb{A},\mathbb{A}')}(r)$ is defined to be the intersection of the images of these maps, taken over all r-perturbations of f:

$$\mathsf{U}_{(\mathbb{A},\mathbb{A}')}(r) = \bigcap_{\|h-f\|_{\mathcal{P}} \leqslant r} \operatorname{im} j_h.$$

When a distinction is needed, we will refer to the previous notion of well groups as *absolute* well groups.

3. Combinatorics of Homology

In this section, we present the first half of our point calculus, showing how to read the homology of a level or interlevel set from the extended persistence diagram. The crucial technical concept is that of a basis of the pyramid of zigzag modules, which we establish by strengthening the Pyramid Theorem in [**6**].

3.0.4. Flipping a basis

We construct a basis for the pyramid one step at a time, by flipping the basis of one zigzag module to the next. For this purpose, we consider two zigzag modules that differ at one position, and we assume that there is a Mayer-Vietoris diamond serving as a connecting bridge between the two modules at that position. Drawing the diamond with the intersection at the bottom and the union at the top, as in (5), we say the diamond connects the *lower* module with the *upper* module. Given a basis of the lower module, we can show that we can construct a basis of the upper module so that the two bases agree on the overlap. We refer to this operation as *flipping* the first basis to the second.

Lemma 1. *Given two zigzag modules that differ by a single Mayer-Vietoris diamond, we can flip any basis of the lower module to a basis of the upper module.*

Proof. We give a proof by construction. Writing $\{e^i_k\}$ for the basis of the lower zigzag module, we describe a basis $\{v^i_k\}$ of the upper zigzag module that differs from the lower one only at the position j at which the modules differ; as in (5). We thus at once set $v^i_k = e^i_k$ for all $k \neq j$, and the main task is then the construction of the v^i_j. Put briefly, our rule will be that $v^i_j \neq 0$ iff an odd number of e^i_{j-1}, e^i_j, e^i_{j+1} are non-zero. We give more specifics via a case analysis. The cases are labeled pictorially, with black dots denoting non-zero classes, showing only the positions $j-1, j, j+1$.

CASE 1 (→): We have $e^i_{j-1} \neq 0$ and $e^i_j = e^i_{j+1} = 0$, and define v^i_j as well as the advancing map using the Mayer-Vietoris diamond, namely $v^i_j = a_{j-1}(e^i_{j-1})$, which is non-zero by exactness and because $e^i_j = 0$.

CASE 2 (→): Again we set $v^i_j = a_{j-1}(e^i_{j-1})$, which is zero by exactness and because $e^i_j \neq 0$.

CASE 3 (→): We set $v^i_j = a_{j-1}(e^i_{j-1}) = b_j(e^i_{j+1})$, which in this case is non-zero. Indeed, if it were zero, then, by exactness, the pair $(e^i_{j-1}, 0)$ would be in the image of the map $b_{j-1} \oplus a_j$. By the direct-sum decomposition of the maps in the basis, this would imply that $a_j(e_j) = 0$, a contradiction.

CASE 4 (→): We have $e^i_j \neq 0$ and $e^i_{j-1} = e^i_{j+1} = 0$. If there are $\ell \geqslant 0$ indices i of this kind, then the orthogonal complement to the image of the map c_j, defined below, has rank ℓ, as we prove shortly. We pick ℓ classes v^i_j that span this complement. Since v^i_j maps to e^i_j via the connecting homomorphism of the Mayer-Vietoris sequence, the homological dimension of v^i_j is one higher than that of e^i_j.

CASE 5 ($\vee \to \wedge$): This is symmetric to Case 2, and we set $v_j^i = b_j(e_{j+1}^i) = 0$.
CASE 6 ($\vee \to \wedge$): This is symmetric to Case 1, and we set $v_j^i = b_j(e_{j+1}^i) \neq 0$.

Note first that we now have interval modules $\{e_j^i\}$ in the lower zigzag module, and interval modules $\{v_j^i\}$ in the upper zigzag module. To show that the latter are indeed summands, we only need to verify that the non-zero classes v_j^i form a basis of $\mathsf{H}(\mathbb{V}, \mathbb{V}')$, the new group in the upper zigzag module. Using the notation in (5), we let E denote the vector space spanned by the pairs (e_{j-1}^i, e_{j+1}^i), noting that E is a subspace of $\mathsf{H}(\mathbb{C}, \mathbb{C}') \oplus \mathsf{H}(\mathbb{D}, \mathbb{D}')$, but because of Case 3 it is not necessarily the entire direct sum. We consider the subspaces E_N of E spanned by the pairs (e_{j-1}^i, e_{j+1}^i) in each Case N, for $1 \leqslant N \leqslant 6$. These subspaces are independent and span the entire space E. In other words, zero is the only element common to any two of the subspaces, and the ranks of the subspaces add up to the rank of E.

The case analysis suggests a map $c_j \colon \mathsf{E} \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$ with $c_j((e_{j-1}^i, e_{j+1}^i)) = v_j^i$, if $(e_{j-1}^i, e_{j+1}^i) \neq (0, 0)$, and zero otherwise. Since $\mathsf{E}_4 = 0$, this map is zero on E_4, but it is also zero on E_2 and E_5. Furthermore, c_j is injective when restricted to E_1, E_3, and E_6. We proceed to show that the images of these latter three vector spaces under c_j are independent of one another. To derive a contradiction, we first suppose that $c_j(\mathsf{E}_1) \cap c_j(\mathsf{E}_6)$ contains a non-zero class. Then there must exist $(\alpha, 0) \in \mathsf{E}_1$ and $(0, \beta) \in \mathsf{E}_6$ with $a_{j-1}(\alpha) = b_j(\beta) \neq 0$. Hence, $(\alpha, \beta) \in \ker(a_{j-1} \oplus b_j)$, which, by exactness, tells us that $\alpha \in \operatorname{im} b_{j-1}$. But this contradicts the direct-sum decomposition of the map b_{j-1}. Next, suppose that $c_j(\mathsf{E}_1) \cap c_j(\mathsf{E}_3)$ contains a non-zero class, which means there exists $(\alpha, 0) \in \mathsf{E}_1$ and $(\gamma, \beta) \in \mathsf{E}_3$ such that $a_{j-1}(\alpha) = b_j(\beta) \neq 0$. As above, this implies that $(\alpha, \beta) \in \ker(a_{j-1} \oplus b_j)$, and we reach the same contradiction. Finally, a symmetric argument gives $c_j(\mathsf{E}_3) \cap c_j(\mathsf{E}_6) = 0$. We conclude that $c_j(\mathsf{E}_1)$, $c_j(\mathsf{E}_3)$, and $c_j(\mathsf{E}_6)$ are independent subspaces of $\mathsf{H}(\mathbb{V}, \mathbb{V}')$. In Case 4, we picked a basis for the orthogonal complement to their span; all together, we have a basis of $\mathsf{H}(\mathbb{V}, \mathbb{V}')$, as required. □

3.0.5. Establishing a basis

The Pyramid Theorem in [**6**] establishes an explicit bijection between the interval modules that arise in the decomposition of any two zigzags within the pyramid. We strengthen this result by establishing bases on all the zigzag modules in such a way that the basis elements correspond to the intervals and respect the same bijections. We call this a *basis* of the pyramid. To construct it, we note that the paths in the pyramid are connected by Mayer-Vietoris diamonds. We can therefore flip a basis of the level set zigzag upwards through the entire pyramid via repeated application of Lemma 1.

Theorem 1. *A basis of the level set zigzag module extends to a basis of the entire pyramid.*

We now give an explicit description of how the interval modules of the various paths in the pyramid relate to each other. A convenient reference in this description is the extended filtration (2), which follows the upward slope through the middle of the pyramid. Its first half is parameterized from $-\infty$ to ∞, and its second half from ∞ back to $-\infty$. Let now x and y be two points along the upward slope, with x to the left of y. We distinguish between the ordinary case ($x < y$, both in the first half),

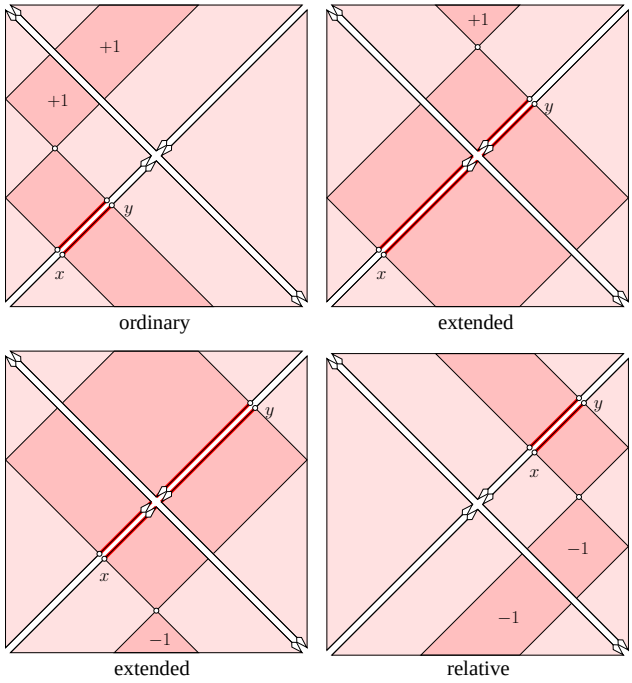

Figure 4: The basis element that corresponds to the interval from x to y along the upward slope maps to all spaces between the paths of its two endpoints. The four squares show the pattern for the four different types of intervals.

the relative case ($y < x$, both in the second half), and the two extended cases ($x < y$ and $y < x$, with x in the first half and y in the second half). For each case, we sketch how the basis element of the interval corresponds to basis elements of other homology groups in Figure 4. As a general pattern, the two points trace out two curves consisting of segments with slopes $\pm 45°$ that reflect before they hit the vertical sides and end at the horizontal sides of the square. The reason for the slopes are Cases 1, 2, 5, and 6 in the proof of the Lemma 1, and the reason for the reflection is the local change in the zigzag structure caused by moving the terminal zero group up. The two curves cross at one point inside the square, and the location of that point is characteristic for the case (the triangular region on the left in the ordinary case, at the top and at the bottom in the two extended cases, and on the right in the relative case). The crossing is caused by Case 4, in which the correspondence between the basis elements is constructed via the connecting homomorphism of the Mayer-Vietoris sequence and therefore comes with a shift in homological dimension.

3.0.6. Turning the table

The regions in Figure 4 show all the spaces represented by points in the pyramid to which the basis element corresponding to the interval $[x, y]$ is relevant. We are now interested in the inverse question: which basis elements are relevant to a given space? More specifically: which intervals in the decomposition of the extended filtration (2)

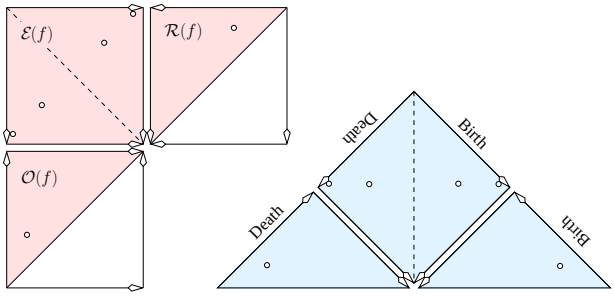

Figure 5: The three overlaid subdiagrams in the standard extended persistence diagram are unfolded by flipping pages: keeping $\mathcal{O}(f)$ fixed, $\mathcal{E}(f)$ flips up, followed by $\mathcal{R}(f)$ which flips up and then to the right. Finally, we clip the ordinary and relative subdiagrams along the diagonal and rotate the entire design by 45 degrees so it rests on its long side. The arrows of the diagram go from negative to positive infinity.

map to the basis of the homology group of the space represented by a point with coordinates a and b? We answer this question by considering the following subregions of the p-dimensional persistence diagram:

$$\begin{aligned}
\lambda_p[a,b] &= \{(x,y) \in \mathcal{O}_p(f) \mid x < b < y\} \sqcup \{(x,y) \in \mathcal{E}_p(f) \mid x < b, a < y\},\\
\varrho_p[a,b] &= \{(x,y) \in \mathcal{E}_p(f) \mid b < x, y < a\} \sqcup \{(x,y) \in \mathcal{R}_p(f) \mid y < a < x\},\\
\lambda_p[a,b) &= \{(x,y) \in \mathcal{E}_p(f) \mid a < y < b\} \sqcup \{(x,y) \in \mathcal{R}_p(f) \mid a < y < b < x\},\\
\varrho_p[a,b) &= \{(x,y) \in \mathcal{R}_p(f) \mid y < a < x < b\},\\
\lambda_p(a,b] &= \{(x,y) \in \mathcal{O}_p(f) \mid x < a < y < b\},\\
\varrho_p(a,b] &= \{(x,y) \in \mathcal{O}_p(f) \mid a < x < b < y\} \sqcup \{(x,y) \in \mathcal{E}_p(f) \mid a < x < b\},\\
\lambda_p(a,b) &= \{(x,y) \in \mathcal{O}_p(f) \mid x < a < y\} \sqcup \{(x,y) \in \mathcal{E}_p(f) \mid x < a, b < y\},\\
\varrho_p(a,b) &= \{(x,y) \in \mathcal{E}_p(f) \mid a < x, y < b\} \sqcup \{(x,y) \in \mathcal{R}_p(f) \mid y < b < x\},
\end{aligned}$$

where we assume that a and b are both homological regular values. To display these multisets, we first introduce a new, and for our purposes more convenient, way of drawing the extended persistence diagram. See Figure 5 for a guide. We glue the domains of the three sub-diagrams and draw the result as a right-angled triangle. In this triangle, the birth and death axes go from $-\infty$ up to $+\infty$ and then continue on back to $-\infty$. In other words, we flip the extended subdiagram upside down and glue its (formerly) upper side to the upper side of the ordinary subdiagram. Similarly, we rotate the relative subdiagram by 180 degrees and glue its (formerly) right side to the right side of the (flipped) extended subdiagram. After gluing the three domains, we rotate the design by -45 degrees so the triangle rests on its longest side, consisting of the diagonals in the ordinary and relative subdiagrams. The diagonal of the extended subdiagram is now the vertical symmetry axis passing through the middle of the triangle. These changes having been made, the multisets referenced above are displayed in Figure 6.

Remark 2. There is a straightforward translation of this triangular design to the representation of persistence advocated in [**4**]. Namely, draw a isosceles right-angled

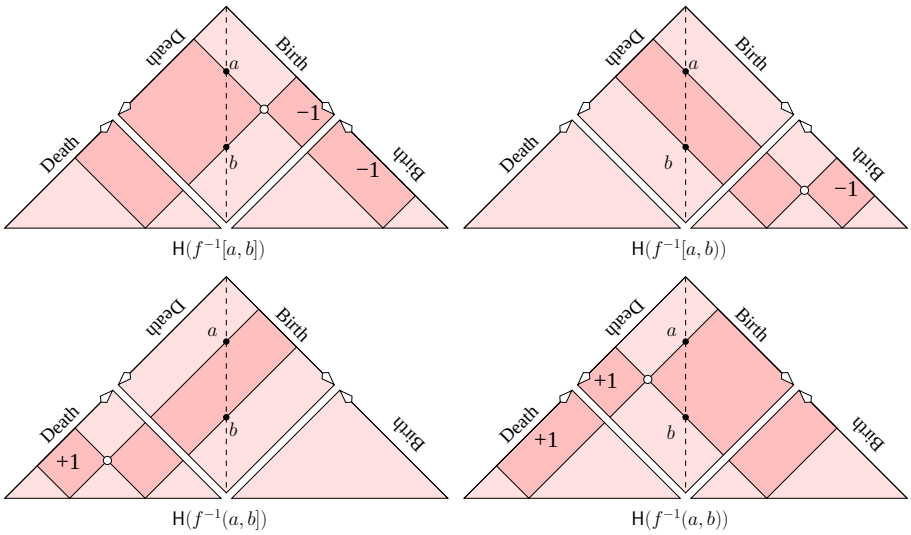

Figure 6: The triangle design of the persistence diagram showing the regions λ and ϱ for the four types of intervals in darker shading. When we collect the points to compute the rank of the p-th homology group, we shift the homological dimension of classes as shown.

triangle downward from each point in the multiset and call the horizontal lower edge the corresponding *bar*. The *barcode* is the multiset of bars, one for each point in the diagram. Similarly, we can translate the triangular design into the square design of the pyramid by cutting along the vertical axis, turning the right triangle upside-down, and gluing the two triangles along their hypotenuses.

3.0.7. Reading interlevel sets

The purpose of the multisets defined above is to offer a convenient way to read the absolute or relative homology of an interlevel set from the extended persistence diagram. We need some definitions to combine all four types into one. First, we let $\mathcal{B}$ be the collection of interval modules in the decomposition of the extended filtration (2). As mentioned earlier, this collection is in bijective correspondence with the points in $\mathcal{D}(f)$. We write $\mathsf{V} = \langle \mathcal{B} \rangle$ for the abstract vector space spanned by $\mathcal{B}$, and we let $\mathcal{V} = \{\langle \mathcal{B}' \rangle \mid \mathcal{B}' \subseteq \mathcal{B}\}$ be the collection of vector spaces spanned by subsets of this basis. Second, we write

$$\mathcal{W}_p(I) = \begin{cases} \lambda_p[a,b] \sqcup \varrho_{p+1}[a,b] & \text{if } I = [a,b], \\ \lambda_p[a,b) \sqcup \varrho_{p+1}[a,b) & \text{if } I = [a,b), \\ \lambda_{p-1}(a,b] \sqcup \varrho_p(a,b] & \text{if } I = (a,b], \\ \lambda_{p-1}(a,b) \sqcup \varrho_p(a,b) & \text{if } I = (a,b), \end{cases}$$

for the region of points in the persistence diagram that correspond to the basis elements of $\mathsf{H}_p(f^{-1}(I))$, and call it a *pair of wings*. With these concepts, we have the following result, which implies that the rank of $\mathsf{H}_p(f^{-1}(I))$ is the number of points in $\mathcal{W}_p(I)$:

Theorem 2. *For each dimension p and each interval I whose endpoints are homological regular values, there exists an isomorphism that takes $\mathsf{H}_p(f^{-1}(I))$ to the vector space $\mathsf{G}_p(I) \in \mathcal{V}$ spanned by the basis vectors corresponding to the points in $\mathcal{W}_p(I)$.*

Proof. Write $\mathcal{B} = \{e^i\}$ and let $\{v^i\}$ be the basis of the group $\mathsf{H}_p(f^{-1}(I))$, where I is an interval with endpoints $a \leqslant b$ that can be closed, closed-open, open-closed, or open. The claimed isomorphism is then the linear map $\gamma\colon \mathsf{H}_p(f^{-1}(I)) \to \mathsf{V}$ defined by $\gamma(v^i) = \{e^i\}$ for all non-zero v^i.

To understand why the image of γ consists of the intervals that correspond to the points in $\mathcal{W}_p(I)$, we need to recall the transformation rules sketched in Figure 4. Consider for example the closed interval case, $I = [a, b]$, for which $\mathcal{W}_p(I) = \lambda_p[a, b] \sqcup \varrho_{p+1}[a, b]$. Since the interval is closed, the homology group is represented by the point (a, b) in the lower triangular region. To lie in the dark shaded region, this point must satisfy the constraint $x < b < y$ in the ordinary case, $x < b$ and $a < y$ in the first extended case, and $x < b$ and $a < y$ without dimension shift in the second extended case. These inequalities define $\lambda_p[a, b]$. Furthermore, we get $b < x$ and $y < a$ with dimension shift in the second extended case, and $y < a < x$, again with dimension shift, in the relative case. These inequalities define $\varrho_{p+1}[a, b]$, which completes the proof in the closed case. For a proof of the closed-open, open-closed, and open cases, note that the points representing $\mathsf{H}_p(f^{-1}(I))$ are found in the right, left, and top triangular region of the pyramid, and then argue in a similar fashion. □

4. Combinatorics of Robustness

The definition of well group given in Section 2 involves an uncountable number of perturbations, which give rise to the intersection of a potentially large number of subgroups, and as such does not seem amenable to computation. In this section, we show that the situation in the real-valued case is simpler, and that we are able to read the absolute and relative well groups directly from the extended persistence diagram. We begin with a consequence of the Mayer-Vietoris sequence, which provides the main technical ingredient of our proofs.

4.0.8. A corollary of Mayer-Vietoris

For convenience, we establish the following notational convention, wherein we reuse the same letter in different fonts. We will need it for absolute and for relative homology groups. To avoid repetition, we state it now for the more general relative case. Letting $\mathbb{U}' \subseteq \mathbb{U}$ and $\mathbb{V}' \subseteq \mathbb{V}$ be pairs of topological spaces, we write $(\mathbb{U}, \mathbb{U}') \hookrightarrow (\mathbb{V}, \mathbb{V}')$ if $\mathbb{U} \subseteq \mathbb{V}$ and $\mathbb{U}' \subseteq \mathbb{V}'$. This inclusion of pairs induces a map $\mathsf{u}\colon \mathsf{H}(\mathbb{U}, \mathbb{U}') \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$ on homology groups, and we write $\mathsf{U} = \operatorname{im} \mathsf{u}$ for the image of this map. Note that U is always a subgroup of $\mathsf{H}(\mathbb{V}, \mathbb{V}')$, namely the subgroup of homology classes that have a chain representative carried by $(\mathbb{U}, \mathbb{U}')$. Note also that the rank of U can never exceed the rank of $\mathsf{H}(\mathbb{V}, \mathbb{V}')$. Suppose that, furthermore, $(\mathbb{T}, \mathbb{T}') \hookrightarrow (\mathbb{U}, \mathbb{U}')$. Then, from the sequence of maps $\mathsf{H}(\mathbb{T}, \mathbb{T}') \to \mathsf{H}(\mathbb{U}, \mathbb{U}') \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$, we see that T, the image of $\mathsf{H}(\mathbb{T}, \mathbb{T}')$ in $\mathsf{H}(\mathbb{V}, \mathbb{V}')$, must be a subgroup of U. The following lemma is a direct consequence of the exactness of the Mayer-Vietoris sequence. However, we will use it often enough that it seems reasonable to state and prove it formally.

Lemma 2. *Suppose the pair of topological spaces $\mathbb{V}' \subseteq \mathbb{V}$ can be decomposed as $\mathbb{V} = \mathbb{C} \cup \mathbb{D}$ and $\mathbb{V}' = \mathbb{C}' \cup \mathbb{D}'$, where $\mathbb{C}' \subseteq \mathbb{C}$ and $\mathbb{D}' \subseteq \mathbb{D}$. Set $(\mathbb{E}, \mathbb{E}') = (\mathbb{C} \cap \mathbb{D}, \mathbb{C}' \cap \mathbb{D}')$. If a class $\alpha \in \mathsf{H}(\mathbb{V}, \mathbb{V}')$ belongs to C and to D, then α also belongs to E.*

Proof. Following our convention, we use the notation $\mathsf{c}\colon \mathsf{H}(\mathbb{C}, \mathbb{C}') \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$ for the map on homology induced by the inclusion of $(\mathbb{C}, \mathbb{C}')$ in $(\mathbb{V}, \mathbb{V}')$. Similarly, we write $\mathsf{d}\colon \mathsf{H}(\mathbb{D}, \mathbb{D}') \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$ and $\mathsf{e}\colon \mathsf{H}(\mathbb{E}, \mathbb{E}') \to \mathsf{H}(\mathbb{V}, \mathbb{V}')$, as well as $\mathsf{e}_c\colon \mathsf{H}(\mathbb{E}, \mathbb{E}') \to \mathsf{H}(\mathbb{C}, \mathbb{C}')$ and $\mathsf{e}_d\colon \mathsf{H}(\mathbb{E}, \mathbb{E}') \to \mathsf{H}(\mathbb{D}, \mathbb{D}')$. Note that $\mathsf{C} = \operatorname{im}\mathsf{c}$, $\mathsf{D} = \operatorname{im}\mathsf{d}$, and $\mathsf{E} = \operatorname{im}\mathsf{e}$. Consider now the relevant portion of the Mayer-Vietoris sequence for $(\mathbb{V}, \mathbb{V}')$:

$$\mathsf{H}(\mathbb{E}, \mathbb{E}') \xrightarrow{(\mathsf{e}_c, \mathsf{e}_d)} \mathsf{H}(\mathbb{C}, \mathbb{C}') \oplus \mathsf{H}(\mathbb{D}, \mathbb{D}') \xrightarrow{\mathsf{c}-\mathsf{d}} \mathsf{H}(\mathbb{V}, \mathbb{V}').$$

By assumption, $\alpha \in \mathsf{C}$, so there exists some $\alpha_c \in \mathsf{H}(\mathbb{C}, \mathbb{C}')$ such that $\mathsf{c}(\alpha_c) = \alpha$. Similarly, there exists an $\alpha_d \in \mathsf{H}(\mathbb{D}, \mathbb{D}')$ such that $\mathsf{d}(\alpha_d) = \alpha$. This implies that the pair (α_c, α_d) belongs to the kernel of $\mathsf{c} - \mathsf{d}$, and thus also, by exactness of the sequence, belongs to the image of $(\mathsf{e}_c, \mathsf{e}_d)$. Hence, there exists $\alpha_e \in \mathsf{H}(\mathbb{E}, \mathbb{E}')$ with $\mathsf{e}_c(\alpha_e) = \alpha_c$ and $\mathsf{e}_d(\alpha_e) = \alpha_d$. In particular, since $\mathsf{e} = \mathsf{c} \circ \mathsf{e}_c$, we have $\mathsf{e}(\alpha_e) = \alpha$, and therefore $\alpha \in \mathsf{E}$ as claimed. □

In the typical application of Lemma 2, we will construct further pairs $(\mathbb{T}, \mathbb{T}') \hookrightarrow (\mathbb{C}, \mathbb{C}')$ and $(\mathbb{B}, \mathbb{B}') \hookrightarrow (\mathbb{D}, \mathbb{D}')$ such that $\alpha \in \mathsf{T} \cap \mathsf{B}$. From the remark above, we know that $\mathsf{T} \subseteq \mathsf{C}$ and $\mathsf{B} \subseteq \mathsf{D}$. The lemma then applies and we can conclude that $\alpha \in \mathsf{E}$, as before.

4.0.9. The well group of a level set

As a warm-up exercise, we first consider the case in which $\mathbb{A}$ is a single point. More specifically, we suppose that we have a compact topological space $\mathbb{X}$ and a function $f\colon \mathbb{X} \to \mathbb{R}$, and we find the well groups $\mathsf{U}(r) = \mathsf{U}_{\mathbb{A}}(r)$, where $\mathbb{A} = \{a\}$ is some point on the real line. In this case, $\mathbb{X}_r(f_{\mathbb{A}}) = f_{\mathbb{A}}^{-1}[0, r] = f^{-1}[a - r, a + r]$. To state the formula, we distinguish two particular subspaces of $\mathbb{X}_r = \mathbb{X}_r(f_{\mathbb{A}})$, namely the *top level set*, $\mathbb{T}_r = f^{-1}(a + r)$, and the *bottom level set*, $\mathbb{B}_r = f^{-1}(a - r)$. Using the convention from before, we write T_r and B_r for the images of $\mathsf{H}(\mathbb{T}_r)$ and $\mathsf{H}(\mathbb{B}_r)$ in $\mathsf{H}(\mathbb{X}_r)$.

Theorem 3. $\mathsf{U}(r) = \mathsf{T}_r \cap \mathsf{B}_r$, *for every* $r \geqslant 0$.

Proof. We prove equality by establishing the two inclusions in turn. To show $\mathsf{U}(r) \subseteq \mathsf{T}_r \cap \mathsf{B}_r$, consider an arbitrary class $\alpha \in \mathsf{U}(r)$. We define $h_{\text{top}} = f - r$ and $h_{\text{bot}} = f + r$ and note that they are r-perturbations of f, with $h_{\text{top}}^{-1}(a) = \mathbb{T}_r$ and $h_{\text{bot}}^{-1}(a) = \mathbb{B}_r$. By definition of the well group, α is supported by every r-perturbation of f, and therefore by h_{top} and by h_{bot}. It follows that $\alpha \in \mathsf{T}_r \cap \mathsf{B}_r$. To show $\mathsf{T}_r \cap \mathsf{B}_r \subseteq \mathsf{U}(r)$, we consider an arbitrary class $\alpha \in \mathsf{T}_r \cap \mathsf{B}_r$ and let h be an arbitrary r-perturbation of f. To finish the proof, we need to show that α is supported by h. We define $\mathbb{C}_r = h^{-1}[a, \infty) \cap \mathbb{X}_r$ and $\mathbb{D}_r = h^{-1}(-\infty, a] \cap \mathbb{X}_r$. Note that $\mathbb{C}_r \cup \mathbb{D}_r = \mathbb{X}_r$ while $\mathbb{C}_r \cap \mathbb{D}_r = h^{-1}(a)$. Furthermore, the inequality $\|h - f\|_\infty \leqslant r$ implies that $\mathbb{T}_r \subseteq \mathbb{C}_r$ and $\mathbb{B}_r \subseteq \mathbb{D}_r$. By Lemma 2, α is supported by $h^{-1}(a)$, as required. □

Remark 3. Theorem 3 implies that the well group for a Morse function f can change only at critical values of the function $f_{\mathbb{A}}$, where $\mathbb{A} = \{a\}$. In other words, terminal critical values are, in this simple context, just ordinary critical values. Indeed, if

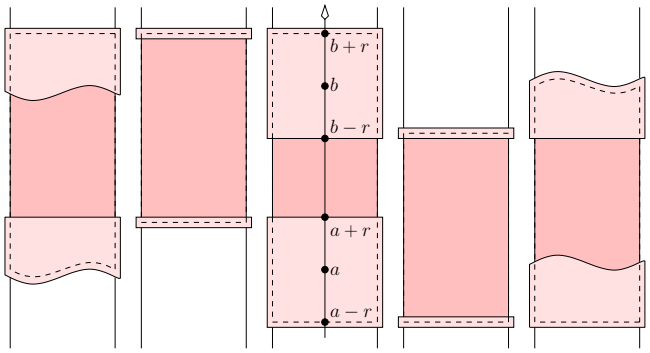

Figure 7: Each vertical strip represents $\mathbb{X}$, and the shaded portions mark $(\mathbb{C}_r, \mathbb{C}'_r)$ and $(\mathbb{T}_r, \mathbb{T}'_r)$ on the left, $(\mathbb{X}_r, \mathbb{X}'_r)$ in the middle, and $(\mathbb{B}_r, \mathbb{B}'_r)$ and $(\mathbb{D}_r, \mathbb{D}'_r)$ on the right.

$[r, s]$ is an interval that contains no critical values of $f_{\mathbb{A}}$, then there is a deformation retraction $\mathbb{X}_s(f_{\mathbb{A}}) \to \mathbb{X}_r(f_{\mathbb{A}})$ providing an isomorphism $\mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}})) \to \mathsf{H}(\mathbb{X}_r(f_{\mathbb{A}}))$. Furthermore, this retraction maps $\mathbb{T}_s$ onto $\mathbb{T}_r$, in such a way that that the images of $\mathsf{H}(\mathbb{T}_r)$ and $\mathsf{H}(\mathbb{T}_s)$ in $\mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}}))$ are identical. Similarly, the images of $\mathsf{H}(\mathbb{B}_r)$ and $\mathsf{H}(\mathbb{B}_s)$ in $\mathsf{H}(\mathbb{X}_s(f_{\mathbb{A}}))$ are identical. Hence the well groups $\mathsf{U}(r)$ and $\mathsf{U}(s)$ are isomorphic.

4.0.10. The well group of an interlevel set
We generalize from a point to an interval, which can be closed, closed-open, open-closed, or open. To that end, we define the spaces and maps so that the formula for the well group is the same in all four cases, and indeed the same as in Theorem 3 above. Assume $a < b$, set $\mathbb{A} = [a, b]$, and let $\mathbb{A}' \subseteq \{a, b\}$. We thus get $\mathbb{X}_r = \mathbb{X}_r(f_{\mathbb{A}}) = f^{-1}[a - r, b + r]$ and $\mathbb{X}'_r = \mathbb{X}_r(f_{\mathbb{A}'})$, which is the empty set, $f^{-1}[b - r, b + r]$, $f^{-1}[a - r, a + r]$, or the union of these two interlevel sets. Correspondingly, we define the *top* and *bottom interlevel sets*:

$$\begin{aligned} \mathbb{T}_r &= f^{-1}[a + r, b + r], \quad \mathbb{T}'_r \subseteq \{f^{-1}(a + r), f^{-1}(b + r)\}, \\ \mathbb{B}_r &= f^{-1}[a - r, b - r], \quad \mathbb{B}'_r \subseteq \{f^{-1}(a - r), f^{-1}(b - r)\}; \end{aligned}$$

see Figure 7. The pairs $(\mathbb{T}_r, \mathbb{T}'_r)$ and $(\mathbb{B}_r, \mathbb{B}'_r)$ include into $(\mathbb{X}_r, \mathbb{X}'_r)$ in all four cases. Still using the notational convention from above, we write T_r and B_r for the images of $\mathsf{H}(\mathbb{T}_r, \mathbb{T}'_r)$ and $\mathsf{H}(\mathbb{B}_r, \mathbb{B}'_r)$ in $\mathsf{H}(\mathbb{X}_r, \mathbb{X}'_r)$. The formula for the well group, $\mathsf{U}(r) = \mathsf{U}_{(\mathbb{A},\mathbb{A}')}(r)$, is then, unsurprisingly:

Theorem 4. *$\mathsf{U}(r) = \mathsf{T}_r \cap \mathsf{B}_r$, for every $r \geqslant 0$.*

Proof. We give the argument for the most complicated of the four cases, when $\mathbb{A}' = \{a, b\}$. The proofs of the other three cases are simpler versions of the same argument. We may assume $a + r < b - r$, else $\mathbb{X}_r = \mathbb{X}'_r$, which implies that all groups in the claimed formula are zero and so we are done. To prove the inclusion $\mathsf{U}(r) \subseteq \mathsf{T}_r \cap \mathsf{B}_r$, we consider the two r-perturbations $h_{\text{top}} = f - r$ and $h_{\text{bot}} = f + r$, as before. Note that $(\mathbb{T}_r, \mathbb{T}'_r) = h^{-1}_{\text{top}}(a, b)$ and $(\mathbb{B}_r, \mathbb{B}'_r) = h^{-1}_{\text{bot}}(a, b)$, and the desired inclusion follows from the definition of relative well groups. To prove $\mathsf{T}_r \cap \mathsf{B}_r \subseteq \mathsf{U}(r)$, we choose an

arbitrary class $\alpha \in \mathsf{T}_r \cap \mathsf{B}_r$ and an r-perturbation h of f. Furthermore, we introduce the following pairs of subspaces:

$$\begin{aligned}
\mathbb{C}_r &= h^{-1}[a,\infty) \cap f^{-1}(-\infty, b+r],\\
\mathbb{C}'_r &= (h^{-1}[a,\infty) \cap f^{-1}(-\infty, a+r]) \cup (h^{-1}[b,\infty) \cap f^{-1}(-\infty, b+r]),\\
\mathbb{D}_r &= h^{-1}(-\infty, b] \cap f^{-1}[a-r,\infty),\\
\mathbb{D}'_r &= (h^{-1}(-\infty, a] \cap f^{-1}[a-r,\infty)) \cup (h^{-1}(-\infty, b] \cap f^{-1}[b-r,\infty));
\end{aligned}$$

see Figure 7 for a depiction of the open case. Since h is an r-perturbation, we have $(\mathbb{T}_r, \mathbb{T}'_r) \hookrightarrow (\mathbb{C}_r, \mathbb{C}'_r)$ and similarly $(\mathbb{B}_r, \mathbb{B}'_r) \hookrightarrow (\mathbb{D}_r, \mathbb{D}'_r)$. This implies $\mathsf{T}_r \subseteq \mathsf{C}_r$ and $\mathsf{B}_r \subseteq \mathsf{D}_r$, and therefore $\alpha \in \mathsf{C}_r \cap \mathsf{D}_r$. It is easy to see that $(\mathbb{C}_r \cup \mathbb{D}_r, \mathbb{C}'_r \cup \mathbb{D}'_r) = (\mathbb{X}_r, \mathbb{X}'_r)$, and also that $(\mathbb{C}_r \cap \mathbb{D}_r, \mathbb{C}'_r \cap \mathbb{D}'_r) = (h^{-1}(\mathbb{A}), h^{-1}(\mathbb{A}'))$. Lemma 2 thus implies $\alpha \in (h^{-1}(\mathbb{A}), h^{-1}(\mathbb{A}'))$. Since this is true for all r-perturbations h, we have $\alpha \in \mathsf{U}(r)$, as required. □

4.0.11. Including intervals

We again need some definitions to unify the four cases into one. Given two intervals I and J of the same type, we say I *includes* into J, denoted as $I \hookrightarrow J$, if $f^{-1}(I)$ includes as a pair in $f^{-1}(J)$. Unfolding the definition of the four types and assuming $a \leqslant b \leqslant c \leqslant d$, we have $[b,c] \hookrightarrow [a,d]$, $[b,d) \hookrightarrow [a,c)$, $(a,c] \hookrightarrow (b,d]$, and $(a,d) \hookrightarrow (b,c)$; compare this with the Mayer-Vietoris diamonds in Figure 2. Suppose now that we have intervals $I \hookrightarrow J$, both of the same type. By Theorem 2, there are isomorphisms that take $\mathsf{H}_p(f^{-1}(I))$ and $\mathsf{H}_p(f^{-1}(J))$ to groups $\mathsf{G}_p(I)$ and $\mathsf{G}_p(J)$ in $\mathcal{V}$. The inclusion induces a map on homology, which composes with these isomorphisms to give $\mathsf{g}\colon \mathsf{G}_p(I) \to \mathsf{G}_p(J)$. On the other hand, since the two groups are members of $\mathcal{V}$, there is also a natural map from $\mathsf{G}_p(I)$ to $\mathsf{G}_p(J)$, namely the one that restricts to the identity on the span of their shared vectors and is zero otherwise. Not surprisingly, g is exactly that map. We formalize this claim and give a proof.

Theorem 5. *Let $I \hookrightarrow J$ and let $\mathsf{G}_p(I)$, $\mathsf{G}_p(J)$ be the corresponding p-dimensional groups in $\mathcal{V}$. Then the image of $\mathsf{g}\colon \mathsf{G}_p(I) \to \mathsf{G}_p(J)$ is a vector space in $\mathcal{V}$, and its basis is in bijection with the multiset $\mathcal{W}_p(I) \cap \mathcal{W}_p(J)$.*

Proof. To restate the theorem, we consider the diagram defined by the homology groups of the preimages of the including intervals, $I \hookrightarrow J$, and the corresponding vector spaces in $\mathcal{V}$:

$$\begin{array}{ccc}
\mathsf{H}_p(f^{-1}(I)) & \xrightarrow{\mathsf{h}} & \mathsf{H}_p(f^{-1}(J)) \\
\uparrow & & \downarrow \\
\mathsf{G}_p(I) & \xrightarrow{\mathsf{g}} & \mathsf{G}_p(J).
\end{array}$$

The vertical maps are isomorphisms given by Theorem 2. The map h is induced by inclusion, and g maps a basis vector of $\mathsf{G}_p(I)$ to the same basis vector of $\mathsf{G}_p(J)$, if it exists, and to zero, otherwise. Hence, the basis of $\operatorname{im}\mathsf{g}$ consists of the vectors that are common to the bases of $\mathsf{G}_p(I)$ and $\mathsf{G}_p(J)$. This theorem states that we can get g by composing h with the two isomorphisms. Equivalently, the diagram commutes. To prove commutativity, we consider again the zigzag modules drawn as monotonic paths in the square; see Figure 2. Since $I \hookrightarrow J$, we can find two non-crossing modules, one containing $\mathsf{H}_p(f^{-1}(I))$ and the other containing $\mathsf{H}_p(f^{-1}(J))$. To get a basis for

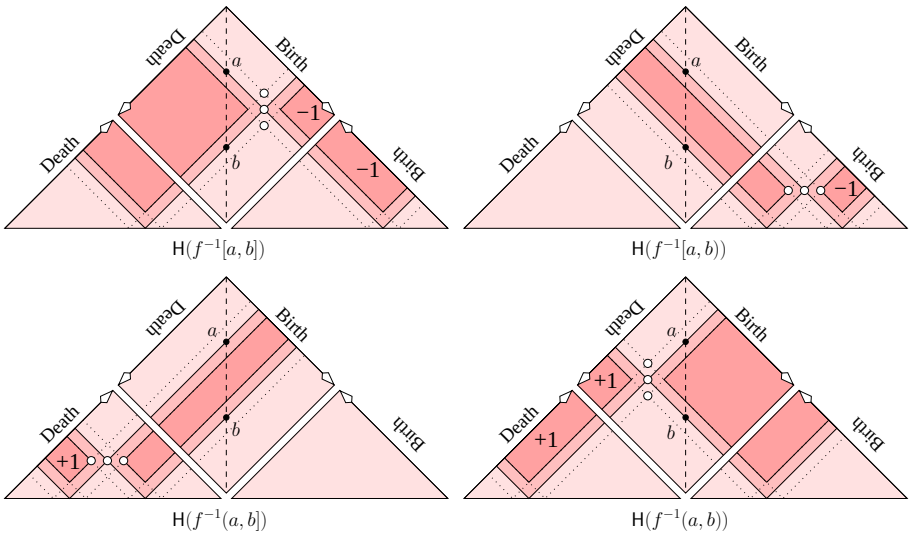

Figure 8: Reading the robust homology in the four different cases. The shaded region gives the basis of $\mathsf{H}_p(f^{-1}(I))$, while the dark shaded region gives the basis of the well subgroup, $\mathsf{U}_p(r)$.

$\operatorname{im}\mathsf{h}$, we translate intervals from one path to the other, keeping only the ones that cover both $\mathsf{H}_p(f^{-1}(I))$ and $\mathsf{H}_p(f^{-1}(J))$. Further translating these intervals to the hypotenuse gives the corresponding points in the persistence diagram. These points are precisely the ones shared by $\mathcal{W}_p(I)$ and $\mathcal{W}_p(J)$. In other words, $\operatorname{im}\mathsf{g}$ in $\mathcal{V}$ is isomorphic to $\operatorname{im}\mathsf{h}$, as desired. □

4.0.12. Reading robustness

Theorem 5 allows us to compute the well groups and the well diagram associated to a single interval, $I = (\mathbb{A}, \mathbb{A}')$. The homology of $f^{-1}(I)$ can be read off the persistence diagram of f, as stated in Theorem 2. Similarly, the homology of $(\mathbb{X}_r, \mathbb{X}'_r)$, where $\mathbb{X}_r = \mathbb{X}_r(f_{\mathbb{A}})$ and $\mathbb{X}'_r = \mathbb{X}_r(f'_{\mathbb{A}})$, can be read off the same diagram, as we now explain. By Theorem 4, the well group for r is the intersection of the images of the maps $\mathsf{t}_r \colon \mathsf{H}_p(\mathbb{T}_r, \mathbb{T}'_r) \to \mathsf{H}_p(\mathbb{X}_r, \mathbb{X}'_r)$ and $\mathsf{b}_r \colon \mathsf{H}_p(\mathbb{B}_r, \mathbb{B}'_r) \to \mathsf{H}_p(\mathbb{X}_r, \mathbb{X}'_r)$ induced by the inclusions. By Theorem 5, this intersection corresponds to a pair of rectangles within the region of $f^{-1}(I)$; see the intersection between $\mathcal{W}_p(I)$ and the dotted rectangles in Figure 8. In the closed case, this intersection gradually recedes to infinity, while in the two half-open cases, the intersection disappears when r reaches half the length of the interval. Correspondingly, the well group shrinks gradually in the closed case, while it vanishes at or before $r = (b-a)/2$ in the half-open cases. Similarly, the well group vanishes when r reaches $(b-a)/2$ in the open case. However, here it vanishes abruptly. More precisely, the range of the maps t_r and b_r, which is $\mathsf{H}_p(f^{-1}(a+r, b-r))$, approaches the homology group of the suspension of the level set at $(a+b)/2$, when r goes toward $(b-a)/2$, before it suddenly becomes zero when r reaches that limit.

In all four cases, a point contributes to the well group until r reaches a value at which the shrinking intersection no longer contains the point. Finding this value of r is easy since both rectangles shrink uniformly along all of their sides. Consider for example the case $I = [a, b]$ illustrated by the upper left design in Figure 8. For a point $(x, y) \in \mathcal{D}(f)$, the value of r at which the point drops out of the relevant region is

$$\begin{aligned}
&\min\{b - x, y - b\} \text{if} (x, y) \in \mathcal{O}(f) \cap \lambda[a, b],\\
&\min\{b - x, y - a\} \text{if} (x, y) \in \mathcal{E}(f) \cap \lambda[a, b],\\
&\min\{x - b, a - y\} \text{if} (x, y) \in \mathcal{E}(f) \cap \varrho[a, b],\\
&\min\{x - a, a - y\} \text{if} (x, y) \in \mathcal{R}(f) \cap \varrho[a, b].
\end{aligned}$$

The well diagram is the multiset of the values we get from the points in the persistence diagram.

4.0.13. Measuring the difference
We can interpret the rank of the well group as a measure of the similarity between the image of the map $\mathsf{t}_r \colon (\mathbb{T}_r, \mathbb{T}'_r) \to (\mathbb{X}_r, \mathbb{X}'_r)$ and the image of the map $\mathsf{b}_r \colon (\mathbb{B}_r, \mathbb{B}'_r) \to (\mathbb{X}_r, \mathbb{X}'_r)$. Alternatively, we could use the cokernels of these two maps to measure their difference. Indeed, it is not difficult to prove counterparts of Theorem 5 for cokernels as well as for kernels.

Theorem 6. *Let $I \hookrightarrow J$ and let $\mathsf{G}_p(I)$, $\mathsf{G}_p(J)$ be the corresponding p-dimensional groups in $\mathcal{V}$. Then the kernel and cokernel of $\mathsf{g} \colon \mathsf{G}_p(I) \to \mathsf{G}_p(J)$ are vector spaces in $\mathcal{V}$, the basis of $\ker \mathsf{g}$ is in bijection with $\mathcal{W}_p(I) - \mathcal{W}_p(J)$, and the basis of $\operatorname{cok} \mathsf{g}$ is in bijection with $\mathcal{W}_p(J) - \mathcal{W}_p(I)$.*

To measure the difference, we would therefore take the (algebraic) sum of the two cokernels. Consider for example the open case. By the above lemma, we get a basis of $\operatorname{cok} \mathsf{t}_r$ and $\operatorname{cok} \mathsf{b}_r$ by setting $J = (a + r, b - r)$ and first setting I to $I_1 = (a + r, b + r)$ and second to $I_2 = (a - r, b - r)$. The basis of the sum, $\operatorname{cok} \mathsf{t}_r + \operatorname{cok} \mathsf{b}_r$, is in bijection with the union of the two multisets of points, which is $\mathcal{W}_p(J) - \mathcal{W}_p(I_1) - \mathcal{W}_p(I_2)$.

5. Discussion

The main contribution of this paper is the introduction of the point calculus for homology computations of level and interlevel sets. This comprises interlevel sets defined by closed, half-open, and open intervals, images, kernels, and cokernels of maps induced by inclusions, and the robustness of homology as defined by well groups. The point calculus provides a compact interface to a wealth of homological information that can be useful to researchers with and without background in algebraic topology. For the expert, it provides a compact summary of information that may be used to formulate conjectures about the topology of spaces and of functions. For the non-expert, the interface offers an intuitive approach to understand the topology of datasets that by-passes the introduction of algebraic topology foundations. It is directly applicable to data in the form of continuous functions, which is common in medical imaging and in scientific visualization.

We conclude by formulating an open question aimed at casting light on two- and higher-dimensional notions of robustness. This paper provides a solution to computing

robustness when $\mathbb{Y} = \mathbb{R}$ and perturbations are measured using the L_∞-metric, and [**3**] shows that our results also hold for a broader class of metric function spaces. In [**9**], the authors give an algorithm when $\mathbb{X}$ is an orientable 2-manifold, $\mathbb{Y} = \mathbb{R}^2$, and $\mathbb{A}$ is a point. Algorithms for other cases are not yet known.

References

[**1**] C. L. Bajaj, V. Pascucci and D. R. Schikore. The contour spectrum. *In* "Proc. 8th IEEE Conf. Visualization, 1997", 167–173.

[**2**] P. Bendich, H. Edelsbrunner and M. Kerber. Computing persistence and robustness for images. *IEEE Trans. Visual. Comput. Graphics* **16** (2010), 1251-1260.

[**3**] P. Bendich, H. Edelsbrunner, M. Kerber and A. Patel Persistent homology under non-uniform error. *In* "Proc. 35th Internat. Sympos. Math. Found. Comput. Sci. 2010," Springer-Verlag, LNCS **6281**, 12–23.

[**4**] G. Carlsson, A. Collins, L. J. Guibas and Z. Zomorodian. Persistence barcodes for shapes. *Internat. J. Shape Modeling* **11** (2005), 149–187.

[**5**] G. Carlsson and V. de Silva. Zigzag persistence. *Found. Comput. Math.* **10** (2010), 367–405.

[**6**] G. Carlsson and V. de Silva and D. Morozov. Zigzag persistent homology and real-valued functions. *In* "Proc. 25th Ann. Sympos. Comput. Geom., 2009", 247–256.

[**7**] D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79-103.

[**8**] H. Edelsbrunner and J. Harer. Persistent homology — a survey. *Surveys on Discrete and Computational Geometry. Twenty Years Later*, eds. J. E. Goodman, J. Pach and R. Pollack, Contemporary Mathematics **453**, 257–282, Amer. Math. Soc., Providence, Rhode Island, 2008.

[**9**] H. Edelsbrunner, D. Morozov and A. Patel. The stability of the apparent contour of an orientable 2-manifold. In *Topological Methods in Data Analysis and Visualization*, 27–42, eds. V. Pascucci, X. Tricoche, H. Hagen and J. Tierny, Springer-Verlag, Heidelberg, Germany, 2011.

[**10**] H. Edelsbrunner, D. Morozov and A. Patel. Quantifying transversality by measuring the robustness of intersections. *Found. Comput. Math.* **11** (2011), 345-361.

[**11**] I. Fujishiro, Y. Maeda and H. Sato. Interval volume: a solid fitting technique for volumetric data display and analysis. *In* "Proc. 6th IEEE Conf. Visualization, 1995", 151–158.

[**12**] M. van Kreveld, R. van Oostrum, C. L. Bajaj, V. Pascucci and D. R. Schikore. Contour trees and small seed sets for isosurface traversal. *In* "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 212–220.

[**13**] J. R. Munkres. *Elements of Algebraic Topology.* Perseus, Cambridge, Massachusetts, 1984.

Paul Bendich `bendich@math.duke.edu`

Mathematics Department, Duke University, Box 90320 Durham, NC 27708-0320, USA, and IST Austria (Institute of Science and Technology Austria), Am Campus 1, A-3400, Klosterneuburg, AUSTRIA

Herbert Edelsbrunner edels@ist.ac.at

IST Austria (Institute of Science and Technology Austria), Am Campus 1, A-3400, Klosterneuburg, AUSTRIA, Departments of Computer Science and of Mathematics, Duke University, Box 90320, Durham, NC 27708-0320, USA, and Geomagic, 430 Davis Drive, Suite 300 Morrisville, NC 27560, USA

Dmitriy Morozov dmitriy@mrzv.org

Departments of Computer Science and of Mathematics, Stanford University, 353 Serra Mall Stanford, CA 94305-9025, USA

Amit Patel amit@akpatel.org

IST Austria (Institute of Science and Technology Austria), Am Campus 1, A-3400, Klosterneuburg, AUSTRIA. and Duke University, Department of Computer Science, LSRC Building D101, 308 Research Drive, Campus Box 90129 Durham, NC 27708, USA