The Robustness of Level Sets ⋆

Paul Bendich[1,2,3], Herbert Edelsbrunner[1,2,3,5], Dmitriy Morozov[4], and Amit Patel[1,2]

[1] IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
[2] Dept. Comput. Sci., Duke Univ., Durham, North Carolina
[3] Dept. Mathematics, Duke Univ. Durham, North Carolina
[4] Depts. Comput. Sci. and Math., Stanford Univ., Stanford, California
[5] Geomagic, Research Triangle Park, North Carolina.

Abstract. We define the robustness of a level set homology class of a function $f : \mathbb{X} \to \mathbb{R}$ as the magnitude of a perturbation necessary to kill the class. Casting this notion into a group theoretic framework, we compute the robustness for each class, using a connection to extended persistent homology. The special case $\mathbb{X} = \mathbb{R}^3$ has ramifications in medical imaging and scientific visualization.

Keywords. Topological spaces, continuous functions, level sets, perturbations, homology, extended persistence, well groups, well diagrams, robustness.

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 between topological spaces. Assuming a metric space of perturbations, we have such a group for each subspace $\mathbb{A} \subseteq \mathbb{Y}$, each bound $r \geq 0$ on the magnitude of the perturbation, and each dimension p. These groups 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 [8] and prove the stability of the apparent contour of a mapping from an orientable 2-manifold to $\mathbb{R}^2$ [7]. In this paper, we contribute to the general understanding of well groups by studying the real-valued case. Specifically,

I. we characterize the well group of $f : \mathbb{X} \to \mathbb{R}$ when the space $\mathbb{A}$ is a single point;
II. we give an algorithm relating the well diagram of f and $\mathbb{A}$ with the extended persistence diagram of f.

In the full version of this paper, we extend these results to the case when $\mathbb{A}$ is a finite union of points and intervals. Applications of this theoretical work can be found in scientific visualization, where data in the form of real-valued functions is common. To mention one example, the magnetic resonance image of a person's brain results in a

⋆ This research is partially supported by the Defense Advanced Research Projects Agency (DARPA), under grants HR0011-05-1-0057 and HR0011-09-0065, as well as the National Science Foundation (NSF), under grant DBI-0820624.

3-dimensional array of intensity values, best viewed as a function from the unit cube to the real numbers. Generically, the preimage of a value $a \in \mathbb{R}$ is a 2-manifold, referred to as a *contour* or an *isosurface* [10, 12]. We contribute to the state of the art by

III. explaining how the homology of an isosurface can be read off the extended persistence diagram of the function;
IV. describing how the robustness of features in isosurfaces can be read off the same diagram.

We believe that these results warrant the development of extended persistence diagrams as a new interface tool to efficiently select interesting collections of isosurfaces. 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 sequence of level sets. The most novel aspect of this new tool is the robustness information, which is readily available through subregions of the diagram.

Outline. In Section 2, we review necessary background on persistence diagrams and well groups. In Section 3, we present our main result: the characterization of the well groups. In Section 4, we explain the relationship between the points in the persistence diagram and the homology of level sets, extending it to robust homology in Section 5. Finally, Section 6 concludes the paper with a brief discussion of further research directions.

2 Background

This section begins with a review of persistence and persistence diagrams. Then we give a description of the 1-parameter family of well groups, and conclude with an example that illustrates the given definitions.

Persistence. The persistence of homology classes along a filtration of a topological space can be defined in a quite general context [5]. 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 $r \leq s$, 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 $\mathbb{Z}/2\mathbb{Z}$ coefficients. Often we will suppress the dimension from our notation, writing $\mathsf{H}(\mathbb{X}_r(f)) = \bigoplus_p \mathsf{H}_p(\mathbb{X}_r(f))$; in this case, a map $\mathsf{H}(\mathbb{X}_r(f)) \to \mathsf{H}(\mathbb{X}_s(f))$ will of course decompose into maps on each factor. A real value r is called a *homological regular value* of f if there exists an $\varepsilon > 0$ such that the inclusion $\mathbb{X}_{r-\delta}(f) \hookrightarrow \mathbb{X}_{r+\delta}(f)$ induces an isomorphism between homology groups for all $\delta < \varepsilon$. 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

$\mathbb{X}_i = \mathbb{X}_{s_i}(f)$. We have $\mathbb{X}_0 = \emptyset$ and $\mathbb{X}_n = \mathbb{X}$, by compactness. The inclusions $\mathbb{X}_i \hookrightarrow \mathbb{X}_j$ induce maps $\mathsf{f}^{i,j} : \mathsf{H}(\mathbb{X}_i) \to \mathsf{H}(\mathbb{X}_j)$ for $0 \leq i \leq j \leq n$ and give the following filtration:

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

Given a class $\alpha \in \mathsf{H}(\mathbb{X}_i)$, we say that α 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.

Extended persistence. Note that the filtration in (1) begins with the zero group but ends with a potentially nonzero group. Hence, 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 [4] and extend (1) using relative homology groups. More precisely, we consider for each i the *superlevel set* $\mathbb{X}^i = f^{-1}[s_{n-i}, \infty)$. For $i \leq j$, the inclusion $\mathbb{X}^i \hookrightarrow \mathbb{X}^j$ induces a map on relative homology $\mathsf{H}(\mathbb{X}, \mathbb{X}^i) \to \mathsf{H}(\mathbb{X}, \mathbb{X}^j)$. We have $\mathbb{X}^0 = \emptyset$ and $\mathbb{X}^n = \mathbb{X}$ by compactness. These maps therefore lead to the extended filtration:

$$\begin{aligned} 0 &= \mathsf{H}(\mathbb{X}_0) \to \mathsf{H}(\mathbb{X}_1) \to \ldots \to \mathsf{H}(\mathbb{X}_n) = \mathsf{H}(\mathbb{X}) \\ &= \mathsf{H}(\mathbb{X}, \mathbb{X}^0) \to \mathsf{H}(\mathbb{X}, \mathbb{X}^1) \ldots \to \mathsf{H}(\mathbb{X}, \mathbb{X}^n) = 0. \end{aligned} \tag{2}$$

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.

The information contained within the extended filtration (2) can be compactly represented by *persistence diagrams* $\mathrm{Dgm}_p(f)$, one for each dimension p in homology; see Figure 1. Each such diagram is a multiset of points in the plane: it contains 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*, $\mathrm{Ord}_p(f)$, represents classes that are born and die during the first half of (2). The *relative subdiagram*, $\mathrm{Rel}_p(f)$, represents classes that are born and die during the second half. Finally, the *extended subdiagram*, $\mathrm{Ext}_p(f)$, represents classes that are born during the first half and die during the second half of the extended filtration. Note that points in $\mathrm{Ord}_p(f)$ all lie above the main diagonal while points in $\mathrm{Rel}_p(f)$ all lie below. On the other hand, $\mathrm{Ext}_p(f)$ may contain points on either side of the main diagonal. By $\mathrm{Dgm}(f)$, we mean the points of all diagrams in all dimensions, overlaid as one multiset of points.

Note that the number of points in $\mathrm{Ext}_p(f)$ is precisely the rank of the p-th homology group of $\mathbb{X}$. A similar formula holds for the sublevel set $\mathbb{X}_r(f)$. Using levelset zigzag modules introduced in [3], we will see that this way of reading the rank of homology groups can be extended to level sets and, more generally, to sets of the form $f^{-1}[a, b]$.

Well groups. Given a continuous function $f : \mathbb{X} \to \mathbb{R}$ and a value $a \in \mathbb{R}$, we review the definition of the well groups $\mathsf{U}_p(a, r)$ for each radius $r \geq 0$ and each dimension p.

Since a will be fixed, we usually drop it from the notation and simply write $\mathsf{U}(r)$, by which we mean the direct sum of groups $\mathsf{U}_p(a, r)$, over all homology dimensions p. We will need the assumption that $f^{-1}(a)$ has homology groups of finite rank.

To begin, we define the *radius function* $f_a : \mathbb{X} \to \mathbb{R}$ by mapping each point x to $f_a(x) = |f(x) - a|$. Using this real-valued function, we filter $\mathbb{X}$ via sublevel sets: $\mathbb{X}_r(f_a) = f_a^{-1}[0, r]$. For $r \leq s$, there is a map $\mathsf{f}^{r,s} : \mathsf{H}(\mathbb{X}_r(f_a)) \to \mathsf{H}(\mathbb{X}_s(f_a))$. By an *$r$-perturbation* h of f, we mean a function $h : \mathbb{X} \to \mathbb{R}$ such that $\|h - f\|_\infty = \sup_{x \in \mathbb{X}} |h(x) - f(x)| \leq r$. The preimage of a under any such h will obviously be a subset of $\mathbb{X}_r(f_a)$, and hence there is a map on homology, $\mathsf{j}_h : \mathsf{H}(h^{-1}(a)) \to \mathsf{H}(\mathbb{X}_r(f_a))$. Given a class $\alpha \in \mathsf{H}(\mathbb{X}_r(f_a))$, we say that α is *supported* by h if $\alpha \in \operatorname{im} \mathsf{j}_h$. Equivalently, $h^{-1}(a)$ carries a chain representative of α. The *well group* $\mathsf{U}(r) \subseteq \mathsf{H}(\mathbb{X}_r(f_a))$ is then defined to consist of the classes that are supported by all r-perturbations of f:

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

For $r \leq s$, the map $\mathsf{f}^{r,s}$ restricts to $\mathsf{U}(r) \to \mathsf{H}(\mathbb{X}_s(f_a))$. On the other hand, $\mathsf{H}(\mathbb{X}_s(f_a))$ contains $\mathsf{U}(s)$ as a subgroup. It can be shown that $\mathsf{U}(s) \subseteq \mathsf{f}^{r,s}(\mathsf{U}(r))$ whenever $r \leq s$; see [8]. In other words, the rank of the well group can only decrease as the radius increases.

We call a value of r at which the rank of the well group decreases a *terminal critical value*. The *well diagram* of f and a is then the multiset of terminal critical values of f_a, taking a value k times if the rank of the well group drops by k at the value. Here we note that well groups can be defined in a more general context [8], given a mapping $f : \mathbb{X} \to \mathbb{Y}$, a subspace $\mathbb{A} \subseteq \mathbb{Y}$, and a metric space of perturbations. In this general setting, the relationship between the terminal critical values and the homological critical values of f_a is not completely understood. However, for $\mathbb{Y} = \mathbb{R}$ and $\mathbb{A} = \{a\}$, we will see shortly that the former is a subset of the latter.

Example. Consider the torus $\mathbb{X}$, as shown in Figure 1, along with the vertical height function f and a value $a \in \mathbb{R}$. The preimage of a, $f^{-1}(a) = f_a^{-1}(0)$, consists of two disjoint circles on the torus; hence there are two components and two independent 1-cycles, all belonging to the well group. For small values of r, $\mathbb{X}_r(f_a)$ 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 group.

Now consider the value of r shown in Figure 1. For this r, the sublevel set $\mathbb{X}_r(f_a)$ consists of two pair-of-pants glued together along two common circles. We note that $\mathsf{H}_0(\mathbb{X}_r(f_a))$ has dropped in rank by one, while the rank of $\mathsf{H}_1(\mathbb{X}_r(f_a))$ has grown to three. In contrast, the rank of $\mathsf{U}_1(r)$ is less than or equal to one. Indeed, the function $h : \mathbb{X} \to \mathbb{R}$, defined by $h = f - r$, is an r-perturbation of f and its level set at a, $h^{-1}(a) = f^{-1}(a + r)$, 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 the next section.

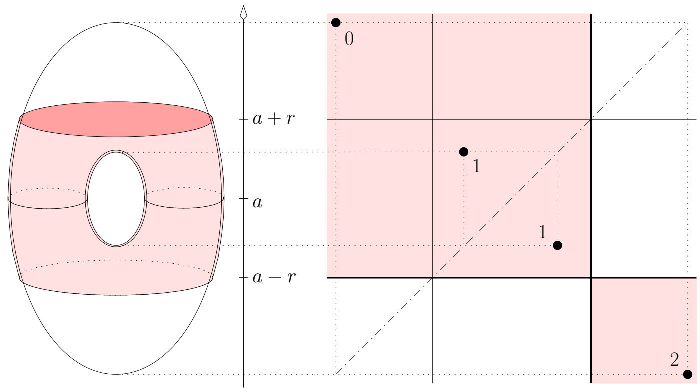

Fig. 1: Left: the torus and the preimage of the interval $[a-r, a+r]$. Right: the persistence diagram of the vertical height function. Each point is labeled by the dimension of the corresponding homology class.

3 Characterization

In this section, we characterize the well groups. We begin with a consequence of the exactness of the Mayer-Vietoris sequence, see eg. [11], which will provide the main technical ingredient of our proof.

Mayer-Vietoris sequence. For convenience, we establish the following notational convention, wherein we reuse the same letter in different fonts. If $\mathbb{X} \subseteq \mathbb{Y}$ are topological spaces, then inclusion induces a map $\mathsf{x} : \mathsf{H}(\mathbb{X}) \to \mathsf{H}(\mathbb{Y})$ on homology groups and we write $\mathsf{X} = \operatorname{im} \mathsf{x}$ for the image of this map. Note that X is always a subgroup of $\mathsf{H}(\mathbb{Y})$, namely the subgroup of homology classes that have a chain representative carried by $\mathbb{X}$. Note also that the rank of X can never exceed the rank of $\mathsf{H}(\mathbb{X})$. Suppose that $\mathbb{W} \subseteq \mathbb{X}$ are two subspaces of $\mathbb{Y}$. Then, from the chain of maps $\mathsf{H}(\mathbb{W}) \to \mathsf{H}(\mathbb{X}) \to \mathsf{H}(\mathbb{Y})$, we see that W must be a subgroup of X. The following lemma is a direct consequence of the exactness of the Mayer-Vietoris sequence. However, we will use it often so it seems reasonable to state it formally.

1 (Mayer-Vietoris Lemma) *Suppose that we can write a topological space $\mathbb{Y}$ as $\mathbb{Y} = \mathbb{C} \cup \mathbb{D}$, with $\mathbb{E} = \mathbb{C} \cap \mathbb{D}$. If a class $\alpha \in \mathsf{H}(\mathbb{Y})$ belongs to C as well as to D, then α also belongs to E.*

PROOF. Following our convention, we use the notation $\mathsf{c} : \mathsf{H}(\mathbb{C}) \to \mathsf{H}(\mathbb{Y})$ for the map on homology induced by the inclusion of $\mathbb{C}$ in $\mathbb{Y}$. Similarly, we write $\mathsf{d} : \mathsf{H}(\mathbb{D}) \to \mathsf{H}(\mathbb{Y})$ and $\mathsf{e} : \mathsf{H}(\mathbb{E}) \to \mathsf{H}(\mathbb{Y})$, as well as $\mathsf{e}_c : \mathsf{H}(\mathbb{E}) \to \mathsf{H}(\mathbb{C})$ and $\mathsf{e}_d : \mathsf{H}(\mathbb{E}) \to \mathsf{H}(\mathbb{D})$. 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 the union $\mathbb{Y} = \mathbb{C} \cup \mathbb{D}$:

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

By assumption, $\alpha \in \mathsf{C}$, so there exists some $\alpha_c \in \mathsf{H}(\mathbb{C})$ such that $\mathsf{c}(\alpha_c) = \alpha$. Similarly, there exists an $\alpha_d \in \mathsf{H}(\mathbb{D})$ 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 $\alpha_e \in \mathsf{H}(\mathbb{E})$ 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 the Mayer-Vietoris Lemma, we will construct subspaces $\mathbb{B}_0 \subseteq \mathbb{C}$ and $\mathbb{B}_1 \subseteq \mathbb{D}$ such that $\alpha \in \mathsf{B}_0 \cap \mathsf{B}_1$. From the remark above, we know that $\mathsf{B}_0 \subseteq \mathsf{C}$ and $\mathsf{B}_1 \subseteq \mathsf{D}$. The lemma then applies and we can conclude that $\alpha \in \mathsf{E}$, as before.

One-point case. We now suppose that we have a topological space $\mathbb{X}$ and a function $f : \mathbb{X} \to \mathbb{R}$, and we find the well groups $\mathsf{U}(a, r) = \mathsf{U}(r)$. Recall that $\mathbb{X}_r(f_a) = f_a^{-1}[0, r] = f^{-1}[a - r, a + r]$. To state the formula, we distinguish two particular subspaces of $\mathbb{X}_r(f_a)$, namely the top level set, $\mathbb{B}_{0,r} = f^{-1}(a + r)$, and the bottom level set, $\mathbb{B}_{1,r} = f^{-1}(a - r)$. Using the convention from before, we write $\mathsf{B}_{0,r}$ and $\mathsf{B}_{1,r}$ for the images of $\mathsf{H}(\mathbb{B}_{0,r})$ and $\mathsf{H}(\mathbb{B}_{1,r})$ in $\mathsf{H}(\mathbb{X}_r(f_a))$.

2 (One-Point Formula) $\mathsf{U}(r) = \mathsf{B}_{0,r} \cap \mathsf{B}_{1,r}$, *for every* $r \geq 0$.

PROOF. We simplify notation by fixing r and dropping it from our notation. We prove equality by proving the two inclusions in turn. To show $\mathsf{U} \subseteq \mathsf{B}_0 \cap \mathsf{B}_1$, consider a class $\alpha \in \mathsf{U}$. We define $h_0 = f - r$ and $h_1 = f + r$ and note that they are r-perturbations of f, with $h_0^{-1}(a) = \mathbb{B}_0$ and $h_1^{-1}(a) = \mathbb{B}_1$. By definition of the well group, α is supported by every r-perturbation of f, and therefore by h_0 and by h_1. It follows that $\alpha \in \mathsf{B}_0 \cap \mathsf{B}_1$.

To show $\mathsf{B}_0 \cap \mathsf{B}_1 \subseteq \mathsf{U}$, we consider a class $\alpha \in \mathsf{B}_0 \cap \mathsf{B}_1$ and let h be an arbitrary r-perturbation of f. To finish the proof, we just need to show that α is supported by h. We define $\mathbb{C} = \{x \in \mathbb{X}_r(f_a) \mid h(x) \geq a\}$ and $\mathbb{D} = \{x \in \mathbb{X}_r(f_a) \mid h(x) \leq a\}$. Note that $\mathbb{C} \cup \mathbb{D} = \mathbb{X}_r(f_a)$ while $\mathbb{C} \cap \mathbb{D} = h^{-1}(a)$. Furthermore, the inequality $\|h - f\|_\infty \leq r$ implies that $\mathbb{B}_0 \subseteq \mathbb{C}$ and $\mathbb{B}_1 \subseteq \mathbb{D}$. By the Mayer-Vietoris Lemma, α is supported by $h^{-1}(a)$, as required. ◻

We note that the One-Point Formula implies that the well group for a Morse function f can change only at critical values of the function f_a. In other words, terminal critical values are, in this simple context, just ordinary critical values. Indeed, if $[r, s]$ is an interval that contains no critical values of f_a, then there is a deformation retraction $\mathbb{X}_s(f_a) \to \mathbb{X}_r(f_a)$ providing an isomorphism $\mathsf{H}(\mathbb{X}_r(f_a)) \to \mathsf{H}(\mathbb{X}_s(f_a))$. Furthermore, this retraction maps $\mathbb{B}_{0,s}$ onto $\mathbb{B}_{0,r}$, in such a way that that the images of $\mathsf{H}(\mathbb{B}_{0,r})$ and $\mathsf{H}(\mathbb{B}_{0,s})$ in $\mathsf{H}(\mathbb{X}_s(f_a))$ are identical. Similarly, the images of $\mathsf{H}(\mathbb{B}_{1,r})$ and $\mathsf{H}(\mathbb{B}_{1,s})$ in $\mathsf{H}(\mathbb{X}_s(f_a))$ are identical. Hence the well groups $\mathsf{U}(r)$ and $\mathsf{U}(s)$ are isomorphic.

4 Combinatorics of Homology

We note that the groups relevant to the One-Point Formula are all groups of a very particular type. Namely, each is the image, under a map induced by inclusion, of the homology of a level set of f. In this section, we describe a relationship between the points in the extended persistence diagram of f and the homology groups of any level set. More generally, we show how the homology of the preimage of any interval can be read from the extended persistence diagram. We also give a similar relationship for the image induced by the inclusion of a smaller interval into a bigger one.

Common basis. The main idea here, stated intuitively, is that each point in the persistence diagram corresponds to a unique basis vector of an abstract vector space in such a way that the points in certain subregions of the diagram give crucial information. Slightly more precisely, we choose a persistence module basis $\mathcal{B}$ for the extended filtration (2), one which results from transforming a basis of the levelset zigzag module in the way described by the Pyramid Basis Theorem; for complete precision, we refer the reader to the full version of this paper. Following [6] and [3], we note that these basis vectors correspond bijectively to the points in the extended persistence diagram. We then define an abstract vector space $\mathsf{V} = \langle \mathcal{B} \rangle$. By $\mathcal{V}$, we will mean the collection of those particular vector subspaces of V that have a basis consisting of vectors chosen from $\mathcal{B}$; in other words, $\mathcal{V} = \{\langle \mathcal{B}' \rangle \mid \mathcal{B}' \subseteq \mathcal{B}\}$.

Now suppose that we have a pair of real numbers $a \leq b$ and consider the homology of $f^{-1}[a, b]$, the *interlevel set* defined by $[a, b]$. For convenience, we assume that a and b are different from all coordinates of points in $\mathrm{Dgm}(f)$. We will demonstrate shortly that a basis for $\mathsf{H}(f^{-1}[a, b])$ can be read directly off the extended persistence diagram. To formulate this claim, we define two multisets of points:

$$\mathcal{L}_p[a, b] = \{(x, y) \in \mathrm{Ord}_p(f) \mid x < b, y > b\} \sqcup \{(x, y) \in \mathrm{Ext}_p(f) \mid x < b, y > a\},$$
$$\mathcal{R}_p[a, b] = \{(x, y) \in \mathrm{Ext}_p(f) \mid x > b, y < a\} \sqcup \{(x, y) \in \mathrm{Rel}_p(f) \mid x > a, y < a\},$$

for every dimension p; see Figure 2. It will be convenient to glue the domains of the three subdiagrams and draw the result as a right-angled triangle, as in Figure 3. In this

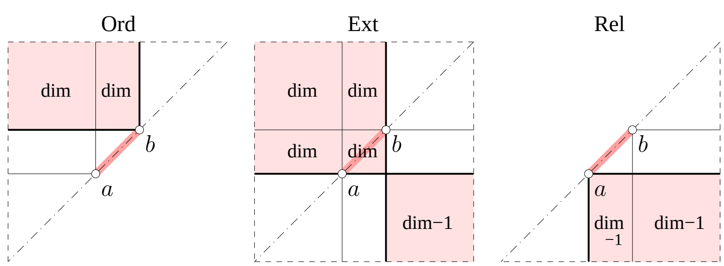

Fig. 2: From left to right: the shaded regions in the ordinary, extended, and relative subdiagrams in which the points correspond to the basis of the homology of the interlevel set defined by $[a, b]$.

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 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. We note that there is a straightforward translation of this triangular design to the representation of persistence advocated in [2]. Namely, draw a symmetric right-angled 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.

Reading homology. The purpose of the multisets $\mathcal{L}_p[a,b]$ and $\mathcal{R}_p[a,b]$ is to offer a convenient way to read the homology of a level set or an interlevel set from the persistence diagram. We make this statement precise in the following lemma, which is a corollary of the Pyramid Basis Theorem given in the full version of this paper.

3 (Interlevel Set Lemma) *For each dimension p and each pair of real numbers $a \leq b$, there exists an isomorphism taking $\mathsf{H}_p(f^{-1}[a,b])$ onto the vector space $\mathsf{G} \in \mathcal{V}$ spanned by the basis vectors corresponding to the points in $\mathcal{L}_p[a,b] \cup \mathcal{R}_{p+1}[a,b]$*

Recall that the points in $\mathrm{Ext}_p(f)$ determine the homology of $\mathbb{X}$. This is a special case of the lemma. To get $f^{-1}[a,b] = \mathbb{X}$, we choose a smaller than the minimum function value and b larger than the maximum function value. Hence, $\mathcal{L}_p[a,b] = \mathrm{Ext}_p(f)$ and $\mathcal{R}_p[a,b] = \emptyset$ for all dimensions p, as required. Of course the homology of a level set $f^{-1}(a)$ can also be read off via the Interlevel Set Lemma; one simply sets $a = b$ and makes the necessary adaptations to the formula.

Now suppose we have a pair of nested intervals $[a,b] \subseteq [c,d]$. By the Interlevel Set Lemma, there are isomorphisms that take the homology groups $\mathsf{H}(f^{-1}[a,b])$ and $\mathsf{H}(f^{-1}[c,d])$ onto groups $\mathsf{G}, \mathsf{G}' \in \mathcal{V}$, respectively. The inclusion of the smaller into the larger interval induces a map on homology, which composes with the isomorphisms obtained from the Interlevel Set Lemma to give $\mathrm{g} : \mathsf{G} \to \mathsf{G}'$. Since the two groups are members of $\mathcal{V}$, there is a natural map from G to G', 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 give the proof of this result in the full version of this paper.

4 (Interval Mapping Lemma) *Let $[a,b] \subseteq [c,d]$ and let G, G' be the corresponding groups in $\mathcal{V}$. Then the image of the map $\mathrm{g} : \mathsf{G} \to \mathsf{G}'$ is in $\mathcal{V}$, with basis $\mathcal{B}(\mathrm{im}\, \mathrm{g})$ in bijection with the multiset $(\mathcal{L}_p[a,b] \cap \mathcal{L}_p[c,d]) \cup (\mathcal{R}_{p+1}[a,b] \cap \mathcal{R}_{p+1}[c,d])$.*

5 Combinatorics of Robustness

This section gives a procedure for reading the well diagrams from the persistence diagram for f. The homology of $\mathbb{X}_0(f_a) = f^{-1}(a)$ can be read off the persistence diagram of f, as stated in the Interlevel Set Lemma. Specifically, $\mathsf{H}_p(\mathbb{X}_0(f_a))$ is isomorphic to the vector space whose basis corresponds to $\mathcal{L}_p[a,a] \cup \mathcal{R}_{p+1}[a,a]$. Similarly, the homology of $\mathbb{X}_r(f_a) = f^{-1}[a-r, a+r]$ can be read off the same diagram. By the One-Point Formula, the well group for r is the intersection of the images of the homology maps induced by the inclusions of $f^{-1}(a-r)$ and $f^{-1}(a+r)$ in $\mathbb{X}_r(f_a)$. By the

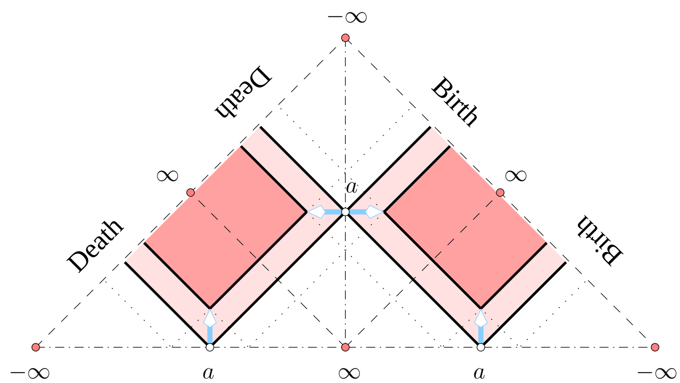

Fig. 3: The triangle design of the extended persistence diagram. The shaded region gives the basis of $\mathsf{H}(f^{-1}(a))$, while the dark shaded region gives the basis of $\mathsf{U}(a, r)$.

Interval Mapping Lemma, this intersection corresponds to a pair of rectangles within the region of $\mathbb{X}_0(f_a)$; see Figure 3.

A point contributes to the well group until r reaches a value at which the pair of rectangles no longer contains the point. For a point $(x, y) \in \mathcal{L}_p[a, a]$, this value of r is $\min\{a - x, y - a\}$, and for $(x, y) \in \mathcal{R}_{p+1}[a, a]$, this value is $\min\{x - a, a - y\}$. The well diagram is the multiset of the values we get from the points in the persistence diagram.

6 Discussion

The main contribution of this paper is a characterization of the well groups of real-valued functions and a recipe for deriving their well diagrams from the extended persistence diagram of the function. These results have ramifi cations in scientifi c visualization, in particular in the selection and display of isosurfaces. We conclude this paper by formulating two directions for further research.

The general problem of well group computation remains wide open. One way to think about this is the following. If we have a mapping $f : \mathbb{X} \to \mathbb{Y}$ from an m-dimensional topological space $\mathbb{X}$ to an n-dimensional topological space $\mathbb{Y}$, and a submanifold $\mathbb{A} \subseteq \mathbb{Y}$ of dimension k, we call the computation of the well groups a variant of the (m, n, k) *problem*. The full version of this paper provides a complete solution for $(m, 1, 0)$ and $(m, 1, 1)$, when $\mathbb{Y} = \mathbb{R}$. In [7], the authors give an algorithm for $(2, 2, 0)$, when $\mathbb{X}$ is an orientable 2-manifold and $\mathbb{Y} = \mathbb{R}^2$. Their algorithm extends to $(m, n, n - m)$. Everything else is as yet unsolved.

The use of well diagrams to provide local measures of robustness for isosurfaces is a promising research direction in scientifi c visualization. From the extended persistence diagram drawn as in Figure 3, we obtain a compact representation of all homology groups and the robustness of their classes. Can this rich representation of information be effectively used to design transfer functions [9, 13] for highlighting important features in 3-dimensional data sets?

References

1. C. L. BAJAJ, V. PASCUCCI AND D. R. SCHIKORE. The contour spectrum. *In* "Proc. 8th IEEE Conf. Visualization, 1997", 167–173.
2. G. CARLSSON, A. COLLINS, L. J. GUIBAS AND Z. ZOMORODIAN. Persistence barcodes for shapes. *Internat. J. Shape Modeling* **11** (2005), 149–187.
3. 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.
4. D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79-103.
5. 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.
6. H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511-533.
7. 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: Theory, Algorithms, and Applications*, eds. V. Pascucci and J. Tierny, Springer-Verlag, to appear.
8. H. EDELSBRUNNER, D. MOROZOV AND A. PATEL. Quantifying transversality by measuring the robustness of intersections. Manuscript, Dept. Comput. Sci., Duke Univ., Durham, North Carolina, 2009.
9. S. FANG, T. BIDDLECOME AND M. TUCERYAN. Image-based transfer function design for data exploration in volume visualization. *In* "Proc. 9th IEEE Conf. Visualization, 1998", 319–326.
10. M. VAN KREFELD, 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.
11. J. R. MUNKRES. *Elements of Algebraic Topology.* Perseus, Cambridge, Massachusetts, 1984.
12. T. S. NEWMAN AND H. YI. A survey of the marching cube algorithm. *Computers and Graphics* **30** (2006), 854–879.
13. C. M. WITTENBRINK, T. MALZBENDER AND M. E. GOSS. Opacity-weighted color interpolation for volume sampling. *In* "Proc. IEEE Proc. Volume Visualization, 1998", 135–142.