

ON ε APPROXIMATIONS OF PERSISTENCE DIAGRAMS

JONATHAN JAQUETTE AND MIROSLAV KRAMÁR

ABSTRACT. Biological and physical systems often exhibit distinct structures at different spatial/temporal scales. Persistent homology is an algebraic tool that provides a mathematical framework for analyzing the multi-scale structures frequently observed in nature. In this paper a theoretical framework for the algorithmic computation of an arbitrarily good approximation of the persistent homology is developed. We study the filtrations generated by sub-level sets of a function $f: X \rightarrow \mathbb{R}$, where X is a CW-complex. In the special case $X = [0, 1]^N$, $N \in \mathbb{N}$, we discuss implementation of the proposed algorithms. We also investigate *a priori* and *a posteriori* bounds of the approximation error introduced by our method.

1. INTRODUCTION

The formation of complex patterns can be encountered in almost all sciences. Many patterns observed in nature are extremely complicated and it is usually impossible to completely understand them using analytical methods. Often a coarser but computationally tractable description is needed. In recent years computational topology [8, 14] has become widely recognized as an important tool for quantifying complex structures. Computational homology has been used to quantitatively study spatio-temporal chaos [9] as well as to analyze complicated flow patterns generated by Rayleigh-Bénard convection [19, 20], time evolution of the isotropically compressed granular material [15] and complex microstructures generated in binary metal alloys through a process called spinodal decomposition [10, 21]. Unfortunately homology groups can be extremely sensitive to small perturbations [17]. Persistent homology [7] overcomes this problem [4] and provides a more powerful tool for studying evolution of complex patterns in the granular media [16, 17] and fluid dynamics [18].

In the applications mentioned above the pattern is given by a sub- or super-level set of a real valued function f . Therefore analyzing the patterns usually involves numerical study based on a suitable discretization. It is important to know when the homology groups of a sub-level set of f can be inferred from a discrete approximation of this set. Probabilistic results, for a function $f: [a, b] \subset \mathbb{R} \rightarrow \mathbb{R}$, are presented in [24]. An alternative approach is to construct a cubical approximation of the sub-level set with the same homology groups [5]. The algorithm presented in [5] can handle functions defined on a unit square and uses a uniform grid on $[0, 1]^2$.

Received by the editor December 4, 2014 and, in revised form, September 15, 2015 and December 27, 2015.

2010 *Mathematics Subject Classification*. Primary 55-04, 55N99.

The first author's research was funded in part by AFOSR Grant FA9550-09-1-0148 and NSF Grant DMS-0915019.

The second author's research was funded in part by NSF Grants DMS-1125174 and DMS-0835621.

If the resolution of a grid is fine enough, then the sub-level set can be approximated by a collection of the grid elements that intersect the sub-level set. This often leads to an unnecessarily large complex. The size of the complex can be reduced by using a randomized version of the algorithm [3]. Using a regular CW-complex to construct the nodal approximation [6] simplifies the homology computations.

In this paper we develop a method for approximating the sub-level sets of a continuous function $f : X \rightarrow \mathbb{R}$, where X is a regular CW-complex. Note that the sub-level set $X_t = f^{-1}(-\infty, t] = \{x \in X : f(x) \leq t\}$ does not have to be a CW-complex. The following definition provides conditions on a CW structure \mathcal{E} of X under which the singular homology groups of the CW-complex

$$\mathcal{X}_t = \bigcup \{e \in \mathcal{E} : f(\text{cl}(e)) \leq t\}$$

are isomorphic to the homology groups $H_*(X_t)$. Throughout the paper, we use the notation “ $f(\text{cl}(e)) \leq t$ ” to mean “ $f(x) \leq t$ for all $x \in \text{cl}(e)$ ”.

Definition 1.1. Let X be a finite regular complex with a CW structure \mathcal{E} . Suppose that $f : X \rightarrow \mathbb{R}$ is a continuous function and $t \in \mathbb{R}$. We say that the CW structure \mathcal{E} is *compatible* with X_t if every cell $e \in \mathcal{E}$ satisfies one of the following conditions:

- (1) $f(\text{cl}(e)) > t$,
- (2) $f(\text{cl}(e)) \leq t$,
- (3) there exists a deformation retraction $h(\mathbf{x}, s) : X_t \cap \text{cl}(e) \times [0, 1] \rightarrow X_t \cap \text{cl}(e)$ of the set $X_t \cap \text{cl}(e)$ onto $X_t \cap \text{bd}(e)$.

According to the following theorem, proved in Section 3, the singular homology groups $H_*(X_t)$ can be obtained by computing the cellular homology groups of the complex \mathcal{X}_t .

Theorem 1.2. *Let X be a finite regular CW-complex. if the CW decomposition \mathcal{E} of X is compatible with X_t , then the map $i_* : H_*(\mathcal{X}_t) \rightarrow H_*(X_t)$, induced by the inclusion $i : \mathcal{X}_t \rightarrow X_t$, is an isomorphism.*

Due to computational complexity of computing homology, it is extremely important that the number of cells in the complex \mathcal{X}_t be as small as possible. To achieve this goal we use a multi-scale CW structure \mathcal{E} . The diameter (size) of the cells in \mathcal{E} varies with local complexity of the pattern. Larger cells are used at places with smaller complexity, while small cells might be necessary at the places where spacial scales of the pattern are small. To produce a multi-scale CW structure we start from a coarse CW structure of X and keep refining the cells that do not satisfy any of the conditions in Definition 1.1. For $X = [0, 1]^2$ the number of cells in \mathcal{X}_t is smaller than the number of cells in the cubical complex used in [5, 6].

To facilitate the refinement process we introduce self similar grids. The size of the grid elements might vary, but all the grid elements of \mathcal{G} are homeomorphic to X . The resolution of the self similar grid \mathcal{G} can be dynamically adjusted at different parts of X . To increase the resolution, each grid element can be refined into a union of smaller complexes homeomorphic to X . We represent a self similar grid by a tree structure. The root of this tree corresponds to the whole complex X . The i -th level nodes represent the complexes obtained by i refinements of X . Finally, the leaves correspond to the grid elements of \mathcal{G} .

The tree representation of \mathcal{G} provides a basis for a memory efficient data structure for storing the complex \mathcal{X}_t and its boundary operator. We stress that the

boundary operator is not stored as a large matrix. Instead it is computed using the boundary operator $\partial: H_*(X) \rightarrow H_{*-1}(X)$ corresponding to the coarse (unrefined) CW structure, the tree structure representing the self similar grid and geometric information about intersections of the grid elements.

While constructing a compatible CW structure for an arbitrary regular CW-complex X might be challenging, it is straightforward when $X = [0, 1]^N$, where $N \in \mathbb{N}$. The coarse CW structure of $[0, 1]^N$ employed in this paper is given by standard decomposition of a cube into vertices, edges and higher-dimensional faces (cells). A self similar grid \mathcal{G} on $[0, 1]^N$ consists of dyadic cubes of different sizes. If every cube (grid element) in \mathcal{G} satisfies the following definition, then every cell in the CW structure \mathcal{E} generated by the grid \mathcal{G} satisfies one of the conditions in Definition 1.1. Moreover, if a cell $e \in \mathcal{E}_{s_i}$ satisfies condition (3) in Definition 1.1, then the function $f \circ h(\mathbf{x}, s)$ can be taken to be non-increasing in s .

Definition 1.3. Let $f \in C^1([0, 1]^N, \mathbb{R})$ and C be an N -dimensional dyadic cube. We say that C is (f, t) -verified if $0 \notin \frac{\partial f}{\partial x_{i_i}}(C)$ for some coordinate vectors x_{i_1}, \dots, x_{i_n} , where $0 \leq n \leq N$, and for every $(N - n)$ -dimensional cell e of the cube C which is orthogonal to the vectors x_{i_1}, \dots, x_{i_n} , either $f(\text{cl}(e)) > t$ or $f(\text{cl}(e)) \leq t$.

A cell e is considered to be orthogonal to a collection of vectors v_1, \dots, v_n if the projection of e onto $\text{span}\{v_1, \dots, v_n\}$ is zero dimensional. The condition in Definition 1.3 can be checked using interval arithmetic. Therefore the homology groups $H_*(X_t)$ can be evaluated on a computer. However, if t is close to a critical value, then using interval arithmetic to guarantee that the cubes surrounding the critical point satisfy Definition 1.3 might require an extremely fine resolution or might not be possible at all. Another problem with this approach is that the topology of the set X_t depends on the choice of the threshold t , and small changes of t can lead to large changes of the topology as demonstrated in [17]. So the particular choice of the threshold t might be hard to justify. Both of these problems can be overcome by using the persistent homology [7, 8].

Persistent homology relates the homology groups of X_t for different values of t . It reduces the function f to a collection of points in the plane. This collection of points is called a *persistence diagram* and is denoted by $\text{PD}(f)$. Each point in the persistence diagram encodes a well defined geometric feature of f . Hence the persistent homology provides a finer description of the pattern than the homology groups of any one sub-level set. The set of all persistence diagrams can be turned into a complete metric space by using a *Bottleneck distance* d_B between the diagrams [4, 22].

The main contribution of this paper is a theoretical framework for constructing an ε approximation of the persistence diagram $\text{PD}(f)$. For $\varepsilon > 0$ and a continuous function $f: X \rightarrow \mathbb{R}$ we define an ε approximate filtration of f as follows.

Definition 1.4. Let X be a finite regular CW-complex and $f: X \rightarrow \mathbb{R}$ a continuous function. Suppose that a sequence of real numbers $\{s_i\}_{i=-1}^{m+1}$ has the following properties:

- (1) $s_{-1} = -\infty$ and $s_0 < \min_{x \in X} f(x)$,
- (2) $s_m > \max_{x \in X} f(x)$ and $s_{m+1} = \infty$,
- (3) $|s_i - s_{i-1}| < \varepsilon$ for $1 \leq i \leq m$.

Then the collection of sub-level sets $\{X_{s_i}\}_{i=-1}^{m+1}$ is called an ε approximate filtration of f .

It follows from the stability results [4] that the persistence diagram PD' of any ε approximate filtration of f is close to $\text{PD}(f)$. More precisely

$$(1.1) \quad d_B(\text{PD}(f), \text{PD}') < \varepsilon.$$

Construction of the persistence diagram PD' requires determining homology groups of the sub-level sets in some ε approximate filtration of f . Freedom in choosing an ε approximate filtration of f allows us to stay away from the critical values. In order to compute the persistence diagram of an ε approximate filtration $\{X_{s_i}\}_{i=-1}^{m+1}$ one might be tempted to replace this filtration by its cellular counterpart $\{\mathcal{X}_{s_i}\}_{i=-1}^{m+1}$. However there is no guarantee that $\mathcal{X}_{s_i} \subseteq \mathcal{X}_{s_{i+1}}$ and so $\{\mathcal{X}_{s_i}\}_{i=-1}^{m+1}$ may not be a filtration. (See Figure 2 for a counterexample.) In order to construct a cellular filtration we have to require that the CW structures $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ are commensurable.

Definition 1.5. Let $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ be a collection of CW structures on X . We say that they are *commensurable* if for every $e, e' \in \mathcal{E}' := \bigcup \mathcal{E}_{s_i}$ such that $e \cap e' \neq \emptyset$ then either $e \subseteq e'$ or $e' \subseteq e$.

The CW structures $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ constructed using a self similar grid are always commensurable. The CW-complexes $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ defined by $\bar{\mathcal{X}}_{s_i} = \bigcap_{j \geq i} \mathcal{X}_{s_j}$ form an ε approximate cellular filtration of f . This filtration has the following property.

Theorem 1.6. Let $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ be an ε approximate cellular filtration of f corresponding to $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$. Suppose that for every cell $e \in \mathcal{E}_{s_i}$ which satisfies condition (3) in Definition 1.1, the function $f \circ h(\mathbf{x}, s)$ is non-increasing in s . Then the persistence diagrams of the filtrations $\{X_{s_i}\}_{i=-1}^{m+1}$ and $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ are equal.

Corollary 1.7. Let $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ be an ε approximate cellular filtration of f . If PD' is the persistence diagram of this filtration, then $d_B(\text{PD}(f), \text{PD}') < \varepsilon$.

The persistence diagram PD' can be computed even if we fail to construct CW structures compatible with X_{s_i} for some of the thresholds s_i . In this case the distance $d_B(\text{PD}(f), \text{PD}')$ can be evaluated *a posteriori*. The required resources for computing the persistence homology can be further reduced if zig-zag persistence [1] is used.

The remainder of this paper is organized as follows. In Section 2 we recall the definition of a CW-complex, homology and persistent homology. Theorem 1.2 is proven in Section 3. In Section 4 we address the inequality (1.1). A precise definition of a self similar grid is given in Section 5. These grids are employed in Section 6 to build a compatible CW structure. A memory efficient data structure for storing the CW-complex \mathcal{X}_t and its boundary operator is presented in Section 7. In the last section we use our method to compute the approximate persistence diagram of a two-dimensional Fourier series with 5 modes and discuss the *a posteriori* error estimates.

2. BACKGROUND

In this section we recall definitions of a CW-complex [12], homology [14] and persistent homology [7]. Notation introduced in this section is used throughout the paper. We start by introducing a regular CW-complex.

Let X be a Hausdorff space. A (finite) regular CW decomposition of X is a (finite) set \mathcal{E} of subspaces of X with the following properties:

- (1) \mathcal{E} is a partition of X , that is, $X = \bigcup_{e \in \mathcal{E}} e$ and $e \neq e' \Rightarrow e \cap e' = \emptyset$.
- (2) Every $e \in \mathcal{E}$ is homeomorphic to some euclidean space $\mathbb{R}^{\dim(e)}$.
- (3) For every n -cell $e \in \mathcal{E}$ there exists a continuous embedding $\Phi_e(B^n, S^{n-1}) \rightarrow (X^{n-1} \cup e, X^{n-1})$ such that $\Phi_e(B^n) = \text{cl}(e)$ where $\text{cl}(e)$ is the closure of e , $B^n = \{x \in \mathbb{R}^n : \|x\| \leq 1\}$ and $S^n = \{x \in B^{n+1} : \|x\| = 1\}$.

The number $\dim(e)$ is well determined and is called the dimension of e . The sets $e \in \mathcal{E}$ which are homeomorphic to \mathbb{R}^n are the n -cells. The boundary of a cell e is defined by $\text{bd}(e) := \text{cl}(e) \setminus e$. The union $X^n = \bigcup_{\dim(e) \leq n} e$ is the n -skeleton of the CW decomposition. The map Φ_e is called a characteristic map for e and $\varphi_e = \Phi_e|_{S^{n-1}} : S^{n-1} \rightarrow X^{n-1}$ an attaching map. The space X is called a regular CW-complex if there exists a regular CW decomposition of X .

Cellular homology is useful for computing the homology groups of CW-complexes. For a regular CW-complex X the n -chains are defined by

$$C_n(X) = H_n(X^n, X^{n-1}).$$

There is a one-to-one correspondence between the basis elements of $C_n(X)$ and the n -cells of X . For every n -dimensional cell e the map Φ_e generates a monomorphism $(\Phi_e)_* : H_n(B^n, S^{n-1}) \rightarrow C_n(X)$. For a fixed generator $a \in H_n(B^n, S^{n-1})$ the set

$$\Lambda^n = \{(\Phi_e)_*(a) : e \in \mathcal{E} \text{ and } \dim(e) = n\}$$

is a basis of $C_n(X)$. For every $\lambda \in \Lambda = \bigcup_{n \in \mathbb{Z}} \Lambda^n$ we denote its corresponding geometric cell by $|\lambda|$.

The boundary operator $\partial_n : C_n(X) \rightarrow C_{n-1}(X)$ is defined by the composition

$$H_n(X^n, X^{n-1}) \xrightarrow{d_n} H_{n-1}(X^{n-1}) \xrightarrow{j_{n-1}} H_{n-1}(X^{n-1}, X^{n-2})$$

where d_n is the connecting homomorphism from the long exact sequence of the pair (X^n, X^{n-1}) and j_{n-1} is the quotient map. In the above chosen basis of $C_n(X)$, the boundary operator $\partial_n : C_n(X) \rightarrow C_{n-1}(X)$ is completely determined by the values on the basis elements Λ^n and can be uniquely expressed as

$$\partial_n(\lambda) = \sum_{\mu \in \Lambda^{n-1}} [\lambda : \mu] \mu.$$

The coefficient $[\lambda : \mu]$ is called an incidence number of the cells. Regularity of the complex X implies that the incidence number is non-zero only for the pairs of cells $\lambda \in \Lambda^n, \tau \in \Lambda^{n-1}$ such that $\text{bd}(|\lambda|) \cap |\tau| \neq \emptyset$. Moreover $[\lambda : \tau] = \pm 1$ in this case.

The cellular complex corresponding to X is defined by $C(X) = \{C_n(X), \partial_n\}_{n \in \mathbb{Z}}$. The singular homology $H_*(X)$ of X is isomorphic to the cellular homology of $C(X)$, which is given by

$$H_*(C(X)) = \frac{\ker(\partial_*(C_*(X)))}{\text{im}(\partial_{*+1}(C_{*+1}(X)))}.$$

We close this section with a short overview of persistent homology, following the presentation given in [7]. Let $f : X \rightarrow \mathbb{R}$ be a function defined on a CW-complex X . We denote its sub-level sets by $X_t = f^{-1}(-\infty, t]$. The function f is called tame if the homology groups of every sub-level set have finite ranks and there are only finitely many values t across which the homology groups are not isomorphic. Let $t_1 < t_2 \dots < t_m$ be the values for which the homology of $X_{t_i - \varepsilon}$ and $X_{t_i + \varepsilon}$ differ for

all $0 < \epsilon \ll 1$. We define an interleaved sequence $\{s_i\}_{i=-1}^{m+1}$ such that $s_{-1} = -\infty$, $s_{m+1} = \infty$ and $s_{i-1} < t_i < s_i$, for $1 \leq i \leq m$. For each $-1 \leq i \leq j \leq m+1$ there is a natural inclusion of the set X_{s_i} into X_{s_j} , and we denote the induced homomorphism between the corresponding homology groups by

$$\mathbf{f}_n^{i,j} : H_n(X_{s_i}) \rightarrow H_n(X_{s_j}).$$

Persistent homology makes use of the above-mentioned maps to compare the homology groups of X_t for different values of t . We say that the homology class $\alpha \in H_n(X_{s_i})$ is *born* at t_i if it does not come from a class in $H_n(X_{s_{i-1}})$, that is, $\alpha \notin \text{im}(\mathbf{f}_n^{i-1,i})$. The class α *dies* at t_j if α is in $\text{im}(\mathbf{f}_n^{i,j-1})$ but not in $\text{im}(\mathbf{f}_n^{i,j})$. We denote the birth-death pair for α by (b_α, d_α) . The n -th persistence diagram $\text{PD}_n(f)$ is defined to be the multiset of the points containing:

- (1) one point for each pair (b_α, d_α) ,
- (2) infinitely many copies of the points (s, s) at the diagonal.

If X has dimension n , then the persistence diagram for the function $f : X \rightarrow \mathbb{R}$ is the collection $\text{PD}(f) = \{\text{PD}_i(f)\}_{i=0}^n$. Actually, persistence diagrams can be defined for continuous functions $f : X \rightarrow \mathbb{R}$, where X is the realization of a simplicial complex as a topological space [2]. In this case the persistence diagrams can contain accumulation points at the diagonal. In other words there might be a sequence of off diagonal points converging to the diagonal. Using the same reasoning as in [2] the same is true if X is a finite regular CW-complex.

Differences between the persistence diagrams can be assessed using a bottleneck distance defined as follows.

Definition 2.1. Let $\text{PD} = \{\text{PD}_i\}_{i=0}^n$ and $\text{PD}' = \{\text{PD}'_i\}_{i=0}^n$ be persistence diagrams. The *bottleneck distance* between PD and PD' is defined to be

$$d_B(\text{PD}, \text{PD}') = \max_i \inf_{\gamma : \text{PD}_i \rightarrow \text{PD}'_i} \sup_{p \in \text{PD}_i} \|p - \gamma(p)\|_\infty,$$

where $\|(a_0, b_0) - (a_1, b_1)\|_\infty := \max\{|a_0 - a_1|, |b_0 - b_1|\}$ and γ ranges over all bijections.

3. HOMOLOGY OF SUB-LEVEL SETS

In this section we present a method for computing the homology of sub-level sets. Let X be a CW-complex and $f : X \rightarrow \mathbb{R}$ a continuous function; we are interested in the homology of the sub-level set $X_t := f^{-1}(-\infty, t]$. Suppose that \mathcal{E} is a CW structure on X ; then we define the *cellular approximation* to be

$$(3.1) \quad \mathcal{X}_t := \bigcup \{e \in \mathcal{E} : f(\text{cl}(e)) \leq t\}.$$

Note that \mathcal{X}_t is a closed subset of X_t , and since \mathcal{X}_t is a finite union of cells, then it is a sub-complex of X . For the function f shown in Figure 1 the set X_0 is homologous to \mathcal{X}_0 . According to Theorem 1.2, the homology groups of X_t and \mathcal{X}_t are isomorphic if the CW structure \mathcal{E} is compatible with X_t . We close this section by proving Theorem 1.2.

Proof. If $H_*(X_t, \mathcal{X}_t) \cong 0$ (using singular homology), then the fact that i_* is an isomorphism follows from the long exact sequence of the pair. We prove that $H_*(X_t, \mathcal{X}_t) \cong 0$ by induction on the dimension of X .

For a zero-dimensional complex X , the set X_t is equal to \mathcal{X}_t and so trivially $H_*(X_t, \mathcal{X}_t) \cong 0$. Now we show that $H_*(X_t, \mathcal{X}_t) \cong 0$ for an n -dimensional complex

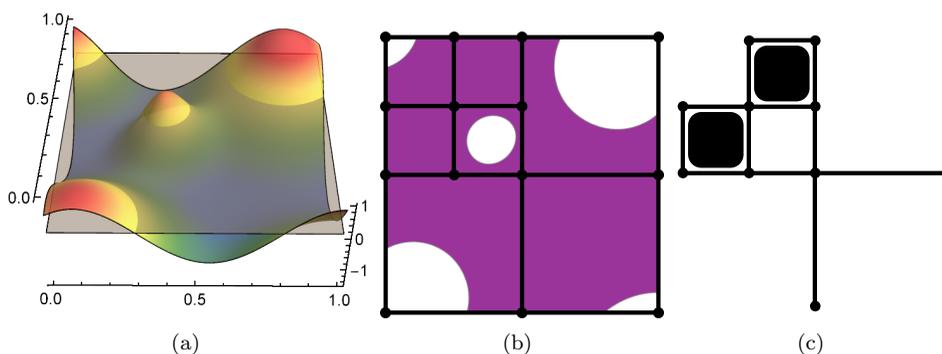


FIGURE 1. (a) A simple function $f : [0, 1]^2 \rightarrow \mathbb{R}$. The sub-level set $X_0 = f^{-1}((-\infty, 0])$ corresponds to the part of $[0, 1]^2$ at which the value of f is below the semitransparent plane. (b) The set X_0 is shaded. The black vertices, edges and squares bounded by the edges define a CW structure \mathcal{E} on $[0, 1]^2$. (c) The set \mathcal{X}_0 corresponding to the decomposition \mathcal{E} .

X under the assumption that the statement is true for any complex with dimension less than n . Let us first suppose that X contains a single n -dimensional cell e . Then $X = X^{n-1} \cup e$. Define \mathcal{X}_t^{n-1} to be the $n-1$ skeleton of \mathcal{X}_t . While X_t is not a CW-complex in general, we define an analog of its $n-1$ skeleton by $X_t^{n-1} := X_t \cap X^{n-1}$. By the inductive hypothesis, the homology groups $H_*(X_t^{n-1}, \mathcal{X}_t^{n-1})$ are trivial.

The CW structure \mathcal{E} is compatible with X_t , so the cell e satisfies one of the conditions in Definition 1.1. If $f(\text{cl}(e)) > t$, then $X_t = X_t^{n-1}$ and $\mathcal{X}_t = \mathcal{X}_t^{n-1}$. Hence we obtain the set equality $H_*(X_t, \mathcal{X}_t) = H_*(X_t^{n-1}, \mathcal{X}_t^{n-1})$, and so by the inductive assumption the relative homology groups stay trivial.

Now suppose that $f(\text{cl}(e)) \leq t$. In this case both $\text{cl}(e) \subseteq X_t$ and $\text{cl}(e) \subseteq \mathcal{X}_t$. In general, we have the equality $(X_t \setminus e, \mathcal{X}_t \setminus e) = (X_t^{n-1}, \mathcal{X}_t^{n-1})$, so it follows from the inductive assumption that $H_*(X_t \setminus e, \mathcal{X}_t \setminus e) \cong 0$. We would like to apply the excision theorem to obtain $H_*(X_t, \mathcal{X}_t) \cong H_*(X_t \setminus e, \mathcal{X}_t \setminus e) \cong 0$. However, the assumptions of the singular excision theorem are not actually satisfied because $\text{cl}(e)$ is not in the interior of \mathcal{X}_t . This problem can be overcome by expanding X_t and \mathcal{X}_t inside e using a small ε neighborhood of $\text{bd}(e)$. Note that for ε small enough the homology groups of X_t and \mathcal{X}_t do not change.

If e satisfies condition (3) in Definition 1.1, then there exists a deformation retraction $h(\mathbf{x}, s)$ of the set $X_t \cap \text{cl}(e)$ onto $X_t \cap \text{bd}(e)$. Note that h is a continuous map on the closed set $X_t \cap \text{cl}(e) \times [0, 1]$, and the identity map on the closed set $X_t \setminus e \times [0, 1]$ is also continuous. Since these two maps agree on the intersection of their domains, we can extend h to a deformation retraction of X_t onto X_t^{n-1} . Existence of h also implies that there exists some $t' > t$ such that $t' \in f(\text{cl}(e))$, hence $e \cap \mathcal{X}_t = \emptyset$ and $\mathcal{X}_t = \mathcal{X}_t^{n-1}$. By induction $H_*(X_t, \mathcal{X}_t) \cong 0$ again.

Finally suppose that X is an arbitrary n -dimensional complex. Let e_1, \dots, e_m be its n -dimensional cells. We repeat the above process m times. At each step we add a cell e_k to $X^n \cup e_1 \cup \dots \cup e_{k-1}$. \square

4. PERSISTENT HOMOLOGY

In order to exactly compute the persistence diagrams $\text{PD}(f)$ of a continuous bounded function, one must identify all thresholds t_i for which the homology of $X_{t_i-\epsilon}$ and $X_{t_i+\epsilon}$ differs for all $0 < \epsilon \ll 1$. If f is a Morse function, then the thresholds t_i are the critical values of f , which one may identify after locating every point where the Jacobian Df vanishes. Such a root finding problem is frequently studied in nonlinear optimization [11]. In applications though, it may not always be feasible to rigorously find every root of Df . In such a situation, it would be desirable to still be able to compute the persistence diagrams of f , even at the cost of losing some accuracy.

In this section we present a framework for computing an approximation of the persistence diagram $\text{PD}(f)$. For a given $\epsilon > 0$ we construct a cellular approximation of the filtration corresponding to f . The persistence diagram PD' of this approximation is ϵ close to $\text{PD}(f)$, that is, $d_B(\text{PD}(f), \text{PD}') < \epsilon$. We start by choosing an approximate filtration (cf. Definition 1.4). The following lemma shows that the bottleneck distance between the persistent diagram for the ϵ approximate filtration of f and the persistence diagram of f is at most ϵ .

Lemma 4.1. *Let X be a CW-complex and $f : X \rightarrow \mathbb{R}$ a continuous bounded function. If PD' is a persistence diagram of any ϵ approximate filtration of f , then $d_B(\text{PD}(f), \text{PD}') < \epsilon$.*

Proof. To prove the lemma we have to show that for every n -th persistence diagram $\text{PD}_n(f)$ there exists a bijection γ from $\text{PD}_n(f)$ to PD'_n such that $\|p - \gamma(p)\|_\infty < \epsilon$ for every $p \in \text{PD}(f)$.

If $p = (b, d)$ and $d - b < \epsilon$, then we define $\gamma(p) = (b, b)$. Clearly $\|p - \gamma(p)\|_\infty < \epsilon$ in this case. The collection of sets $\{X_{s_i}\}_{i=-1}^{m+1}$ is an ϵ filtration. Therefore, if $d - b \geq \epsilon$, then there exist indices i and j such that $s_{i-1} < b \leq s_i \leq s_{j-1} < d \leq s_j$. So the topological feature corresponding to $p = (b, d)$ is not present in the sub-level set $X_{s_{i-1}}$, but it can be found in X_{s_i} . Hence for the ϵ approximate filtration there exists a unique topological feature q , corresponding to p , born at s_i . By the same reasoning the feature q dies at s_j , and we define $\gamma(p) = (s_i, s_j)$. Again the distance $\|p - \gamma(p)\|_\infty < \epsilon$, and the number of points in PD' is smaller than or equal to the number of points in $\text{PD}(f)$. So γ is a bijection. \square

It is appealing to construct a cellular approximate filtration by replacing the sets X_{s_i} with their cellular surrogates \mathcal{X}_{s_i} introduced in the previous section. However the collection $\{\mathcal{X}_{s_i}\}_{i=-1}^{m+1}$ might not be a filtration, as shown in Figure 2. Namely, the set \mathcal{X}_0 is not a subset of \mathcal{X}_1 because of different choices of the compatible CW structures for the sub-level sets. To avoid this problem we construct a common refinement of the CW structures. Our refinement process relies on the fact that the CW structures \mathcal{E}_{s_i} are commensurable (cf. Definition 1.5).

Definition 4.2. Let $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ be a commensurable collection of CW structures on X . Then the set

$$\mathcal{E} := \{e \in \mathcal{E}' : \text{either } e \cap e' = \emptyset \text{ or } e \subseteq e' \text{ for every } e' \in \mathcal{E}'\}$$

is called the *common refinement* of the CW structures $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$.

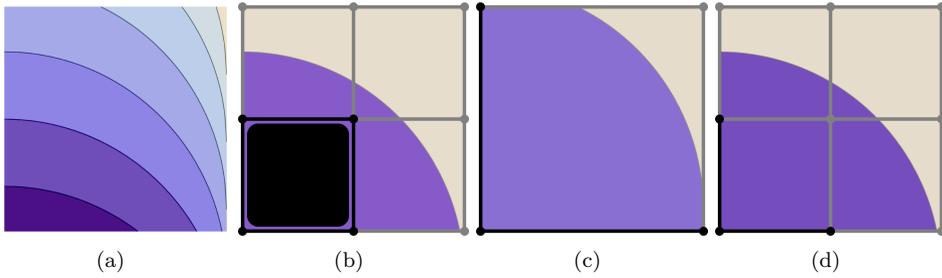


FIGURE 2. (a) A contour plot of a function $f : [0, 1]^2 \rightarrow \mathbb{R}$. (b) The sub-level set X_0 and its cellular approximation \mathcal{X}_0 obtained by the compatible CW structure with four squares (2-dimensional cells). (c) The sub-level set X_1 and its cellular approximation \mathcal{X}_1 obtained by the compatible CW structure with a single 2-dimensional cell. (d) The sub-level set X_0 with the refined approximation $\bar{\mathcal{X}}_0$, which is equal to $\mathcal{X}_0 \cap \mathcal{X}_1$.

The CW structures \mathcal{E}_0 and \mathcal{E}_1 shown in Figure 2 (b-c) are commensurable, and their common refinement is \mathcal{E}_0 . In this paper we will use the common refinement of all the structures $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ to compute the persistent homology. If zig-zag persistence [1] is used, then it suffices to consider the common refinements of consecutive CW structures \mathcal{E}_{s_i} and $\mathcal{E}_{s_{i+1}}$. We stress that the common refinement does not need to be compatible with any of the sub-level sets X_{s_i} . Hence we need to develop a more sophisticated approach for building a cellular representation of the approximate filtration. Before introducing a cellular representation of the filtration $\{X_{s_i}\}_{i=-1}^{m+1}$ let us prove that the common refinement \mathcal{E} is a regular CW structure on X .

Lemma 4.3. *Let \mathcal{E} be a common refinement of the commensurable CW structures $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ on X . Then \mathcal{E} is a CW structure on X .*

Proof. First we show that every point $x \in X$ is contained in a unique cell $e \in \mathcal{E}$. Note that the set $C(x) = \{e \in \mathcal{E}' : x \in e\}$ is not empty. If $e, e' \in C(x)$, then by Definition 1.5 either $e \subseteq e'$ or $e' \subseteq e$. This implies that the cell $e_x := \bigcap \{e \in C(x)\}$ is an element of \mathcal{E}' . Moreover for any $e \in C(x)$ the cell $e_x \subseteq e$ and $e \cap e_x = \emptyset$ if $e \notin C(x)$. Therefore $e_x \in \mathcal{E}$. Suppose by contradiction that e_x is not a unique cell in \mathcal{E} containing x . Then there exist $e, e' \in \mathcal{E}$ such that $x \in e \cap e'$. It follows from Definition 4.2 that $e \subseteq e'$ and $e' \subseteq e$. So $e = e'$.

We showed that \mathcal{E} is a disjoint cellular decomposition of X . Naturally, every cell $e \in \mathcal{E} \subseteq \mathcal{E}'$ is homeomorphic to some euclidean space $\mathbb{R}^{\dim(e)}$. Now we need to define the characteristic maps for the cells. Every cell in $e \in \mathcal{E}$ belongs to at least one \mathcal{E}_{s_i} . We identify the characteristic map for $e \in \mathcal{E}$ with the characteristic map given by one (arbitrary but fixed) CW structure \mathcal{E}_{s_i} . Since every cell in \mathcal{E}' is covered by a collection of smaller cells in \mathcal{E} , the boundary of a cell $e \in \mathcal{E}$ is contained in the union of cells in \mathcal{E} of dimension less than $\dim(e)$. Hence the set of cells \mathcal{E} equipped with these maps is a CW structure on X . \square

Since the common refinement \mathcal{E} does not have to be compatible with any of the sub-level sets, we cannot use it to generate a cellular approximation of the filtration $\{X_{s_i}\}_{i=-1}^{m+1}$. Instead we approximate the sub-level sets X_{s_i} by

$$(4.1) \quad \bar{\mathcal{X}}_{s_i} = \bigcap_{j \geq i} \mathcal{X}_{s_j}.$$

Definition 4.4. Let $\{X_{s_i}\}_{i=-1}^{m+1}$ be an ε approximate filtration of a continuous function f and $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ a collection of commensurable CW structures, where \mathcal{E}_{s_i} is compatible with X_{s_i} for every s_i . Then the collection of sets $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ is called an ε approximate cellular filtration of f corresponding to $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$.

Remark 4.5. Every set $\bar{\mathcal{X}}_{s_i}$ is a sub-complex of X , where the CW structure of X is given by the common refinement \mathcal{E} . Moreover $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ is a filtration, that is,

$$\emptyset = \bar{\mathcal{X}}_{s_{-1}} \subseteq \bar{\mathcal{X}}_{s_0} \subseteq \dots \subseteq \bar{\mathcal{X}}_{s_{m+1}} = X.$$

Our goal is to show that the persistence diagrams of the filtrations $\{X_{s_i}\}_{i=-1}^{m+1}$ and $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ are equal. This is true if the diagram shown in Figure 3, with the maps induced by inclusion, commutes and the maps j_*^i are isomorphisms. It follows from Remark 4.5 and the functorality of homology that the diagram commutes. To prove that j_*^i are isomorphisms we will have to strengthen the definition of a compatible CW structure.

$$\begin{array}{ccccccc} \dots & \longrightarrow & H_*(X_{s_{i-1}}) & \longrightarrow & H_*(X_{s_i}) & \longrightarrow & H_*(X_{s_{i+1}}) & \longrightarrow & \dots \\ & & \uparrow j_*^{i-1} & & \uparrow j_*^i & & \uparrow j_*^{i+1} & & \\ \dots & \longrightarrow & H_*(\bar{\mathcal{X}}_{s_{i-1}}) & \longrightarrow & H_*(\bar{\mathcal{X}}_{s_i}) & \longrightarrow & H_*(\bar{\mathcal{X}}_{s_{i+1}}) & \longrightarrow & \dots \end{array}$$

FIGURE 3. The persistence diagrams of the filtrations $\{X_{s_i}\}_{i=-1}^{m+1}$ and $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ are equal if the above diagram commutes and the maps j_*^i , induced by inclusion, are isomorphisms.

First, we illustrate the idea on a simple example shown in Figure 2. The cells $e \in \mathcal{X}_0$ that are not in $\bar{\mathcal{X}}_0$ are inside the single 2-dimensional cell $e' := (0, 1)^2 \in \mathcal{E}_1$ and additionally $\bar{\mathcal{X}}_0 = \mathcal{X}_0 \cap \text{bd}(e')$. Furthermore $e' \cap \mathcal{X}_1 = \emptyset$ because the image $f(\text{cl}(e'))$ contains both values above and below 1. The fact that the CW structure \mathcal{E}_1 is compatible with X_1 implies that there is a deformation retraction $h(\mathbf{x}, s)$ from $X_1 \cap \text{cl}(e')$ onto $X_1 \cap \text{bd}(e')$. If $f \circ h(\mathbf{x}, s)$ is non-increasing in s , then it can be restricted to a deformation retraction from $X_0 \cap \text{cl}(e')$ onto $X_0 \cap \text{bd}(e')$. By the same argument as in the proof of Theorem 1.2 the inclusion map from $\bar{\mathcal{X}}_0 = \mathcal{X}_0 \cap \text{bd}(e')$ to X_0 induces an isomorphism on the homology level. We now prove Theorem 1.6, which treats the general case.

Proof. The persistence diagrams of the filtrations $\{X_{s_i}\}_{i=-1}^{m+1}$ and $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ are equal if the diagram in Figure 3 commutes and the maps j_* , induced by inclusion, are isomorphisms. We already showed that the diagram commutes. Now we have

to prove that every $j_*^i: H_*(\bar{\mathcal{X}}_{s_i}) \rightarrow H_*(X_{s_i})$ is an isomorphism. According to Theorem 1.2, the inclusion of each \mathcal{X}_{s_i} into X_{s_i} induces an isomorphism. Hence it suffices to show that the inclusion from $\bar{\mathcal{X}}_{s_i}$ into \mathcal{X}_{s_i} also induces an isomorphism in homology.

The set $\bar{\mathcal{X}}_{s_i}$ is a sub-complex of \mathcal{X}_{s_i} , so

$$\bar{\mathcal{X}}_{s_i} = \mathcal{X}_{s_i} \setminus \{e_1 \cup e_2 \cup \dots \cup e_N\},$$

for some cells $e_a \in \mathcal{E}$. We denote $E = \{e_1, e_2, \dots, e_N\}$. First we divide the cells in E into disjoint groups. Then we show that by removing all of the cells in the first group we obtain a CW sub-complex of \mathcal{X}_{s_i} and the map induced by the inclusion of this complex into \mathcal{X}_{s_i} is an isomorphism. We prove that the maps $\bar{j}_*^i: H_*(\bar{\mathcal{X}}_{s_i}) \rightarrow H_*(\mathcal{X}_{s_i})$ are isomorphisms by inductively removing the cells in all the other groups.

Before grouping the cells into disjoint classes let us show that for every cell $e_a \in E$ there exists $e'_a \in \mathcal{E}_{s_i}$ and $e''_a \in \mathcal{E}_{s_j}$, where $j > i$, such that:

- (1) $e_a \subseteq e'_a \subsetneq e''_a$,
- (2) $e''_a \cap \mathcal{X}_{s_j} = \emptyset$.

By the construction of \mathcal{E} , each cell e_a is a subset of a unique cell $e'_a \in \mathcal{E}_{s_i}$ and $e'_a \subseteq \mathcal{X}_{s_i}$. The fact that $e_a \cap \mathcal{X}_{s_i} = \emptyset$ implies $e_a \cap \mathcal{X}_{s_j} = \emptyset$ for some $j > i$. It follows from the construction of the cellular approximation \mathcal{X}_{s_j} that $e'_a \notin \mathcal{E}_{s_j}$. Otherwise we would have $e_a \subseteq e'_a \subseteq \mathcal{X}_{s_j}$. Therefore there exists a cell $e''_a \in \mathcal{E}_{s_j}$ which satisfies the properties (1) and (2).

Let $\{e''_1, \dots, e''_N\}$ be the cells with the properties (1) and (2) corresponding to the cells in E and let $E'' = \{e''_1, \dots, e''_M\} \subseteq \{e''_1, \dots, e''_N\}$ be a set of the maximal cells. That is, a cell e''_i is removed from the set $\{e''_1, \dots, e''_N\}$ if there is another cell e''_j such that $e''_i \subseteq e''_j$. We order the cells in E'' by dimension so that $\dim(e''_i) \geq \dim(e''_j)$ for $i \leq j$. We say that the cells $e_i, e_j \in E$ are in the same group if there exists some l such that both e_i and e_j are subsets of e''_l .

We will remove the cells $e \in E$ from \mathcal{X}_{s_i} in M steps. During the l -th step all of the cells $e \in E$ such that $e \subset e''_l$ are removed. Note that $\bar{\mathcal{X}}_{s_i} = \mathcal{X}_{s_i} \setminus \bigcup_{a=1}^M e''_a$. Due to the ordering of the cells in E'' we get a CW-complex at each step. To show this let $e \in E$ be contained in some $e''_l \in E''$. We show that e is not in the boundary of any cell in $\mathcal{X}_{s_i} \setminus \bigcup_{a=1}^l e''_a$. The cell $e''_l \in \mathcal{E}_{s_j}$ for some $j > i$ and $e''_l \not\subseteq \mathcal{X}_{s_j}$. It follows from the construction of \mathcal{X}_{s_j} that every cell $c \in \mathcal{E}_{s_j}$ such that $e''_l \cap \text{cl}(c) \neq \emptyset$ is not in \mathcal{X}_{s_j} . So every cell in \mathcal{X}_{s_i} whose closure intersects the cell e is contained in some cell $e''_k \in E''$ such that $\dim(e''_k) > \dim(e''_l)$ and was removed before e . Therefore at each step we get a closed subset of the cells from \mathcal{X}_{s_i} which is a CW-complex. Moreover, at each step we can express the complex as a finite union of cells from \mathcal{E}_{s_i} , so it also has a CW structure induced by \mathcal{E}_{s_i} .

First we suppose that $l = 1$ and prove that the map $\ell_*: H_*(\mathcal{X}_{s_i} \setminus e''_1) \rightarrow H_*(\mathcal{X}_{s_i})$, induced by inclusion, is an isomorphism. To do so we apply the five lemma to the following diagram:

$$\begin{array}{ccccccc}
 \longrightarrow & H_*(\mathcal{X}_{s_i} \cap \text{bd}(e''_1)) & \longrightarrow & H_*(\mathcal{X}_{s_i} \setminus e''_1) \oplus H_*(\mathcal{X}_{s_i} \cap \text{cl}(e''_1)) & \longrightarrow & H_*(\mathcal{X}_{s_i}) & \longrightarrow \\
 & \uparrow & & \uparrow & & \uparrow \ell_* & \\
 \longrightarrow & H_*(\mathcal{X}_{s_i} \cap \text{bd}(e''_1)) & \longrightarrow & H_*(\mathcal{X}_{s_i} \setminus e''_1) \oplus H_*(\mathcal{X}_{s_i} \cap \text{bd}(e''_1)) & \longrightarrow & H_*(\mathcal{X}_{s_i} \setminus e''_1) & \longrightarrow
 \end{array}$$

The horizontal maps are given by the cellular Mayer-Vietoris sequences for the CW-complexes given by the middle terms. The vertical maps are induced by inclusion. The first two vertical maps are induced by identity maps and are thereby isomorphisms. To satisfy the assumptions of the five lemma we have to show that the map $k_* : H_*(\mathcal{X}_{s_i} \cap \text{bd}(e''_1)) \rightarrow H_*(\mathcal{X}_{s_i} \cap \text{cl}(e''_1))$ is also an isomorphism.

We recall that the cell $e''_1 \in \mathcal{E}_{s_j}$ and $e''_1 \cap \mathcal{X}_{s_j} = \emptyset$ for some $j > i$. Therefore the image $f(e_1)$ contains points both above and below s_j and there has to exist a deformation retraction $h(\mathbf{x}, s) : X_{s_j} \cap \text{cl}(e''_1) \rightarrow X_{s_j} \cap \text{bd}(e''_1)$. We have assumed that $f \circ h(\mathbf{x}, s)$ is non-increasing in s . So h restricts to a deformation retraction from $X_{s_i} \cap \text{cl}(e''_1)$ onto $X_{s_i} \cap \text{bd}(e''_1)$. Hence the top horizontal map in the diagram below is an isomorphism:

$$\begin{CD} H_*(X_{s_i} \cap \text{bd}(e''_1)) @>\cong>> H_*(X_{s_i} \cap \text{cl}(e''_1)) \\ @V\cong VV @VV\cong V \\ H_*(\mathcal{X}_{s_i} \cap \text{bd}(e''_1)) @>k_*>> H_*(\mathcal{X}_{s_i} \cap \text{cl}(e''_1)) \end{CD}$$

There exist regular CW decompositions of $\text{bd}(e''_1)$ and $\text{cl}(e''_1)$ comprised of cells from \mathcal{E}_{s_i} . It follows from Theorem 1.2 that the vertical arrows in the diagram above are isomorphisms. By the commutativity of the diagram, it follows that the map k_* is an isomorphism. So by the five lemma, it follows that the map $\ell_* : H_*(\mathcal{X}_{s_i} \setminus e''_1) \rightarrow H_*(\mathcal{X}_{s_i} \setminus e''_1)$ is an isomorphism.

Using mathematical induction we suppose that the inclusion map from $\mathcal{X}_{s_i} \setminus \bigcup_{a=1}^l e''_a$ to \mathcal{X}_{s_i} induces an isomorphism on the homology level. To show that this is also true for $l+1$ we repeat the above process replacing the set \mathcal{X}_{s_i} by $\mathcal{X}_{s_i} \setminus \bigcup_{a=1}^l e''_a$ and e''_1 by e''_{l+1} . □

5. CELLULAR COMPLEXES ON GRID

In order to build a cellular approximation \mathcal{X}_t of X_t we need to construct comensurable CW structures on the space X . We do this in an adaptive manner, starting with a coarse CW structure \mathcal{E} on X . To obtain a cellular approximation which is compatible with a given sub-level set, the cells that do not satisfy any of the conditions in Definition 1.1 are repeatedly refined. If the critical values of f are bounded away from t , then this process terminates in a finite number of steps.

To formalize this process we introduce the notion of a self similar grid. Using this type of grid also facilitates a memory efficient method for storing the complex and its boundary operator; see Section 7. At the end of this section we show how to compute the boundary operator for the refined CW structure on X in terms of the incidence numbers for the basis Λ given by \mathcal{E} .

Before defining a self similar grid in full generality we consider the case of a unit cube $[0, 1]^N$. We choose \mathcal{E} to be the standard CW structure on the cube consisting of the vertices, edges, higher-dimensional faces and a single N -dimensional cell. A self similar grid for the unit cube is built out of dyadic cubes defined as follows.

Definition 5.1. Let n be an arbitrary natural number. For every sequence $\{m_i\}_{i=1}^N$, where $0 \leq m_i < 2^n$, we define an N -dimensional dyadic cube as a closed set

$$C := \prod_{i=1}^N \left[\frac{m_i}{2^n}, \frac{m_i + 1}{2^n} \right].$$

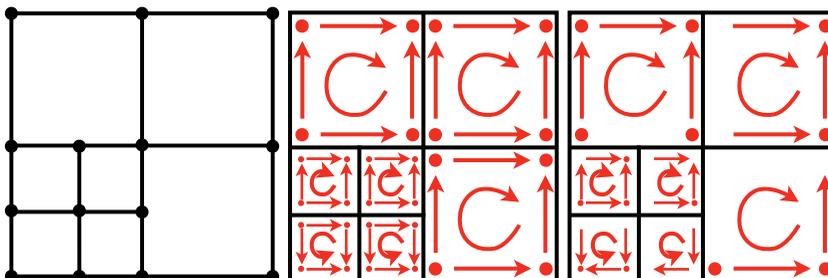


FIGURE 4. (a) A dyadic grid \mathcal{G} on a unit square. (b) CW structures on the dyadic cubes. (c) CW structure $\mathcal{E}_{\mathcal{G}}$ induced by \mathcal{G} .

Every dyadic cube is homeomorphic to $[0, 1]^N$. Compositions of this homeomorphism with the characteristic maps for the cells $e \in \mathcal{E}$ define a CW structure on the dyadic cubes. A dyadic grid is a particular self similar grid consisting of dyadic cubes.

Definition 5.2. Let $\mathcal{G} = \{C_i\}_{i=0}^M$ be a finite collection of N -dimensional dyadic cubes. Then \mathcal{G} is a *dyadic grid* if the following conditions are satisfied:

- (1) $[0, 1]^N = \bigcup_{i=1}^n C_i$,
- (2) $\text{int}(C_i) \cap \text{int}(C_j) = \emptyset$ if $i \neq j$.

Figure 4(a) shows a simple example of a dyadic grid on a unit square. The CW structures of the dyadic cubes are depicted in Figure 4(b). Note that orientation of a cell depends on the precise choice of the homeomorphism $h_i: [0, 1]^N \rightarrow C_i$. In our example we reversed orientation of the cells in two cubes at the bottom left corner.

We generalize the idea of dyadic grids to CW-complexes which are homeomorphic to a closed ball. In the rest of this paper we assume that there is a regular CW structure \mathcal{E} on X with a single N -dimensional cell e such that $X = \text{cl}(e)$, thus implying that X is homeomorphic to an N -dimensional closed ball. The building blocks of a self similar grid are homeomorphic to X , and their CW structure is again induced by these homeomorphisms. They also have to cover X in a similar manner as the dyadic cubes cover the unit cube. A self similar grid on X is a special collection of the building blocks.

Definition 5.3. Let X be an N -dimensional regular CW complex with a single N -dimensional cell e such that $X = \text{cl}(e)$. The set \mathcal{B} is called a *collection of building blocks* of X if it satisfies the following conditions:

- (1) An element $C \in \mathcal{B}$ is a CW-complex homeomorphic to X , and $C \subseteq X$.
- (2) Let $C_1, C_2 \in \mathcal{B}$ and $e_i \subset C_i$ for $i = 1, 2$. If $e_1 \cap e_2 \neq \emptyset$, then either $e_1 \subseteq e_2$ or $e_2 \subseteq e_1$.
- (3) $X \in \mathcal{B}$.

Definition 5.4. Let \mathcal{B} be a collection of building blocks of X . Then $\mathcal{G} \subseteq \mathcal{B}$ is a *self similar grid* on X if the following conditions are satisfied:

- (1) $X = \bigcup_{C \in \mathcal{G}} C$.
- (2) For all $C_1, C_2 \in \mathcal{G}$ the intersection $\text{int}(C_1) \cap \text{int}(C_2) = \emptyset$ if $C_1 \neq C_2$.

Using the same arguments as in the proof of Lemma 4.3, one can show that by removing the redundant cells, a self similar grid \mathcal{G} on X generates a CW structure $\mathcal{E}_{\mathcal{G}}$. For our simple example of a dyadic grid, this structure is shown in Figure 4(c). If we fix a collection of building blocks \mathcal{B} , then a collection of CW structures $\{\mathcal{E}_{\mathcal{G}_i}\}$ generated by the self similar grids $\{\mathcal{G}_i\}$ is always commensurable. The following two lemmas formalize the above statements.

Lemma 5.5. Let $\mathcal{G} = \{C_i\}_{i=0}^M$ be a self similar grid on X and

$$\mathcal{E}'_{\mathcal{G}} := \{e : e \text{ is a cell of } C_i \text{ for some } i = 0, \dots, M\}.$$

Then

$$\mathcal{E}_{\mathcal{G}} := \{e \in \mathcal{E}'_{\mathcal{G}} : \text{either } e \cap e' = \emptyset \text{ or } e \subseteq e' \text{ for every } e' \in \mathcal{E}'_{\mathcal{G}}\}$$

is a CW structure on X generated by \mathcal{G} .

Lemma 5.6. Let \mathcal{B} be a collection of building blocks of X and $\{\mathcal{G}_i\}_{i=-1}^{m+1}$ a collection of self similar grids on X such that $\mathcal{G}_i \subseteq \mathcal{B}$ for $i \in \{-1, \dots, m+1\}$. Then the CW structures $\{\mathcal{E}_{\mathcal{G}_i}\}_{i=-1}^{m+1}$, generated by the grids $\{\mathcal{G}_i\}_{i=-1}^{m+1}$, are commensurable.

In the second part of this section we analyze the boundary operator associated with the CW structure $\mathcal{E}_{\mathcal{G}}$. Our goal is to express this operator using the incidence numbers given by the basis Λ corresponding to \mathcal{E} . We start by building a basis for the chain complex $C_*(X)$ induced by the cellular structure $\mathcal{E}_{\mathcal{G}}$. Every cell $e \in \mathcal{E}_{\mathcal{G}}$ belongs to some $C_i \in \mathcal{G}$. We associate this cell with one of these (arbitrary but fixed) building blocks C_i . Let h_i be the homeomorphism from X to C_i . Then there is a cell $e' \in \mathcal{E}$ such that $e = h_i(e')$. Now the characteristic map for e is given by $h_i \circ \Phi_{e'}$ where $\Phi_{e'}$ is the characteristic map for e' . The map h_i induces a chain isomorphism between the chain complexes, $(h_i)_* : C_*(X) \rightarrow C_*(C_i)$. Fix $\lambda \in \Lambda$ to be the basis element associated with the cell $e' \in \mathcal{E}$ and $\ell_i : C_i \rightarrow X$ the inclusion map. We choose $(\ell_i \circ h_i)_*(\lambda)$ to be a basis element in $C_*(X)$ with the cellular structure $\mathcal{E}_{\mathcal{G}}$, corresponding to the geometric cell e . We denote $(\ell_i \circ h_i)_*(\lambda)$ by (i, λ) and the corresponding basis by $\Lambda_{\mathcal{G}}$. As before $|i, \lambda|$ is the topological cell associated with (i, λ) .

To define a boundary operator in the basis $\Lambda_{\mathcal{G}}$ using the incidence numbers for Λ , we need to understand the relationship between the cells belonging to different building blocks. Namely, we need a method for comparing the orientations of two intersecting cells with the same dimension, and we do so using their local homology groups.

Let us consider cells $e_i = |i, \lambda| \subset C_i$ and $e_j = |j, \mu| \subset C_j$ such that $n = \dim(e_i) = \dim(e_j)$ and $e_i \cap e_j \neq \emptyset$. Without loss of generality we can suppose that $e_i \subseteq e_j$ and select a point $q \in e_i$. Since $\text{bd}(e_i)$ and $\text{bd}(e_j)$ are deformation retracts of $\text{cl}(e_i) \setminus q$ and $\text{cl}(e_j) \setminus q$ respectively, it follows that the maps induced by inclusion,

$$\begin{aligned} r_* &: H_*(\text{cl}(e_i), \text{bd}(e_i)) \rightarrow H_*(\text{cl}(e_i), \text{cl}(e_i) \setminus q), \\ s_* &: H_*(\text{cl}(e_j), \text{bd}(e_j)) \rightarrow H_*(\text{cl}(e_j), \text{cl}(e_j) \setminus q), \end{aligned}$$

are isomorphisms. By excision the map $t_* : H_*(\text{cl}(e_i), \text{cl}(e_i) \setminus q) \rightarrow H_*(\text{cl}(e_j), \text{cl}(e_j) \setminus q)$ induced by inclusion is also an isomorphism. Thereby we obtain the following commutative diagram where all of the maps are isomorphisms:

$$\begin{array}{ccccc}
 & & H_*(\text{cl}(e_i), \text{bd}(e_i)) & \xrightarrow{r_*} & H_*(\text{cl}(e_i), \text{cl}(e_i) \setminus q) \\
 & \nearrow^{(h_i \circ \Phi_{|\lambda|})_*} & & & \downarrow t_* \\
 H_*(B^n, S^{n-1}) & & & & \\
 & \searrow_{(h_j \circ \Phi_{|\mu|})_*} & & & \\
 & & H_*(\text{cl}(e_j), \text{bd}(e_j)) & \xrightarrow{s_*} & H_*(\text{cl}(e_j), \text{cl}(e_j) \setminus q)
 \end{array}$$

Hence for a fixed generator $a \in H_*(B^n, S^{n-1})$ then

$$(5.1) \quad r_*^{-1} \circ t_*^{-1} \circ s_* \circ (h_j \circ \Phi_{|\mu|})_*(a) = \{(i, \lambda) : (j, \mu)\} (h_i \circ \Phi_{|\lambda|})_*(a)$$

where $\{(i, \lambda) : (j, \mu)\}$ is a unit. If we are working with \mathbb{Z} coefficients, then

$$\{(i, \lambda) : (j, \mu)\} = \pm 1,$$

and we say that e_i and e_j have the same orientation if $\{(i, \lambda) : (j, \mu)\} = 1$ and the opposite otherwise.

First we evaluate a boundary for some cells in the complex shown in Figure 4(c). The boundary of the two-dimensional cell in the top left square C consists of five one-dimensional cells. Three of them belong to C . However the bottom edge was replaced by two smaller edges. Orientation of the removed edge agrees with the smaller edges. Hence we can simply replace this edge by those smaller edges. For the bottom right square the problem is a little more complicated. The missing left edge is divided into two smaller edges, but orientation of the bottom smaller edge is opposite. Therefore we need to reverse orientation of this cell by multiplying it by -1 . The following lemma formalizes this procedure.

Lemma 5.7. *Let $(i, \lambda) \in \Lambda_{\mathcal{G}}^n$. Then*

$$(5.2) \quad \partial_n(i, \lambda) = \sum_{\mu \in \Lambda^{n-1}} [\lambda : \mu] \sum_{(j, \tau) \in \mathcal{N}(i, \mu)} \{(i, \mu) : (j, \tau)\} (j, \tau),$$

where $\mathcal{N}(i, \mu) = \{(j, \tau) \in \Lambda_{\mathcal{G}}^{n-1} : |i, \mu| \cap |j, \tau| \neq \emptyset\}$.

Proof. First we consider the boundary of (i, λ) inside C_i equipped with the basis $\{(i, \lambda) : \lambda \in \Lambda\}$. This boundary can be written as $\sum_{\mu \in \Lambda^{n-1}} [\lambda : \mu](i, \mu)$. However this is not necessarily the boundary of (i, λ) in the basis $\Lambda_{\mathcal{G}}$. This is because some elements (i, μ) might be represented by basis elements (j, τ) corresponding to a different building block, or they can be broken up to smaller pieces also from different building blocks. The set $\mathcal{N}(i, \mu) \subset \Lambda_{\mathcal{G}}$ is a representation of (i, μ) in the basis $\Lambda_{\mathcal{G}}$. Geometrically, this means that $\text{cl}(|i, \mu|) = \bigcup_{(j, \tau) \in \mathcal{N}(i, \mu)} \text{cl}(|j, \tau|)$. To obtain $\partial_n(i, \lambda)$ we have to replace each (i, μ) by the elements in $\mathcal{N}(i, \mu)$. In order to accommodate possible differences in orientation, the sign of the incidence number $[\lambda : \mu]$ has to be changed if $(j, \tau) \in \mathcal{N}(i, \mu)$ and (i, μ) have opposite orientation. Therefore each term contained in the second sum is multiplied by $\{(i, \mu), (j, \tau)\}$. \square

6. CONSTRUCTION OF COMPATIBLE CW STRUCTURES ON A UNIT CUBE

Until now the results presented in this paper have been mostly theoretical. In this section we restrict our attention to functions $f \in C^1([0, 1]^N, \mathbb{R})$ and present

an algorithm for constructing a CW structure compatible with the sub-level set $X_t = f^{-1}((-\infty, t])$. It is complicated to check the conditions in Definition 1.1 algorithmically, especially condition (3). Instead we introduce analytical conditions for dyadic cubes that force their cells to satisfy the conditions in Definition 1.1. These conditions can be checked using interval arithmetic, and the algorithm can be carried out by a computer. We divide this section into two parts. In the first part we present the conditions posed on the dyadic cubes. The second part describes the algorithm for constructing compatible CW structures on a unit cube.

6.1. Verified cubes. The dyadic grid \mathcal{G} is said to be (f, t) -verified if every dyadic cube $C \in \mathcal{G}$ is (f, t) -verified (cf. Definition 1.3). We will show that if \mathcal{G} is (f, t) -verified, then the associated CW-complex $\mathcal{E}_{\mathcal{G}}$ is compatible with X_t . Note that every cell in $\mathcal{E}_{\mathcal{G}}$ belongs to at least one dyadic cube $C \in \mathcal{G}$. Hence it suffices to show that every cell in an (f, t) -verified cube $C \in \mathcal{G}$ satisfies one of the three conditions in Definition 1.1. The following lemma allows us to prove this for just the top-dimensional cell.

Lemma 6.1. *If a cube C is (f, t) -verified, then for every cell $e \subseteq C$, the cube $\text{cl}(e)$ is (f, t) -verified.*

Proof. Fix vectors x_{i_1}, \dots, x_{i_n} such that $0 \notin \frac{\partial f}{\partial x_{i_i}}(C)$ for all $i \leq n$, with $0 \leq n \leq N$. Select any cell $e \subseteq C$. There exists a $k \leq n$ such that if we reorder the sequence $\{i_i\}$, then $0 \notin \frac{\partial f}{\partial x_{i_i}}(\text{cl}(e))$ for all $i \leq k$ and $\text{cl}(e)$ is orthogonal to x_{i_i} for any $i > k$. Select any $(\dim(e) - k)$ -dimensional cell $e' \subseteq \text{cl}(e)$ such that e' is orthogonal to the vectors x_{i_1}, \dots, x_{i_k} . Since C is a cube, there exists an $N - n$ cell $e'' \subseteq C$ such that $e' \subseteq e''$ and e'' is orthogonal to the vectors x_{i_1}, \dots, x_{i_n} . Since C is (f, t) -verified, either $f(\text{cl}(e'')) > t$ or $f(\text{cl}(e'')) \leq t$. Thereby either $f(\text{cl}(e')) > t$ or $f(\text{cl}(e')) \leq t$, thus showing that $\text{cl}(e)$ is (f, t) -verified. \square

There are three special cases of a cube C being (f, t) -verified corresponding to the cases when $n = 0$ and $n = N$:

- (1) $f(C) > t$,
- (2) $f(C) \leq t$,
- (3) $0 \notin \frac{\partial f}{\partial x_{i_i}}(C)$ for $1 \leq i \leq N$.

If an (f, t) -verified cube C satisfies one of conditions (1) or (2) above, then all of its cells satisfy condition (1) or (2) in Definition 1.1 respectively. On the other hand, if a cube C satisfies neither condition (1) nor (2) above, then we must construct a deformation retraction required by condition (3) in Definition 1.1. Using rescaling we can further restrict our attention to a unit cube.

Before constructing the deformation retraction in full generality, we demonstrate the basic ideas in two examples shown in Figure 5. The unit cube depicted in Figure 5(a) satisfies condition (3) above. In this case both partial derivatives are negative. For $\bar{\mathbf{x}} \in \text{int}([0, 1]^2 \setminus X_t)$ we define a vector field $\xi(\mathbf{x}) = \mathbf{x} - \bar{\mathbf{x}}$. The set X_t can be retracted onto $X_t \cap \partial[0, 1]^2$ along the flow lines, indicated by the arrows in Figure 5(a), of the flow generated by $\xi(\mathbf{x})$. To ensure that the function $f \circ h(\mathbf{x}, s)$ is non-increasing in s we have to choose $\bar{\mathbf{x}}$ sufficiently close to the origin.

For the example shown in Figure 5(b) the partial derivative $\frac{\partial f}{\partial x_1}([0, 1]^2) < 0$ and $\frac{\partial f}{\partial x_2}$ is both positive and negative on $[0, 1]^2$. In this case the cells $e_1 = \{0\} \times (0, 1)$ and $e_2 = \{1\} \times (0, 1)$ are orthogonal to the coordinate vector x_1 and satisfy $f(\text{cl}(e_1)) > t$

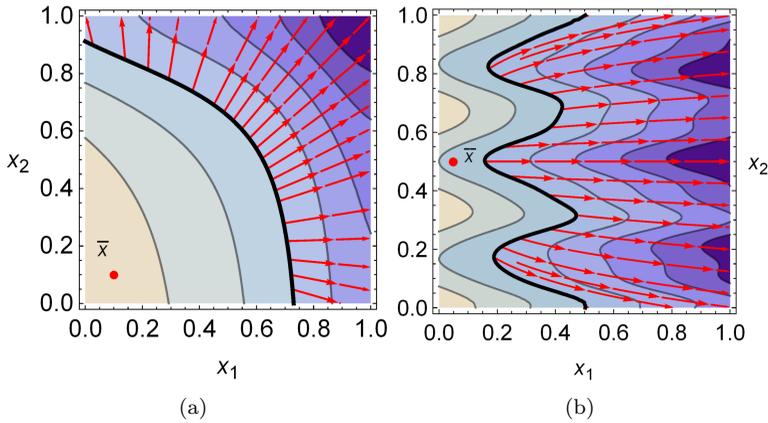


FIGURE 5. (a) A deformation retract of a sub-level set to the boundary when the cube satisfies condition (3). (b) A deformation retract of a sub-level set to the boundary for a generic (f, t) -verified cube.

and $f(\text{cl}(e_2)) \leq t$. Hence cube $[0, 1]^2$ is (f, t) -verified. In this case we define the vector field $\xi(x_1, x_2) = (x_1 - \bar{x}_1, \varepsilon(x_2 - \frac{1}{2}))$ for some point $(\bar{x}_1, \frac{1}{2}) \in \text{int}(X_t \setminus [0, 1]^2)$ and ε small enough. Again we retract the set X_t along the flow lines indicated by the arrows. The first coordinate of the vector field ξ is similar to the one used in the previous example. Part of the set $X_t \cap \partial[0, 1]^2$ is contained inside the horizontal edges of the cube. The second coordinate of the vector field ξ ensures that some flow lines reach this part of the set $X_t \cap \partial[0, 1]^2$. The following two lemmas formalize the above discussion.

Lemma 6.2. *If $0 \notin \frac{\partial f}{\partial x_i}([0, 1]^N)$ for $1 \leq i \leq N$ and f passes through t , then there exists a deformation retraction h from X_t onto $X_t \cap \partial[0, 1]^N$. Moreover $f \circ h(\mathbf{x}, s)$ is non-increasing in s .*

Proof. Without loss of generality we can suppose that $\frac{\partial f}{\partial x_i}(\mathbf{x}) < 0$ for all $\mathbf{x} \in [0, 1]^N$ and $1 \leq i \leq N$. Thereby the value of f at the origin has to be greater than t . To prove the existence of a deformation retraction $h : X_t \rightarrow X_t \cap \partial[0, 1]^N$ we use the Ważewski principal. We will show that X_t is a Ważewski set for a flow φ corresponding to an appropriate linear vector field $\xi(\mathbf{x})$ on \mathbb{R}^N .

The vector field is given by $\xi(\mathbf{x}) = \mathbf{x} - \bar{\mathbf{x}}$ where $\bar{\mathbf{x}} \in \text{int}([0, 1]^N \setminus X_t)$ such that $(\mathbf{x} - \bar{\mathbf{x}}) \cdot \nabla f(\mathbf{x}) < 0$ for all $\mathbf{x} \in X_t$. To show that there exists such $\bar{\mathbf{x}}$ we analyze the continuous function $\mathbf{x} \cdot \nabla f(\mathbf{x})$. For every $\mathbf{x} \in [0, 1]^N$ the function $\mathbf{x} \cdot \nabla f(\mathbf{x}) \leq 0$ and it attains zero only at the origin.

So there exists $\varepsilon > 0$ such that $\mathbf{x} \cdot \nabla f(\mathbf{x}) \leq -\varepsilon$ on X_t . Moreover for $\delta > 0$ sufficiently small $\mathbf{x} \cdot \nabla f(\mathbf{x}) > \frac{-\varepsilon}{2}$ for all $\mathbf{x} \in B_\delta := \{\mathbf{y} \in [0, 1]^N : \|\mathbf{y}\| < \delta\}$. We choose $\bar{\mathbf{x}}$ to be an arbitrary but fixed point in $B_\delta \cap \text{int}([0, 1]^N)$. Clearly $\bar{\mathbf{x}} \notin X_t$ and $(\mathbf{x} - \bar{\mathbf{x}}) \cdot \nabla f(\mathbf{x}) < 0$ for all $\mathbf{x} \in X_t$.

Let φ be a flow given by the vector field $\xi(\mathbf{x})$. We define the exit set and the immediate exit set for X_t as follows:

$$\begin{aligned} W^0 &= \{ \mathbf{x} \in X_t : \exists s > 0 \text{ such that } \varphi(\mathbf{x}, s) \notin X_t \}, \\ W^- &= \{ \mathbf{x} \in X_t : \forall s > 0 \text{ the set } \varphi(\mathbf{x}, [0, s]) \not\subset X_t \}. \end{aligned}$$

The unstable fixed point of the linear system $\mathbf{x}' = \xi(\mathbf{x})$ is not in X_t . Therefore all the orbits have to leave X_t and $W^0 = X_t$. By the chain rule $\frac{d}{ds} f \circ \varphi(\mathbf{x}, s) = (\mathbf{x} - \bar{\mathbf{x}}) \cdot \nabla f(\mathbf{x})$, which is negative for all $\mathbf{x} \in X_t$, and $f \circ \varphi(\mathbf{x}, s)$ is strictly decreasing with respect to s for all $\mathbf{x} \in X_t$. Hence for every $\mathbf{x} \in X_t \cap (0, 1)^N$ there exists $s_{\mathbf{x}} > 0$ such that $\varphi(\mathbf{x}, [0, s_{\mathbf{x}}]) \subseteq X_t$ and the immediate exit set satisfies $W^- \subseteq X_t \cap \partial[0, 1]^N$. The point $\bar{\mathbf{x}}$ is in $\text{int}([0, 1]^N)$. Thus for every $\mathbf{x} \in \partial[0, 1]^N$ the vector field $\xi(\mathbf{x})$ points out of the unit cube and $\varphi(\mathbf{x}, s) \notin [0, 1]^N$ for $s > 0$. Thereby $W^- = X_t \cap \partial[0, 1]^N$.

The sets X_t and $X_t \cap \partial[0, 1]^N$ are closed, so it follows that X_t is a Ważewski set. By the Ważewski principal there exists a deformation retraction $h(\mathbf{x}, s)$ from $W^0 = X_t$ onto $W^- = X_t \cap \partial[0, 1]^N$. The set X_t is retracted along the flow lines of φ . So $f \circ h(\mathbf{x}, s)$ is non-increasing in s . \square

Lemma 6.3. *If $[0, 1]^N$ is (f, t) -verified and f passes through t , then there exists a deformation retraction h from X_t onto $X_t \cap \partial[0, 1]^N$. Moreover $f \circ h(\mathbf{x}, s)$ is non-increasing in s .*

Proof. Without loss of generality we can suppose that $\frac{\partial f}{\partial x_i}(\mathbf{x}) < 0$ for all $\mathbf{x} \in [0, 1]^N$ and $1 \leq i \leq k$. We denote by π_A projection on the first k coordinates and by π_B projection on the last $N - k$ coordinates. The cube $B = \pi_B([0, 1]^N)$ is the closure of a cell in $[0, 1]^N$ which is orthogonal to the vectors x_1, \dots, x_k , so either $f(B) > t$ or $f(B) \leq t$. Since $t \in f([0, 1]^N)$ and $\frac{\partial f}{\partial x_i}([0, 1]^N) < 0$ for $1 \leq i \leq k$ it follows that $f(B) > t$.

We apply the Ważewski principal to a linear flow φ radiating from some point $\bar{\mathbf{x}}$. By the same argument as in the previous proof one can show that for some $\varepsilon > 0$ there exists $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_k, 1/2, \dots, 1/2) \in \text{int}([0, 1]^N \setminus X_t)$ such that $\pi_A(\bar{\mathbf{x}} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) < -\varepsilon$ for all $\mathbf{x} \in X_t$. We now define the vector field

$$\xi(\mathbf{x}) = \pi_A(\mathbf{x} - \bar{\mathbf{x}}) + \frac{\varepsilon}{1 + \eta} \pi_B(\mathbf{x} - \bar{\mathbf{x}})$$

where $\eta = \sup_{\mathbf{x} \in X_t} |\pi_B(\mathbf{x} - \bar{\mathbf{x}}) \cdot \nabla f(\mathbf{x})|$. The value η has to be finite because it is supremum of a continuous function on the compact set X_t . Define φ to be the flow generated by the vector field ξ .

The function f is non-increasing along the flow lines of φ as corroborated by the following calculation:

$$\begin{aligned} \frac{d}{dt} f \circ \varphi(\mathbf{x}, s) &= \left(\pi_A(\mathbf{x} - \bar{\mathbf{x}}) + \frac{\varepsilon}{1 + \eta} \pi_B(\mathbf{x} - \bar{\mathbf{x}}) \right) \cdot \nabla f(\varphi(\mathbf{x}, s)) \\ &< -\varepsilon + \varepsilon \frac{\eta}{1 + \eta} < 0. \end{aligned}$$

Applying the Ważewski principal to the same sets W^0 and W^- as in the proof of the previous lemma guarantees that there exists a deformation retraction h from X_t onto $X_t \cap \partial[0, 1]^N$ and $f \circ h(\mathbf{x}, s)$ is non-increasing in s . \square

Finally we are in a position to show that an (f, t) -verified grid \mathcal{G} induces a CW structure $\mathcal{E}_{\mathcal{G}}$ compatible with X_t .

Theorem 6.4. *Let $f \in C^1([0, 1]^N, \mathbb{R})$. If the dyadic grid \mathcal{G} on $[0, 1]^N$ is (f, t) -verified, then the CW structure $\mathcal{E}_{\mathcal{G}}$ is compatible with X_t .*

Proof. If e is a cell in $\mathcal{E}_{\mathcal{G}}$, then there exists some (f, t) -verified cube $C \in \mathcal{G}$ which contains e . By Lemma 6.1 it follows that $\text{cl}(e)$ is (f, t) -verified. If either $f(\text{cl}(e)) > t$ or $f(\text{cl}(e)) \leq t$, then $\text{cl}(e)$ will satisfy condition (1) or (2) in Definition 1.1, respectively. Otherwise f passes through t , and so by Lemma 6.3 there exists a deformation retraction h of the set $X_t \cap \text{cl}(e)$ onto $X_t \cap \text{bd}(e)$. Moreover the function $f \circ h(\mathbf{x}, s)$ is non-increasing in s . Thereby $\mathcal{E}_{\mathcal{G}}$ is compatible with X_t . \square

Remark 6.5. The fact that the function $f \circ h(\mathbf{x}, s)$ is non-increasing in s implies that the sets \mathcal{X}_t obtained using the dyadic grid can be used to produce a cellular filtration (see Theorem 1.6).

6.2. Algorithm for constructing a verified grid. Recall that a dyadic grid \mathcal{G} is (f, t) -verified if every cube $C \in \mathcal{G}$ is (f, t) -verified. To construct a verified grid we start with a grid \mathcal{G} containing a single cube $[0, 1]^N$. Then we iteratively refine the cubes $C \in \mathcal{G}$ that are not (f, t) -verified. If C is not (f, t) -verified, then we replace it by 2^N dyadic cubes inside C whose edge length equals half of the edge length of C . We call these smaller cubes the offspring of C . If the following algorithm terminates, then the returned grid \mathcal{G} is (f, t) -verified.

```

 $\mathcal{G} = \text{ConstructVerifiedGrid}(f, t)$ 
 $\mathcal{G} = \{[0, 1]^N\}$ 
while there exists a cube  $C \in \mathcal{G}$  which is not  $(f, t)$ -verified do
    Remove  $C$  from  $\mathcal{G}$ 
    Add the offspring of  $C$  to  $\mathcal{G}$ 
end while
return  $\mathcal{G}$ 
    
```

7. DATA STRUCTURES AND ALGORITHMS

In order to compute the homology of the sub-level set $X_t \subseteq X$ we need to develop data structures for storing the basis $C_*(\mathcal{X}_t)$ and the boundary operator. In Section 7.1 we introduce a representation of the basis $\Lambda_{\mathcal{G}}$ associated with the grid \mathcal{G} . A special structure of self similar grid is exploited to devise a memory efficient data structure. We do not store the boundary operator for the basis $\Lambda_{\mathcal{G}}$, which is usually a large (sparse) matrix and requires a lot of memory to be stored. Only the incidence numbers for the coarse CW structure $\mathcal{E}_{\mathcal{G}}$ need to be stored. In the case of a unit cube the incidence numbers can be computed, and the memory requirements are further reduced.

In Section 7.2, the boundary operator $\partial_n: \Lambda_{\mathcal{G}}^n \rightarrow \Lambda_{\mathcal{G}}^{n-1}$ is computed using formula (5.2). Evaluating the terms in (5.2) involves manipulating geometric objects. This might be complicated for a general self similar grid. However for a dyadic grid it can be performed using a special representation of the cells, as explained in Section 7.3.

7.1. Representation of $\Lambda_{\mathcal{G}}$. A collection of building blocks \mathcal{B} , given by Definition 5.3, can be partially ordered by inclusion. For $C_i, C_j \in \mathcal{B}$ we say that C_i precedes C_j and write $C_i \prec C_j$ if $C_j \subseteq C_i$. The fact that $X \prec C$ for all $C \in \mathcal{B}$ implies that the graph representing the partial ordering is a tree $T_{\mathcal{B}}$ with the root

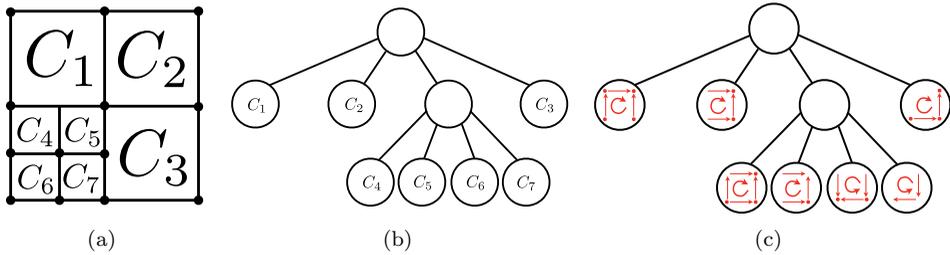


FIGURE 6. (a) A dyadic cubical grid \mathcal{G} with seven building blocks. (b) Tree representation of \mathcal{G} . A building block $C_i \in \mathcal{G}$ corresponds to the leaf with the same label. (c) Basis $\Lambda_{\mathcal{G}}$ chosen in Figure 4(c) is stored in the leaves of the tree corresponding to \mathcal{G} .

corresponding to the set X . Therefore a self similar grid $\mathcal{G} = \{C_i\}_{i=0}^M$ can be represented by a subtree $T_{\mathcal{G}}$ of $T_{\mathcal{B}}$ containing the vertices corresponding to the building blocks $C_i \in \mathcal{G}$ and their predecessors.

Figure 6 shows an example of a self similar grid with seven building blocks and its tree representation $T_{\mathcal{G}}$. There is a one-to-one correspondence between the building blocks $C_i \in \mathcal{G}$ and the leaves of the tree $T_{\mathcal{G}}$. The nodes of the tree correspond to the building blocks that were refined in the process of creating the grid. In our example the root corresponds to the unit square and the unlabeled node in the middle layer to the subdivided cube at the bottom left corner.

Definition 7.1. For every node (leaf) \mathbf{s} in $T_{\mathcal{G}}$ we denote the corresponding building block by $B(\mathbf{s})$. If \mathbf{s} is not a leaf, then $B(\mathbf{s}) \notin \mathcal{G}$. For a leaf \mathbf{s} we define $\Lambda^n(\mathbf{s}) := \{(i, \lambda) \in \Lambda_{\mathcal{G}}^n \text{ stored in the leaf } \mathbf{s}\}$.

The basis $\Lambda_{\mathcal{G}}$ is stored in the leaves of the tree $T_{\mathcal{G}}$. Figure 6(c) shows how the basis of the CW structure chosen in Figure 4(c) is stored. In general, the leaf corresponding to the building block $C_i \in \mathcal{G}$ stores the basis elements $(i, \lambda) \in \Lambda_{\mathcal{G}}$.

7.2. Boundary operator. Evaluation of formula (5.2) requires construction of the set $\mathcal{N}(i, \mu)$. In this section we present a tree traversing algorithm for constructing $\mathcal{N}(i, \mu)$. The algorithm consists of two parts. In the first part, the algorithm moves up the tree. It stops at the node \mathbf{s} such that $|i, \mu| \subseteq \text{int}(B(\mathbf{s}))$. We will prove that the elements of $\mathcal{N}(i, \mu)$ are stored in the leaves that can be reached by descending down from the node \mathbf{s} . In the second part, the algorithm descends to the leaves that can contain elements of $\mathcal{N}(i, \mu)$ and retrieves them. Before stating the algorithm we recall some basic terminology.

Definition 7.2. Let T be a tree. Suppose that \mathbf{s}, \mathbf{s}' are two nodes of T connected by an edge. We say that \mathbf{s} is a predecessor (child) of \mathbf{s}' if the path from \mathbf{s} to the root is shorter (longer) than the path from \mathbf{s}' to the root.

$\mathcal{N} = \text{ConstructN}(i, \mu)$

Let \mathbf{s} be a leaf corresponding to $C_i \in \mathcal{G}$
while $|i, \mu| \not\subseteq \text{int}(B(\mathbf{s}))$ and \mathbf{s} is not the root **do**
 Set \mathbf{s} to the predecessor of \mathbf{s}
end while

```

Put  $\mathbf{s}$  to the stack  $\mathcal{H}$ 
while  $\mathcal{H} \neq \emptyset$  do
  Remove one element  $\mathbf{s}$  from  $\mathcal{H}$ 
  if  $\mathbf{s}$  is a leaf then
    for every  $(j, \tau) \in \Lambda^{\dim(\mu)}(\mathbf{s})$  such that  $|i, \mu| \cap |j, \tau| \neq \emptyset$  do
      Put  $(j, \tau)$  to  $\mathcal{N}(i, \mu)$ 
    end for
  else
    for every child  $\mathbf{s}'$  of  $\mathbf{s}$  such that  $|i, \mu| \cap B(\mathbf{s}') \neq \emptyset$  do
      Add  $\mathbf{s}'$  to  $\mathcal{H}$ 
    end for
  end if
end while
return  $\mathcal{N}$ 

```

The following lemma guarantees that the set \mathcal{N} retrieved by $\text{ConstructN}(i, \mu)$ is equal to the set $\mathcal{N}(i, \mu)$ given by (5.2).

Lemma 7.3. *For every basis element $(i, \mu) \in \Lambda_{\mathcal{G}}$ the algorithm $\text{ConstructN}(i, \mu)$ returns the set $\mathcal{N}(i, \mu)$.*

Proof. Let \mathbf{s} be the node reached at the end of the ascending part (first **while** cycle) of the algorithm. We show that every leaf containing an element of $\mathcal{N}(i, \mu)$ can be reached by descending from the node \mathbf{s} . If \mathbf{s} is the root, then this is trivial. If \mathbf{s} is not the root, then $|i, \mu| \subseteq \text{int}(B(\mathbf{s}))$. For every $(j, \tau) \in \mathcal{N}(i, \mu)$ the cell $|j, \tau| \subseteq |i, \mu| \subseteq \text{int}(B(\mathbf{s}))$. So $\text{int}(C_j) \cap \text{int}(B(\mathbf{s})) \neq \emptyset$ for every building block $C_j \in \mathcal{G}$, such that $|j, \tau| \subseteq C_j$. The interior of every building block is formed by a single N -dimensional cell. Hence by condition (2) of Definition 5.3 either $\text{int}(C_j) \subseteq \text{int}(B(\mathbf{s}))$ or $\text{int}(B(\mathbf{s})) \subseteq \text{int}(C_j)$. If $\text{int}(B(\mathbf{s})) \subseteq \text{int}(C_j)$, then there exists a building block $C_k \in \mathcal{G}$, corresponding to some leaf below \mathbf{s} , such that $\text{int}(C_j) \cap \text{int}(C_k) \neq \emptyset$. This violates condition (2) of Definition 5.4. Therefore $C_j \subseteq B(\mathbf{s})$ and the leaf containing (j, τ) can be reached by descending from the node \mathbf{s} .

In the descending part, the algorithm reaches every leaf \mathbf{s} corresponding to a building block C_j that intersects $|i, \mu|$. All the elements of $\Lambda^{\dim(\mu)}(\mathbf{s})$ which intersect $|i, \mu|$ are added to \mathcal{N} . Hence $\mathcal{N} = \mathcal{N}(i, \mu)$. \square

7.3. Dyadic grid. In this section we deal with implementation of the algorithm for a dyadic grid. First we label the dyadic cubes. Our labeling schema is motivated by the refinement process applied to the unit cube. During a single refinement the cube $[0, 1]^N$ is divided into 2^N cubes with edge lengths equal to $1/2$. Each of these cubes contains exactly one vertex of the original cube. Let C be a cube containing the vertex $(s^0, s^1, \dots, s^{N-1})$. The coordinates of the vertex can be interpreted as a binary number. The number $s = \sum_{n=0}^{N-1} s^n 2^n$ is the index of the cube C . Note that the index contains complete information about the position of the cube. Coordinates of the smallest vertex (in the lexicographical order) of C are given by $x_i = \frac{s^i}{2}$; see Figure 7.

Further refinement of the cube C , considered above, produces another collection of 2^N dyadic cubes. Each of these cubes contains exactly one vertex of C . Again, the position of each cube inside C is encoded by a binary number with N digits. Therefore every cube obtained by two consecutive refinements of $[0, 1]^N$ is indexed

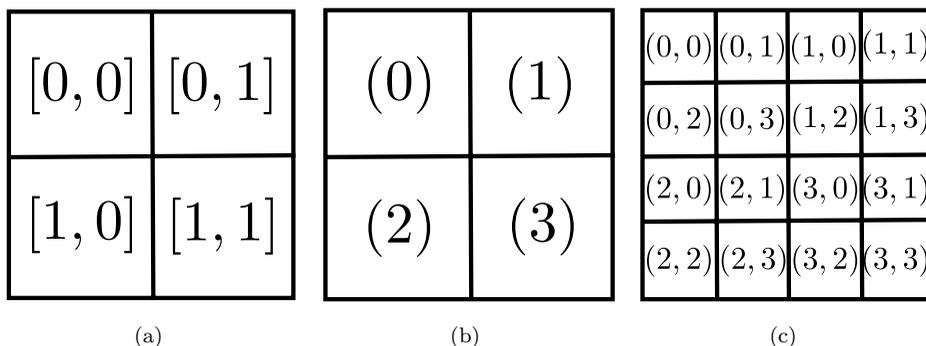


FIGURE 7. (a) Binary, (b) decimal indices for the dyadic cubes obtained by one subdivision of a unit square. (c) Binary indices for the dyadic cubes obtained by two subdivisions of a unit square.

by a sequence (s, s') ; see Figure 7(c). In general a dyadic cube obtained by n subdivisions is uniquely indexed by $\mathbf{s} = (s_1, s_2, \dots, s_n)$ where $s_i \in \{0, 1, \dots, 2^N - 1\}$. The position of the cube corresponding to \mathbf{s} is given by

$$(7.1) \quad C_{\mathbf{s}} = \prod_{i=0}^{N-1} \left[\sum_{j=1}^n 2^{-j} s_j^i, \sum_{j=1}^n 2^{-j} (1 + s_j^i) \right],$$

where s_j^i is the i -th binary digit of s_j .

We choose a CW structure for the unit cube to be the standard cellular structure on $[0, 1]^N$. Every cell e is a product of N generalized intervals and can be uniquely indexed by $\lambda = (\lambda^0, \lambda^1, \dots, \lambda^{N-1}) \in \mathbb{Z}_3^N$. In particular the cell corresponding to the index λ is $|\lambda| = I_0 \times I_1 \times \dots \times I_{N-1}$ where

$$I_i = \begin{cases} 0 & \text{if } \lambda^i = 0, \\ 1 & \text{if } \lambda^i = 1, \\ (0, 1) & \text{if } \lambda^i = 2. \end{cases}$$

If the cell e' is a translation of the cell e , then $\Phi_{e'} = T \circ \Phi_e$, where T is the translation map. The characteristic map of the cell e whose closure contains the origin can be chosen arbitrarily. The orientation of the cells e' that do not contain the origin in their closure is given by $\Phi_{e'}$.

To define a CW structure on $C_{\mathbf{s}}$ we use the homeomorphism $h_{\mathbf{s}}: [0, 1]^N \rightarrow C_{\mathbf{s}}$ defined coordinate-wise by

$$(7.2) \quad h_i(\mathbf{x}) = 2^{-n} x_i + \sum_{j=1}^n 2^{-j} s_j^i,$$

where $h_{\mathbf{s}} = (h_0(\mathbf{x}), h_1(\mathbf{x}), \dots, h_{N-1}(\mathbf{x}))$. Our choice of orientation and homeomorphisms $h_{\mathbf{s}}$ guarantees that two intersecting cells have the same orientation. So formula (5.2) reduces to

$$\partial_n(i, \lambda) = \sum_{\mu \in \Lambda^{n-1}} [\lambda : \mu] \sum_{(j, \tau) \in \mathcal{N}(i, \mu)} (j, \tau).$$

Implementation of the algorithm `ConstructN`(i, μ) requires evaluation of intersections of the cells. For the dyadic grid the cells are cubes and their intersection can be computed from the cell index (s, λ) using (7.1) and (7.2). In practice we do not compute the intersections. Instead we formally manipulate indices of the cells to obtain the index of the cell corresponding to the intersection.

8. APPLICATION

8.1. Error estimates for persistence diagrams. The algorithm we have described for computing persistence diagrams largely depends on constructing (f, s_i) -verified grids at a variety of thresholds. Depending on how many (f, s_i) -verified grids we are able to construct, the following corollary provides an *a posteriori* bound on the bottleneck distance between our calculated persistence diagrams and the true persistent diagrams for our function.

Corollary 8.1. *Fix a sequence of real numbers $\{s_i\}_{i=-1}^{m+1}$ with the following properties:*

- (1) $s_{-1} = -\infty$ and $s_0 < \min_{x \in X} f(x)$,
- (2) $s_{m+1} = \infty$ and $s_m > \max_{x \in X} f(x)$.

Suppose there exists a sequence of grids $\{\mathcal{G}_{s_i}\}_{i=-1}^{m+1}$ such that each grid \mathcal{G}_{s_i} is (f, s_i) -verified. Define $\{\mathcal{E}_{s_i}\}_{i=-1}^{m+1}$ to be the associated CW-complexes, define $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ to be the corresponding approximate filtration, and define PD' to be the persistence diagram associated with $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$. Then

$$d_B(PD(f), PD') \leq \varepsilon$$

where $\varepsilon = \max_{1 \leq i \leq m} |s_i - s_{i-1}|$.

Proof. By Theorem 6.4 it follows that each CW structure \mathcal{E}_{s_i} is compatible with X_{s_i} . Moreover, for every cell $e \in \mathcal{E}_{s_i}$ which satisfies condition (3) in Definition 1.1, the function $f \circ h(\mathbf{x}, s)$ is non-increasing in s . It follows from Theorem 1.6 that $\{\bar{\mathcal{X}}_{s_i}\}_{i=-1}^{m+1}$ and $\{X_{s_i}\}_{i=-1}^{m+1}$ have the same persistence diagrams. For every $\delta > 0$ the sequence $\{X_{s_i}\}_{i=-1}^{m+1}$ is an $(\varepsilon + \delta)$ -approximate filtration. Taking the limit as δ goes to zero, it follows from Lemma 4.1 that $d_B(PD(f), PD') \leq \varepsilon$. \square

In practice, we do not know *a priori* the thresholds for which we can construct an (f, s_i) -verified grid. Consequently, we can only give *a posteriori* error estimates for our persistence diagrams. For example, fix a sequence of real numbers $\{s_i\}_{i=-1}^{m+1}$ as above and suppose that it is uniformly spaced, that is, $\Delta = |s_i - s_{i-1}|$ for $1 \leq i \leq m$. If we attempt to create (f, s_i) -verified grids corresponding to $\{s_i\}_{i=-1}^{m+1}$, then our resulting persistence diagram will satisfy

$$(8.1) \quad d_B(PD(f), PD') \leq \Delta(F + 1)$$

where F is the largest number of consecutive thresholds s_i for which we failed to construct (f, s_i) -verified grids. While we have *a priori* control on choosing Δ , the final error estimate can be determined only after finishing the calculation.

The algorithm we employ for creating an (f, t) -verified grid will either terminate successfully or halt after reaching one of two failure conditions. The first failure condition is reached if the algorithm has exceeded a predetermined maximum resolution of the domain. This condition is imposed to prevent the algorithm from subdividing grid elements *ad infinitum*. The second failure condition occurs if, by using interval arithmetic, a vertex has an image which is neither strictly greater

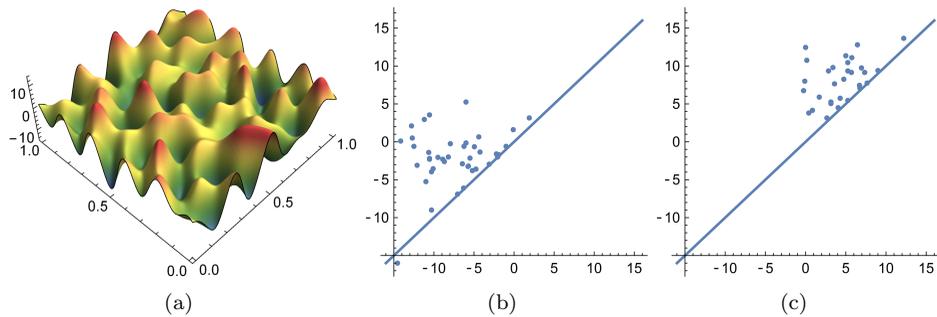


FIGURE 8. (a) Plot of a trigonometric polynomial with random coefficients. (b) β_0 persistence diagram and (c) β_1 persistence diagram.

than nor strictly less than the threshold t . If this condition is satisfied, then we are unable to decide whether or not to include that vertex in the approximation.

In general, we only need to worry about the first failure condition and regions of the sub-level set which are geometrically complex. In particular, if a cube C contains a critical point of the function f , then it can only be (f, t) -verified if it satisfies either $t > f(C)$ or $t \leq f(C)$. Consequently, if f has a critical value very close to the threshold t , then a dyadic grid will need to have a very fine resolution in order to be (f, t) -verified. Additionally, it is sometimes the case that the offspring of a cube which is (f, t) -verified are themselves not (f, t) -verified. For these reasons, having very precise interval arithmetic can help prevent runaway subdivisions of the domain.

8.2. Computational results. We implemented our algorithm for studying the sub-level sets of functions defined on the unit square $[0, 1]^2$. The source code can be downloaded from [13]. In particular, we studied random Fourier series of the form

$$f(x, y) := \sum_{i=1}^N \sum_{j=1}^N a_{i,j,1} \sin(2\pi ix) \sin(2\pi jy) + a_{i,j,2} \sin(2\pi ix) \cos(2\pi jy) \\ + a_{i,j,3} \cos(2\pi ix) \sin(2\pi jy) + a_{i,j,4} \cos(2\pi ix) \cos(2\pi jy)$$

where the coefficients $a_{i,j,k}$ were taken from a normal distribution of mean 0 and standard deviation 1. Depicted in Figure 8 is one such computation, where we used $N = 5$ as the maximum number of modes.

For the filtration we tried to construct (f, s_i) -verified grids for thresholds between -15 and 15 using spacings of 0.05 . Out of the 601 thresholds, we failed to construct (f, s_i) -verified grids for 11 thresholds. The largest number of consecutive failures occurred when both -5.2 and -5.15 failed to be verified. Hence the largest gap between any two verified thresholds in our filtration was 0.15 . Therefore the persistent diagrams we calculated are within a bottleneck distance of 0.15 of the true persistent diagrams for our given function.

The computation took 193 hours and used 16 MB of memory, and was performed on a single Intel i7-2600 processor. The algorithm created a filtered CW structure on $[0, 1]^2$ which contained $145,721$ cells, which was then processed by the *Perseus*

software program to compute its persistent homology [23]. The ϵ -approximate filtration was constructed sequentially by computing $\bar{\mathcal{X}}_{s_i} := \mathcal{X}_{s_i} \cap \bar{\mathcal{X}}_{s_{i+1}}$, so at most three CW-complexes were stored in memory at a given time. At the cost of using more memory, one could reduce the computation time by computing the cellular approximations $\{\mathcal{X}_{s_i}\}$ in parallel and subsequently creating the filtration $\{\bar{\mathcal{X}}_{s_i}\}$. Additionally, the one point of departure our implementation took from the algorithm described here is that the cells were represented using a lookup table, as opposed to the product of generalized intervals.

The reason our algorithm failed to verify the 11 thresholds is because when attempting to verify them, it surpassed 25 subdivisions of the domain. In general, the size of geometric features in these sub-level sets was much larger than 2^{-25} . However, the error bounds produced by interval arithmetic have a risk of creating a cascade of cubes which fail to become (f, t) -verified. In order to compute interval arithmetic more precisely, whenever we computed a function on a cube, we subdivided the cube seven times and computed the function using interval arithmetic on each piece.

When we repeated the computation but instead used 6 subdivisions, the number of thresholds we failed to verify increased 4.5 fold and the accuracy of our computed persistence diagram was only 0.4 in the bottleneck distance (it failed to verify seven consecutive thresholds). Using 6 subdivisions also increased the number of cells in the filtered complex by 10.5%; however the entire computation finished 3.7 times faster. We refer the reader to [3, 5] for more efficient techniques to compute precise interval arithmetic bounds needed by this algorithm.

ACKNOWLEDGMENT

The authors would like to thank Konstantin Mischaikow for his insightful comments and conversations.

REFERENCES

- [1] G. Carlsson, V. de Silva, and D. Morozov, *Zigzag persistent homology and real-valued functions*, Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry (New York, NY, USA), SCG '09, ACM, 2009, pp. 247–256.
- [2] F. Chazal, V. de Silva, M. Glisse, and S. Oudot, *The Structure and Stability of Persistence Modules*, SpringerBriefs in Mathematics, Springer International Publishing, Switzerland, 2016. DOI 10.1007/978-3-319-42545-0.
- [3] G. S. Cochran, T. Wanner, and P. Dłotko, *A randomized subdivision algorithm for determining the topology of nodal sets*, SIAM J. Sci. Comput. **35** (2013), no. 5, B1034–B1054, DOI 10.1137/120903154. MR3106488
- [4] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, *Stability of persistence diagrams*, Discrete Comput. Geom. **37** (2007), no. 1, 103–120, DOI 10.1007/s00454-006-1276-5. MR2279866 (2008i:68130)
- [5] S. Day, W. D. Kalies, and T. Wanner, *Verified homology computations for nodal domains*, Multiscale Model. Simul. **7** (2009), no. 4, 1695–1726, DOI 10.1137/080735722. MR2539195 (2010i:60156)
- [6] P. Dłotko, T. Kaczynski, M. Mrozek, and T. Wanner, *Coreduction homology algorithm for regular CW-complexes*, Discrete Comput. Geom. **46** (2011), no. 2, 361–388, DOI 10.1007/s00454-010-9303-y. MR2812514 (2012j:57001)
- [7] H. Edelsbrunner and J. Harer, *Persistent homology—a survey*, Surveys on Discrete and Computational Geometry, Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 257–282, DOI 10.1090/conm/453/08802. MR2405684 (2009h:55003)
- [8] H. Edelsbrunner and J. L. Harer, *Computational Topology: An Introduction*, American Mathematical Society, Providence, RI, 2010. MR2572029 (2011e:00001)

- [9] M. Gameiro, K. Mischaikow, and W. Kalies, *Topological characterization of spatial-temporal chaos*, Phys. Rev. E (3) **70** (2004), no. 3, 035203, 4, DOI 10.1103/PhysRevE.70.035203. MR2129999 (2005k:37063)
- [10] M. Gameiro, K. Mischaikow, and T. Wanner, *Evolution of pattern complexity in the Cahn-Hilliard theory of phase separation*, Acta Materialia **53** (2005).
- [11] E. Hansen and G. W. Walster, *Global Optimization Using Interval Analysis*, Monographs and Textbooks in Pure and Applied Mathematics, vol. 264, Marcel Dekker, Inc., New York, 2004. Second edition, revised and expanded; With a foreword by Ramon Moore. MR2025041 (2004j:90004)
- [12] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002. MR1867354 (2002k:55001)
- [13] J. Jaquette and M. Kramar, C++ code available at: http://chomp.rutgers.edu/Projects/Rigorous_Computational_Dynamics/Approximating_Persistence_Diagrams.html.
- [14] T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational Homology*, Applied Mathematical Sciences, vol. 157, Springer-Verlag, New York, 2004. MR2028588 (2005g:55001)
- [15] L. Kondic, A. Goulet, C. S. O'Hern, M. Kramar, K. Mischaikow, and R. P. Behringer, *Topology of force networks in compressed granular media*, Europhys. Lett. **97** (2012), 54001.
- [16] M. Kramar, A. Goulet, L. Kondic, and K. Mischaikow, *Persistence of force networks in compressed granular media*, Phys. Rev. E **87** (2013), 042207.
- [17] M. Kramar, A. Goulet, L. Kondic, and K. Mischaikow, *Quantifying force networks in particulate systems*, Phys. D **283** (2014), 37–55, DOI 10.1016/j.physd.2014.05.009. MR3228988
- [18] M. Kramár, R. Levanger, J. Tithof, B. Suri, M. Xu, M. Paul, M. F. Schatz, and K. Mischaikow, *Analysis of Kolmogorov flow and Rayleigh–Bénard convection using persistent homology*, Phys. D **334** (2016), 82–98, DOI 10.1016/j.physd.2016.02.003. MR3545971
- [19] K. Krishan, H. Kurtuldu, M. Schatz, M. Gameiro, K. Mischaikow, and S. Madruga, *Homology and symmetry breaking in Rayleigh-Benard convection: Experiments and simulations*, Physics of Fluids **19** (2007).
- [20] H. Kurtuldu, K. Mischaikow, and M. F. Schatz, *Measuring the departures from the Boussinesq approximation in Rayleigh-Benard convection experiments*, Journal of Fluid Mechanics **682** (2011), 543–557.
- [21] R. Mendoza, K. Thornton, I. Savin, and P. W. Voorhees, *The evolution of interfacial topology during coarsening*, Acta Materialia **54** (2006), 743–750.
- [22] Y. Mileyko, S. Mukherjee, and J. Harer, *Probability measures on the space of persistence diagrams*, Inverse Problems **27** (2011), no. 12, 124007, 22, DOI 10.1088/0266-5611/27/12/124007. MR2854323 (2012k:60008)
- [23] K. Mischaikow and V. Nanda, *Morse theory for filtrations and efficient computation of persistent homology*, Discrete Comput. Geom. **50** (2013), no. 2, 330–353, DOI 10.1007/s00454-013-9529-6. MR3090522
- [24] K. Mischaikow and T. Wanner, *Topology-guided sampling of nonhomogeneous random processes*, Ann. Appl. Probab. **20** (2010), no. 3, 1068–1097, DOI 10.1214/09-AAP652. MR2680558 (2011g:60068)

DEPARTMENT OF MATHEMATICS, HILL CENTER-BUSCH CAMPUS, RUTGERS UNIVERSITY, 110 FRELINGHUSEN ROAD, PISCATAWAY, NEW JERSEY 08854-8019

E-mail address: jaquette@math.rutgers.edu

DEPARTMENT OF MATHEMATICS, HILL CENTER-BUSCH CAMPUS, RUTGERS UNIVERSITY, 110 FRELINGHUSEN ROAD, PISCATAWAY, NEW JERSEY 08854-8019

Current address: Advanced Institute for Material Research, Tohoku University, 2-1-1 Katahira, Aoba-ku, Sendai, 980-8577 Japan

E-mail address: kramar.miroslav.e1@tohoku.ac.jp