

The Impact of Changes in Resolution on the Persistent Homology of Images

Teresa Heiss*, Sarah Tymochko†, Brittany Story‡, Adélie Garin§, Hoa Bui¶, Bea Bleile||, and Vanessa Robins**

*Institute of Science and Technology (IST) Austria, Klosterneuburg, Austria

†Michigan State University, Department of Computational Mathematics, Science and Engineering, East Lansing, MI, USA

‡Colorado State University, Department of Mathematics, Fort Collins, CO, USA

§ École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

¶ School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University, Australia

|| School of Science and Technology, University of New England, Armidale, Australia

** Research School of Physics, The Australian National University, Canberra, Australia

Abstract—Digital images enable quantitative analysis of material properties at micro and macro length scales, but choosing an appropriate resolution when acquiring the image is challenging. A high resolution means longer image acquisition and larger data requirements for a given sample, but if the resolution is too low, significant information may be lost. This paper studies the impact of changes in resolution on persistent homology, a tool from topological data analysis that provides a signature of structure in an image across all length scales. Given prior information about a function, the geometry of an object, or its density distribution at a given resolution, we provide methods to select the coarsest resolution yielding results within an acceptable tolerance. We present numerical case studies for an illustrative synthetic example and samples from porous materials where the theoretical bounds are unknown.

Index Terms—image processing, image resolution, persistent homology

I. INTRODUCTION

Engineers know that macro-scale properties, such as the strength of a composite material and the rate of fluid flow through a porous solid, depend crucially on the microscopic geometric and topological structure of the material. In practice, micro-CT x-ray images of porous materials provide detailed three-dimensional maps of spatial structure for the different phases present — typically summarised as solid grains and fluid-filled pores [1]. The size, shape, and connectivity of the pore space are known to dictate physical properties of fluid flow such as permeability and trapping capacity [2], [3]. This structure is exactly the type of information captured by a key tool of Topological Data Analysis (TDA), namely persistent homology.

The authors thank the Mathematical Sciences Institute at ANU, the US National Science Foundation through the award CCF-1841455, the Australian Mathematical Sciences Institute and the Association for Women in Mathematics for funding the second Workshop for Women in Computational Topology in July 2019 where this project began. The project reached completion thanks to funding from the MSRI Summer Research in Mathematics program awarded in 2020. Teresa Heiss has received funding from the ERC under the Horizon 2020 programme (No. 788183). We are also grateful to Kritika Singhal for her input to the first part of the project. The authors want to thank Mathijs Wintraecken for the idea of tightening the bounds when the Hausdorff distance is below the reach, i.e. the idea for Lemma B.1 in the appendix.

Simply put, persistent homology is the study of connected components and holes in a sequence of nested shapes [4]. In image analysis, these nested shapes are obtained by thresholding a real-valued function defined on the image domain. For example, to characterise the geometry of a porous material an appropriate function is the signed Euclidean distance transform derived from the binary map of solid and pore [5], [6]. The persistence diagrams computed from such a distance transform then quantify the sizes, shapes, and connectivity of all the individual pores and grains. A key property of persistent homology is its stability with respect to perturbations in the input function [7]. Topological features with small persistence can be filtered out and may correlate with noise (although this depends on the type of noise and how the input function is defined and discretized). This stability does not hold for the regular homology of a shape because every topological feature is counted with equal ‘weight’ no matter its size.

Capturing both the micro- and macro-scale properties of a material using micro-CT is a trade-off between accuracy at the micro-scale and the computational expense associated with the large image dimensions required to represent a macro-scale sample. The main focus of this paper is to establish both theoretical and numerical results that bound the difference between persistence diagrams computed from images of an object taken at different resolutions. With prior knowledge of the object these bounds can guide the selection of the coarsest resolution that will yield a persistence diagram within an acceptable distance of the ground-truth. When high resolution images lead to “big data” problems these can be scaled back to a more manageable size by reducing the image resolution or subsampling to a lower resolution. As persistent homology is very expensive to compute on high resolution imagery, studying large datasets of imagery is nearly impossible; being able to select a coarser resolution can alleviate this issue and allow for more data to be analyzed.

After covering the mathematical background in Section II, we state results in Section III comparing real-valued functions (not necessarily continuous) with digital approximations for different resolutions. Since we define the digital approximation of f by taking average values within each voxel, these results

involve fairly straightforward bounds on the variation of f and an application of the stability theorem for persistence diagrams.

In Section IV we focus on a particular application where the ground-truth is the continuous signed Euclidean distance transform (CSEDT) derived from an object $X \subset \mathbb{R}^d$. We compare this function with the discrete signed Euclidean distance transform (DSEDT) of a digital approximation to X . This is more complex than simply applying the results of Section III to the CSEDT of X , because the digital approximation to X may result in small components being lost, and this in turn may lead to a large difference between the CSEDT and the DSEDT. The results and bounds obtained in this section are illustrated with simple examples to show their relevance.

In Section V-A we present a numerical case study of a synthetic image with structure on different length scales to illustrate circumstances where the theoretical bounds are unknown. Finally, we study some porous materials in Section V-B. These experiments show that the actual data behaves even better than the bounds derived in Section IV suggest.

A. Related Work

There are several results on point-cloud approximations of manifolds in the context of geometric triangulations built from randomly sampled points near a manifold embedded in \mathbb{R}^d , [8]–[12]. Most of these results start with the assumption that the point-cloud approximation is close to the original object X in the Hausdorff distance. In contrast, the main challenge in Section IV is to establish that the digital approximation is close to the object X in the Hausdorff distance. Once this is achieved, we apply Lemma IV.6 — an analogous result to those for point-clouds, but in the setting of digital images and extended to the signed distance transforms.

An earlier paper working with the persistent homology of digital images [13] uses an adaptive grid derived from an oct-tree data structure to reduce data size. Those authors bound distance between persistence diagrams of the original high-resolution image and their oct-tree approximation, similar in spirit to the results in our Section III and the numerical case study in Section V-A, but using the adaptive grid in the approximation rather than a single coarser voxel size.

Related work in [14] starts with a known continuous function, f , builds an adaptive rectangular subdivision of the domain and uses rigorous computer arithmetic to guarantee their piecewise constant approximation is within ε of f thus guaranteeing the same bound on the bottleneck distances between persistence diagrams.

Papers that have studied how estimates of material properties from micro-CT images change with image resolution include [15], [16].

II. MATHEMATICAL BACKGROUND

A. Persistent Homology

A summary of the mathematical formalism and results of persistent homology is available in [4], and we very briefly cover the basic idea and necessary notation in this section.

Persistent homology is a tool from TDA that quantifies shape features that appear and then disappear during a *filtration*, i.e., a sequential growth process indexed by a real parameter. The filtrations in this paper are obtained as sublevel sets of real-valued functions. Given $f: \mathbb{R}^d \rightarrow \mathbb{R}$, these sublevel sets are $X_\delta = f^{-1}(-\infty, \delta]$. In three-dimensions, the shape features of X_δ are the numbers of components, tunnels, and voids, i.e., its homology classes in dimensions $k = 0, 1, 2$. Each feature has an associated interval $[b, d]$ that defines the range of filtration parameter values for which that feature persists. The persistence intervals for a filtration defined by f are collected in *persistence diagrams* which are multisets of points for each dimension of homology, $\text{PD}_k(f) = \{(b_i, d_i)\}_{i=1}^{n_k}$.

Let f, f' be two functions with persistence diagrams $\text{PD}_k(f)$ and $\text{PD}_k(f')$. The *bottleneck distance* between $\text{PD}_k(f)$ and $\text{PD}_k(f')$ is defined using matchings γ on points in the two diagrams,

$$d_B(\text{PD}_k(f), \text{PD}_k(f')) = \min_{\gamma} \max_x \|x - \gamma(x)\|_{\infty}.$$

Note that a matching, γ , is permitted to pair any point (b, d) in either diagram with a point on the diagonal in the other diagram. The stability theorem of persistent homology [7] now tells us that if f and f' are tame functions that are point-wise close then their persistence diagrams are also:

$$d_B(\text{PD}_k(f), \text{PD}_k(f')) \leq \|f - f'\|_{\infty}.$$

B. Digital Images as Discretized Functions

There is a variety of different methods to define digital approximations to algebraic functions whose domain is \mathbb{R}^d . Here we consider a tame μ -integrable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ (where μ denotes Lebesgue measure) and take local averages over each voxel.

Let r be the desired spacing for the digital grid and define a voxel σ to be an open d -cube of side length r , and volume $\mu(\sigma) = r^d$. Voxels are indexed by the integer grid \mathbb{Z}^d so that $\sigma(i) = \sigma(i_1, \dots, i_d) \subset \mathbb{R}^d$ is the product of d open intervals $(i_k r, (i_k + 1)r)$. The *digital approximation* to f is the function

$$f_r: \mathbb{R}^d \rightarrow \mathbb{R},$$

which is a piecewise constant function defined on voxels $\sigma(i)$ with

$$f_r(x) = r^{-d} \int_{\sigma(i)} f d\mu, \quad \text{if } x \in \sigma(i).$$

On the voxel faces, we define $f_r(x)$ to be the minimum value taken on voxels σ with $x \in \partial\sigma$.

If we restrict f_r to the voxel centers indexed by integers $i \in \mathbb{Z}^d$, this discrete function is usually referred to as a *grayscale digital image*. The above definition of f_r then agrees with the *T*-construction for digital images defined in [17]. The cubical complex and its sublevel sets are equivalent to using the indirect adjacency of digital grids.

In Section III we use a bound on the differences $|f(x) - f_r(x)|$ to obtain a corresponding bound on the bottleneck distance of their persistence diagrams.

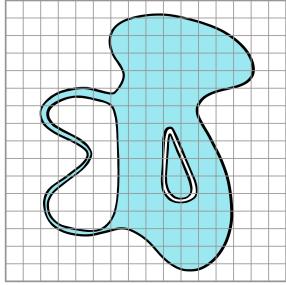


Fig. 1. An object X in blue, with a voxel grid of side r overlaid in gray. There is no t that makes $X(r, t)$ have the same homology as X . If t is close to 1, the handle-shaped part of X is lost. If t is small, then the narrow annulus in the middle will be filled in.

C. Digital Images of Objects and their Distance Transforms

Let $X \subset \mathbb{R}^d$ be a compact subset representing the solid object we approximate with a digital image. By a solid object, we mean $\text{cl}(\text{int}(X)) = X$. Typically $d = 2, 3$, and we can shift X so that it is a subset of a rectangular prism, R .

A *digital approximation* to X is defined using a discrete function $\rho_r : \mathbb{Z}^d \rightarrow [0, 1]$ that quantifies the proportion of space occupied by X in each voxel $\sigma(i)$ with

$$\rho_r(i) = \frac{\mu(X \cap \sigma(i))}{\mu(\sigma(i))}.$$

Equivalently, ρ_r is the grayscale digital image obtained from the digital approximation to the indicator function for X .

The digital approximation $X(r, t)$ is then the union of closed voxels of size r , with ρ_r -values at least threshold t with $0 < t \leq 1$,

$$X(r, t) = \bigcup_{\rho_r(i) \geq t} \text{cl}(\sigma(i)).$$

This definition of digital approximation is a simplified model of segmenting an x-ray CT image. If the material consists of just two components with a high degree of x-ray contrast, (e.g., silica and air), then the CT image measures the average x-ray density of each voxel-sized patch of the sample. In this case, x-ray density encodes the proportion of space occupied by the higher-density material. Segmenting converts the grayscale image to a binary image that approximates the distribution in space of the higher-density material. The simplest method of segmentation is to use a single threshold value, as we do here. In the presence of noise and imaging artefacts this method is unsatisfactory, and there is a considerable literature describing practical methods for achieving good segmentations [1].

Note that changing the choice of threshold t can cause significant changes to the regular homology of $X(r, t)$. Fig. 1 shows an example where there is no choice of t for which $X(r, t)$ and X have the same homology. This is why we use persistent homology to compare X and $X(r, t)$.

To quantify the size of the holes and other topological and geometric features of X , a suitable function to use is the signed

Euclidean distance. The *continuous signed Euclidean distance transform* (CSEDT) of X is defined as

$$d_X^{\mp} : \mathbb{R}^d \rightarrow \mathbb{R}; y \mapsto \begin{cases} -d(y, \partial X) & \text{if } y \in X \\ d(y, \partial X) & \text{if } y \in X^c, \end{cases}$$

where ∂X denotes the boundary of X and X^c denotes the complement of X in \mathbb{R}^d .

In practice, since we want to compare d_X^{\mp} and its discretized version, we restrict the distance transforms to a rectangular prism R that contains X . The discrete analog of the CSEDT is the discrete signed Euclidean distance transform (DSEDT)

$$D[X(r, t)] : R \rightarrow \mathbb{R}$$

This is a piecewise constant function defined on voxels $\sigma(i) \subseteq X(r, t)$ and $\sigma(j) \subseteq X(r, t)^c$ by

$$x \mapsto \begin{cases} -\min_{\sigma(j) \subseteq X(r, t)^c} r d(i, j) & \text{if } x \in \sigma(i) \subseteq X(r, t) \\ \min_{\sigma(i) \subseteq X(r, t)} r d(i, j) & \text{if } x \in \sigma(j) \subseteq X(r, t)^c. \end{cases}$$

On voxel faces, $D[X(r, t)]$ takes the minimum value over all voxels adjacent to the given face. Again, this corresponds to the T -construction mentioned earlier.

In Section IV, we investigate conditions on the geometry of X and the voxel size r necessary to control the distance between the persistence diagrams of the CSEDT of X and the DSEDT of $X(r, t)$. Even though the DSEDT depends on the threshold t , the bounds we obtain do not depend on t . In other words, even though a wise choice of t might help to achieve a better approximation at a larger voxel size, we can guarantee a good approximation for any value of $t > 0$ once r is chosen small enough compared to geometric characteristics of X , such as the *reach*.

Intuitively, the reach of a closed subset $A \subseteq \mathbb{R}^d$ encodes the minimum distance at which two or more fire fronts meet after A is set on fire [18]. Mathematically, the reach of a closed set $A \subseteq \mathbb{R}^d$ is the largest ε (possibly ∞) such that for every $p \in \mathbb{R}^d$ with distance $d(p, A) < \varepsilon$, A contains a unique point, $\xi_A(p)$, nearest to p , i.e. $d(p, A) = d(p, \xi_A(p))$ [19]. In this paper we use just one property of the reach proven by Federer in [19, Theorem 4.8 (12)]. For completeness, we give this lemma in Appendix A.

Note that the $\text{reach}(A)$ characterizes the geometry of its complement. Since we work with the signed Euclidean distance transform of a solid object X in this paper, we will use

$$\text{reach}(\partial X) = \min\{\text{reach}(X), \text{reach}(\text{cl}(X^c))\}$$

to characterize the geometry of both X and its complement. The geometric attributes that determine the reach of ∂X are its radii of curvature and distances to the generalised critical points of the signed Euclidean distance transform. It is known [20], that if ∂X is a closed C^2 submanifold of \mathbb{R}^d , then $\text{reach}(\partial X) > 0$.

In Section IV we establish a bound on the difference between the CSEDT of X and the DSEDT of $X(r, t)$ in terms of a geometric quantity we call the *leash*. We denote by A_δ ,

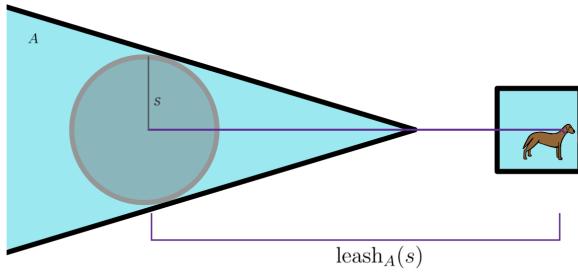


Fig. 2. Illustration of the $\text{leash}_A(s)$: The ball of radius s gets stuck inside part of A but with a long enough leash the dog can reach every part of A .

the set of all points within distance δ of a set A . Recall now, that the *erosion* of X by balls of radius δ is given by

$$E_\delta(X) = ((X^c)_\delta)^c = \{x \in X \mid B_\delta(x) \subseteq X\}.$$

The *leash* then measures the Hausdorff distance between X and $E_\delta(X)$:

$$\text{leash}_X(\delta) = d_H(X, E_\delta(X)) = \sup_{x \in X} d(x, E_\delta(X)).$$

The name *leash* can be best understood by imagining an infinitesimally small dog connected by a leash to the center of a ball of radius s , see Figure 2. The ball must stay fully inside the set A , while the dog is free to visit the whole of \mathbb{R}^d . The minimal length of the leash such that the dog can access every point of A is then $\text{leash}_A(s)$.

In Lemma IV.8, we show that if the reach is strictly positive and $0 < s < \text{reach}(\text{cl}(X)^c)$, then $\text{leash}_X(s) = s$. This is equivalent to saying that the closed δ -neighborhood of $E_\delta(X)$ exactly recovers X .

As for the reach, we need to characterize both the geometry of X and its complement X^c , so we will use the two-sided leash, defined as

$$l_X(\delta) = \max\{\text{leash}_X(\delta), \text{leash}_{\text{cl}(X^c)}(\delta)\}.$$

III. PERSISTENT HOMOLOGY OF GRayscale IMAGES

Using the definitions and tools outlined in Section II, we compare the persistence diagrams of a real-valued function to the persistence diagrams of its digital approximation at a given voxel size. This is achieved by bounding the difference between the function and its approximation and invoking the stability theorem for persistence diagrams in Proposition III.1. This result has a number of simple corollaries giving us bounds on the differences between persistence diagrams of digital approximations for different voxel sizes, and for the case that the ground-truth function is Lipschitz continuous.

Proposition III.1. *Let f be a tame μ -integrable function, and f_r its digital approximation on a grid with spacing r . Suppose there is a positive number $M_r > 0$ such that for every closed voxel $\text{cl}(\sigma(i))$ in the domain of f , the difference*

$$\sup_{x \in \text{cl}(\sigma(i))} f(x) - \inf_{x \in \text{cl}(\sigma(i))} f(x) \leq M_r.$$

Then $d_B(PD(f), PD(f_r)) \leq M_r$.

Proof. For any x in the domain of f , there is at least one voxel $\sigma(i)$ with side-length r such that $x \in \text{cl}(\sigma(i))$. If x belongs to more than one closed voxel, choose the one for which $f_r(x) = f_r(y)$ for $y \in \sigma(i)$. Let

$$M_x := \sup_{y \in \text{cl}(\sigma(i))} f(y), \quad \text{and} \quad m_x := \inf_{y \in \text{cl}(\sigma(i))} f(y).$$

We have $M_r \geq M_x - m_x$. The definition of f_r and choice of $\sigma(i)$ ensures that $m_x \leq f_r(x) \leq M_x$. It follows that

$$\begin{aligned} |f(x) - f_r(x)| &\leq \max\{f(x) - m_x, M_x - f(x)\} \\ &\leq M_x - m_x \\ &\leq M_r. \end{aligned}$$

We conclude that $\|f - f_r\|_\infty \leq M_r$. Thus, by the stability theorem [7] the bottleneck distance between the persistence diagrams is also bounded by M_r . \square

We now compare the persistence diagrams of the digital approximations f_r at different resolutions. Given two grid spacings $r_1 > 0$ and $r_2 > 0$, suppose r_2 is divisible by r_1 , i.e., $a := r_2/r_1 \in \mathbb{N}$, then any voxel of size r_2 contains a^d voxels of size r_1 . Since the measure of the voxel faces is zero, f_{r_2} is the r_2 digital approximation of f_{r_1} . Therefore, Proposition III.1, applied to f_{r_1} , implies an upper bound for the distance between the persistence diagrams of the two digital approximations f_{r_1} and f_{r_2} as follows.

Corollary III.2. *Suppose $r_2 > r_1 > 0$, $a = r_2/r_1 \in \mathbb{N}$, and choose $M \geq 0$ such that for all d -voxels $\sigma(i)$ of size r_2 , the difference*

$$\max_{\sigma' \subseteq \sigma(i)} \{f_{r_1}(x) : x \in \text{cl}(\sigma')\} - \min_{\sigma' \subseteq \sigma(i)} \{f_{r_1}(x) : x \in \text{cl}(\sigma')\} \leq M,$$

where σ' refers to voxels of size r_1 . Then $d_B(PD(f_{r_1}), PD(f_{r_2})) \leq M$.

For the special case when f is Lipschitz continuous, we can estimate the constant M in Proposition III.1 by the Lipschitz constant of f . Recall that a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is Lipschitz continuous with constant $L > 0$ if $|f(x) - f(y)| \leq L\|x - y\|$. For brevity, we say f is L -Lipschitz continuous. Note that the CSEDT d_X^+ is Lipschitz continuous with constant 1.

Corollary III.3. *Suppose $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is L -Lipschitz continuous. Then $d_B(PD(f), PD(f_r)) \leq Lr\sqrt{d}$.*

Proof. For any voxel $\sigma(i)$, since the function f is continuous and $\text{cl}(\sigma(i))$ is compact, there are $x_1, x_2 \in \text{cl}(\sigma(i))$ such that

$$f(x_1) = \max_{y \in \text{cl}(\sigma(i))} f(y), \quad \text{and} \quad f(x_2) = \min_{y \in \text{cl}(\sigma(i))} f(y).$$

It follows that

$$\begin{aligned} \max_{y \in \text{cl}(\sigma(i))} f(y) - \min_{y \in \text{cl}(\sigma(i))} f(y) &= |f(x_1) - f(x_2)| \\ &\leq L\|x_1 - x_2\| \leq Lr\sqrt{d}. \end{aligned}$$

From Proposition III.1, we see that the bottleneck distance between the persistence diagrams is bounded by $Lr\sqrt{d}$. \square

Under the Lipschitz continuity assumption, the bottleneck distance between persistence diagrams of two digital approximations is bounded as follows.

Corollary III.4. *Suppose $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is L -Lipschitz continuous. Then, $d_B(PD(f_{r_1}), PD(f_{r_2})) \leq Lr_2\sqrt{d}$, when $r_2 > r_1 > 0$ and r_2 is divisible by r_1 .*

When r_2 is not an integer multiple of r_1 , it is still possible to bound the bottleneck distance between the persistent diagrams of f_{r_1} and f_{r_2} using Proposition III.1 and triangle inequality. However, in this case, the bound is not as tight as the bound in Corollary III.4.

Corollary III.5. *Suppose $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is L -Lipschitz continuous. Then, $d_B(PD(f_{r_1}), PD(f_{r_2})) \leq L(r_1 + r_2)\sqrt{d}$, for all $r_1, r_2 > 0$.*

Proof. From the proof of Proposition III.1, triangle inequality and Corollary III.3, we have

$$\begin{aligned} \|f_{r_1} - f_{r_2}\|_\infty &\leq \|f_{r_1} - f\|_\infty + \|f_{r_2} - f\|_\infty \\ &\leq Lr_1\sqrt{d} + Lr_2\sqrt{d} = L(r_1 + r_2)\sqrt{d}. \end{aligned}$$

By the stability theorem, the bottleneck distance between the persistence diagrams is also bounded by $L(r_1 + r_2)\sqrt{d}$. \square

IV. PERSISTENT HOMOLOGY OF THE DISTANCE TRANSFORM OF BINARY IMAGES

In this section, we consider $X \subseteq \mathbb{R}^d$ as a solid object that is imaged with voxel spacing r , yielding its digital approximation $X(r, t)$ for some choice of $0 < t \leq 1$. We derive bounds on the bottleneck distance between the persistence diagram of the DSEDT of $X(r, t)$, and the CSEDT of X , by first comparing both of them to the CSEDT of $X(r, t)$, as outlined in Fig. 3.

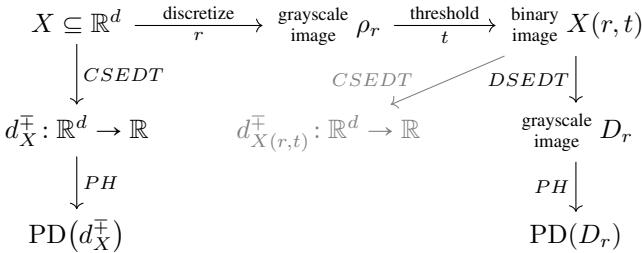


Fig. 3. Diagram of our model for the digital approximation of a solid object X , and its geometric characterisation using signed Euclidean distance transforms and persistent homology. The top row is a simplified version of CT-imaging and segmentation. In gray: In the proof of Thm. IV.1, instead of comparing the continuous distance transform d_X^\mp to the discrete distance transform D_r directly, we compare both to the continuous distance transform $d_{X(r, t)}^\mp$ of the discrete object.

A. Bounds using the reach or leash of X

The most relevant result of this section is Corollary IV.9, which states that when discretizing with voxel length $r < \frac{1}{\sqrt{d}}\text{reach}(\partial X)$, the bottleneck distance is bounded by a constant times r . This follows directly from Theorem IV.1, which

bounds the bottleneck distance in terms of the two-sided leash $l_X(\sqrt{d}r)$ when discretizing with any voxel length r , not necessarily smaller than the reach. Note that the bounds we give here do not depend on the threshold t . The results could be slightly improved by incorporating t but we do not do so for the sake of brevity.

Theorem IV.1. *Given $X \subseteq \mathbb{R}^d$ with $X = \text{cl}(\text{int}(X))$. Let d_X^\mp be the CSEDT of X , $r > 0$, $t \in (0, 1]$, and $X(r, t)$ be its digital approximation. Let $D_r = D[X(r, t)]: \mathbb{R}^d \rightarrow \mathbb{R}$ denote the discrete signed Euclidean distance transform. Then,*

$$d_B(PD(d_X^\mp), PD(D_r)) \leq l_X(\sqrt{d}r) + 2\sqrt{d}r$$

The proof requires the following lemmas. The first lemma shows that the DSEDT of $X(r, t)$ is a good approximation to the CSEDT of $X(r, t)$.

Lemma IV.2. *With the notation of Theorem IV.1,*

$$\|d_{X(r, t)}^\mp - D_r\|_\infty \leq \sqrt{d}r.$$

Proof. Let $p \in \mathbb{R}^d$ and assume $p \in X(r, t)$; the proof for $p \in X(r, t)^c$ works analogously. Let y be a closest point of $\partial X(r, t)$ to p , i.e. $d_{X(r, t)}^\mp(p) = -d(p, \partial X(r, t)) = -d(p, y)$, see Fig. 4. As $y \in \partial X(r, t)$, there exists a voxel $\sigma(i) \subseteq$

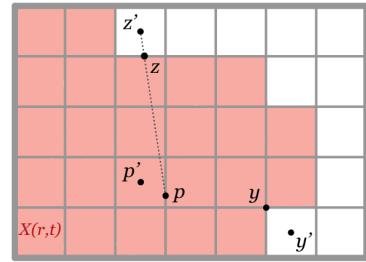


Fig. 4. Illustration of the variables appearing in the proof of Lemma IV.2.

$X(r, t)^c$ with voxel center y' of distance $d(y, y') \leq \frac{1}{2}\sqrt{d}r$. Let $\sigma(j)$ be the voxel with p in its closure with minimal filtration value (in case there are several), i.e. the voxel giving filtration value $D_r(p) = D_r(\sigma(j))$ to p , and let p' be its center. Thus, $d(p, p') \leq \frac{1}{2}\sqrt{d}r$. Hence,

$$\begin{aligned} -D_r(p) &= -D_r(\sigma(j)) \\ &= \min\{d(p', c) \mid c \text{ voxel center of } \sigma(k) \subseteq X(r, t)^c\} \\ &\leq d(p', y') \leq d(p', p) + d(p, y) + d(y, y') \\ &\leq \frac{1}{2}\sqrt{d}r - d_{X(r, t)}^\mp(p) + \frac{1}{2}\sqrt{d}r, \end{aligned}$$

yielding $d_{X(r, t)}^\mp(p) - D_r(p) \leq \sqrt{d}r$.

Let z' be the voxel center minimizing

$$\min\{d(p', c) \mid c \text{ voxel center of } \sigma(k) \subseteq X(r, t)^c\} = -D_r(p).$$

As $z' \in X(r, t)^c$ and $p \in X(r, t)$, there exists a $z \in \partial X(r, t)$ on the straight line segment between z' and p . Hence,

$$\begin{aligned} -d_{X(r, t)}^\mp(p) &= d(p, \partial X(r, t)) \leq d(p, z) \leq d(p, z') \\ &\leq d(p, p') + d(p', z') \leq \frac{1}{2}\sqrt{d}r - D_r(p), \end{aligned}$$

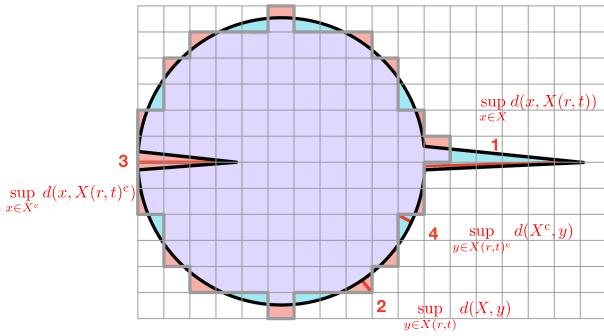


Fig. 5. The four different suprema in the bound of Lemma IV.3 can all have different values. In the proof of Theorem IV.1 we show that suprema 1 and 3 are bounded by the two-sided leash $l_X(\sqrt{dr})$, while suprema 2 and 4 are bounded by the voxel diameter \sqrt{dr} .

yielding $-(d_{X(r,t)}^\mp(p) - D_r(p)) \leq \frac{1}{2}\sqrt{dr}$. \square

The authors of [7] mention that, by definition, the Hausdorff distance between two sets equals the L_∞ -distance between the (unsigned) Euclidean distance transforms of these two sets. For the proof of Theorem IV.1 an extension of this to the *signed* Euclidean distance transform is needed, see Lemma IV.3. Note that the bound is between the maximum and the sum of the two terms $d_H(A, B)$ and $d_H(A^c, B^c)$.

Lemma IV.3. *Let $A, B \subseteq \mathbb{R}^d$, then*

$$\begin{aligned} \|d_A^\mp - d_B^\mp\|_\infty &\leq \max\{\sup_{b \in B^c} d(A^c, b) + \sup_{a \in A} d(a, B), \\ &\quad \sup_{b \in B} d(A, b) + \sup_{a \in A^c} d(a, B^c)\} \end{aligned}$$

Proof. Let d_A denote the unsigned Euclidean distance transform defined as $d_A(p) = d(p, A)$, for any $p \in \mathbb{R}^d$. As mentioned above, we have by definition that $\|d_A - d_B\|_\infty = d_H(A, B)$. To bound $\|d_A^\mp - d_B^\mp\|_\infty$, we consider four cases for the point $p \in \mathbb{R}^d$.

Case 1: $p \in A$, $p \in B$. Then $|d_A^\mp(p) - d_B^\mp(p)| = |-d_{A^c}(p) + d_{B^c}(p)| \leq \|d_{B^c} - d_{A^c}\|_\infty = d_H(A^c, B^c) = \max\{\sup_{b \in B^c} d(A^c, b), \sup_{a \in A^c} d(a, B^c)\}$.

Case 2: $p \notin A$, $p \notin B$. Analogously $|d_A^\mp(p) - d_B^\mp(p)| \leq \max\{\sup_{a \in A} d(a, B), \sup_{b \in B} d(A, b)\}$.

Case 3: $p \in A$, $p \notin B$. Then $|d_A^\mp(p) - d_B^\mp(p)| = |-d(A^c, p) - d(p, B)| = d(A^c, p) + d(p, B) \leq \sup_{b \in B^c} d(A^c, b) + \sup_{a \in A} d(a, B)$.

Case 4: $p \notin A$, $p \in B$. Analogously $|d_A^\mp(p) - d_B^\mp(p)| \leq \sup_{b \in B} d(A, b) + \sup_{a \in A^c} d(a, B^c)$. \square

Remark IV.4. In the special case that $\max\{d_H(A, B), d_H(A^c, B^c)\} < \text{reach}(\partial A)$, Lemma IV.3 can be tightened to $\|d_A^\mp - d_B^\mp\|_\infty \leq \max\{d_H(A, B), d_H(A^c, B^c)\}$, see Appendix B.

However, when the sets differ by more than their reach, the bound in Lemma IV.3 can be tight, as Example IV.5 shows.

Example IV.5. Let $A = [-1.5, 1.5] \times \mathbb{R}$ and $B = [-1.5, 1.5] \times ((-\infty, -1] \cup [1, \infty))$. Then, the bound given

in Lemma IV.3 is tight: $\|d_A^\mp - d_B^\mp\|_\infty = |d_A^\mp(0) - d_B^\mp(0)| = |-1.5 - 1| = 1.5 + 1 = \sup_{b \in B^c} d(A^c, b) + \sup_{a \in A} d(a, B)$.

Lemma IV.6. *With the notation of Theorem IV.1,*

$$\begin{aligned} d_B(\text{PD}(d_X^\mp), \text{PD}(D_r)) \\ \leq \max\{\sup_{y \in X(r,t)^c} d(X^c, y) + \sup_{x \in X} d(x, X(r, t)), \\ \sup_{y \in X(r,t)} d(X, y) + \sup_{x \in X^c} d(x, X(r, t)^c)\} + \sqrt{dr} \end{aligned}$$

Proof. The proof is a combination of the stability theorem of persistent homology [7], the triangle inequality, and Lemmas IV.3, IV.2:

$$\begin{aligned} d_B(\text{PD}(d_X^\mp), \text{PD}(D_r)) \\ \leq \|d_X^\mp - D_r\|_\infty \\ \leq \|d_X^\mp - d_{X(r,t)}^\mp\|_\infty + \|d_{X(r,t)}^\mp - D_r\|_\infty \\ \leq \max\{\sup_{y \in X(r,t)^c} d(X^c, y) + \sup_{x \in X} d(x, X(r, t)), \\ \sup_{y \in X(r,t)} d(X, y) + \sup_{x \in X^c} d(x, X(r, t)^c)\} + \sqrt{dr}. \end{aligned}$$

\square

Lemma IV.7. *With the notation of Theorem IV.1,*

- 1) $\sup_{x \in X} d(x, X(r, t)) \leq \text{leash}_X(\sqrt{dr})$
- 2) $\sup_{x \in X^c} d(x, X(r, t)^c) \leq \text{leash}_{\text{cl}(X^c)}(\sqrt{dr})$
- 3) $\sup_{y \in X(r, t)} d(X, y) \leq \sqrt{dr}$
- 4) $\sup_{y \in X(r, t)^c} d(X^c, y) \leq \sqrt{dr}$

Proof. 1) To prove $\sup_{x \in X} d(x, X(r, t)) \leq \text{leash}_X(\sqrt{dr})$, let $x \in X$ be arbitrary. By the definition of $\text{leash}_X(\sqrt{dr})$, there exists a point $a \in ((X^c)_{\sqrt{dr}})^c$ with $d(x, a) \leq \text{leash}_X(\sqrt{dr})$. As a is at least \sqrt{dr} far from X^c , the open ball $B_{\sqrt{dr}}(a)$ is fully inside X . Let σ be a voxel containing a in its closure. As the diameter of the voxel is $\text{diam}(\sigma) = \sqrt{dr}$, the voxel $\sigma \subseteq B_{\sqrt{dr}}(a) \subseteq X$ has density $\rho_r(\sigma) = 1 \geq t$, for all $t \in (0, 1]$. Hence, $a \in \text{cl}(\sigma) \subseteq X(r, t)$ and $d(x, X(r, t)) \leq d(x, a) \leq \text{leash}_X(\sqrt{dr})$.

2) The proof for $\sup_{x \in X^c} d(x, X(r, t)^c) \leq \text{leash}_{\text{cl}(X^c)}(\sqrt{dr})$ is analogous.

3) To prove $\sup_{y \in X(r, t)} d(X, y) \leq \sqrt{dr}$, let $y \in X(r, t)$ be arbitrary. Let $\sigma \subseteq X(r, t)$ be such that $y \in \text{cl}(\sigma)$. As $\rho_r(\sigma) > t > 0$, there exists a point $x \in X \cap \sigma$. Hence, $d(X, y) \leq d(x, y) \leq \text{diam}(\sigma) = \sqrt{dr}$.

4) The proof for $\sup_{y \in X(r, t)^c} d(X^c, y) \leq \sqrt{dr}$ is analogous. \square

Proof of Theorem IV.1. Lemma IV.7 shows that both the terms $\sup_{y \in X(r, t)^c} d(X^c, y) + \sup_{x \in x} d(x, X(r, t))$ and $\sup_{y \in X(r, t)} d(X, y) + \sup_{x \in X^c} d(x, X(r, t)^c)$ are bounded by $\sqrt{dr} + l_X(\sqrt{dr})$. In combination with Lemma IV.6 this finishes the proof. \square

In the special case when the voxel diameter \sqrt{dr} is smaller than $\text{reach}(\partial X)$, we can calculate the term $l_X(\sqrt{dr})$ explicitly using the following Lemma.

Lemma IV.8. Let A be a subset of \mathbb{R}^d with $A = \text{cl}(\text{int}(A))$ and $\text{reach}(\text{cl}(A^c)) > 0$. Let $s < \text{reach}(\text{cl}(A^c))$, then $\text{leash}_A(s) = s$.

Proof. To show that $\text{leash}_A(s) = \sup_{a \in A} d(a, E_s(A))$ is at most s , we need to find for every $a \in A$ an element $b \in E_s(A)$ with distance $d(a, b)$ at most s . Let $a \in A$ be arbitrary. If $a \in E_s(A)$, then $b = a$ fulfills $d(a, b) = 0 \leq s$. Therefore $a \notin E_s(A)$, i.e. $a \in (A^c)_s$. Let $x \in \partial \text{cl}(A^c)$ be s.t. $d(a, x) = d(a, \text{cl}(A^c))$. If $x \neq a$, then $n_x = \frac{a-x}{\|a-x\|}$ is a unit normal vector of $\text{cl}(A^c)$ at x (for a rigorous definition of normal vector, see Definition A.2). If $x = a$, let n_x be any unit normal vector of $\text{cl}(A^c)$ at x . Define $b = x + sn_x$. Since $s < \text{reach}(\text{cl}(A^c))$, Federer's Lemma A.3 yields, $d(A^c, b) = d(\text{cl}(A^c), b) = d(x, b) = s$. Hence, $b \notin (A^c)_s$ and thus $b \in ((A^c)_s)^c = E_s(A)$. The points x , a , and b lie on a straight line by construction. The point a has distance $d(a, x) = d(a, \text{cl}(A^c)) = d(a, A^c)$ to x , which is less than s as $a \in (A^c)_s$. Hence, a lies on the line segment between x and b , and thus $d(a, b) \leq d(x, b) = s$, concluding the proof of $\text{leash}_A(s) \leq s$.

For any $a \in \partial A$, we get $d(a, E_s(A)) \geq s$, yielding $\text{leash}_A(s) = s$. \square

Corollary IV.9. Using the notation of Theorem IV.1. If $r < \frac{1}{\sqrt{d}} \text{reach}(\partial X)$, then

$$d_B(\text{PD}(d_X^\mp), \text{PD}(D_r)) \leq 3\sqrt{dr}$$

Proof. Apply Theorem IV.1 and Lemma IV.8. \square

Remark IV.10. Using Remark IV.4, the constant 3 in Corollary IV.9 can be tightened to 2, see Appendix B.

B. Bounds using the density ρ_r

Suppose we start with ρ_r , the function that tells us the proportion of X in each voxel defined in Section II-C. As mentioned there, this is a reasonable model of the information contained in an x-ray CT image of a two-phase material with high x-ray contrast. Theorem IV.12 bounds the bottleneck distance between the persistence diagrams of d_X^\mp and D_r in terms of ρ_r . The proof reuses Lemma IV.6, but we need to bound the suprema by quantities other than the leash this time.

Lemma IV.11. Using the notation of Theorem IV.1.

- 1) $\sup_{x \in X} d(x, X(r, t)) \leq \max_{\rho_r(i) > 0} D_r(\sigma(i))$
- 2) $\sup_{x \in X^c} d(x, X(r, t)^c) \leq -\min_{\rho_r(i) < 1} D_r(\sigma(i))$.

Proof. We prove the first statement here; the second can be shown similarly. Let $x \in X$, and assume $x \notin X(r, t)$ (otherwise the statement follows trivially). As $x \in X = \text{cl}(\text{int}(X))$, any voxel $\sigma(j)$ containing x in its closure has $\rho_r(j) > 0$. Let x' be the center of $\sigma(j)$. As $x \notin X(r, t)$, its voxel $\sigma(j)$ is in $X(r, t)^c$. Let y' be the voxel center minimizing

$$\min\{d(x', c) \mid c \text{ voxel center } \in X(r, t)\} = D_r(\sigma(j)).$$

As $x - x'$ is a vector pointing from a voxel center to a point in the closure of the same voxel, adding this vector to the voxel

center y' yields a point in the closure of the same voxel, and thus a point in $X(r, t)$. Hence,

$$\begin{aligned} d(x, X(r, t)) &\leq d(x, y' + (x - x')) \\ &= d(x - (x - x'), y') \\ &= d(x', y') \\ &= \min\{d(x', c) \mid c \text{ voxel center } \in X(r, t)\} \\ &= D_r(\sigma(j)) \leq \max_{\rho_r(i) > 0} D_r(\sigma(i)). \end{aligned}$$

\square

Define m_{ρ_r} as the maximum of $\max_{\rho_r(i) > 0} D_r(\sigma(i))$ and $-\min_{\rho_r(i) < 1} D_r(\sigma(i))$. This measures how far $X(r, t)$ is from the two most extreme ways of thresholding.

Theorem IV.12. With the notation of Theorem IV.1,

$$d_B(\text{PD}(d_X^\mp), \text{PD}(D_r)) \leq m_{\rho_r} + 2\sqrt{dr}$$

Proof. Combining Lemma IV.6, Lemma IV.11, and items 3 and 4 of Lemma IV.7. \square

C. Examples

We describe an example that illustrates many of the bounds derived above. The ground-truth object X consists of an array of small circular spots lying inside a large disk together with the outside of this disk; see the image at top left of Fig. 6. In total this image is 2048×2048 pixels and the large white disk is a circle of radius $R_2 = 510$ pixels. The small black disks have radius $R_1 = 5$ pixels and centers $w = 85$ pixels apart.

The reach of ∂X is the radius R_1 of the small spots, so the image with 2048^2 pixels has $1 = r < \text{reach}(\partial X)/\sqrt{2} = 3.54$. For any resolution with $r < 3.54$, (no. pixels $> 580^2$) Remark IV.10 guarantees that the bottleneck distance between the PDs for d_X^\mp and $\text{D}[X(r, t)]$ is no larger than $2\sqrt{2}r$.

At coarser resolutions, Theorem IV.1 still applies with $l_X(\sqrt{dr}) = R_2 - \frac{w}{\sqrt{2}} + R_1 + \sqrt{2}r = 454.9 + \sqrt{2}r$ when $\text{reach}(\partial X) \leq \sqrt{2}r \leq \frac{w}{\sqrt{2}} - R_1$. So we see that

$$d_B(\text{PD}(d_X^\mp), \text{PD}(\text{D}[X(r, t)])) \leq 454.9 + 3\sqrt{2}r,$$

for $3.55 \leq r \leq 39.0$.

For this synthetic example, we can also compute ρ_r for all choices of r . Let $\epsilon > 0$ be a small tolerance to accommodate noise. The pixels where $\epsilon < \rho_r < 1 - \epsilon$ are those that have non-empty intersection with ∂X . At high resolutions, i.e., pixel sizes $r < \text{reach}(\partial X)/\sqrt{2}$, this means the Hausdorff distances between $X(r, t)$ and $X(r, 1 - \epsilon)$ or $X(r, \epsilon)$ are small.

At coarser resolutions, for example the image with 64×64 pixels depicted at lower left of Fig. 6, each spot is smaller than a pixel and all pixels intersecting the large disk have $\epsilon < \rho_r < 1 - \epsilon$. In this situation, setting $t = 1 - \epsilon$ means the set $X(r, 1 - \epsilon)$ is contained in the outside of the large disk, while $X(r, \epsilon)$ is the entire square.

Fig. 6 shows that the red bound from Section IV-B has the advantage of being tighter than the green bound from Section IV-A at lower resolutions (i.e., pixel sizes larger than the reach). However, at high resolutions the bound given by Remark IV.10 is better.

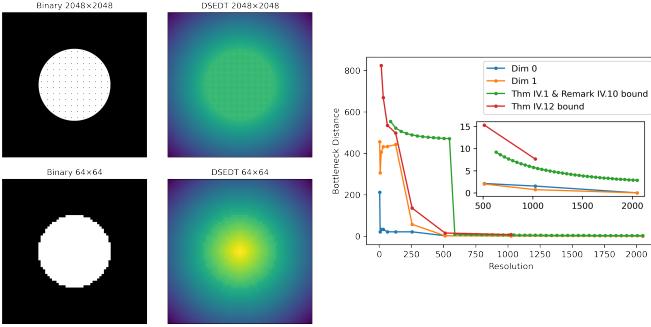


Fig. 6. Left two columns: The top row shows a high resolution binary image with corresponding DSEDT and the bottom row shows a lower resolution version obtained by downsampling with the averaging method and thresholding with $t = 0.5$ and its corresponding DSEDT. Right column: Bottleneck distances comparing persistence diagrams for each resolution with that for the highest resolution. The distances for dimensions 0 and 1 are shown along with the bounds provided by Theorem IV.1 with Remark IV.10 (green) and Theorem IV.12 (red).

V. APPLICATIONS

In any real-world situation, the reach and the leash for the object being imaged are likely to be unknown, and/or computationally expensive to estimate. Similarly, the density ρ_r at a given pixel length r might be known, but noisy. So in this section, we explore a different approach to determining an adequate digital resolution using the persistence diagrams of DSEDTs computed for a succession of larger voxel sizes, i.e., decreasing image dimensions.

In our case studies, we begin with the highest resolution image and downsample to lower resolution images as follows. For a d -dimensional binary image with grid spacing r , and dimensions $n_1 \times \dots \times n_d$, we compute the averages in blocks of a^d voxels, where a is called the kernel size. Based on these average values, we create the downsampled binary image with grid spacing $r' = ar$ and dimensions $n_1/a \times \dots \times n_d/a$ by thresholding at value $t = 0.5$, similar to the process described in Section II-C. From the new image, we compute the DSEDT, with the new grid spacing r' used as a scaling factor.

All examples in this section are square, so the image size is n^d and we use n to quantify resolution. The voxel size is specified with arbitrary units so that $r = 1$ for the highest-resolution image in each case. To avoid issues caused by boundary effects, we only choose kernel sizes that fit evenly within the original high-resolution image dimensions. Thus kernel sizes are always integer divisors of the image size. This gives us a collection of images of different resolutions that are approximations of the highest resolution image. Then we compute the bottleneck distance between the persistence diagram for each resolution with the persistence diagram from the high resolution image. This allows us to quantify how much information is lost when we downsample.

We compute the persistence diagrams and bottleneck distance, using the Giotto-TDA python package [21]. Code for this project, including code to generate the synthetic examples, can be found on Github [22].

A. Structure at Different Length Scales

Our goal in this section is to explore how the persistent homology of a DSEDT changes at different image resolutions when approximating a particular object $X \subset \mathbb{R}^d$, with structure at different length scales, illustrated in Fig. 7. Although the overall trend is that bottleneck distance decreases with resolution, this decrease is not monotonic. The distinct length scales mean the distances show a succession of plateaus as each feature is resolved and the persistence diagrams remain relatively stable over an interval of resolutions.

We now define a Δ - ε plateau to capture the change in bottleneck distance between persistence diagrams as the image resolution changes. Recall that the (linear) image size n is used to quantify resolution, and, for the remainder of this section, denote the DSEDT $D[X(r, t)]$ by D^n . Let N denote the highest resolution available and take an interval $\Delta = [\ell, m]$ of image resolutions and $\varepsilon > 0$. We then say there is a Δ - ε plateau if for all $j, k \in [\ell, m]$,

$$|d_B(\text{PD}(D^j), \text{PD}(D^N)) - d_B(\text{PD}(D^k), \text{PD}(D^N))| < \varepsilon.$$

By the triangle inequality $d_B(\text{PD}(D^j), \text{PD}(D^k)) < \varepsilon$ is a sufficient condition for a Δ - ε plateau.

As an example, consider the image in Figure 7 with three nested rings of different thicknesses. The original high-resolution image, X_N , has $r = 1$ and $N^2 = 5040^2$ pixels. We downsample the original image to create 57 different images, where each new image has a kernel of size a where a is a proper divisor of 5040, so that the resolution in each case is $n = 5040/a$, and the voxel size at this resolution is $r = a$. The DSEDT and persistence diagrams are computed for each image, and then we find the bottleneck distance $d_B(\text{PD}(D^n), \text{PD}(D^{5040}))$ comparing each lower-resolution image with the original one. The plot of these results in Figure 7 shows three specific behaviors we would like to emphasize; spikes, plateaus, and final plateaus.

As we increase the resolution, some fluctuations in the bottleneck distance are to be expected, even within a plateau. These fluctuations are on the order of magnitude of the pixel diameter. For the example in Fig. 7 it makes a difference whether the center of the image is the center of a pixel (for n odd) or the center is in the closure of 4 pixels (for n even). These effects can cause the bottleneck distance to spike, i.e. to increase and then immediately decrease. The lower the resolution, the more pronounced these spikes usually appear. The blue zero-dimensional bottleneck distance curve in Fig. 7 has a spike at $n = 21$, the only odd number in this range of resolutions, with an increase by 366 from $n = 20$ to $n = 21$ and a decrease by 333 from $n = 21$ to $n = 24$. The magnitude of this spike should be compared with the pixel diameter for $n = 21$, which is $r\sqrt{2} = 240\sqrt{2} = 339.4$.

Next, we observe the plateau behavior outlined in the definition of a Δ - ε plateau. This is caused by the introduction of a new topological feature, such as a grain of sand or a small pore. In the nested ring example, the plot of one-dimensional bottleneck distances shows three plateaus. The

zero-dimensional bottleneck distances mimics this behavior, but the first one or two plateaus are overshadowed by noise, like spikes, explained above. The three plateaus are due to the three different ring widths. After each of these is resolved, an increase in resolution does little to change the persistence diagram. As shown in Fig. 7, we see just one ring at resolution 18, two rings at $n = 40$, the third ring starts being resolved at $n = 105$ (where it only consists of 4 pixels) and gets fully resolved at $n = 210$. Specifically, the first Δ - ε plateau is at $\Delta = [18, 24]$ with $\varepsilon = 36.96$ in dimension 1 and $\varepsilon = 402.51$ in dimension 0. The second Δ - ε plateau is at $\Delta = [40, 90]$ with $\varepsilon = 17.31$ in dimension 1 and $\varepsilon = 104.49$ in dimension 0. The third Δ - ε plateau is at $\Delta = [210, 5040]$ with $\varepsilon = 13.43$ in dimension 1 and $\varepsilon = 15.80$ in dimension 0.

The existence of structure at just a few distinct, well separated length scales in the synthetic example means there is a sequence of plateaus as each structure is resolved. This does not happen for the porous materials discussed in Section V-B. However, we note that Corollary IV.9, and Remark IV.10 guarantee the existence of a final plateau for any object with positive reach. Specifically, let M be the lowest resolution required for the corresponding voxel size $r_M < \frac{1}{\sqrt{d}}\text{reach}(\partial X)$, and assume that the highest resolution image has $N > M$. Then the bottleneck distances will have a Δ - ε plateau with $\Delta = [M, N]$ and $\varepsilon = 4\sqrt{d}r_M \leq 4\text{reach}(\partial X)$. In practice, we often see a final plateau even when the reach is zero, such as when the leash bounding the bottleneck distance converges to zero. The glass bead packing in the following section is such an example.

Although a final plateau suggests the image resolution is sufficient to capture the actual underlying structure of the imaged object, we can never definitively know whether a plateau is final or not, because we do not know what structure exists at length-scales finer than the voxel size. Hence, the plateau serves only as a guide for resolution choice.

B. Application to Porous Materials

It is common for material science examples to have zero reach. For example, packings of spherical or elliptical grains have $\text{reach}(X) = 0$, because two grains touch at a single point, while other porous materials, such as a metal foam, have sharp corners. Here we present three examples of micro-CT images of porous materials segmented into the two phases of solid and void [23]. The images are subregions from a packing of spherical glass beads, a sandstone (from Castlegate) and an unconsolidated sandpack, each with 512^3 voxels.

In Fig. 8 we depict slices through each 3D binary image and their signed Euclidean distance transforms, followed by plots of the bottleneck distances between persistence diagrams computed at different resolutions. In each case, we subsample the binary image to a lower resolution using the averaging technique described at the start of Section V with a threshold $t = 0.5$. We only use kernel sizes that fit evenly within the image; with $N = 512$ we must use powers of two, $a = 2^k$. As the persistence diagrams for these examples have so many points, we use an approximation algorithm for the bottleneck

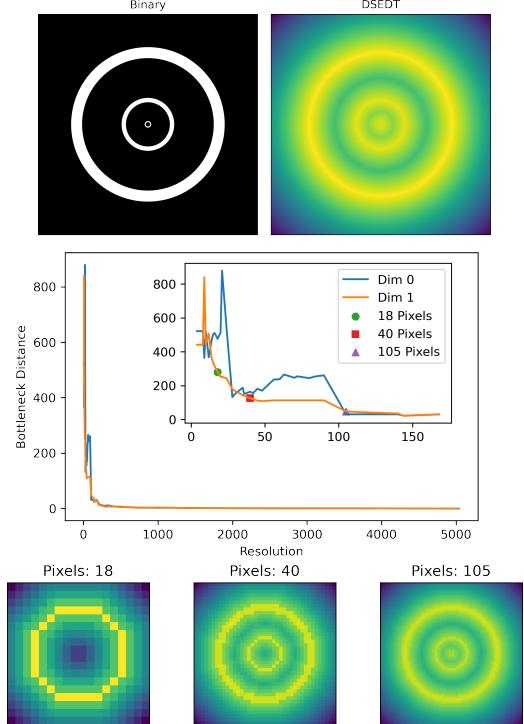


Fig. 7. Top, the digital image we use as the ground truth with the corresponding DSEDT. Middle, the plot of the bottleneck distances between 1-dimensional persistence diagrams from each resolution and the highest resolution. Below, the DSEDT of the downsampled image at three different resolutions.

distances, as implemented in Giotto-TDA. Specifically, we use an approximation value of $\delta = 0.1$ for the glass bead packing, and $\delta = 0.5$ for the other two to make these computations feasible.

The green curve shown on the plots of bottleneck distances in Fig. 8 is the function $2\sqrt{3}r = 2\sqrt{3}(512/n)$, where $r = 1$ is the voxel spacing for the highest-resolution image, and n is the resolution as measured by the number of voxels along each side of the cube. This is the bound on bottleneck distance we derived in Section IV for the case that the voxel size $r < \text{reach}(\partial X)/\sqrt{d}$. As already argued, the reach is zero for the glass bead pack, and likely to be zero for the sand pack, so the fact that this bound holds suggests that our estimation results are too generous, and/or that $l_X(s)$ is approximately s despite the reach being zero for these examples.

We note that the glass bead example has significantly larger bottleneck distances between its 1-dimensional persistence diagrams compared to the 0-dimensional distances for the resolutions $n = 32, 64, 128, 256$. Inspection of the dimension one $\text{PD}_1(D^n)$ diagrams and the original image shows that the larger values of $d_B(\text{PD}_1(D^n), \text{PD}_1(D^{512}))$ are due to the presence of just a couple of high-persistence cycles near the boundary of the image, each cycle due to a bead that intersects two faces of the boundary. The sensitivity of the bottleneck distance to outliers is well known and is the reason Wasserstein distances between diagrams are often preferred.

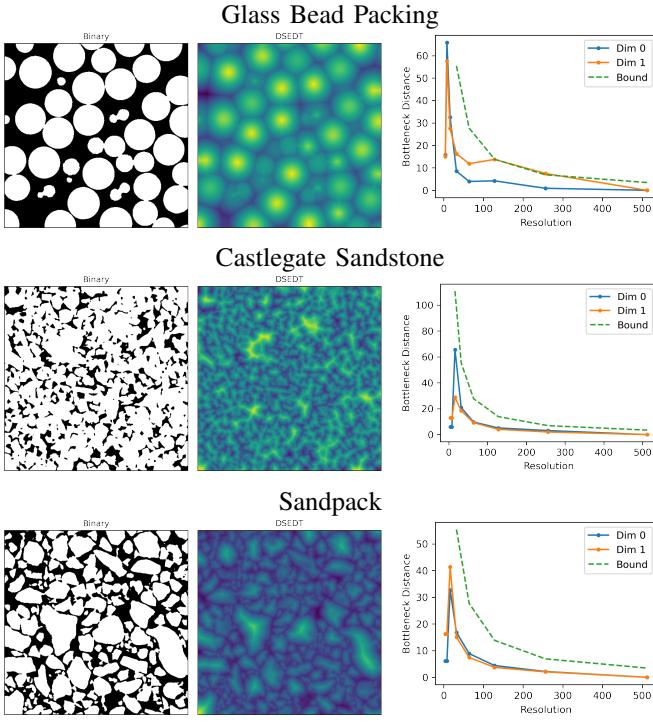


Fig. 8. 2D slices of the binary and DSEDT images for the glass bead packing, Castlegate sandstone, and sand packing samples. The colormap for the DSEDT is scaled to the max and min values in each case. The plots in the right column show the bottleneck distance between persistence diagrams for each downsampled resolution and the highest resolution. The green dashed curve is the function $2\sqrt{3}/(512/n)$, where n is the image resolution.

An important physical parameter associated with porous materials is the percolation threshold, l_c . This is the radius of the largest sphere that can pass through the pore space from one side of the image to the opposite. The distribution of points in $\text{PD}_0(D^n)$ shows a clear signature of this critical length scale; see [24] for details. As can be seen in the persistence diagrams in Appendix D, this signature yields the same estimate of l_c for image resolutions $n = 512, 256, 128$ in the three example materials. This supports the designation of a Δ - ε plateau in the 0-dimensional bottleneck distances with $\Delta = [64, 512]$, and $\varepsilon = 10$ for all three samples.

VI. CONCLUSION

This paper has presented two sets of results about the digital approximation of functions, and of solid objects. Section III makes explicit how far points in a persistence diagram can move when working with locally averaged digital approximations to a function at different resolutions.

Section IV has considered a more subtle question of how close are the persistence diagrams of a solid object X and a digital approximation to it, $X(r, t)$, when they are filtered by the continuous and discrete signed Euclidean distance transforms respectively. These results are the analogues of seminal work for point-cloud approximations of manifolds, which do not translate easily to the digital image setting.

In general, the difference between the continuous and discrete distance transforms is given in terms of the voxel diameter plus the leash (Theorem IV.1). The leash may be large, for example, if X is a material that has regions of micro-porosity or extended structures with geometric detail below the voxel size r . When the voxel diameter $\sqrt{dr} < \text{reach}(\partial X)$, we show that the leash, $l_X(\sqrt{dr}) = \sqrt{dr}$, so that the bottleneck distance between persistence diagrams is bounded by $2\sqrt{dr}$ (Corollary IV.9 and Remark IV.10).

The practical consequences of these results are that we expect the persistent homology to converge with increasing image resolution, but the error may not be monotonic, especially when considering images at low resolutions. When there is no prior information about the critical length scales of the object X , the x-ray density function given by the CT-scan can be interpreted as a (noisy) approximation to the density function ρ_r , and the Hausdorff distance between the two threshold choices $X(r, 1 - \varepsilon)$ and $X(r, \varepsilon)$ provides an estimate of the possible error in the persistence diagrams (Theorem IV.12).

Ultimately, our results provide guidance to practitioners on how to balance the time, cost, and processing power required for image acquisition and persistent homology computations against the desired level of accuracy in their results.

REFERENCES

- [1] D. Wildenschild and A. P. Sheppard, "X-ray imaging and analysis techniques for quantifying pore-scale structure and processes in subsurface porous medium systems," *Advances in Water Resources*, vol. 51, pp. 217–246, 2013.
- [2] A. L. Herring, E. J. Harper, L. Andersson, A. Sheppard, B. K. Bay, and D. Wildenschild, "Effect of fluid topology on residual nonwetting phase trapping: Implications for geologic CO₂ sequestration," *Advances in Water Resources*, vol. 62, Part A, pp. 47–58, Dec. 2013.
- [3] A. L. Herring, V. Robins, and A. P. Sheppard, "Topological Persistence for Relating Microstructure and Capillary Fluid Trapping in Sandstones," *Water Resources Research*, vol. 55, no. 1, pp. 555–573, Jan. 2019.
- [4] H. Edelsbrunner and J. Harer, "Persistent homology—a survey," in *Surveys on discrete and computational geometry*, ser. Contemp. Math. Amer. Math. Soc., Providence, RI, 2008, vol. 453, pp. 257–282.
- [5] V. Robins, P. J. Wood, and A. P. Sheppard, "Theory and algorithms for constructing discrete Morse complexes from grayscale digital images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 1646–1658, 05 2011.
- [6] O. Delgado-Friedrichs, V. Robins, and A. Sheppard, "Morse theory and persistent homology for topological analysis of 3D images of complex materials," in *2014 IEEE International Conference on Image Processing (ICIP)*, Oct. 2014, pp. 4872–4876.
- [7] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, "Stability of persistence diagrams," *Discrete & computational geometry*, vol. 37, no. 1, pp. 103–120, 2007.
- [8] N. Amenta, S. Choi, T. K. Dey, and N. Leekha, "A Simple Algorithm for Homeomorphic Surface Reconstruction," in *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, ser. SCG '00. ACM, 2000, pp. 213–222. [Online]. Available: <http://doi.acm.org/10.1145/336154.336207>
- [9] P. Niyogi, S. Smale, and S. Weinberger, "Finding the Homology of Submanifolds with High Confidence from Random Samples," *Discrete Comput Geom*, vol. 39, no. 1–3, pp. 419–441, Mar. 2008. [Online]. Available: <http://link.springer.com/10.1007/s00454-008-9053-2>
- [10] F. Chazal, D. Cohen-Steiner, and A. Lieutier, "A Sampling Theory for Compact Sets in Euclidean Space," *Discrete Comput Geom*, vol. 41, no. 3, pp. 461–479, Apr. 2009. [Online]. Available: <http://link.springer.com/10.1007/s00454-009-9144-8>
- [11] D. Attali, A. Lieutier, and D. Salinas, "Vietoris-rips complexes also provide topologically correct reconstructions of sampled shapes," *27th Annual Symposium on Computational Geometry*, vol. 46, 05 2013.

- [12] J. Kim, J. Shin, F. Chazal, A. Rinaldo, and L. Wasserman, “Homotopy reconstruction via the cech complex and the veltoris-rips complex,” in *36th International Symposium on Computational Geometry (SoCG 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [13] P. Bendich, H. Edelsbrunner, and M. Kerber, “Computing robustness and persistence for images,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, pp. 1251–1260, 2010.
- [14] P. Dłotko and T. Wanner, “Rigorous cubical approximation and persistent homology of continuous functions,” *Computers & Mathematics with Applications*, vol. 75, no. 5, pp. 1648–1666, Mar. 2018.
- [15] P. W. Botha and A. P. Sheppard, “Mapping permeability in low-resolution micro-CT images: A multiscale statistical approach,” *Water Resources Research*, vol. 52, no. 6, pp. 4377–4398, Jun. 2016.
- [16] R. Huang, A. L. Herring, and A. Sheppard, “Effect of Saturation and Image Resolution on Representative Elementary Volume and Topological Quantification: An Experimental Study on Bentheimer Sandstone Using Micro-CT,” *Transport in Porous Media*, vol. 137, no. 3, pp. 489–518, Apr. 2021.
- [17] B. Bleile, A. Garin, T. Heiss, K. Maggs, and V. Robins, “The persistent homology of dual digital image constructions,” <https://arxiv.org/abs/2102.11397v1>, 2021.
- [18] H. Blum, “A transformation for extracting descriptors of shape,” in *Models for the Perception of Speech and Visual Forms*. MIT Press, 1967, pp. 362–380.
- [19] H. Federer, “Curvature measures,” *Transactions of the American Mathematical Society*, vol. 93, no. 3, pp. 418–491, 1959.
- [20] C. Thäle, “50 Years of sets with positive reach, a survey.” *Surveys in Mathematics and its Applications*, vol. 3, pp. 123–165, 2008.
- [21] G. Tauzin, U. Lupo, L. Tunstall, J. B. Pérez, M. Caorsi, A. Medina-Mardones, A. Dassatti, and K. Hess, “giotto-tda: A topological data analysis toolkit for machine learning and data exploration,” 2020.
- [22] “PH-of-Images GitHub Repository,” <https://github.com/sarahtymochko/PH-of-Images/>.
- [23] A. Sheppard and M. Prodanovic, “Network generation comparison forum,” <http://www.digitalrocksportal.org/projects/16>, 2015.
- [24] V. Robins, M. Saadatfar, O. Delgado-Friedrichs, and A. P. Sheppard, “Percolating length scales from topological persistence analysis of micro-CT images of porous materials,” *Water Resources Research*, vol. 52, no. 1, pp. 315–329, Jan. 2016.

APPENDIX A

FEDERER’S LEMMA ABOUT THE REACH

This section describes a lemma about the reach, proven by Federer in [19, Theorem 4.8 (12)], stating that when walking from a point $a \in \partial A$ orthogonally away from a closed set A for distance $r < \text{reach}(A)$, then the closest point of A is still the starting point a , and thus the distance to A is r . Before we can state this rigorously in Lemma A.3, we need formal definitions [19, Definitions 4.3 and 4.4] of tangent vectors and normal vectors:

Definition A.1 (Tangent vector). Let $A \subseteq \mathbb{R}^d$ be closed and $a \in \partial A$. Then $u \in \mathbb{R}^d$ is a *tangent vector of A at a* if either $u = 0$ or for every $\varepsilon > 0$ there exists a point $b \in A$ with

$$0 < \|b - a\| < \varepsilon \text{ and } \left\| \frac{b - a}{\|b - a\|} - \frac{u}{\|u\|} \right\| < \varepsilon.$$

Definition A.2 (Normal vector). Let $A \subseteq \mathbb{R}^d$ be closed and $a \in \partial A$. Then $v \in \mathbb{R}^d$ is a *normal vector of A at a* if for every tangent vector u of A at a , the scalar product $v \cdot u$ is non-positive.

Lemma A.3 (Federer). Let $A \subseteq \mathbb{R}^d$ be closed and $a \in \partial A$. Let $\text{reach}(A) > r > 0$. Let v be a normal vector of A at a . Then,

$$d(a + r \frac{v}{\|v\|}, A) = d(a + r \frac{v}{\|v\|}, a) = r.$$

APPENDIX B TIGHTER BOUNDS

As mentioned in Remark IV.10, the bound on the bottleneck distance in Corollary IV.9 can be tightened from $3\sqrt{dr}$ to $2\sqrt{dr}$. To prove this, we first need to prove the tighter version of Lemma IV.3 mentioned in Remark IV.4:

Lemma B.1. Let $A \subseteq \mathbb{R}^d$ have boundary with positive reach, and let $B \subseteq \mathbb{R}^d$.

If $\max\{d_H(A, B), d_H(A^c, B^c)\} < \text{reach}(\partial A)$, then

$$\|d_A^\mp - d_B^\mp\|_\infty \leq \max\{d_H(A, B), d_H(A^c, B^c)\}.$$

Proof. Similarly to the proof of Lemma IV.3, we distinguish between 4 different cases for an arbitrary point p . Cases 1 and 2 stay unchanged.

Case 3: $p \in A$, $p \notin B$. As in the proof of Lemma IV.3, we need to bound $|d_A^\mp(p) - d_B^\mp(p)| = |-d(A^c, p) - d(p, B)| = d(A^c, p) + d(p, B)$. To prove Case 3, we distinguish between two sub-cases.

Case 3a: $p \in \partial A \cap A$, $p \notin B$. As $p \in \partial A$, there is a sequence in A^c converging to p , proving $d(A^c, p) = 0$. Thus,

$$\begin{aligned} |d_A^\mp(p) - d_B^\mp(p)| &= d(A^c, p) + d(p, B) \leq 0 + d_H(A, B) \\ &\leq \max\{d_H(A, B), d_H(A^c, B^c)\}. \end{aligned}$$

Case 3b: $p \in \text{int}(A)$, $p \notin B$. Let $a \in \partial \text{cl}(A^c)$ be such that $d(A^c, p) = d(\text{cl}(A^c), p) = d(a, p)$. As $a \neq p$, the vector $p - a$ has non-zero length. From $d(\text{cl}(A^c), p) = d(a, p)$ follows that $n_a = \frac{p - a}{\|p - a\|}$ is a unit normal vector of $\text{cl}(A^c)$ at x (for a rigorous definition of normal vector, see Definition A.2). Let $\varepsilon \in (0, \text{reach}(\partial A) - d_H(A^c, B^c))$. Let $x = a + (d_H(A^c, B^c) + \varepsilon)n_a$. As $d_H(A^c, B^c) + \varepsilon < \text{reach}(\partial A) \leq \text{reach}(\text{cl}(A^c))$, Lemma A.3 yields

$$\begin{aligned} d(A^c, x) &= d(\text{cl}(A^c), x) = d(a, x) = d_H(A^c, B^c) + \varepsilon \\ &> d_H(A^c, B^c). \end{aligned}$$

Hence, by the definition of Hausdorff distance, x cannot be in B^c and thus $x \in B$. The points a , p , and x lie on a straight line by construction. To prove $d(a, p) < d(a, x)$, let us assume $d(a, p) \geq d(a, x)$ which gives a contradiction between $d(A^c, p) = d(a, p) \geq d(a, x) > d_H(A^c, B^c)$ and $p \in B^c$. Therefore, p lies on the line segment between a and x . With this we can bound $d(A^c, p) + d(p, B) \leq d(a, p) + d(p, x) = d(a, x) = d_H(A^c, B^c) + \varepsilon$. As this bound is true for every $\varepsilon \in (0, \text{reach}(A^c) - d_H(A^c, B^c))$, we follow

$$\begin{aligned} |d_A^\mp(p) - d_B^\mp(p)| &= d(A^c, p) + d(p, B) \leq d_H(A^c, B^c) \\ &\leq \max\{d_H(A, B), d_H(A^c, B^c)\}. \end{aligned}$$

Case 4: $p \notin A$, $p \in B$. Analogously $|d_A^\mp(p) - d_B^\mp(p)| \leq \max\{d_H(A, B), d_H(A^c, B^c)\}$. \square

With this we can prove the tighter version of Corollary IV.9, mentioned in Remark IV.10:

Corollary B.2. Using the notation of Theorem IV.1. If $r < \frac{1}{\sqrt{d}}\text{reach}(\partial X)$, then

$$d_B(\text{PD}(d_X^\mp), \text{PD}(D_r)) \leq 2\sqrt{dr}$$

Proof. Note that $\max\{d_{\mathcal{H}}(X, X(r, t)), d_{\mathcal{H}}(X^c, X(r, t)^c)\}$ is the maximum of the 4 suprema from Lemma IV.7. We thus use Lemma IV.7 to bound $\max\{d_{\mathcal{H}}(X, X(r, t)), d_{\mathcal{H}}(X^c, X(r, t)^c)\}$ by $\max\{l_X(\sqrt{dr}), \sqrt{dr}\}$, which is \sqrt{dr} by Lemma IV.8.

Similar to the proof of Lemma IV.6, we combine the stability theorem of persistent homology [7], the triangle inequality, and Lemma IV.2, now with the new Lemma B.1:

$$\begin{aligned} d_B(\text{PD}(d_X^{\mp}), \text{PD}(D_r)) &\leq \|d_X^{\mp} - D_r\|_{\infty} \\ &\leq \|d_X^{\mp} - d_{X(r, t)}^{\mp}\|_{\infty} + \|d_{X(r, t)}^{\mp} - D_r\|_{\infty} \\ &\leq \max\{d_{\mathcal{H}}(X, X(r, t)), d_{\mathcal{H}}(X^c, X(r, t)^c)\} + \sqrt{dr} \\ &\leq 2\sqrt{dr}. \end{aligned}$$

□

APPENDIX C

BOTTLENECK DISTANCES WITH RESPECT TO PIXEL SIZE

For all our examples we present the plots of bottleneck distances vs. image resolution which can exhibit the $\Delta\text{-}\varepsilon$ plateau behavior described in Section V-A. However there is another way of visualizing this, considering bottleneck distance vs. pixel size, defined in Section II-B and denoted by r throughout the paper. Figures 9, 10, and 11 present this alternate perspective on Figures 6, 7, and 8, respectively.

The important thing to note is that in the figures presented in terms of resolution, the x -axis goes from low to high resolution from left to right. However, in the figures presented here in terms of pixel size, the highest resolution image (considered our “ground truth”) has a pixel size of 1, thus the x -axis in these plots go from high to low resolution from left to right. Thus, the plateau behavior we look for occurs from right to left in the plots with respect to pixel size.

APPENDIX D

PERSISTENCE DIAGRAMS FROM MATERIAL SCIENCE EXAMPLES

Section V-B presents three material science examples: glass bead packing, a Castlegate sandstone sample, and a sand packing sample. The original binary images are available from [23]. In Figures 12, 13, and 14, we show the 0- and 1-dimensional persistence diagrams for the 3D images computed using several resolutions for the bead packing, Castlegate sandstone, and sand packing, respectively.

As shown in [24], the percolation threshold, l_c , can be determined from the distribution of points in the zero-dimensional persistence diagram. This threshold is the radius of the largest sphere that can pass through the pore space from one side of the image to the opposite and is an important physical parameter associated with porous materials.

In the three sets of examples shown in Figures 12-14, the distribution of points in $\text{PD}_0(D_n)$ shows a clear signature of this critical length scale. This signature yields the same estimate for l_c for image resolutions $n = 512, 256$, and 128 in the three example materials.

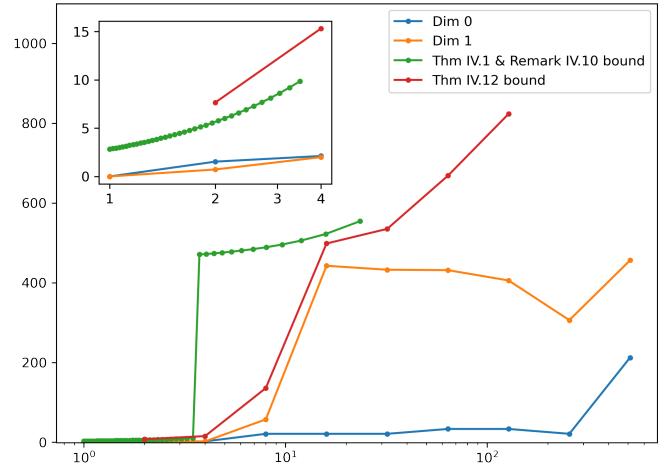


Fig. 9. Bottleneck distances comparing persistence diagrams for each pixel size with the image with pixel size of 1. The distances for dimensions 0 and 1 are shown along with the bounds provided by Theorem IV.1 with Remark IV.10 (green) and Theorem IV.12 (red). Note that the x -axis is on a logarithmic scale. This figure contains the same information as Fig. 6 but presented in terms of pixel size instead of resolution.

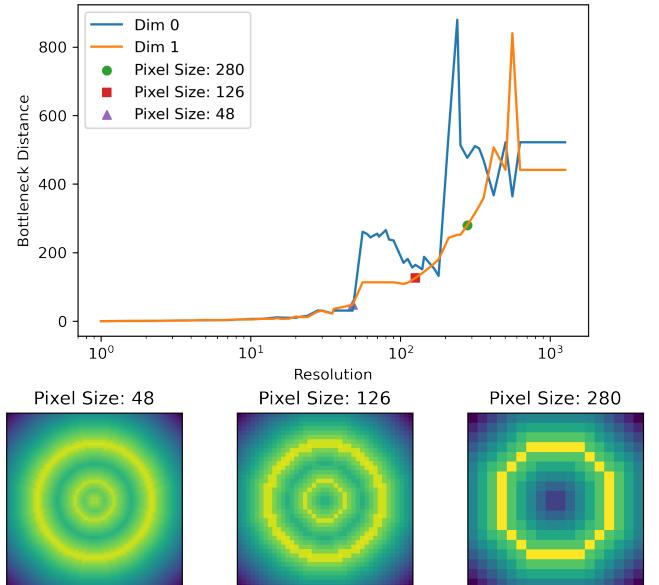


Fig. 10. Top: the plot of the bottleneck distances between 1-dimensional persistence diagrams from each pixel size and the image with pixel size of 1. Note that the x -axis is on a logarithmic scale. Bottom: the DSEDT of the downsampled image at three different pixel sizes. This figure contains the same information as Fig. 7 but presented in terms of pixel size instead of resolution.

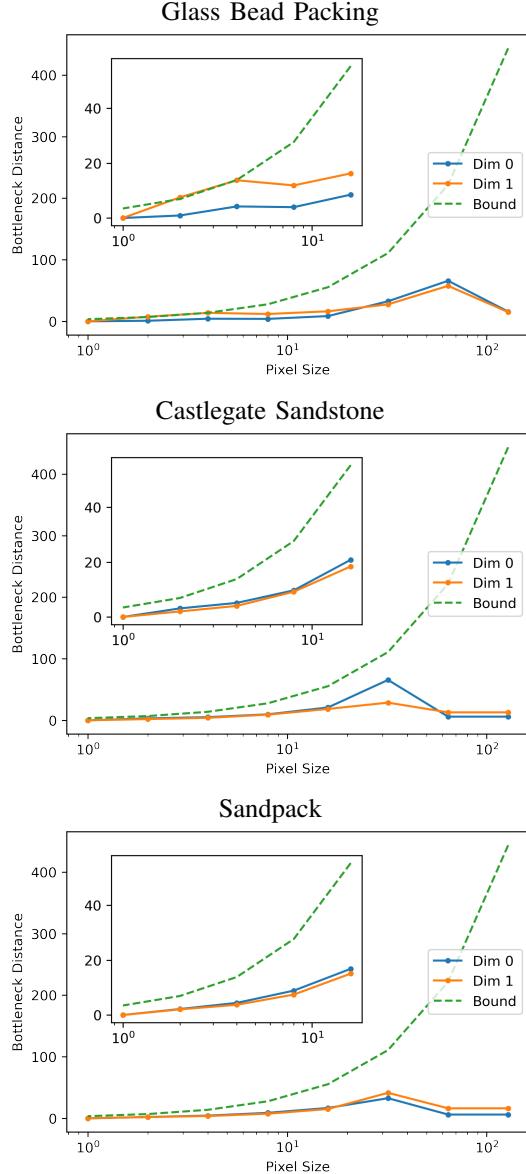


Fig. 11. Bottleneck distance between persistence diagrams for each larger pixel size and the image with pixel size of 1 for each of the three material science examples. The green dashed curve is the function $2r\sqrt{3}$, where r is the pixel size. Note that the x -axes are on a logarithmic scale. This figure contains the same information as Fig. 8 but presented in terms of pixel size instead of resolution.

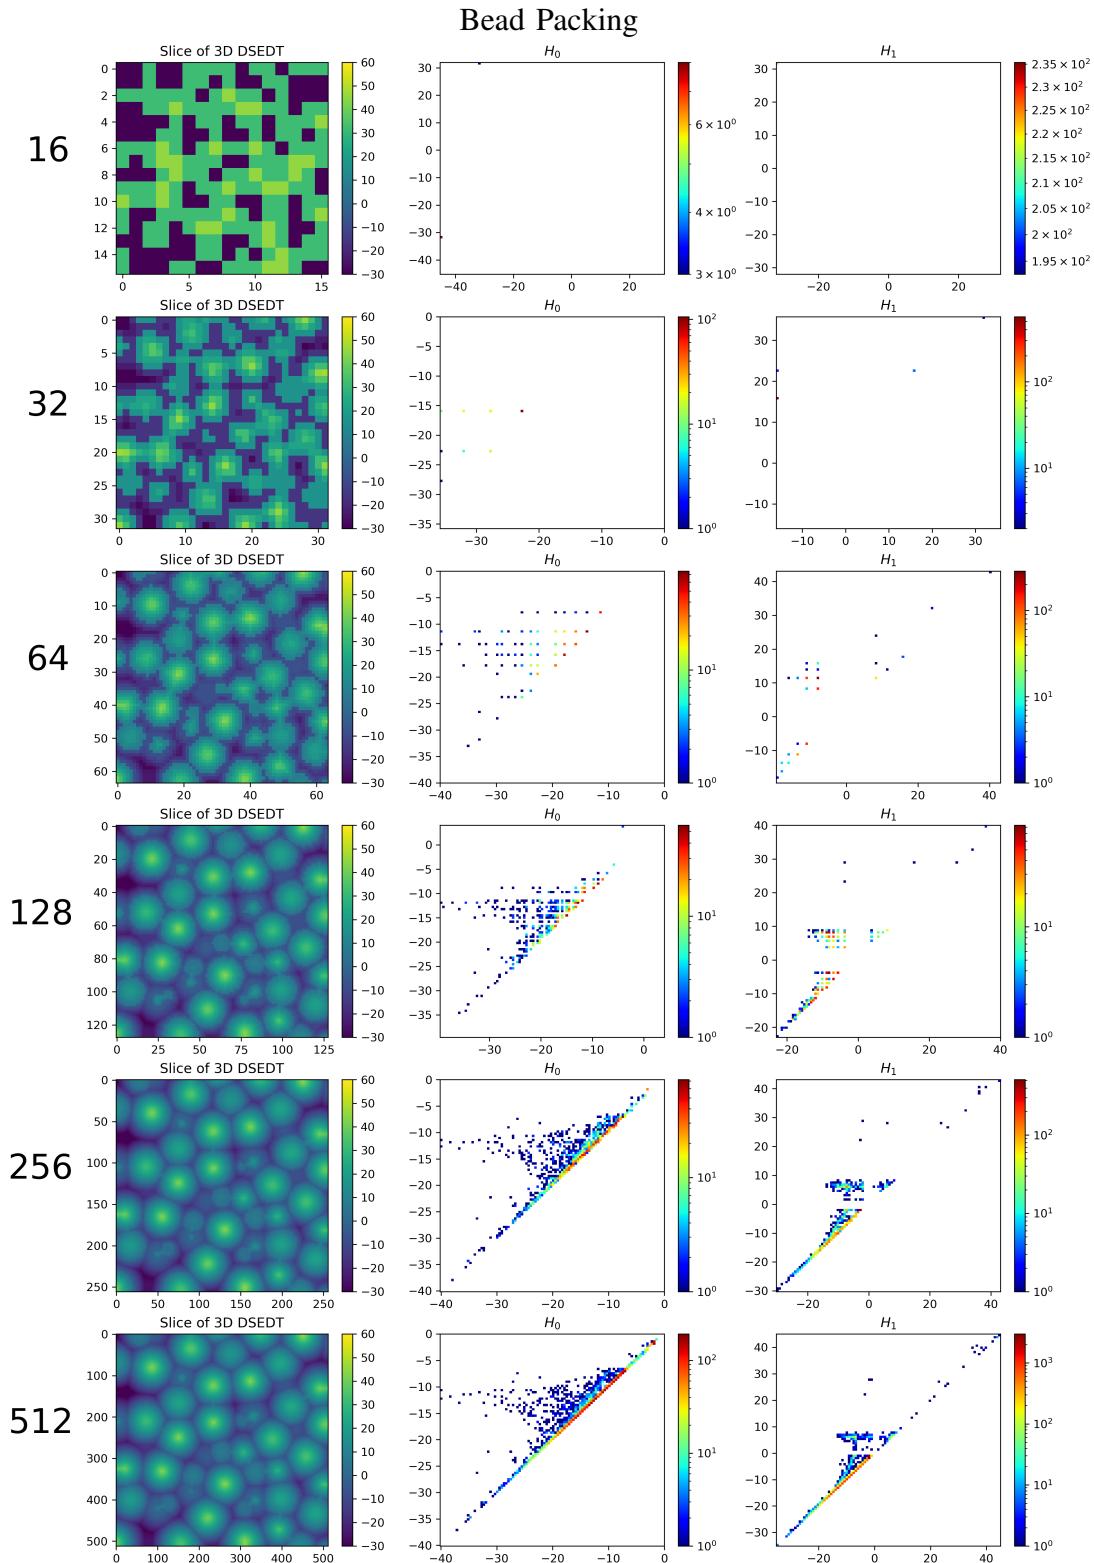


Fig. 12. Each row shows a slice of the 3D SEDT, and the 0 and 1-dimensional persistence diagrams of the bead packing sample at different resolutions.

Castlegate Sandstone

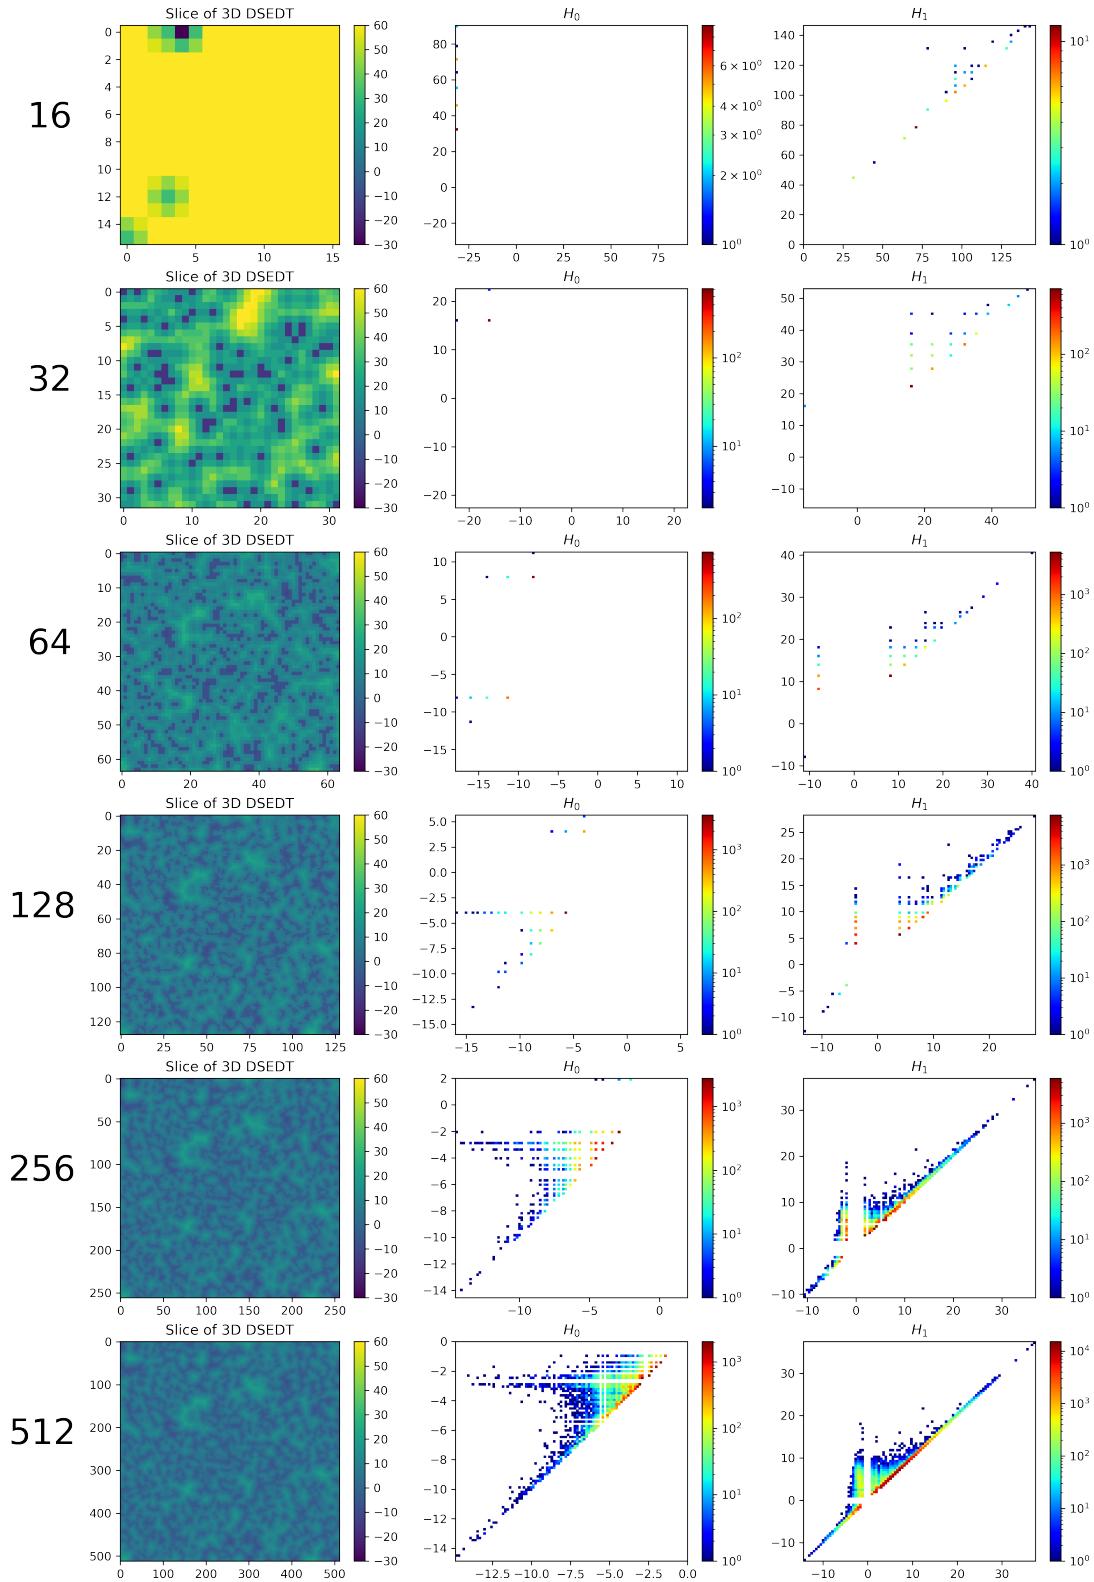


Fig. 13. Each row shows a slice of the 3D SEDT, and the 0 and 1-dimensional persistence diagrams of the Castlegate sandstone sample at different resolutions.

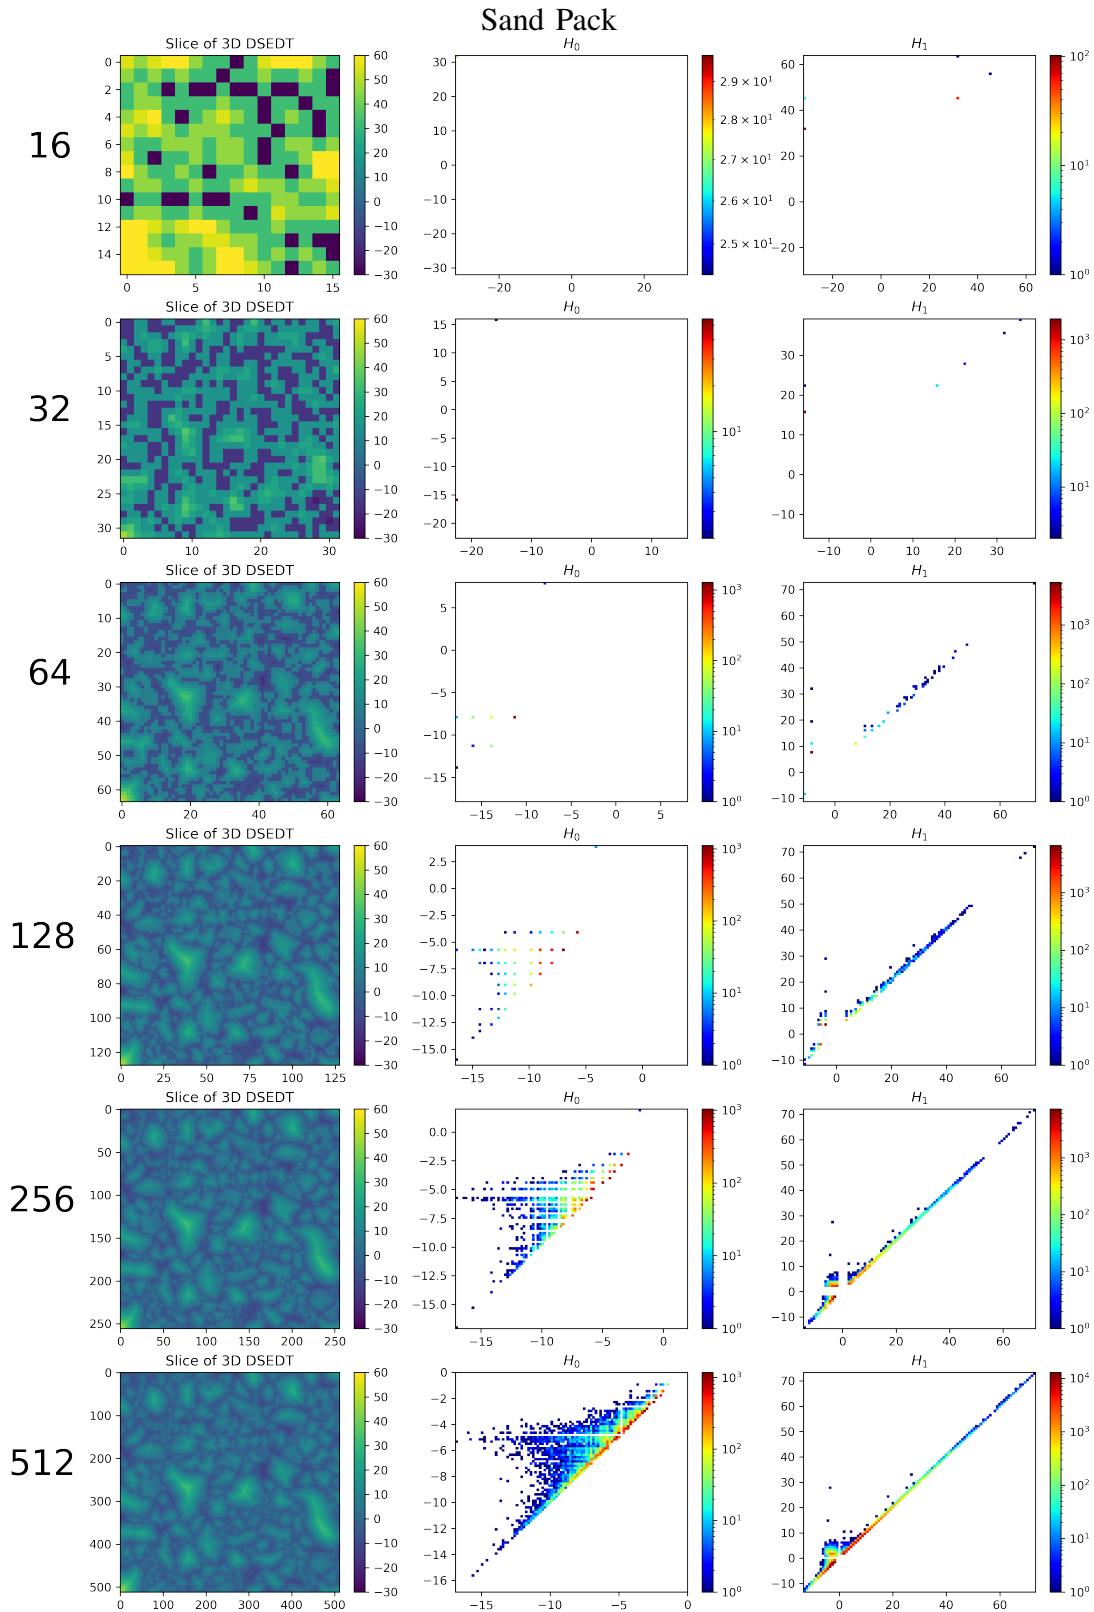


Fig. 14. Each row shows a slice of the 3D SEDT, and the 0- and 1-dimensional persistence diagrams of the sand packing sample at different resolutions.