

Feature Detection and Hypothesis Testing for Extremely Noisy Nanoparticle Images using Topological Data Analysis

Andrew M. Thomas*

Center for Applied Mathematics, Cornell University,

Peter A. Crozier

School for Engineering of Matter, Transport and Energy, ASU,

Yuchen Xu

Department of Statistics and Data Science, Cornell University,

David S. Matteson

Department of Statistics and Data Science, Cornell University

March 31, 2023

*The authors would like to thank the two reviewers and associate editors for their insightful comments which greatly aided in the presentation of the paper. Furthermore, the authors would like to acknowledge Joshua Vincent, Piyush Haluai and Mai Tan at Arizona State University (ASU) for sample preparation and TEM data acquisition. We also thank Ramon Manzorro (University of Cádiz) and Advait Gilankar (ASU) for assistance in imaging simulations and data handling as well as the John M. Cowley Center for High Resolution Electron Microscopy at ASU. The authors gratefully acknowledge funding from NSF grants OAC-1940124, CCF-1934985, DMS-2114143 and specifically CBET-1604971 and OAC-1940263 for the support of sample preparation and data acquisition.

Abstract

We propose a flexible algorithm for feature detection and hypothesis testing in images with ultra-low signal-to-noise ratio using cubical persistent homology. Our main application is in the identification of atomic columns and other features in transmission electron microscopy (TEM). Cubical persistent homology is used to identify local minima and their size in subregions in the frames of nanoparticle videos, which are hypothesized to correspond to relevant atomic features. We compare the performance of our algorithm to other employed methods for the detection of columns and their intensity. Additionally, Monte Carlo goodness-of-fit testing using real-valued summaries of persistence diagrams derived from smoothed images (generated from pixels residing in the vacuum region of an image) is developed and employed to identify whether or not the proposed atomic features generated by our algorithm are due to noise. Using these summaries derived from the generated persistence diagrams, one can produce univariate time series for the nanoparticle videos, thus providing a means for assessing fluxional behavior. A guarantee on the false discovery rate for multiple Monte Carlo testing of identical hypotheses is also established.

Keywords: Cubical persistent homology; ALPS statistic; Catalysis; Persistent entropy; Multiple Monte Carlo testing; Transmission electron microscopy

1 Introduction

Transmission electron microscopy (TEM) has become a critical tool in both physical and life sciences for characterizing materials at the atomic level. Over the last 10 years, recent advances in direct electron detectors have greatly improved sensitivity with detective quantum efficiencies approaching the theoretical maximum of unity (Ruskin et al., 2013; Faruqi and McMullan, 2018; Levin, 2021). As a result, the information content in signals is now limited mostly by the Poisson shot noise associated with the quantum mechanical processes responsible for electron emission and scattering. In principle, for a perfect detector, the fraction of noise in the signal can be made arbitrarily small by counting for longer or by increasing the flux of the incident electron beam. However, increasing the measure-

ment time/electron flux is simply not practical for many materials systems, as they are irreversibly damaged by the electron beam (Egerton et al., 2004; Egerton, 2013, 2019).

The signal-to-noise ratio is also significantly limited when high temporal resolution is required for investigations of dynamic behavior associated with kinetic processes in materials. In such experiments, the exposure time per frame is necessarily short resulting in a high degree of shot noise in each frame. Recent efforts to understand structural dynamics in catalytic nanoparticles have been significantly impacted by the challenges associated with high degrees of noise (Lawrence et al., 2019; Levin et al., 2020; Lawrence et al., 2021; Vincent and Crozier, 2021). Moreover, the large number of noisy image frames required to fully map out the details of the spatio-temporal behavior requires the collection of large image data sets—on the order of terabytes (Lawrence et al., 2019). As noted in Lawrence et al. (2019), there is a continuing need for algorithms to automate the extraction of relevant features from images to facilitate the assessment of dynamics, and to do so in the presence of an immense amount of noise. Denoising methods based on convolutional neural networks trained on simulated nanoparticle configuration images have been developed specifically to deal with such ultra-noisy nanoparticle videos (Mohan et al., 2022). However, it is not always be feasible to implement these methods in practice, frame-by-frame, on the aforementioned terabytes of data.

In this article, we present an algorithm for detecting features in, and hypothesis testing of, severely noisy images based off of topological data analysis (TDA)—more specifically, cubical persistent homology (cPH) (Edelsbrunner et al., 2002; Kaczynski et al., 2006; Mischaikow and Nanda, 2013). Cubical homology is more naturally suited to imaging as it treats images and their connectivity in a natural manner, with no need for triangulation of the inherently pixelated data (Kaczynski et al., 2006). Indeed, cubical persistent homology (Garin and Tauzin, 2019; Rieck et al., 2020; Lawson et al., 2021; Chung and Day, 2018) has been applied to great effect in the statistics, machine learning, and imaging communities in recent years. Additionally, the algorithms which compute cubical persistent homology are often much faster than their simplicial counterparts (Wagner et al., 2012). We will discuss the background behind these concepts and provide intuition for them in Section 2.

Topological methods as used in this paper have the advantage over traditional methods in materials science in that they are isometry invariant. Cubical homology in particular is invariant to translations, thus providing robustness against minor perturbations of atomic columns and ridges across frames. In the previous few years, applications of methods in TDA to materials science (beyond cPH) have seen greater adoption. In [Motta et al. \(2018\)](#), the authors use functionals of persistent homology, such as the variance of H_0 barcode lengths, to characterize the order of a nearly hexagonal planar lattice. In cluster physics, [Chen et al. \(2020\)](#) examined the ability of topological features and machine learning to predict ground-state structure-energy relationships in lithium clusters. In a similar context, [Jiang et al. \(2021\)](#) examined a topological invariant called “atom-specific persistent homology” to predict formation energy of crystal structures. [Nakamura et al. \(2015\)](#) used persistence diagrams to characterize medium-range order in amorphous materials.

Here we use methods from TDA to quantify the fluxional behavior of nanoparticles. As alluded to earlier, we identify atomic columns (see Figure 1) via their estimated *persistence*—quantified within each individual frame as the difference in greyscale threshold at which a given dark region appears and when it merges with another dark region that appeared before it.¹ The appearance of a connected component in our images corresponds to the appearance of a local minima. Though methods such as [Mukherjee et al. \(2020\)](#); [Nord et al. \(2017\)](#) use the local minima of images as initial locations for fitted Gaussians (from which the intensity is estimated), here we estimate the intensity by the “lifetimes” of the local minima, i.e. the concept of persistence that we defined at the start of this paragraph. Thus, the lifetime of a local minimum (its persistence) is defined as the difference in pixel values between the appearance of the local minimum and the pixel value at which it merges with another longer-lived local minimum.

Recently, there have been forays into using the statistics that appear in TDA in hypothesis and goodness-of-fit testing ([Biscio and Møller, 2019](#); [Blumberg et al., 2014](#); [Fasy et al., 2014](#); [Robinson and Turner, 2017](#); [Cericola et al., 2017](#); [Vejdemo-Johansson and Mukherjee, 2022](#)). The present article is the first to attempt this in a cubical setting and the first to

¹Such a paradigm is called the *elder rule* and is described in detail in [Edelsbrunner and Harer \(2010\)](#).

evaluate the efficacy of certain topological summaries to capture relevant topological features in the presence of powerful noise. In particular, both *persistent entropy* (Chintakunta et al., 2015; Rucco et al., 2016) and the *accumulated lifetime persistent survival* (ALPS) statistic—a new topological summary—perform well and evince good statistical power. It is our opinion that they would work well in a litany of tasks in summarizing noisy videos, particularly the ALPS statistic when the number of features and their intensity are both salient. On a final note, we prove a result demonstrating the false discovery rate of a certain multiple Monte Carlo test is bounded by any α almost surely (see Supplementary Material for the result and proof). This yields a theoretically sound—as well as computationally efficient—means of multiple testing in persistent homology, improving on previous studies in the area (Cericola et al., 2017; Vejdemo-Johansson and Mukherjee, 2022).

It is our hope that the algorithm and hypothesis testing framework we have devised here can provide an off-the-shelf method of statistical detection of atomic structure in a flexible manner for those in the material science community and others that deal with necessarily noisy images. Our method performs well in the standard nanoparticle imaging task of determining position and location of atomic columns (Nord et al., 2017; Levin et al., 2020) and also performs well against the state-of-the-art (Manzorro et al., 2022). Additionally, these topological methods do not presume a particular structure to the image data. For example, Gaussian peak fitting and blob detection (ibid.) both assume an elliptical structure to features in the image which are present when individual columns of atoms are well aligned. However, in many cases, the crystal may be tilted and individual columns may not be resolved but planes of atoms may be visible. Moreover, during structural dynamics and in the presence of high concentrations of crystal defect, the structure of the image intensity may be complicated and rapidly changing. To elucidate structural dynamics, it is important to have an image analysis method that can adapt to the changing structure of the image contrast and does not presuppose a particular image form. Finally, between the persistence entropy and the ALPS statistic, we offer a choice in how conservative a practitioner wants to be in determining which atomic features are statistically significant.

In the following, we will discuss the cubical persistence algorithm we devised to process

the extremely noisy videos at high time resolution in Section 3, along with the ability of said algorithm to recover atomic features in simulated datasets before and after the application of noise in Section 4. The main statistical contribution is in Section 5, wherein we describe the parametric assumptions of the noise region of the data, check those assumptions, investigate various topological summaries of persistence diagrams that capture relevant atomic features, detail our Monte Carlo goodness-of-fit test, and see the results on the data.

Before continuing to the description of our algorithm and the rest of the paper, we must first introduce TDA and detail the concepts of cubical homology and persistence.

2 Background

Topological data analysis (TDA) is the field consisting of the methods and tools that directly integrate shape information into the data analysis process. As one can imagine, TDA is most often concerned with topological methods of quantifying shape, rather than explicitly geometric content, such as location and angle. For a good up-to-date overview of these methods, see [Chazal and Michel \(2021\)](#), or [Carlsson and Vejdemo-Johansson \(2022\)](#) for a more in-depth treatment. In the following sections we detail methods for topological analysis of image data.

2.1 Cubical sets

The tool that we use to assess shapes in images in this paper is cubical persistent homology (cPH). One reason for considering a cubical representation of an image is that it is the most natural construction for a 2-dimensional image from the perspective of topology, as detailed in [Kovalevsky \(1989\)](#). Another reason is that cubical persistence algorithms run in linear time in the number of pixels n in the image \mathcal{I} , whereas traditional approaches to persistent homology, such as using the Čech or Vietoris-Rips complex, can only be computed in polynomial time in the number of points in the point cloud ([Wagner et al., 2012](#)). To introduce cPH we must introduce the notion of cubical homology and the objects it acts on: cubical sets. The cubical sets we consider here are collections of unit squares

(2-dimensional elementary cubes) of the form

$$[i, i + 1] \times [j, j + 1],$$

along with all intervals (1-dimensional elementary cubes) and vertices (0-dimensional elementary cubes) on the boundaries, where i and j are integers—i.e. $(i, j) \in \mathbb{Z}^2$. Once we have a cubical set (or *cubical complex*) X , we can calculate *homology*. Homology is an algebraic method of formally characterizing connectivity in various dimensions ([Edelsbrunner and Harer, 2010](#)). Given a cubical set X , we can associate a homology group $H_k(X)$ to it—which captures k -dimensional shape information—for each nonnegative integer k . Of great interest are the dimensions of these homology groups, which are called the *Betti numbers* of X and are denoted $\beta_k(X)$, or β_k when the underlying cubical set is clear from the context. The 0th Betti number β_0 represents the number of connected components in X and β_1 represents the number of loops. As the relevant features we aim to capture are darker than their surroundings, we focus on β_0 . One can see Figure 1 for the calculation of β_0 at various greyscale thresholds. In Figure 1 the black pixels are the ones included in our cubical set—in this sense all cubical sets/complexes that we treat here can be considered as binary images. For more information on cubical homology, one may refer to Chapter 2 of [Kaczynski et al. \(2006\)](#).

2.2 Image model

In this article, a (2-dimensional) *image map* is a function $I : \mathbb{Z}^2 \rightarrow [0, 1]$, where $I(p) = 0$ indicates that p is a black pixel and $I(p) = 1$ indicates that p is a white pixel. We call the smallest rectangle $[k, k+m] \times [l, l+n] \subset \mathbb{R}^2$ which contains all the black pixels—i.e. on which $1 - I > 0$ —the *image set*, which we denote \mathcal{I} . Via appropriate normalization, every image that is bounded or has a finite image set can be modified to have pixels between 0 and 1. We may choose to set the codomain of I to \mathbb{R} (or the first n nonnegative integers) instead. As previously mentioned, for the purposes of cubical homology and cubical persistence we must identify I in some appropriate way with a collection of cubical sets. We do this here by the construction of another function I' on the family of unit squares with integer

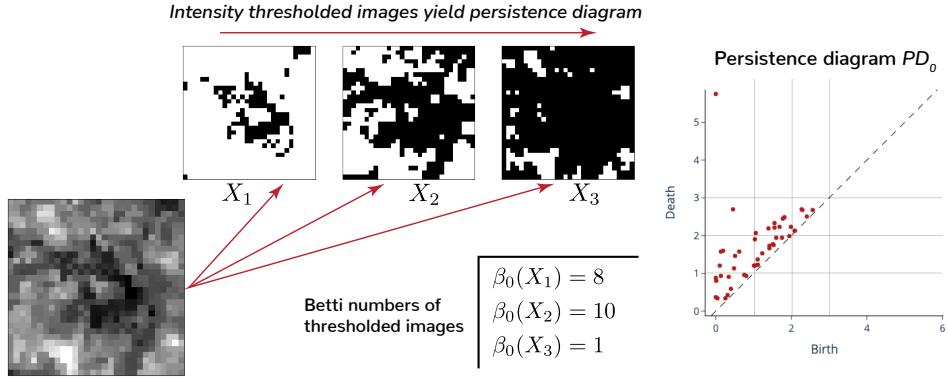


Figure 1: Thresholded images X_t of a single atomic column with greyscale values at most t (where $t = 1, 2, 3$). The persistence diagram PD_0 of the atomic column in the bottom left can be seen to the right. The persistence diagram depicts the pixel values at which a given black-connected region—a feature in H_0 —appears (i.e. is “born”) and when it merges or “dies”. β_0 of X_1 , X_2 , and X_3 can be seen in the table in the center.

vertices. For any such $\tau = [i, i + 1] \times [j, j + 1]$ we define our filtration function

$$I'(\tau) := I(i, j).$$

For lower dimensional elementary cubes τ , such as intervals or vertices, we define the value of I' to be the minimum $I(i, j)$ such that $\tau \subset [i, i + 1] \times [j, j + 1]$. This is consistent with the definition used in the persistent homology software **GUDHI** Python library, which we use for our calculations throughout this article (Dłotko, 2015). We consider the homology of sublevel set filtrations², so darker pixel values will appear first. The cubical complex construction we use here, treating pixels as unit squares (i.e. top-dimensional) is also known as the T -construction and it is dual in some sense to treating pixels as points (or vertices), rather than unit squares (Garin et al., 2020).

Remark 2.1. An important heuristic (see Figures 1, 2, and 3) is that in a binary image *black* pixels $p = (i, j), q = (k, l) \in \mathbb{Z}^2$ are connected if $|i - k| \leq 1$ and $|j - l| \leq 1$. Thus, a

²Other filtrations could be considered here; however, besides the sublevel set filtration, all require choosing a threshold at which to binarize the image (Turkes et al., 2021; Garin and Tausz, 2019)—see also the opening paragraph of Section 4.

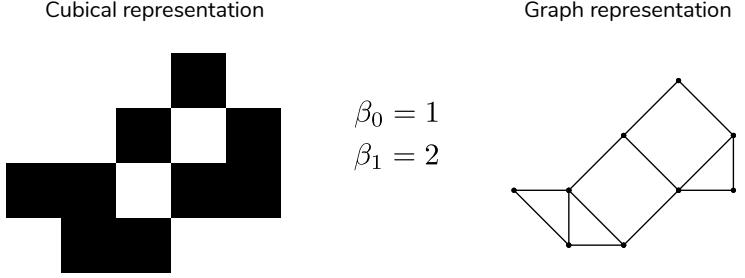


Figure 2: Connectivity information in the cubical set to the left can be conveyed by the graph to the right. Note that there is one connected component and two loops (white regions surrounded by black), because the black region can connect via the diagonals and the white region cannot.

connected component in cubical homology (contributing to β_0) is a *8-connected black region of pixels bordered by either white pixels or the edge of the image*³. The equivalence between the notion of connectivity for the T -construction and 8-connectedness was established in [Kovalevsky \(1989\)](#).

2.3 Persistent homology

Suppose now that we have the collection of cubical complexes $\mathcal{X} = \{X_t\}_{t \in [0,1]}$, where

$$X_t := \bigcup_{(i,j) \in I^{-1}([0,t])} [i, i+1] \times [j, j+1],$$

or alternatively, $X_t = I'^{-1}([0, t])$. It is clear that for $s \leq t$ we have $X_s \subset X_t$ and thus $\mathcal{X} = \{X_t\}_{t \in [0,1]}$ defines a *filtration* of cubical complexes. Given the inclusion maps $\iota_{s,t}$, for $s \leq t$ there exist linear maps $f_k^{s,t} : H_k(X_s) \rightarrow H_k(X_t)$ between all homology groups, which are induced by $\iota_{s,t}$ (see chapter 4 of [Kaczynski et al., 2006](#)). The *persistent homology groups* of the filtered image \mathcal{X} are the vector spaces $\text{im } f_k^{s,t}$ whose elements represent shape features, such as connected components or loops—called *cycles*⁴—that are “born” in or before X_s and that “die” after X_t . The dimensions of these vector spaces are the *persistent*

³Recall that all pixels not in the image (image set) are de facto white pixels in terms of cubical homology.

⁴Technically speaking these are equivalence classes of cycles, which are equivalent modulo a boundary.

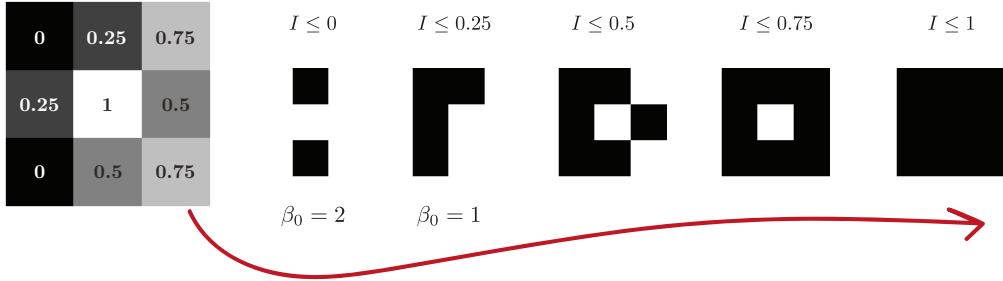


Figure 3: A simple 3×3 image, and the image as it is thresholded by pixel intensity I . Note the appearance of a loop at $I(p) = 0.5$ and the death of the loop entering $I(p) = 1$.

Betti numbers $\beta_k^{s,t}$. The k^{th} persistent homology of \mathcal{X} , denoted PH_k , is the collection of homology groups $H_k(X_t)$ and maps $f_{s,t}^k$, for $0 \leq s \leq t \leq 1$. Persistent homology keeps tracks of the thresholds s, t at which various shapes appear and when they merge with other shapes or disappear. All of the information in the persistent homology groups is contained in a multiset in \mathbb{R}^2 called the *persistence diagram*—see [Boissonnat et al. \(2018\)](#); [Edelsbrunner and Harer \(2010\)](#); these texts also serve as formal, yet accessible introductions to the underlying theory of persistent homology.

The k^{th} persistence diagram of \mathcal{X} , denoted PD_k , consists⁵ of the points (b, d) with multiplicity equal to the number of the cycles that are born at X_b and die entering X_d . For this study, we focus on PD_0 —see Figure 1 for an illustration of the 0^{th} persistence diagram associated to a filtration of a given greyscale image. In this particular setup, if $(b, d) \in PD_0$, this indicates there is a local minimum of the image \mathcal{X} at some pixel p^+ (called a *positive cell*) with $I'(p^+) = b$ and d represents the greyscale threshold at which the connected component containing p^+ merges with a connected component containing a local minimum with birth time *less than* b . Furthermore, we can also find an interval τ^- that kills such a feature, i.e. $I'(\tau^-) = d$ (cf. [Boissonnat et al., 2018](#)).

⁵Often, the diagonal $y = x$ is added to this diagram, but we need not consider this here.

3 The algorithm

In this section, we describe our algorithm for extracting shape, location, and intensity information from ultra-noisy images. To speed computation we may restrict our attention to a rectangular subimage $\mathcal{L} \subset \mathcal{I}$ —see Figure 6 for a depiction. Let us denote the restriction of I to \mathcal{L} by $I_{\mathcal{L}}$. Hence, we process our subimage $p \mapsto I(p)$ according to the following steps:

1. Identify polygonal⁶ nanoparticle region $R \subset \mathbb{R}^2$, which we will use to exclude pixels that lie outside of R .
2. Smooth the image with a Gaussian filter, with smoothing parameter σ .
3. Compute PH_0 for image $I_{\mathcal{L}}$, based off of the filtration function $I'_{\mathcal{L}}$. Note that one should have $R \subset \mathcal{L}$.
4. If the pixel p^+ associated to the creation of connected component is located outside of R , remove point associated to p^+ from PD_0 .
5. (Optional) Remove features with persistence lifetime at or below a threshold $\eta \geq 0$ from PD_0 .

The algorithm is illustrated in Figure 4. For our image I , we denote the output of this algorithm as $A(I)$, which consists of the locations of atomic columns (or other atomic features) as well as their persistence (or, intensity). As such, $A(I)$ may be considered as a finite subset of $R \times [0, \infty)$. Equivalently, we may consider $A(I)$ as a marked point process on R with mark space $[0, \infty)$, as we assumed that our image I is subject to noise. Additionally, we denote the thresholded output as $A_{\eta}(I)$, so that the original output may be considered as $A(I) = A_0(I)$. Formally $A_{\eta}(I) = \{(p, l) \in A(I) : l > \eta\}$, where $l = d - b$ are the *lifetimes* associated to the pixels p^+ . If we preprocess by restricting our image to \mathcal{L} , the algorithm requires only R and σ to be specified. We may choose σ based on the elbow method; see Section S6 of the Supplementary Material.

⁶For specifying polygonal regions and which pixels are contained in them, we use the `Shapely` Python library (Gillies, 2013).

With this in tow, we examine the algorithm in greater detail. For step 1, as cubical persistence does not consider the size of connected component per se, we remove image features (corresponding to atomic columns in our application) outside of some region, which we know either correspond to noise or the structure of which is of no interest. To smooth the image, we consider I convolved with the spherical Gaussian kernel G_σ , which we denote

$$I_\sigma(i, j) := (G_\sigma * I)(i, j) = \frac{1}{C} \sum_{(k, l) \in \mathbb{Z}^2} G_\sigma(i - k, j - l) I(k, l),$$

and

$$G_\sigma(x, y) := \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}},$$

where $C = \sum_{(x, y) \in \mathbb{Z}^2} G_\sigma(x, y)$. By convention let us take $G_0 = \delta_{(0,0)}$ be the Dirac delta function at the origin, i.e. $\delta_{(0,0)}(x, y) = 1$ if and only if $x = y = 0$.

We have chosen the spherical Gaussian kernel because of various desirable properties—see Section S4 in the Supplementary Material. We now compare the performance of the above algorithm with other methods of finding atomic column positions and show how it performs better in certain cases and works well as a method for finding initial positions of atomic features.

4 Noise experiment

In this section, we assess the ability of the algorithm we introduced in Section 3 to perform as well as the combined blob detection/gaussian peak fitting (GPF) method from (Manzorro et al., 2022) in recovering the number, location, and intensity of the atomic columns in a nanoparticle image corrupted by Poisson noise. Images were simulated and Poisson noise was added according to the method described in the section entitled “Generation of the Simulated Dataset” in Manzorro et al. (2022). Persistent homology is purported to be robust to noise; in Skraba and Turner (2020), the authors offer a cubical version of the classical persistence stability theorem (Cohen-Steiner et al., 2007), stating that if two image maps are close, then their persistence diagrams are close as well. There have been other efforts to quantify experimentally the noise robustness of cPH. In Turkes et al. (2021), the

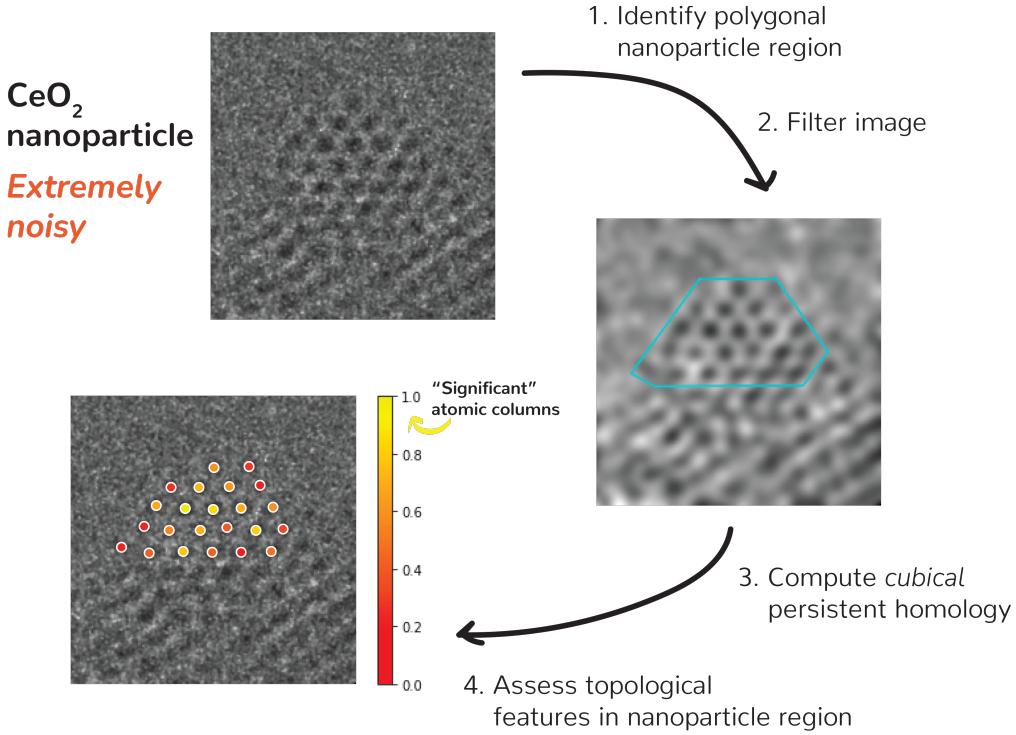


Figure 4: Illustration of our algorithm, applied to a cerium oxide (CeO_2) nanoparticle as described in [Lawrence et al. \(2019\)](#). Persistence in colorbar measured as proportion of longest barcode.

authors demonstrated the empirical robustness of sublevel set (greyscale) filtrations for cPH under affine transformations and additive noise (such as the Poisson noise encountered in our application). Their results buttress our argument that pre-smoothing and thresholding an image can faithfully recover the underlying topology.

Here, we assess the mean and standard error of three statistics related to the recovery of the homology of simulated nanoparticle images based on the output of our algorithm. Let us denote our smoothed noisy *simulated* images as $I_{\sigma,j}^s$, $j = 1, \dots, 10$. Let I_σ be the σ -smoothed version of the noise-free simulated image I . We assess the ability of our cubical homology algorithm applied to the noisy images, to recover atomic column position and intensity that the algorithm outputs on the noise-free image. The three statistics are the

number of columns output by the algorithm; the mean (Pearson) correlation $\hat{\rho}(I_\sigma, I_{\sigma,j}^s)$ of the intensity of the derived columns in the noisy output $A(I_{\sigma,j}^s)$ compared to the output of the noise-free image $A(I_\sigma)$; and the Hausdorff distance

$$d_H(I_\sigma, I_{\sigma,j}^s) := \max \left\{ \max_{(x,l) \in A(I_\sigma)} \min_{(y,l_j) \in A(I_{\sigma,j}^s)} \|x - y\|, \max_{(y,l_j) \in A(I_{\sigma,j}^s)} \min_{(x,l) \in A(I_\sigma)} \|y - x\| \right\}$$

between the locations of the columns in the algorithms output. In this section, if the death time equals $d = \infty$, we set d to be the largest finite death time in the persistence diagram (if we chose d to be the largest pixel value in the image, the death time for the longest barcode is a significant outlier). To assess the performance of each algorithm A (i.e. cPH vs. blob detection/GPF), we assess the output of A with the same parameter set Θ (such as $\Theta = \{\sigma\}$, using our algorithm) on both the noisy and noise-free simulated images. For the combined blob detection/GPF method⁷, the mean correlation was 0.9816 with standard error 0.0046 and the mean Hausdorff distance was 2.319 with standard error 1.088. The corresponding results for our cubical homology algorithm can be seen in Table 1.

For appropriate σ the mean Hausdorff distance was much better using our method, though this could be attributed to the fact that a variety of σ smoothing values (50 different values from 6 to 9) were used to find the best blobs in blob detection, whereas σ remained fixed for the ground truth image as well as the noisy image in our method. As one can see, the mean correlation performs similarly to blob detection, though for lower values of σ the actual intensities of the 25 true atomic columns are retained to a much higher degree, owing to less influence from surrounding pixels attenuating their signal. Setting the threshold η such that we choose the 25 largest persistence values, we achieve correlations of 0.893 (0.042) when $\sigma = 2$ and 0.961 (0.015) when $\sigma = 4$ between the derived intensities of I_σ and $I_{\sigma,j}^s$, but the Hausdorff distances are larger, at 4.46 (0.59) and 2.52 (0.92) respectively.

Besides the blob detection/GPF method described here, there are similar methods in the transmission electron microscopy community for finding atomic locations that iteratively fit Gaussian peaks with initial means often chosen to be local minima/maxima: Atomap

⁷The locations of the blobs/atomic columns was initially calculated using the `blob_log` function in the Python `skimage` library, as in [Manzorro et al. \(2022\)](#). The algorithm was applied in the same fashion as [Manzorro et al. \(2022\)](#) to ensure optimality of parameters chosen and a fair comparison of the methods.

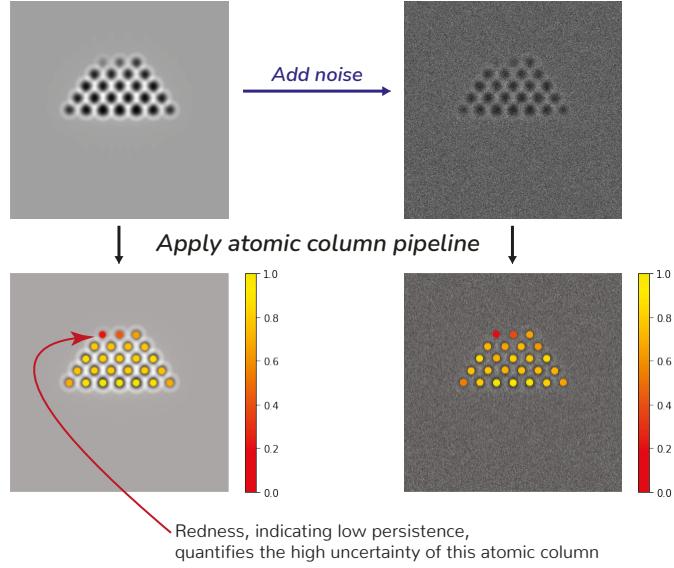


Figure 5: Recovery of the 0^{th} homology (i.e., atomic columns) in the presence of noise. Note that persistence in these images is measured in terms of the proportion of longest barcode.

(Nord et al., 2017), TRACT (Levin et al., 2020) and mpfit (Mukherjee et al., 2020). That local minima and their “intensities”, are used fruitfully in this instance (and stated to have limited value in Levin et al., 2020) is a testament to the efficacy of the global notion of the size of local minima here, rather than a local one.

Comparing these methods with ours yields mean Hausdorff distances of 4.76 (0.94) for the TRACT algorithm and 3.58 (1.39) for Atomap. This is perhaps unsurprising as these algorithms, along the blob detection/GPF approach of Manzorro et al. (2022), yield subpixel precision for the atomic column position. There is perhaps either not enough noise or sufficient smoothing, so that our algorithm does not shift atomic column positions too drastically from the ground truth, which demonstrates a form of spatial stability of the positive cells of persistent homology. The comparison of output intensities of TRACT and blob detection/GPF was done in Manzorro et al. (2022) so we do not replicate it here.

Blob detection using the Laplacian of Gaussian, used as part of the algorithm in Manzorro et al. (2022) and described in Lindeberg (1998) yields images seen in Figures S4a and

σ	0	2	4	6	8	10	12
$N(I_{\sigma,j}^s)$	3055	55	25.2	25	25	25	24.5
#Columns recovered	(22.08)	(5.02)	(0.4)	(0)	(0)	(0)	(0.5)
$d_H(I_\sigma, I_{\sigma,j}^s)$	26.61	23.82	5.16	1.74	2.21	4.19	21.52
Hausdorff distance	(0.35)	(1.02)	(5.68)	(0.88)	(1.15)	(2.39)	(17.26)
$\hat{\rho}(I_\sigma, I_{\sigma,j}^s)$	N/A	N/A	0.961*	0.962	0.929	0.843	0.721*
Pearson correlation			(0.016)	(0.023)	(0.052)	(0.116)	(0.212)

Table 1: In each cell is the mean (standard error) of the summary described in the row label for the smoothing parameter σ seen in the column label, across all 10 noisy frames. An asterisk (*) means that mean/standard error was only taken over the (less than 10) frames where all 25 columns were recovered—respectively 8 and 5 out of 10 instances for $\sigma = 4$ and $\sigma = 12$. N/A indicate that in none of the 10 frames were the correct number of columns identified. The Hausdorff distance is measured here in pixels.

S5 in the Supplementary Material, after tuning parameters optimally. Even in the case of Image $I_{10,281}$ (see below for notation), where there are approximately circular blobs present, the method we present here yields results for atomic features that are very similar to the case of blob detection. More information on this comparison of methods can be seen in the Section S5 of the Supplementary Material.

5 Signal detection with hypothesis testing

As can be seen in the initial image of Figure 4, Figures 7a and 8a and Figures S8–S10 of the Supplementary Material, the images we aim to analyze are extremely noisy. This necessitates a goodness-of-fit test for a pure noise model. Here we use real-valued summaries of cubical persistence as test statistics for this hypothesis for the primary reason that it allows one to easily visualize dynamics as a univariate time series and thus aids interpretability for practitioners.

In many circumstances, such as the initial image of Figure 4 or Figure 7a, there are a natural number of atomic features present in the image that an algorithm should be able to retrieve. However, we are often required to choose a threshold η to return the correct cardinality $|A_\eta(I)|$, which corresponds to the number of atomic features (be they columns or planes). This might bias our result across the whole video due to the myriad different crystal structures that the detector may capture. Hence, we desire real-valued summaries which yield the same statistical conclusion as the number of thresholded columns but which are continuous rather than discrete (so p -values retrieved do not require adjusting for ties) and which do not require any parameter such as η to be tuned.

We concentrated our analysis on a $N = 1124$ frame video of a small area of a catalyst consisting of Pt nanoparticles supported on a larger nanoparticle of CeO_2 ⁸. In the following sections we discuss the relative merits of the various topological summaries and derive time series for our 1124 frame video based on two summaries that seemed to work the best: the persistent entropy and the ALPS statistic. We finish with detailing and applying our hypothesis testing framework to these output time series.

5.1 Setup and assumptions

Let us denote I_1, \dots, I_N as the original image sequence with the same image set \mathcal{I} . For our basic setup, we consider a series of $m \geq 1$ images summed

$$I_{m,\ell} := \sum_{k=0}^{m-1} I_{\ell+k}, \quad \ell = 1, \dots, N - m + 1.$$

Throughout, let us fix a subimage $\mathcal{L} \subset \mathcal{I}$ and let us suppose that our *unsmoothed* pixels take values on the nonnegative integers. We want to test whether or not the output of the algorithm above produces noise or a definitive signal. We assume that in each image I_k , $k = 1, \dots, N$ there is some subset $V_k \subset \mathcal{I}$ that represents the *vacuum*, and as such, is purely constituted of Poisson shot noise, as has been assumed in Levin et al. (2020)—heuristically verified using plotting heuristics for the vacuum region of Pt nanoparticles in

⁸More information on this data and how it was collected can be found in Section S1 of the Supplementary Material.

Mohan et al. (2022). There is still the question of how to choose R . We have done this by choosing a polygonal region R bounding the summed nanoparticle $I_{N,1}$, which is tenable as the nanoparticle experiences minimal movement from its initial location in the video.

In other imaging contexts we could estimate a null hypothesis of i.i.d noise by sampling from the empirical distribution of pixel values within each V_k . However, a Poisson assumption should hold, so we will more rigorously check the Poisson assumption here. If we suppose that $p \in V_{\ell+k}$ for all $k = 1, \dots, m$ so that $I_{\ell+k}(p)$ is Poisson $\mu_{\ell+k}$, then $I_{m,\ell}(p)$ is exactly Poisson with parameter $\sum_{k=0}^{m-1} \mu_{\ell+k}$ if $I_{\ell+k}(p)$ are independent for each k . Thus, we assume in our null model that each value $I_{m,\ell}(p)$ has a Poisson distribution for pixels in the vacuum region. It suffices to show that each value $I_{1,\ell}(p)$ has a Poisson distribution, which we will investigate shortly. Throughout this section, identify R with $R \cap \mathcal{L}$ —adding any 8-connected elementary 2-dimensional cubes lying outside of R , so that we may compute cPH. By convention if $d = \infty$, we set the death time d to be the largest pixel value in the rectangular subimage of I_k , such as in Chung and Day (2018).

To test whether or not the *observed* output $A(I_{m,\ell})$ coincides with the null hypothesis that the probability distribution $P_\ell \equiv P_\ell^m$ which generated the pixel values $\{I_{m,\ell}(p)\}_{p \in R}$ in the polygonal region R is equal to the noise distribution $P_{\ell,0} \equiv P_{\ell,0}^m$, it will help to have some idea of the distribution of $A(I_{m,\ell}^*)$ when the random image $I_{m,\ell}^*$ is generated from the noise distribution, i.e. $P_\ell = P_{\ell,0}$. (Note that P_ℓ and $P_{\ell,0}$ are considered as probability measures on the set \mathbb{N}_0^R of R -tuples of nonnegative integers). When the null hypothesis $H_0 : P_\ell = P_{\ell,0}$ holds, for each $p \in R \subset \mathcal{I}$, $I_{m,\ell}^*(p)$ are sampled i.i.d according to $F_{\ell,0}$.

As we may not know for certain the entire vacuum region (or the boundary of the vacuum region may change) we may select a subregion $U_\ell \subset V_{\ell+k}$, $k = 1, \dots, m$ that we are confident is entirely vacuum for every image $I_{\ell+k}$ —see Figure 6. For simplicity, we suppose that $U_\ell = U$ for all $\ell = 1, \dots, N$. In the following we identify U with $U \cap \mathbb{Z}^2$. We assume $F_{\ell,0}$ to be the Poisson distribution with parameter $\lambda_{m,\ell}$. However, there is no evidence that this mean changes (see Figure S1 of the Supplementary Material), so that $\lambda_{m,\ell} \equiv \lambda_m = m\lambda$ where the maximum likelihood estimate of λ is

$$\hat{\lambda} = \frac{1}{|U|N} \sum_{k=1}^N \sum_{p \in U} I_k(p),$$

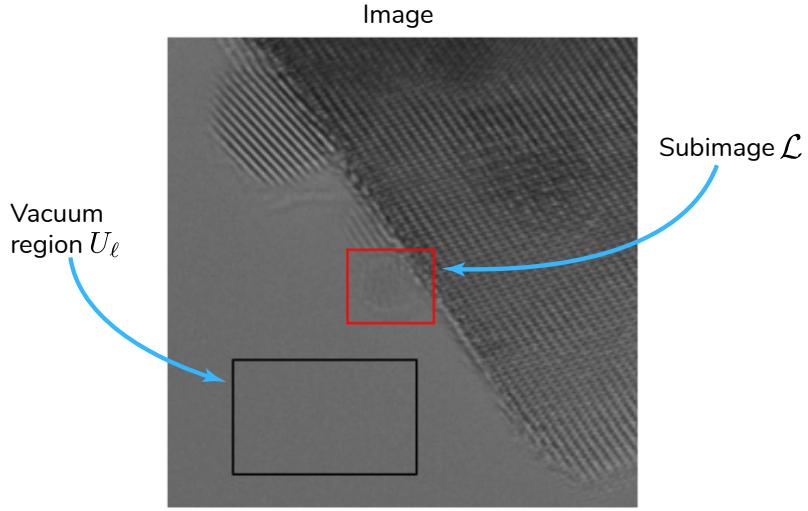


Figure 6: Depiction of selection of a subimage and vacuum region for an image I . Image depicted here is $I_{m,1}$ with $m = 400$. That is, it is the sum of the first 400 frames in a video of *in situ* TEM with 3 Pt nanoparticles. Performing such summing is useful to help identify an important subimage \mathcal{L} and region R which correspond to interesting atomic structural dynamics.

assuming independence across frames. In other words, we assume that $P_{\ell,0} \equiv P_0$ and $F_{\ell,0} \equiv F_0$. It is worth mentioning that we could have chosen to estimate λ_m from the mean intensity in R instead, though the intensity would typically be lower, which would indicate a lower variance given the (Poissonian) nature of the data. Therefore, the average lifetimes in the persistence diagrams would be lower and it would reduce the power of the hypothesis tests below. It is also more natural to estimate vacuum behavior from a known vacuum region, rather than treating the data as if *it were* a vacuum region.

Let us denote \hat{F}_0^m to be the Poisson distribution with parameter $\hat{\lambda}_m = m\hat{\lambda}$ and denote \hat{P}_0^m to be the product measure induced by \hat{F}_0^m on \mathbb{N}_0^R . Denote $F_{\ell,U}$ to be the empirical cdf of the vacuum pixels in frame ℓ and F_U^m to be the empirical cdf of

$$\{I_{m,\ell}(p) : \ell = mk + 1, k = 0, \dots, \lfloor N/m - 1 \rfloor, p \in U\},$$

where $F_U \equiv F_U^1$. Let $\hat{F}_0 := \hat{F}_0^1$ be Poisson with mean $\hat{\lambda}$. In practice, we have taken U to

be the black rectangular region seen in Figure 6, which is 400 pixels by 250 pixels. Here we use the Dvoretzky-Kiefer-Wolfowitz (DKW) inequality (see [DasGupta, 2011](#); [Massart, 1990](#)),

$$\mathbb{P}\left(\sup_{k \in \mathbb{N}_0} |F_U(k) - \hat{F}_0(k)| > \epsilon\right) \leq 2e^{-2|U|\epsilon^2},$$

to test the Poisson assumption in the vacuum region. Noting that $\hat{\lambda}$ is a sufficient statistic and that the Kolmogorov-Smirnov (KS) distance is $h_n = 0.00329$, we apply this inequality with $\epsilon = h_n$ and refute the assertion that the data in the vacuum frame are i.i.d. Poisson random variables.

Though the Poisson assumption may fail to hold precisely over this massive sample, there is little practical and theoretical evidence for doing away with it. Indeed, there are no substantive changes in the Monte Carlo p -values when sampling from the empirical distribution F_U^m (see Tables S1, S2, and Section S3 in the Supplementary Material). Furthermore, were we to have a perfect detector, the rate of arrival of electrons at each pixel in the vacuum would follow a Poisson distribution ([Levin, 2021](#)); thus, we maintain our initial assumption that the marginal distribution for pixels for the noise distribution in each frame is $\text{Poi}(\lambda)$.

Section S2 of the Supplementary Material dives deeper into whether the assumptions mentioned above were satisfied for the vacuum regions in our primary image series. An assumption of independence between frames shows a minor violation, yet only a maximum mean autocorrelation of 10^{-3} across the first 50 lags. Nonetheless, there is robust evidence supporting independence of pixels within frames and stationary and ergodicity across frames. In conclusion, it seems as if the parametric approach we have outlined above is tenable, given the theoretical properties of the materials and the physics, even in spite of the fairly minor violations of the assumptions in practice.

5.2 Empirical results: hypothesis testing and time series

Assume we believe that the number of (thresholded) columns output by our algorithm for the summed image $I_{m,\ell}$ is higher for an image with a strong signal in contrast to an image that is composed entirely of noise. In other words, we are interested in the quantity

$n_\ell := |A_\eta(I_{m,\ell})|$, and in particular the evidence of n_ℓ against H_0 . We follow [Davison and Hinkley \(1997\)](#) in constructing a p -value for our Monte Carlo hypothesis test. First, we generate pixel values in \mathcal{L} (hence R as well) according to \hat{P}_0^m to yield an image $\hat{I}_{m,\ell}^*$. We proceed by generating a number of i.i.d. instances of $\hat{I}_{m,\ell}^*$ denoted $\hat{I}_1, \dots, \hat{I}_n$. There is precedent to the idea of using generated simulated smoothed images (which can be considered as discrete random fields) for hypothesis testing—an example of which can be seen in [Taylor et al. \(2007\)](#). The initial p -value we aim to estimate is

$$\mathbb{P}_0(|A_\eta(I_{m,\ell}^*)| \geq n_\ell), \quad (1)$$

where \mathbb{P}_0 is a probability measure under which the null hypothesis holds. The p -value (1) is valid when we condition on $\hat{\lambda}$. In practice, we estimate the true p -value (1) by the rank of n_ℓ amongst the simulated value $|A(\hat{I}_k)|$. As n_ℓ is discrete (integer-valued), such a rank is not unique, so we utilize the Monte Carlo p -value

$$\frac{1}{n+1} \left(1 + \sum_{i=1}^n \mathbf{1}\{|A_\eta(\hat{I}_k)| \geq n_\ell\} \right).$$

This is reasonable if both $|U_\ell|$ and n are sufficiently large, by (1). For our particular setting, let $m = 10$, $\sigma = 2, 4, 6$, $\eta = t(\sigma)$ and set $n = 9999$, U to be the black region in Figure 6, and R to be the polygon in Figure 7—where $t(\sigma)$ is a function calibrated by the user to recover relevant atomic features in an image. Here we chose $t(2) = 1$, $t(4) = 0.4$, and $t(6) = 0.1$. For $\sigma = 4$ and $\ell = 241$, we estimate that

$$\mathbb{P}_0(|A_\eta(I_{m,\ell}^*)| \geq n_\ell) = 0.0001$$

which yields very strong evidence against H_0 for the pixels in R in Figure 7a, where $n_\ell = 6$. This method worked well because there are many “significant” features in the nanoparticle image; this would also work with any algorithm A which outputs a (marked) point process, such as blob detection ([Lindeberg, 1998](#); [Kong et al., 2013](#)). However, if there is only one or two highly persistent features, this test will be decidedly underpowered. Additionally, there is issue of choosing a threshold η —there are many other values of η that would lead to the anticipated rejection of H_0 ; also, manually inspecting thousands of images to find relevant features is infeasible in practice. In general, we may consider a real-valued functional f of

a marked point process on R , such that larger values of $f(A(I_{m,\ell}))$ for our summed image would lead us to reject H_0 . We can proceed with the exact same framework as the above, but what sort of function would yield useful information? In principle, we could use any method that can be used with a point process, see [Illian et al. \(2008\)](#). But such methods could be used with the output of methods such as blob detection (as in [Manzorro et al., 2022](#)), as well. As such, let us consider a real-valued function of the marks of $A(I)$ —i.e. the lifetimes of the points in the persistence diagram PD_0 —called *persistent entropy* ([Rucco et al., 2016; Atienza et al., 2020](#)). Because we consider output of the algorithm in the case of noise to be associated with more disorder, we actually take the negative of the persistence entropy, or

$$H(A(I)) := \sum_{(p,l) \in A(I)} l/L \log(l/L) = \sum_{(b,d) \in PD_0} ((d-b)/L) \log((d-b)/L),$$

where $L = \sum_{(p,l) \in A(I)} l = \sum_{(b,d) \in PD_0} d - b$. Higher values of $H(A(I))$, i.e. values closer to zero, signify smaller entropy. Using the negative of persistent entropy yields a Monte Carlo p -value of

$$\mathbb{P}_0(H(A(I_{m,\ell}^*)) \geq H_\ell) = 0.0317,$$

where $H_\ell = H(A(I_{m,\ell}))$ is again the observed value for image depicted in Figure 7a. One may also consider the longest barcode (or, the infinity norm), i.e.

$$L(A(I)) := \max\{l : (p, l) \in A(I)\} = \max\{d - b : (b, d) \in PD_0\}$$

as well as the sample mean persistence $E(A(I)) = \sum_{(p,l) \in A(I)} l/|A(I)|$. We also introduce the ALPS (accumulated lifetimes of persistence survival) statistic, $\Delta(A(I))$, defined by

$$\Delta(A(I)) := \int_0^\infty \log U(\eta) d\eta, \tag{2}$$

where $U(\eta) = \sum_{(p,l) \in A(I)} \mathbf{1}\{l > \eta\} = \sum_{(b,d) \in PD_0} \mathbf{1}\{d - b > \eta\}$. We can easily generalize this to any persistence diagram PD_k , just as one can with any of the aforementioned summaries. Let us order the lifetimes of PD_0 as $l_{(1)} \leq \dots \leq l_{(K)}$ where $K = |A(I)|$ is the number of points/barcodes in the persistence diagram. There is another convenient representation of (2) which is worth mentioning.

Test statistic	Number of columns $ A_{t(\sigma)}(I_{10,241}) $	Persistent entropy $H(A(I_{10,241}))$	Longest barcode $L(A(I_{10,241}))$	Mean persistence $E(A(I_{10,241}))$	ALPS statistic $\Delta(A(I_{10,241}))$
$\sigma = 2$	0.0001	0.0001	0.0595	0.2303	0.0008
$\sigma = 4$	0.0001	0.0317	0.0063	0.0013	0.0001
$\sigma = 6$	0.1140	0.7966	0.5573	0.2660	0.1032

Table 2: Monte Carlo p -values p_n for various real-valued topological summaries of $I_{10,241}, \hat{I}_1, \dots, \hat{I}_{9999}$ smoothed with varying values of σ . Images were generated according to $\hat{P}_{\ell,0}$. With at least 95% confidence the true p -value lies in $p_n \pm 0.0137$, truncating at 0 or 1—see Supplementary Material for more information.

Proposition 5.1.

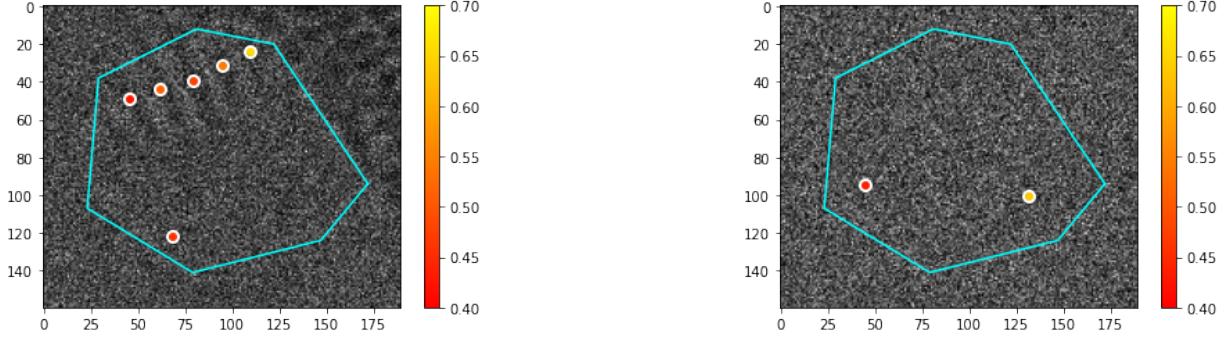
$$\Delta(A(I)) = - \sum_{i=1}^{K-1} l_{(i)} \log \left(1 - \frac{1}{K-i+1} \right),$$

We offer a short proof in Section S7 of the Supplementary Material. Based on Proposition 5.1 there is no need to consider how to treat the infinite barcode, if it is present in R . The ALPS statistic is similar in spirit to the accumulated persistence function (Biscio and Møller, 2019), and aims to balance the information content of the longest barcode with that of the lifetime sum, or 1-norm, $\sum_{i=1}^K l_{(i)}$ of the persistence diagram.

Other summaries contain useful information but fluctuate wildly⁹, such as the sample skewness of the persistence lifetimes, or never yield a significant signal—e.g. the p -norms of persistence (Cohen-Steiner et al., 2010) and the signal-to-noise ratio (of mean lifetime divided by standard deviation of lifetimes). The persistent entropy was found to be most stable to whether or not we were in parametric or nonparametric setting. The ALPS statistic p -values decreased across the board in the nonparametric setting, which is what one would expect as in both images $I_{10,221}$ and $I_{10,241}$ there appears to be some signal. The summary of the Monte Carlo p -values for each of these five test statistics can be seen in Table 2 and 3.

In the image in Figure 7a, as there is an extremely robust nanoparticle structure present,

⁹Tables S3–S6, supporting this conclusion, can be seen in the Supplementary Material.



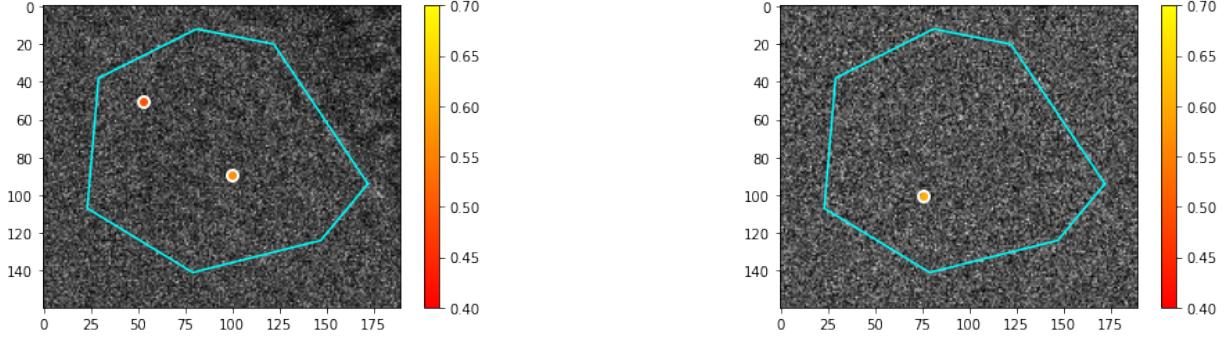
(a) Output $A_{0.4}(I_{10,241})$ (colored points) superimposed over the observed subimage.

(b) Output $A_{0.4}(\hat{I}_k)$ (colored points) superimposed over a pure noise subimage.

Figure 7: The output $A_{0.4}(I)$, $\sigma = 4$, for an observed subimage (left) and a pure noise/simulated subimage (right). Here $I_{10,241}$ can be seen to have a strong signal corresponding to various atomic features.

we would expect to reject the null hypothesis of an the image consisting of i.i.d. Poisson random variables. If we examine a different summed image, as in Figure 8a, we can see upon thresholding the image with $\eta = 0.4$ that there are only two columns. Hence the larger estimated p -value in this case. However, both persistence entropy H and longest barcode L indicate that it is reasonable to reject the null hypothesis that $P_\ell = P_0$ —see Table 3.

Based on the output of the tables above, it seems that the ALPS statistic and persistent entropy are the best metrics of the five that we've considered. The ALPS statistic has the smallest variance across the p -values; it additionally seems to correlate fairly strongly with the p -value associated to the thresholded number of columns. Furthermore, it does not require any tuning of the threshold parameter η . Both the ALPS statistic and the persistent entropy seems to yield the best “separation” between an observed image with clear signal and noisy simulated images—see Figure 9. Persistent entropy also corresponds with known changes in atomic configurations (see Figure 10) and enjoys various stability properties (Atienza et al., 2020). Hence, for smoothing parameters σ in any compact interval bounded away from zero, it can be shown that $\sigma \mapsto I_\sigma$ is a smooth function as well, thus persistent



(a) Output $A_{0.4}(I_{10,221})$ (colored points) superimposed over the observed subimage.

(b) Output $A_{0.4}(\hat{I}_k)$ (colored points) superimposed over a pure noise subimage.

Figure 8: The output $A_{0.4}(I)$, $\sigma = 4$, for an observed subimage (left) and a pure noise/simulated subimage (right). Here $I_{10,221}$ can be seen to have a weak signal corresponding to various atomic features.

entropy does not fluctuate too rapidly for small changes in σ (one can use the stability theorem for cubical persistence from [Skraba and Turner, 2020](#)). Experimentally, it appears this holds with the ALPS statistic too—see Figure 9.

In Figures 10 and 11, we derive time series by calculating the topological summaries persistent entropy and ALPS statistic, respectively, for the two nanoparticle regions seen in the lower right portion of Figure 10. The summaries were calculated framewise. To choose σ we have used the elbow method described in Section S6 of the Supplementary Material. We applied this method and $\sigma = 2$ was chosen as it appeared at the elbow of the plots in Figure S6 and therefore led to a high degree of correspondence with “stable” states, meaning presence of visible nanoparticle structure¹⁰. Such structure began to disappear as σ increased, and the association of the statistic between low values (persistent entropy) or high values (ALPS) and high order reversed when σ was too large. Part of the reason for this is that noisy images with a sufficiently large degree of smoothing will have fewer total output points in $A(I)$, which typically leads to less variation and entropy—see Figure 9.

¹⁰Figure 9 illustrates this for $I_{10,241}$, suggesting that $\sigma \approx 3$ is appropriate for said image.

Test statistic	Number of columns $ A_{t(\sigma)}(I_{10,221}) $	Persistent entropy $H(A(I_{10,221}))$	Longest barcode $L(A(I_{10,221}))$	Mean persistence $E(A(I_{10,221}))$	ALPS statistic $\Delta(A(I_{10,221}))$
$\sigma = 2$	0.0916	0.0604	0.1485	0.3523	0.0851
$\sigma = 4$	0.2015	0.0213	0.0378	0.0051	0.0531
$\sigma = 6$	0.2888	0.2040	0.1620	0.0328	0.1762

Table 3: Monte Carlo p -values p_n for various real-valued topological summaries of $I_{10,221}, \hat{I}_1, \dots, \hat{I}_{9999}$ smoothed with varying values of σ . Images were generated according to $\hat{P}_{\ell,0}$. With at least 95% confidence the true p -value lies in $p_n \pm 0.0137$ (truncating at 0 or 1)—see Supplementary Material for more information.

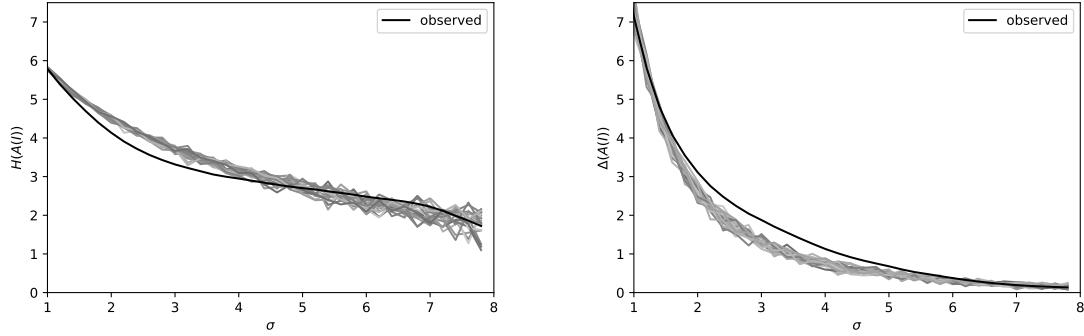


Figure 9: Persistent entropy (left) and ALPS (right) for various smoothing parameters σ for image $I_{10,241}$, compared with same series for 19 simulated images.

5.3 Multiple testing using persistent entropy and the ALPS statistic

Using the extension of the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995; Benjamini and Yekutieli, 2001), we test each of the frames in the summed series $I_{5,5k+1}$, $k = 0, 1, \dots, 223$ versus the null hypothesis that they were generated according to $P_{\ell,0}$, with $m = 5$. We use $m = 5$ here instead of the individual frames, because very few frames (only for the top region, using persistent entropy) are significant. We proceed by generating $n = 9999$ images according to the same recipe outlined at the start of

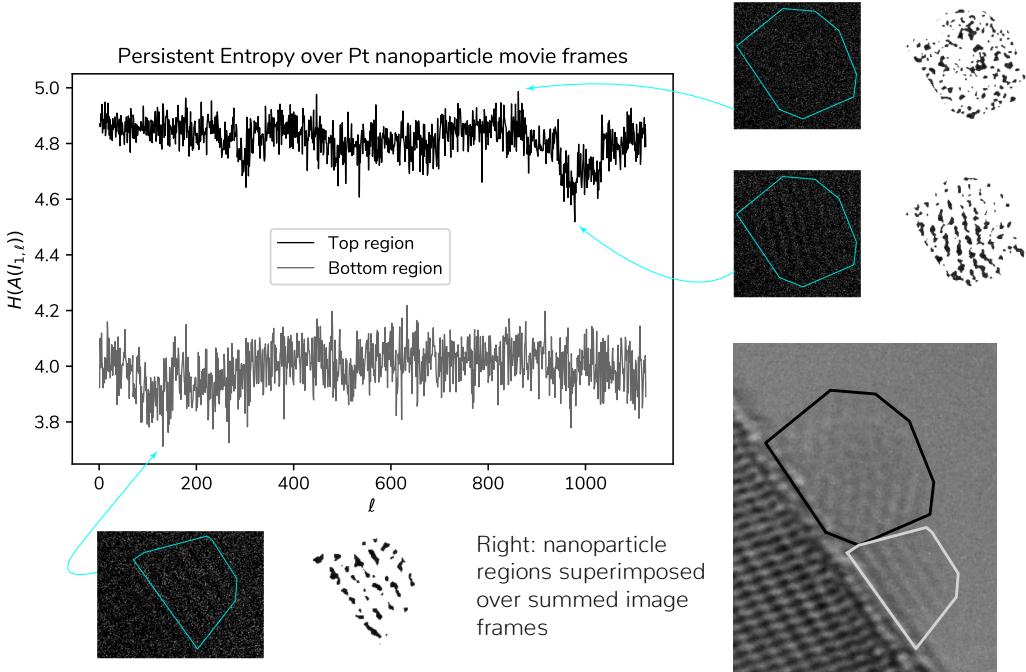


Figure 10: Time series of persistent entropy over all 1124 frames for the nanoparticle regions depicted in lower right. Selected nanoparticle frames corresponding to high/low values of persistent entropy are depicted along with binarized images derived from PD thresholding [Chung and Day \(2018\)](#)—see Supplementary Material for how this was implemented.

Section 5.2. A detailed description of our testing framework can be seen in Section S8 of the Supplementary Material. Again we make the assumption that $P_{\ell,0} = P_0$, or that the same noise hypothesis is shared by each frame. This is not unreasonable based on what we have found so far. Our approach is at least somewhat similar to the framework described for multiple testing using persistent homology described in [Vejdemo-Johansson and Mukherjee \(2022\)](#). However, in said article, they are concerned with point cloud-based rather than cubical homology and potentially different null distributions for each test. Here our hypotheses are $H_0 : P_{5k+1} = P_0$, for each $k = 0, 1, \dots, 223$.

As we generate so few images owing to the stationarity assumption of our time series of topological summaries, the computational costs are significantly less than that of other hypothesis testing settings using persistent homology ([Vejdemo-Johansson and Mukherjee,](#)

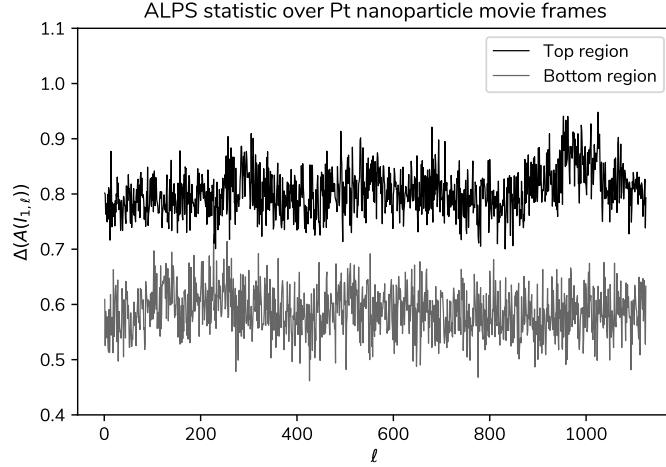


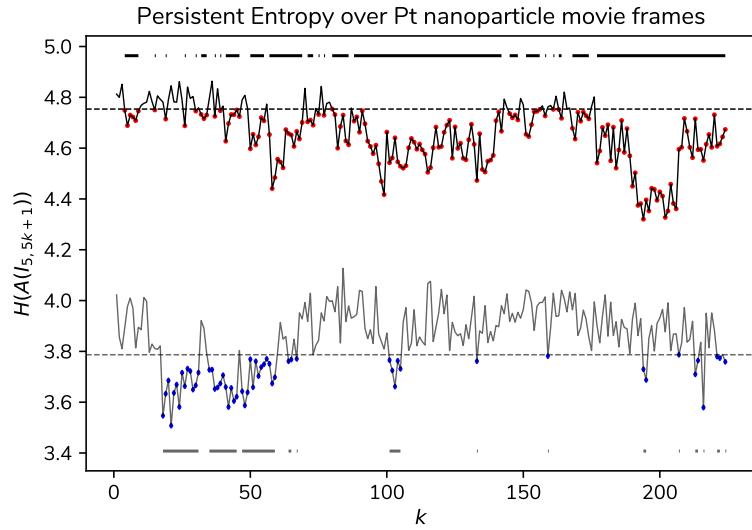
Figure 11: Time series of ALPS statistic over all 1124 frames for nanoparticle regions depicted in lower right of Figure 10.

2022; Robinson and Turner, 2017). We also guarantee that our multiple testing framework has asymptotic false discovery rate less than α , almost surely—see Proposition S8.1 in the Supplementary Material.

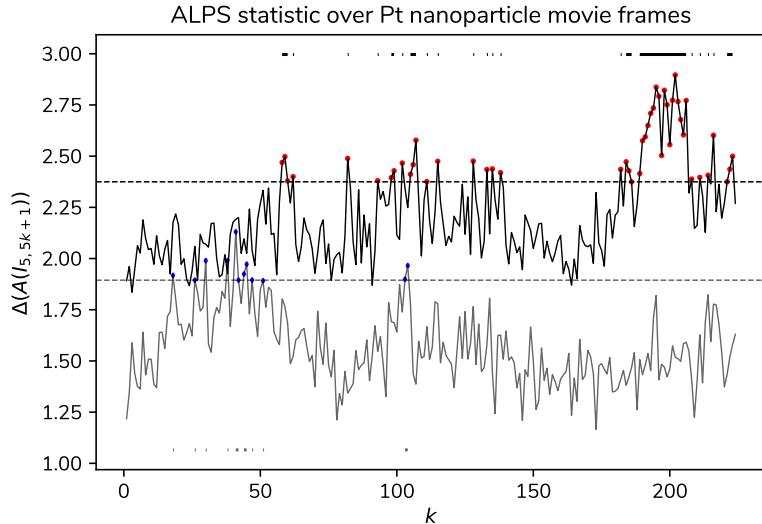
Figures 12a and 12b indicate that the persistent entropy has much greater power than the ALPS statistic, though depending on the application it may be a bit overzealous. For researchers that are more conservative in their desire to detect atomic features, the ALPS statistic may be preferable. It is also possible to normalize the time series and consider the convex combination $\lambda H(A(I)) + (1 - \lambda)\Delta(A(I))$ for $\lambda \in (0, 1)$. One may verify that for the values of k surrounding 200 for the top region and values of k around 25 for the bottom region, there is a significant nanoparticle structure present—see Supplementary Material Section S9, and Figures S8b and S8c.

6 Discussion

In this paper we have discussed a novel means of detecting atomic features in nanoparticle images, which compares favorably to existing methods and is nonparametric in its estimation of intensity. We have also investigated means of hypothesis testing the pres-



(a) Persistent entropy.



(b) ALPS statistic

Figure 12: Persistance summaries $I_{5,5k+1}$, $k = 0, 1, \dots, 223$ image sequence for top and bottom regions with non-noisy frames at the $\alpha = 0.05$ level marked by points. Runs of significant frames for each region depicted above and below the plotted curves.

ence of signals in images with ultra-low signal-to-noise ratio and detailed a useful method for deriving time series from noisy videos. We have also introduced the ALPS statistic, which conveys much of the same information content in a hypothesis testing context as the number of columns present after appropriate thresholding.

As topological data analysis is such a young field, there is no shortage of directions along which the methods in this article could be expanded. For example, one could employ a functional version of $|A_\eta(I)|$; alternatively, some another functional summary of persistence diagrams may hold promise, especially used in conjunction with a global rank envelope test as in [Biscio and Møller \(2019\)](#) and [Myllymäki et al. \(2017\)](#). Furthermore, using persistent homology for point clouds such as the Vietoris-Rips filtration ([Boissonnat et al., 2018](#)), would utilize the location information derived from the above algorithm in an essential way. Using a weighted Vietoris-Rips filtration may even furnish more precise results ([Anai et al., 2020](#)), as the marked point process output of the algorithm naturally yields weights for each point. After deriving an additional persistence diagram from the output $A(I)$, one may choose an appropriate functional summary (see [Berry et al., 2020](#)) and proceed from there.

Furthermore, other choices of filtrations that are more suited to capturing geometric information, may be considered in future studies. Examples include treating the locations of black pixels as embedded into Euclidean space, and then applying a Vietoris-Rips or density-based filtration ([Garin and Tausin, 2019](#); [Turkes et al., 2021](#)). As mentioned earlier, these methods require the choice of a threshold and hence multi-parameter persistent homology may be a more appropriate tool—see [Chung et al. \(2022\)](#). In sum, our hope is that this study can find adoption in the microscopy community and also be used as point of departure for future studies at the intersection of image processing, statistics, and TDA.

Supplementary Material

The Supplementary Material contains additional derivations and simulations pertaining to the main text. Every section of the Supplementary Material is referenced in the text above. Python code and the videos can be found in `DataAndCodeForPaper.zip`.

References

Anai, H., Chazal, F., Glisse, M., Ike, Y., Inakoshi, H., Tinarrage, R., and Umeda, Y. (2020). DTM-based filtrations. In *Topological Data Analysis*, pages 33–66. Springer.

Atienza, N., González-Díaz, R., and Soriano-Trigueros, M. (2020). On the stability of persistent entropy and new summary functions for topological data analysis. *Pattern Recognition*, 107:107509.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165 – 1188.

Berry, E., Chen, Y.-C., Cisewski-Kehe, J., and Fasy, B. T. (2020). Functional summaries of persistence diagrams. *Journal of Applied and Computational Topology*, 4(2):211–262.

Biscio, C. A. and Møller, J. (2019). The accumulated persistence function, a new useful functional summary statistic for topological data analysis, with a view to brain artery trees and spatial point process applications. *Journal of Computational and Graphical Statistics*, 28(3):671–681.

Blumberg, A. J., Gal, I., Mandell, M. A., and Pancia, M. (2014). Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Foundations of Computational Mathematics*, 14(4):745–789.

Boissonnat, J.-D., Chazal, F., and Yvinec, M. (2018). *Geometric and topological inference*, volume 57. Cambridge University Press.

Carlsson, G. and Vejdemo-Johansson, M. (2022). *Topological Data Analysis with Applications*. Cambridge University Press.

Cericola, C., Johnson, I. J., Kiers, J., Krock, M., Purdy, J., and Torrence, J. (2017). Extending hypothesis testing with persistent homology to three or more groups. *Involve, a Journal of Mathematics*, 11(1):27–51.

Chazal, F. and Michel, B. (2021). An introduction to topological data analysis: fundamental and practical aspects for data scientists. *Frontiers in artificial intelligence*, 4:667963.

Chen, X., Chen, D., Weng, M., Jiang, Y., Wei, G. W., and Pan, F. (2020). Topology-Based Machine Learning Strategy for Cluster Structure Prediction. *Journal of Physical Chemistry Letters*, 11(11):4392–4401.

Chintakunta, H., Gentimis, T., Gonzalez-Diaz, R., Jimenez, M.-J., and Krim, H. (2015). An entropy-based persistence barcode. *Pattern Recognition*, 48(2):391–401.

Chung, Y.-M. and Day, S. (2018). Topological fidelity and image thresholding: A persistent homology approach. *Journal of Mathematical Imaging and Vision*, 60(7):1167–1179.

Chung, Y.-M., Day, S., and Hu, C.-S. (2022). A multi-parameter persistence framework for mathematical morphology. *Scientific Reports*, 12(1):6427.

Cohen-Steiner, D., Edelsbrunner, H., and Harer, J. (2007). Stability of persistence diagrams. *Discrete and Computational Geometry*, 37(1):103–120.

Cohen-Steiner, D., Edelsbrunner, H., Harer, J., and Mileyko, Y. (2010). Lipschitz functions have l_p -stable persistence. *Foundations of computational mathematics*, 10(2):127–139.

DasGupta, A. (2011). *Probability for statistics and machine learning: fundamentals and advanced topics*. Springer Science & Business Media.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

Dłotko, P. (2015). Cubical complex. In *GUDHI User and Reference Manual*. GUDHI Editorial Board.

Edelsbrunner, Letscher, and Zomorodian (2002). Topological Persistence and Simplification. *Discrete & Computational Geometry*, 28(4):511–533.

Edelsbrunner, H. and Harer, J. (2010). *Computational topology: an introduction*. American Mathematical Society, Providence, Rhode Island.

Egerton, R. (2013). Control of radiation damage in the TEM. *Ultramicroscopy*, 127:100–108.

Egerton, R. (2019). Radiation damage to organic and inorganic specimens in the TEM. *Micron*, 119:72–87.

Egerton, R., Li, P., and Malac, M. (2004). Radiation damage in the TEM and SEM. *Micron*, 35(6):399–409.

Faruqi, A. and McMullan, G. (2018). Direct imaging detectors for electron microscopy. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 878:180–190.

Fasy, B. T., Lecci, F., Rinaldo, A., Wasserman, L., Balakrishnan, S., and Singh, A. (2014). Confidence sets for persistence diagrams. *The Annals of Statistics*, 42(6):2301–2339.

Garin, A., Heiss, T., Maggs, K., Bleile, B., and Robins, V. (2020). Duality in Persistent Homology of Images. *arXiv preprint arXiv:2005.04597*.

Garin, A. and Tauzin, G. (2019). A topological 'reading' lesson: Classification of MNIST using TDA. *Proceedings - 18th IEEE International Conference on Machine Learning and Applications, ICMLA 2019*, pages 1551–1556.

Gillies, S. (2013). The shapely user manual. URL <https://pypi.org/project/Shapely>.

Illian, J., Penttinen, A., Stoyan, H., and Stoyan, D. (2008). *Statistical analysis and modelling of spatial point patterns*, volume 70. John Wiley & Sons.

Jiang, Y., Chen, D., Chen, X., Li, T., Wei, G. W., and Pan, F. (2021). Topological representations of crystalline compounds for the machine-learning prediction of materials properties. *npj Computational Materials*, 7(1):1–8.

Kaczynski, T., Mischaikow, K., and Mrozek, M. (2006). *Computational homology*, volume 157. Springer Science & Business Media.

Kong, H., Akakin, H. C., and Sarma, S. E. (2013). A generalized laplacian of gaussian filter for blob detection and its applications. *IEEE Transactions on Cybernetics*.

Kovalevsky, V. (1989). Finite topology as applied to image analysis. *Computer Vision, Graphics, and Image Processing*, 46(2):141–161.

Lawrence, E. L., Levin, B. D., Boland, T., Chang, S. L., and Crozier, P. A. (2021). Atomic scale characterization of fluxional cation behavior on nanoparticle surfaces: probing oxygen vacancy creation/annihilation at surface sites. *ACS nano*, 15(2):2624–2634.

Lawrence, E. L., Levin, B. D., Miller, B. K., and Crozier, P. A. (2019). Approaches to Exploring Spatio-Temporal Surface Dynamics in Nanoparticles with in Situ Transmission Electron Microscopy. *Microscopy and Microanalysis*, 2(2020):86–94.

Lawson, A., Hoffman, T., Chung, Y.-M., Keegan, K., and Day, S. (2021). A density-based approach to feature detection in persistence diagrams for firn data. *Foundations of Data Science*.

Levin, B. D. (2021). Direct detectors and their applications in electron microscopy for materials science. *Journal of Physics: Materials*, 4(4):042005.

Levin, B. D., Lawrence, E. L., and Crozier, P. A. (2020). Tracking the picoscale spatial motion of atomic columns during dynamic structural change. *Ultramicroscopy*, 213.

Lindeberg, T. (1998). Feature detection with automatic scale selection. *International journal of computer vision*, 30(2):79–116.

Manzorro, R., Xu, Y., Vincent, J. L., Rivera, R., Matteson, D. S., and Crozier, P. A. (2022). Exploring blob detection to determine atomic column positions and intensities in time-resolved TEM images with ultra-low signal-to-noise. *Microscopy and Microanalysis*, pages 1–14.

Massart, P. (1990). The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pages 1269–1283.

Mischaikow, K. and Nanda, V. (2013). Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, 50(2):330–353.

Mohan, S., Manzorro, R., Vincent, J. L., Tang, B., Sheth, D. Y., Simoncelli, E. P., Matteson, D. S., Crozier, P. A., and Fernandez-Granda, C. (2022). Deep denoising for scientific discovery: A case study in electron microscopy. *IEEE Transactions on Computational Imaging*, 8:585–597.

Motta, F. C., Neville, R., Shipman, P. D., Pearson, D. A., and Bradley, R. M. (2018). Measures of order for nearly hexagonal lattices. *Physica D: Nonlinear Phenomena*, 380–381:17–30.

Mukherjee, D., Miao, L., Stone, G., and Alem, N. (2020). mpfit: a robust method for fitting atomic resolution images with multiple Gaussian peaks. *Advanced Structural and Chemical Imaging*, 6(1):1–12.

Myllymäki, M., Mrkvička, T., Grabarnik, P., Seijo, H., and Hahn, U. (2017). Global envelope tests for spatial processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(2):381–404.

Nakamura, T., Hiraoka, Y., Hirata, A., Escolar, E. G., and Nishiura, Y. (2015). Persistent homology and many-body atomic structure for medium-range order in the glass. *Nanotechnology*, 26(30):304001.

Nord, M., Vullum, P. E., MacLaren, I., Tybell, T., and Holmestad, R. (2017). Atomap: a new software tool for the automated analysis of atomic resolution images using two-dimensional gaussian fitting. *Advanced structural and chemical imaging*, 3(1):1–12.

Rieck, B., Yates, T., Bock, C., Borgwardt, K., Wolf, G., Turk-Browne, N., and Krishnaswamy, S. (2020). Uncovering the topology of time-varying fMRI data using cubical persistence. *Advances in Neural Information Processing Systems (NeurIPS)*, 33(NeurIPS):6900–6912.

Robinson, A. and Turner, K. (2017). Hypothesis testing for topological data analysis. *Journal of Applied and Computational Topology*, 1(2):241–261.

Rucco, M., Castiglione, F., Merelli, E., and Pettini, M. (2016). Characterisation of the idiotypic immune network through persistent entropy. In *Proceedings of ECCS 2014*, pages 117–128. Springer.

Ruskin, R. S., Yu, Z., and Grigorieff, N. (2013). Quantitative characterization of electron detectors for transmission electron microscopy. *Journal of structural biology*, 184(3):385–393.

Skraba, P. and Turner, K. (2020). Wasserstein stability for persistence diagrams.

Taylor, J. E., Worsley, K., and Gosselin, F. (2007). Maxima of discretely sampled random fields, with an application to 'bubbles'. *Biometrika*, 94(1):1–18.

Turkes, R., Nys, J., Verdonck, T., and Latre, S. (2021). Noise robustness of persistent homology on greyscale images, across filtrations and signatures. *PLoS ONE*, 16(9 September):1–26.

Vejdemo-Johansson, M. and Mukherjee, S. (2022). Multiple hypothesis testing with persistent homology. *Foundations of Data Science*, 4(4):667–705.

Vincent, J. L. and Crozier, P. A. (2021). Atomic level fluxional behavior and activity of ceo2-supported pt catalysts for co oxidation. *Nature communications*, 12(1):1–13.

Wagner, H., Chen, C., and Vuçini, E. (2012). Efficient computation of persistent homology for cubical data. In *Topological methods in data analysis and visualization II*, pages 91–106. Springer.

Supplementary Material for “Feature Detection and Hypothesis Testing for Extremely Noisy Nanoparticle Images using Topological Data Analysis”

Andrew M. Thomas

Center for Applied Mathematics, Cornell University,

Peter A. Crozier

School for Engineering of Matter, Transport and Energy, ASU,

Yuchen Xu

Department of Statistics and Data Science, Cornell University,

David S. Matteson

Department of Statistics and Data Science, Cornell University

S1 Details on the collection of the primary image series

By the primary image series, we mean the sequence of 1124 images in Section 5 of the paper. The data is a subset of a large dataset recorded as part of an *in situ* electron microscopy investigation of the behavior of Pt nanoparticles during catalysis, specifically CO oxidation (Vincent and Crozier 2021). The *in situ* TEM experiments explored the behavior of 1–5 nm Pt nanoparticles supported on CeO₂ cubes in a variety of different atmospheres. The experiments were performed on an FEI Titan environmental transmission electron microscope operating at 297 keV and were mostly carried out with an electron fluence

of $600 \text{ e}^-/\text{\AA}^2/\text{s}$. High resolution images were recorded with a Gatan K3 direct electron detector operating in CDS counting mode. Movies were recorded at up to 75 frames per second. In many cases the Pt particles were observed to randomly undergo dynamic structural reconfiguration. For this particular image sequence employed here, the catalyst was exposed to a CO partial pressure of 3×10^{-2} Torr at room temperature. The video is entitled “PtCe02_6.tif”.

S2 Assumptions of the primary image series

It is reasonable to believe that there may exist auto- and/or cross-correlation of pixels between frames, owing to the high temporal resolution. Indeed, the mean autocorrelations over all pixels in the vacuum region $\bar{\rho}(h) = \frac{1}{|U|} \sum_{p \in U} \hat{\rho}_p(h)$ are negative for the first 1093 lags, i.e. $h = 1, \dots, 1093$. If we consider $h \ll N$, under the assumption of independence of each pixel along both spatial and temporal axes, $\hat{\rho}_p(h) \stackrel{\text{i.i.d.}}{\sim} N(0, 1/1124)$ for all $p \in U$. Therefore,

$$\bar{\rho}(h) \simeq N\left(0, \frac{1}{1.124 \times 10^8}\right).$$

Under our hypothesized spatio-temporal independence the probability that *at least one* of the first 50 mean correlations $\bar{\rho}(h)$, $h = 1, \dots, 50$ is less than its observed value has an upper bound of 2.55×10^{-10} , by using the standard probability union bound. Therefore, we can conclude there is strong evidence against temporal independence of the frames in the nanoparticle videos. These negative correlations are small however, with minimum value -0.001003 and maximum value -0.000007 , so summing a small number of frames (such as 10) does not deal a forceful blow to the Poisson assumption of the summed pixels.

The vacuum means for each frame in the primary image series of 1124 images are consonant with a stationary white noise series—see Figure S1. Further reinforcing this notion is that the means of these frames are essentially normally distributed (after standardizing), as displayed in the q-q plot in Figure S2. This follows if independence holds *across* frames.

Setting $m = 1$, the p -values

$$\min \left\{ 1, 2 \exp \left(-2n \sup_{k \in \mathbb{N}_0} |F_{\ell,U}(k) - \hat{F}_0(k)|^2 \right) \right\},$$

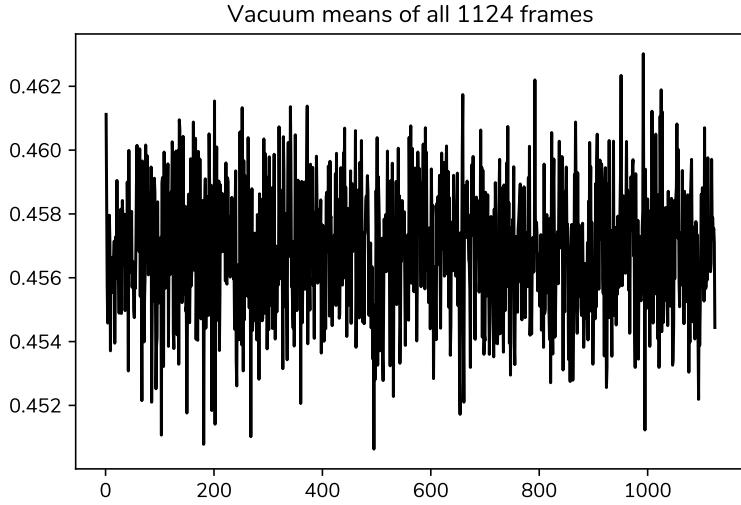


Figure S1: Vacuum means for all 1124 frames.

can be shown to be *valid* for each $\ell = 1, \dots, N$ and thus can be used to construct a level α test. Because of the near independence of the frames, we can use the Bonferroni method; there are 2 frames out of 1124 which reject the null hypothesis that the vacuum pixels in frame ℓ , $\{I_{1,\ell}(p)\}_{p \in U}$, $\ell = 1, \dots, N$ are i.i.d. $\text{Poi}(\lambda)$ when the significance level $\alpha = 0.05$. For practical purposes each frame seems to be identically distributed and very nearly independent, with some distribution that is very close to a Poisson; the KS distances $\sup_{k \in \mathbb{N}_0} |F_{\ell,U}(k) - \hat{F}_{\ell,0}(k)|$ also have mean nearly h_n , and thus satisfy

$$\frac{1}{N} \sum_{\ell=1}^N \sup_{k \in \mathbb{N}_0} |F_{\ell,U}(k) - \hat{F}_0(k)| \approx \sup_{k \in \mathbb{N}_0} |F_U(k) - \hat{F}_0(k)|.$$

This is accordant with $\sup_{k \in \mathbb{N}_0} |F_{\ell,U}(k) - \hat{F}_0(k)|$ being a stationary and ergodic sequence—suggesting the same for the pixels in the vacuum region of each frame.

An analysis of the semivariograms of the vacuum regions (Figure S3) of the two images $I_{10,221}$ and $I_{10,241}$ indicates no violation of independence of the pixels *within* frames. Indeed, the empirical semivariogram [Cressie \[1993\]](#),

$$\hat{\gamma}(l) = \frac{1}{2|N(l)|} \sum_{(p_i, p_j) \in N_l} (I_{m,\ell}(p_i) - I_{m,\ell}(p_j))^2,$$

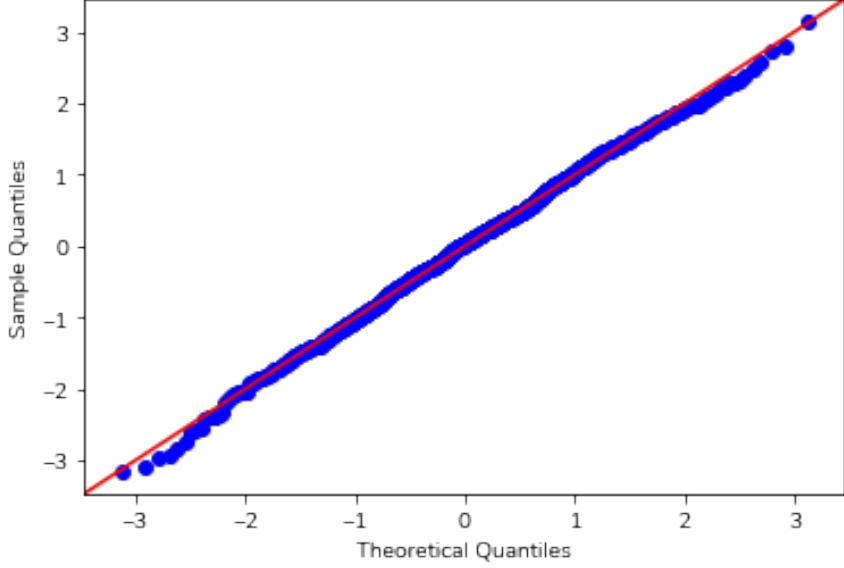


Figure S2: Q-Q Plot of vacuum means for all 1124 frames.

with $l = 0, \dots, 14$ and

$$N(l) := \{(p_i, p_j) \in U \times U : \|p_i - p_j\| \in [4l/3, 4l/3 + 4/3), i \neq j\},$$

indicates a robustly satisfied i.i.d. assumption for the vacuum region in both summed images—see Figure S3 below. At the very least, there is no evidence of correlation. The fact that the plots are nearly indistinguishable lends credibility to the assumption of stationarity across frames as well. Even looking at

$$\hat{\gamma}(1) = \frac{1}{2|N(1)|} \sum_{(p_i, p_j) : \|p_i - p_j\| = 1} (I_{m,\ell}(p_i) - I_{m,\ell}(p_j))^2,$$

where $N(1)$ is the set we sum over, we see that in the case of $I_{10,221}$ the difference between $\hat{\gamma}(1)$ and the variance is $|4.4983 - 4.4743| = 0.024$ and $|4.45 - 4.441| = 0.009$ in the case of $I_{10,241}$.

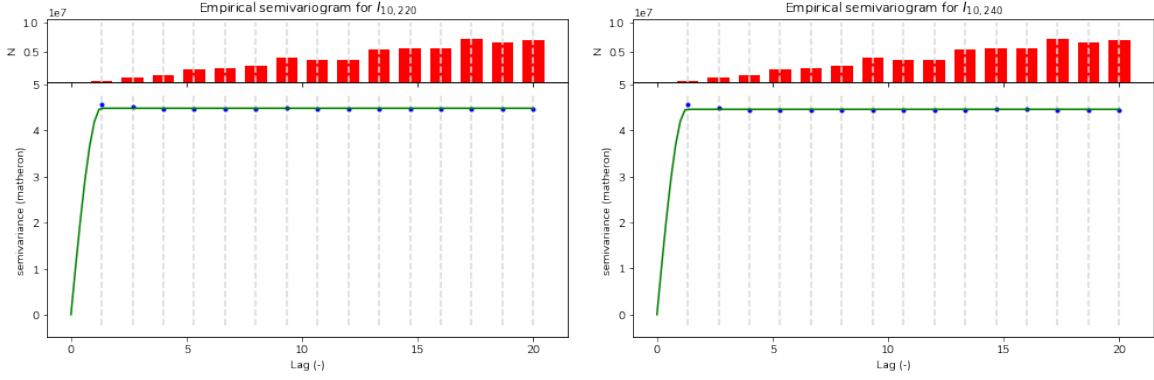


Figure S3: Plots of empirical semivariograms for the vacuum regions of $I_{10,221}$ and $I_{10,241}$, along with plotted estimated spherical model—see [Cressie \[1993\]](#). The histograms are equal as the binning is the same and the collection of pairwise distances for the two 250×400 vacuum regions U are identical.

S3 Further results of the experiments and nonparametric setup

For the p -values seen in the tables we used an earlier image series, consisting of the first 400 frames of the primary image series, and modulo some minor transformations—e.g. rotation, translation, etc. This video was entitled “010402_stk_400fr.tif”. We used F_U^{10} for the empirical cdf of pixel values. We see that there is little difference in Tables S1 and S2 from their parametric counterparts in the main document, especially with regards to the persistent entropy and ALPS statistic.

For completeness, we conduct our Monte Carlo hypothesis test with a few different real-valued summaries of persistence diagrams, i.e. functionals $f(A(I_{m,\ell}))$, from the ones seen in the main document. The first is signal-to-noise ratio (SNR), a summary that appears often in materials science, and defined as

$$\text{SNR}(A(I)) := E(A(I))/\sqrt{M_2(A(I))},$$

where more generally we have

$$M_k(A(I)) = \frac{1}{|A(I)|} \sum_{(p,l) \in A(I)} (l - E(A(I)))^k.$$

With this in mind, we can define the sample skewness as

$$S(A(I)) := \frac{M_3(A(I))}{M_2(A(I))^{3/2}}.$$

It seems that perhaps the negative of SNR and the skewness both perform well for small levels of smoothing, but the association breaks down as σ goes from 2 to 4—see Tables S3–S6.

We also considered two more standard summaries, the lifetime sum (degree-1) and degree-2 total persistence [Cohen-Steiner et al., 2010]. These are generally defined (for $k \geq 1$) as

$$L_k(A(I)) := \sum_{(p,l) \in A(I)} l^k.$$

Both the lifetime sum and the longest barcode were used in Rieck et al. [2020], with the longest barcode performing better in their prediction task.

The functionals L_k have desirable properties, such as stability in the sense that small perturbations of the underlying filtration yield small perturbations of L_k —see Cohen-Steiner et al. [2010]. However, the k -norms $L_k(A(I))^{1/k}$ decay at the rate $O(1/\sigma)$ for $k \gtrsim 5$ [Chen and Edelsbrunner, 2011] which may explain the lack of utility (or consistency) of L^1 and L^2 across a wide range of σ . Furthermore, in the case of a cubical complex, the lifetime sum obeys a central limit theorem (cf. Theorem 2.13 in Hiraoka and Tsunoda 2018). Hence, future studies could investigate multiple testing using the lifetime sum in an ANOVA type setting.

With that being said, the associations are quite sensitive to the smoothing parameter σ , just like the SNR and sample skewness. Hence, they were not considered in the paper. The Monte Carlo p -values using these summaries can be seen in Tables S3–S6.

S4 Justification of the Gaussian kernel

With respect to a suitably large class of kernels and signals, the Gaussian kernel is the only kernel that increases (first-order) local minima in continuous 1-dimensional signals as σ increases—see Babaud et al. [1986] for more details. In other words, the value of a signal

Test statistic	$ A_{t(\sigma)}(I_{10,241}) $	$H(A(I_{10,241}))$	$L(A(I_{10,241}))$	$E(A(I_{10,241}))$	$\Delta(A(I_{10,241}))$
$\sigma = 2$	0.0001	0.0001	0.0587	0.1993	0.0007
$\sigma = 4$	0.0001	0.0328	0.0056	0.0017	0.0001
$\sigma = 6$	0.1017	0.8109	0.5365	0.2476	0.093

Table S1: Monte-Carlo p -values for various real-valued summaries $I_{10,241}, \hat{I}_1, \dots, \hat{I}_{9999}$, with images generated F_U^{10} . That is, generated according to the product empirical pixel distribution.

Test statistic	$ A_{t(\sigma)}(I_{10,221}) $	$H(A(I_{10,221}))$	$L(A(I_{10,221}))$	$E(A(I_{10,221}))$	$\Delta(A(I_{10,221}))$
$\sigma = 2$	0.0782	0.0572	0.1372	0.3118	0.0705
$\sigma = 4$	0.1890	0.0230	0.0329	0.0027	0.0509
$\sigma = 6$	0.2825	0.2104	0.1512	0.0294	0.1681

Table S2: Monte-Carlo p -values for various real-valued summaries $I_{10,221}, \hat{I}_1, \dots, \hat{I}_{9999}$, with images generated according to F_U^{10} . That is, generated according to the product empirical pixel distribution.

Test statistic	$\text{SNR}(I_{10,241})$	$S(A(I_{10,241}))$	$L_1(A(I_{10,241}))$	$L_2(A(I_{10,241}))$
$\sigma = 2$	0.9802	0.0241	0.9999	0.2071
$\sigma = 4$	0.8655	0.3939	0.0038	0.0025
$\sigma = 6$	0.3250	0.8391	0.0838	0.1869

Table S3: Monte Carlo p -values p_n for various real-valued topological summaries of $I_{10,241}, \hat{I}_1, \dots, \hat{I}_{9999}$ smoothed with varying values of σ . Images were generated according to $\hat{P}_{\ell,0}$. With at least 95% confidence the true p -value lies in $p_n \pm 0.0137$ (truncating at 0 or 1). These summaries where not considered in the main document.

Test statistic	$\text{SNR}(I_{10,221})$	$S(A(I_{10,221}))$	$L_1(A(I_{10,221}))$	$L_2(A(I_{10,221}))$
$\sigma = 2$	0.7440	0.0357	0.9006	0.5246
$\sigma = 4$	0.1681	0.4042	0.2817	0.0523
$\sigma = 6$	0.1532	0.6417	0.2733	0.157

Table S4: Monte Carlo p -values p_n for various real-valued topological summaries of $I_{10,221}, \hat{I}_1, \dots, \hat{I}_{9999}$ smoothed with varying values of σ . Images were generated according to $\hat{P}_{\ell,0}$. With at least 95% confidence the true p -value lies in $p_n \pm 0.0137$ (truncating at 0 or 1). These summaries where not considered in the main document.

Test statistic	$\text{SNR}(I_{10,241})$	$S(A(I_{10,241}))$	$L_1(A(I_{10,241}))$	$L_2(A(I_{10,241}))$
$\sigma = 2$	0.9812	0.0231	0.9995	0.1573
$\sigma = 4$	0.8701	0.3795	0.0026	0.0013
$\sigma = 6$	0.3311	0.8375	0.0713	0.1784

Table S5: Monte-Carlo p -values for various real-valued summaries $I_{10,241}, \hat{I}_1, \dots, \hat{I}_{9999}$, with images generated according to F_U^{10} . That is, generated according to the product empirical pixel distribution. These summaries where not considered in the main document.

Test statistic	$\text{SNR}(I_{10,221})$	$S(A(I_{10,221}))$	$L_1(A(I_{10,221}))$	$L_2(A(I_{10,221}))$
$\sigma = 2$	0.7500	0.0368	0.8758	0.4552
$\sigma = 4$	0.1692	0.4069	0.2527	0.0402
$\sigma = 6$	0.1566	0.6462	0.2564	0.1433

Table S6: Monte-Carlo p -values for various real-valued summaries $I_{10,221}, \hat{I}_1, \dots, \hat{I}_{9999}$, with images generated according to F_U^{10} . That is, generated according to the product empirical pixel distribution. These summaries where not considered in the main document.

at local minima increases as smoothing increases—local minima can only be destroyed and not created. This is particularly relevant because the locations of the pixels which create connected components—that correspond to locations of atomic columns—represent local minima [Robins et al., 2011]. Such stability may not strictly be the case for 2-dimensional discrete signals for low-levels of smoothing [Lindeberg, 1990]. However, when $\sigma \geq 1$ (as is the case for all practical settings in this article), the ideal “Lindeberg” kernel (which does not increase local minima) and the discrete Gaussian kernel coincide to a large degree [Getreuer, 2013].

Furthermore, the discretized Gaussian kernel does not increase the number of local minima from the unsmoothed image [Lindeberg, 1990]. Therefore, in theory, tuning the σ parameter appropriately will allow for cubical persistence to recover the precise number of relevant atomic columns—their location as well as an estimate of their intensity. Additionally, the Gaussian kernel has been empirically shown to introduce the fewest “image artifacts” [Levin et al., 2020]. Another option would be to convolve our image with an elliptical Gaussian kernel, such as in Kong et al. [2013] or to apply our cubical homology algorithm directly to a scale-space representation [Lindeberg, 1998]. Decay rates of the norms of persistence diagrams after convolution with a Gaussian kernel have also been established [Chen and Edelsbrunner, 2011] (see Section S3).

S5 Comparison of blob detection versus our algorithm

Before our discussion, we note that PD thresholding [Chung and Day, 2018] was used to threshold our images for visualization. PD thresholding was shown (ibid.) to more accurately represent the topology of the underlying image than traditional histogram-based thresholding methods. Here we chose the t which maximizes the objective function specified at (11) on p. 1172 of Chung and Day. We chose to only weight 0-dimensional topological features and ignore PD_1 . Our work indicates that higher levels of σ , such as $\sigma = 4$, may yield more topologically faithful thresholded images.

Continuing on, in Figure S4a we see the output of LoG blob detection on an image which shows various ridges, or *Miller planes*, in the 2-dimensional nanoparticle image; in

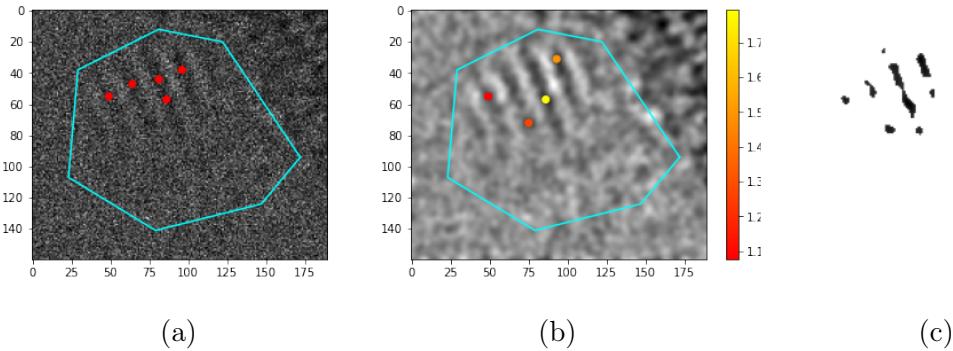


Figure S4: Comparison of the output of Laplacian-of-Gaussian blob detection (Figure S4a) with our algorithm (Figure S4b), along with segmented image using the output of our algorithm (Figure S4c).

Figure S4b we see the same but for our cubical persistence algorithm. In Figure S4c we see the output of PD thresholding (on the unthresholded image). Though LoG blob detection and our algorithm only output points, and Miller planes are visible, there is a one-to-one correspondence between these ridges and points in the case of our algorithm. Thirty values of σ from 4 to 8 and a threshold of 0.4 were chosen for blob detection. The values of $\sigma = 2$ and $\eta = 1$ were chosen for our algorithm in S4b. It did not seem possible to recover 4 planes with blob detection, even after tuning the parameter σ and the threshold.

S6 Elbow method for the selection of σ

The method that we have used for determining the best σ is the elbow method depicted in Figure S6. For a single image (as in $I_{10,221}$ and $I_{10,241}$) we can apply the elbow method to our desired summary to choose the best σ to be located near the point where the derivative of the curve of our topological summary vs. σ stabilizes—that is, our topological summary no longer experiences a rapid decrease¹ (see Figure S7). This method can be made more rigorous than the elbow method in the case of k -means clustering, because of the fact we can choose our mesh of σ to be arbitrarily fine and thus calculate derivatives to a

¹We know these summaries do eventually decrease based off the work of [Chen and Edelsbrunner \[2011\]](#).

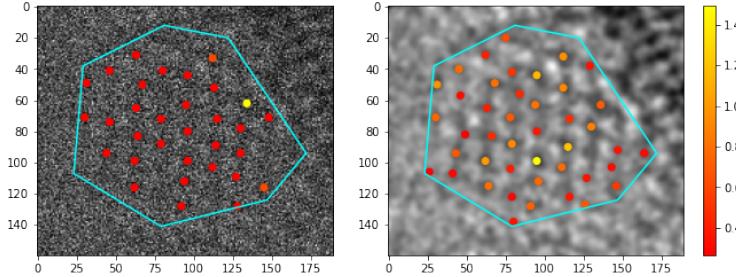


Figure S5: Image $I_{10,281}$ overlaid with output of Laplacian of Gaussian (LoG) blob detection centers (left) and same output of the cubical homology algorithm overlaid over smoothed $\sigma = 2$ version of $I_{10,281}$. For such a high noise regime (see left image), LoG blob detection and our cubical persistence algorithm convey much of the same information. The algorithms (using `blob_log` in scikit-image) run in roughly the same time, though the cubical homology algorithm is marginally faster. Colors in left image correspond to the bandwidth of the convolved Gaussian kernel.

precise degree². For a single image I , the topological summary (e.g. persistent entropy or ALPS) may be calculated on smooth versions of I , in particular for the nanoparticle region R . However, for a large image sequence I_1, \dots, I_N we plot the value of the topological summary in R for smoothed versions of the summed image $I_{N,1}$ in Figure S6. The summed image $I_{N,1}$ can be used to capture the average dynamics of the nanoparticle and can thus be used as a rough estimate of shape. Another idea is to use to some representative subset of $I_{k_m}, m = 1, \dots, M \leq N$ of the image series and use the ensemble of elbow plots as a guide, so as to base the choice of σ off the unaggregated images themselves. Another point is that it seems useful to be a little more conservative about where to identify the elbow (as we have done in Figure S6), as we would like to diminish the noise and capture large-scale features as much as possible.

S7 Proof of Proposition 5.1

He we offer a proof of an alternate representation of the ALPS statistic.

²As mentioned in the main article, $\sigma \mapsto I_\sigma$ can be shown to be smooth for σ bounded away from 0.

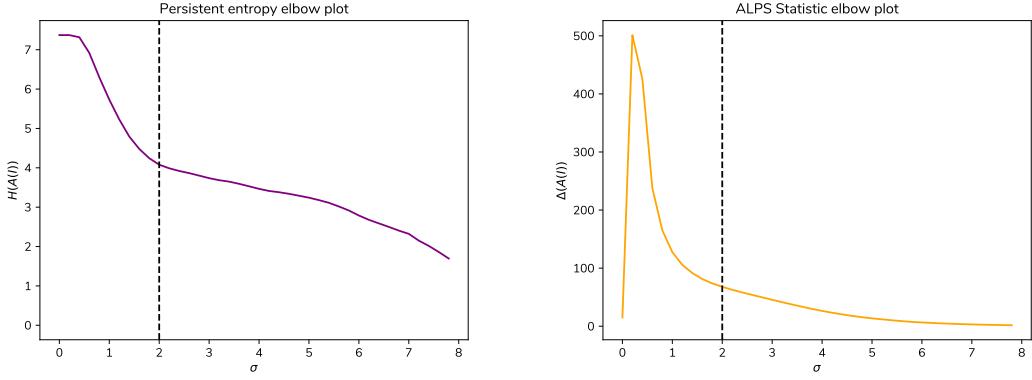


Figure S6: Persistent entropy elbow plot (left) and ALPS statistic elbow plot (right) for summed main primary image series, i.e. $I_{1124,1}$. Values of σ here ranged from 0 to 7.8 in increments of 0.2.

Proof of Proposition 5.1. Let $l_{(0)} \equiv 0$. By definition of Lebesgue integral and Tonelli's theorem, we have

$$\begin{aligned}
\int_0^\infty \log U(\eta) d\eta &= \int_0^\infty \log \left(\sum_{(x,l) \in A(I)} \mathbf{1}\{l > \eta\} \right) d\eta \\
&= \int_0^\infty \sum_{i=0}^{n-1} \log(K-i) \mathbf{1}\{l_{(i)} < \eta \leq l_{(i+1)}\} d\eta \\
&= \sum_{i=0}^{n-1} (l_{(i+1)} - l_{(i)}) \log(K-i) \\
&= \log \left(\prod_{i=0}^{K-1} (K-i)^{l_{(i+1)}} \right) - \log \left(\prod_{i=0}^{K-1} (K-i)^{l_{(i)}} \right).
\end{aligned}$$

As $l_{(0)} = 0$, the product within the logarithm of the second term equals

$$\begin{aligned}
\prod_{i=1}^{K-1} \frac{(K-i)^{l_{(i)}}(K-i+1)^{l_{(i)}}}{(K-i+1)^{l_{(i)}}} &= \prod_{i=1}^{K-1} \left(1 - \frac{1}{K-i+1} \right)^{l_{(i)}} \prod_{i=1}^K (K-i+1)^{l_{(i)}} \\
&= \prod_{i=1}^{K-1} \left(1 - \frac{1}{K-i+1} \right)^{l_{(i)}} \prod_{i=0}^{K-1} (K-i)^{l_{(i+1)}},
\end{aligned}$$

which finishes the proof. \square

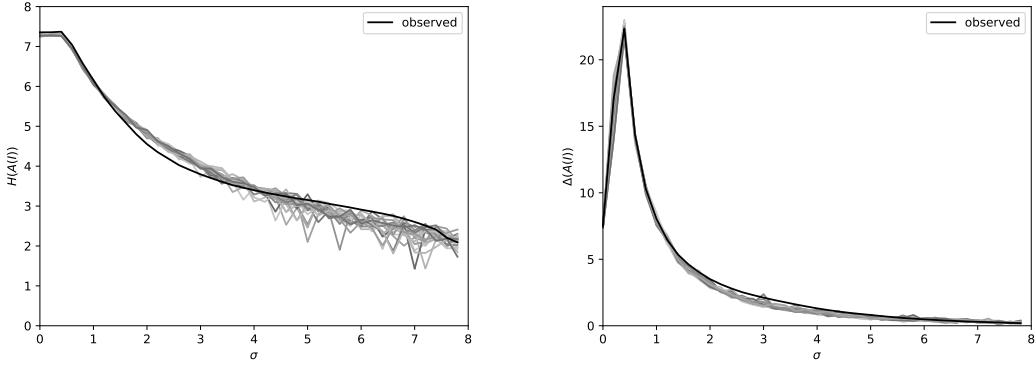


Figure S7: Persistent entropy (left) and ALPS (right) for various smoothing parameters σ for image $I_{10,241}$, compared with same series for 19 simulated images. Same observed curve as in Figure 9 of the main article. As with the Figure 9, values of σ ranged from 0 to 7.8, in increments of 0.2.

S8 Multiple Monte Carlo testing guarantees

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We consider the problem of multiple testing in a Monte Carlo setting where $H_{0,1}, \dots, H_{0,N}$ are null hypotheses under which there exist distribution functions $(F_{0,k})_{k=1}^N$ satisfying $F_{0,k} = F_0$ for all k , for F_0 some continuous cdf on \mathbb{R} . Here N is considered fixed but arbitrary. Denote our (i.i.d.) Monte Carlo sample of test statistics under F_0 as T_1, \dots, T_n , which we assume are also independent from the observed (random) test statistics t_1, \dots, t_N . Let $p_{n,1}, \dots, p_{n,N}$ denote our Monte Carlo p -values and $p_{n,(1)} \leq \dots \leq p_{n,(N)}$ to be their order statistics. We specify $H_{0,(1)}, \dots, H_{0,(N)}$ to be the hypotheses corresponding to the ordered p -values. We suppose that all of our test statistics are nonnegative. Let us define

$$p_{n,k} := \frac{1}{n+1} \left(1 + \sum_{i=1}^n \mathbf{1}\{T_i \geq t_k\} \right),$$

and $r_{n,k}$ to be the rank of $p_{n,k}$ among all of the values of $p_{n,1}, \dots, p_{n,N}$, e.g. $r_{n,(k)} = k$. We construct a test (a Monte Carlo version of the one in [Benjamini and Yekutieli, 2001](#)), based off of the inequalities

$$p_{n,k} \leq \frac{r_{n,k}\alpha}{NC_N}, \tag{1}$$

where $C_N = \sum_{k=1}^N 1/k$. Let FDR_n be the *false discovery rate* for the test based on (1). That is, FDR_n is the expected proportion of falsely rejected hypotheses over the total number of rejected hypotheses due to the criterion in (1) (equal to 0 if no hypotheses are rejected). Note that FDR_n is conditional on $(T_i)_{i=1}^n$. Let FDR denote the false discovery rate based on the true p -values $p_k = \mathbb{P}(T_i \geq t_k)$. We specify the exact nature of test—along with a theoretical guarantee—in the next proposition.

Proposition S8.1. *Suppose T_1, \dots, T_n are test statistics sampled i.i.d. according to continuous cdf F_0 . If the $(t_k)_{k=1}^N$ have continuous distributions under the alternative hypotheses as well, then the test which rejects $H_{0,(1)}, \dots, H_{0,(\ell)}$ if*

$$\ell = \sup \{k = 1, \dots, N : p_{n,(k)} \leq k\alpha/(NC_N)\}$$

satisfies $\text{FDR}_n \xrightarrow{\text{a.s.}} \text{FDR}$, $n \rightarrow \infty$. Thus,

$$\limsup_n \text{FDR}_n \leq \alpha, \quad \text{a.s.}$$

The proof of Proposition S8.1 can be seen below. We are also able to consider the test statistics t_i , $i = 1, \dots, N$ as random as well. Our result is an improvement on Corollary 1 in [Gandy and Hahn \[2014\]](#) for the Benjamini-Hochberg procedure, in that our results holds with probability 1, rather than probability $1 - \epsilon$, for n large enough.

Proof of Proposition S8.1. We let F_N be the empirical cdf of t_1, \dots, t_N —hence we may restate the test corresponding to Proposition A.1 as

$$p_{n,k} \leq \frac{F_N(t_k)\alpha}{C_N}.$$

If F_n is the empirical cdf of the T_i , it is straightforward to show that

$$\sup_{t \in \mathbb{R}} |F_n(t) - F(t)| \xrightarrow{\text{a.s.}} 0, \quad n \rightarrow \infty, \tag{2}$$

implies that

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}(T \geq t) - \frac{1}{n+1} \left(1 + \sum_{i=1}^n \mathbf{1}\{T_i \geq t\} \right) \right| \xrightarrow{\text{a.s.}} 0, \quad n \rightarrow \infty. \tag{3}$$

The equation (2) holds because of the Glivenko-Cantelli theorem. Suppose we denote

$$p_k := \mathbb{P}(T_i \geq t_k).$$

Then (3) implies that

$$\max_k |p_k - p_{n,k}| \xrightarrow{a.s.} 0,$$

as $n \rightarrow \infty$. Additionally,

$$\max_k \frac{|p_k - p_{n,k}|}{F_N(t_k)} \leq N \max_k |p_k - p_{n,k}| \xrightarrow{a.s.} 0, \quad n \rightarrow \infty.$$

As $\mathbb{P}(\bigcup_k \{p_k/F_N(t_k) = \alpha/C_N\}) = 0$, we have that

$$\mathbb{P}\left(\cap_k \left\{ \lim_{n \rightarrow \infty} \mathbf{1}\{p_{n,k}/F_N(t_k) \leq \alpha/C_N\} = \mathbf{1}\{p_k/F_N(t_k) \leq \alpha/C_N\} \right\} = 1\right).$$

This implies that for \mathbb{P} -almost every $\omega \in \Omega$ that there exists some $M = M(\omega)$ such that if $n \geq M$ then

$$\mathbf{1}\{p_{n,k}/F_N(t_k) \leq \alpha/C_N\} = \mathbf{1}\{p_k/F_N(t_k) \leq \alpha/C_N\}, \text{ for all } k = 1, \dots, N.$$

Let FDP be the false discovery proportion³ based on the test using p_k , $k = 1, \dots, N$ and let FDP_n be the corresponding quantity for $p_{n,k}$. Note that the false discovery rate satisfies $\text{FDR} = \mathbb{E}[\text{FDP}]$ (and $\text{FDR}_n = \mathbb{E}[\text{FDP}_n]$). Based on the definition of FDP, for such an ω , $\text{FDP}_n = \text{FDP}$ as the tests are identical. Hence,

$$\text{FDP}_n \xrightarrow{a.s.} \text{FDP}, \quad n \rightarrow \infty.$$

If we let \mathcal{T}_n be the natural filtration with respect to T_1, \dots, T_n we have by the dominated convergence theorem for conditional expectations

$$\text{FDR}_n = \mathbb{E}[\text{FDP}_n | \mathcal{T}_n] \xrightarrow{a.s.} \mathbb{E}[\text{FDP} | \mathcal{T}_n] = \text{FDR}.$$

Therefore based on the results of [Benjamini and Yekutieli \[2001\]](#) for the false discovery rate FDR, we have

$$\limsup_{n \rightarrow \infty} \text{FDR}_n \leq \alpha, \quad \text{a.s.}$$

□

³The ratio of falsely rejected nulls over the total number of rejected nulls. It is defined to be 0 if no null hypotheses are rejected.

We now derive the confidence intervals for the p -values in the Tables S3 and S4 above and seen in Tables 2 and 3 in the main document.

Proof. Let t be our observed quantity of a test statistic T , and T_1, \dots, T_n i.i.d. simulated values of the test statistic T under then null hypothesis. Then if

$$p_n(t) = \frac{1}{n+1} \left(1 + \sum_{i=1}^n \mathbf{1}\{T_i \geq t\} \right),$$

$p_n(t)$ is consistent estimator of the p -value $p(t) := P(T \geq t)$. In the case of a continuous T , we have that the empirical survival function

$$\bar{F}_n(t) := \frac{1}{n} \mathbf{1}\{T_i > t\},$$

is also an unbiased estimator of $p(t)$, and we can use the DKW inequality to see that if

$$q_\alpha = \sqrt{-\frac{1}{2n} \log(\alpha/2)},$$

then

$$\mathbb{P}\left(\sup_{t \in \mathbb{R}} |\bar{F}_n(t) - \bar{F}(t)| \leq q_\alpha\right) \geq 1 - \alpha. \quad (4)$$

Therefore, with *at least* confidence level $1 - \alpha$, we have that

$$\bar{F}_n(t) - q_\alpha \leq p(t) \leq \bar{F}_n(t) + q_\alpha,$$

or

$$\bar{F}_n(t) \pm \sqrt{-\frac{1}{2n} \log(\alpha/2)}.$$

Now, if T is discrete, say on the integers, fix a $\delta \in (0, 1]$ and suppose that t is an integer. Then $\bar{F}_n(t - \delta)$ is a consistent estimator of $p(t)$. Thus, for T discrete on the integers, we have

$$\bar{F}_n(t - \delta) \pm \sqrt{-\frac{1}{2n} \log(\alpha/2)},$$

is an at least $(1 - \alpha)\%$ confidence interval for $p(t)$. It is not difficult to show that

$$\bar{F}_n(t) \leq p_n(t) + 1/n$$

and likewise $\bar{F}_n(t - \delta) \leq p_n(t) + 1/n$ in the discrete case. Similarly, we can establish a lower bound for those terms of $p_n(t) - 1/n$. Hence, with at least $(1 - \alpha)\%$ confidence $p(t)$ lies within the interval

$$\left[\max\{0, p_n(t) - q_{\alpha,n}\}, \min\{1, p_n(t) + q_{\alpha,n}\} \right],$$

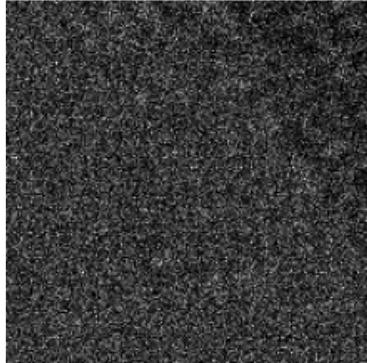
where $q_{\alpha,n} := q_\alpha + 1/n$. The same interval estimate holds in the case that T_i are generated according to some distribution with using a sufficient statistic for a parameter θ instead of θ due to invariance of (4) when conditioning on a sufficient statistic for θ . \square

S9 Frames of interest in multiple testing demonstration

Here we show a few frames corresponding to Figure 12 in the main document. We display the lowest negative entropy and highest ALPS statistic frames with $k = 35$ and $k = 201$, respectively for the top region and the highest negative entropy and lowest ALPS statistic frames for the bottom region with $k = 20$ and $k = 172$, respectively. These figures can be seen in Figure S8.

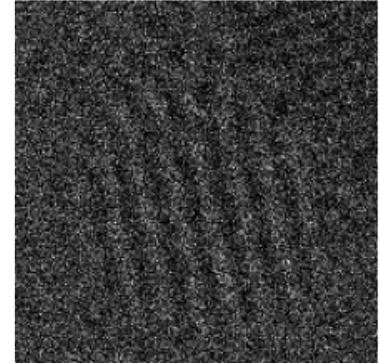
Note that a frame is considered significant by our Benjamini-Hochberg procedure if in the top region the persistent entropy is ≤ 4.7537 or the ALPS statistic is ≥ 2.3747 . The same values corresponding to the bottom region are ≤ 3.7868 and ≥ 1.8943 . We now examine two frames which are significant with respect to persistent entropy but not with respect to the ALPS statistic. The first is $I_{5,761}$ ($k = 152$) for the top region, where $H(A(I_{5,761})) = 4.6924 \leq 4.7537$ and $\Delta(A(I_{5,761})) = 2.1725 \geq 2.3747$, which can be seen in Figure S9.

The second is the image $I_{5,501}$ ($k = 100$) for the bottom nanoparticle region; it is an image which shares significance with respect to persistent entropy $H(A(I_{5,761})) = 3.7654 \leq 3.7868$ and non-significance with respect to the ALPS statistic. However, it is quite closer in absolute terms to the ALPS statistic threshold $\Delta(A(I_{5,501})) = 1.8728 \geq 1.8943$. The intuitive interpretation is perhaps the ALPS statistic is closer to capturing the “intuitive” nature of significance better than persistent entropy. Indeed, there seems to be more signal in Figure S10 than in Figure S9.



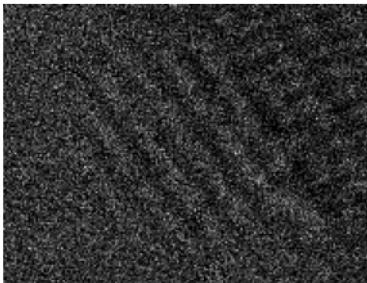
(a) $I_{5,176}$ ($k = 35$), top region.

Lowest negative entropy frame.



(b) $I_{5,1006}$ ($k = 201$), top region.

Highest ALPS statistic frame.



(c) $I_{5,101}$ ($k = 20$), bottom region.

Highest negative entropy frame.



(d) $I_{5,861}$ ($k = 172$), bottom region.

Lowest ALPS statistic frame.

Figure S8: Examples of frames in the image series $I_{5,5k+1}$ that illustrate how the nanoparticles look when they take on low and high values of the ALPS statistic and persistent entropy.

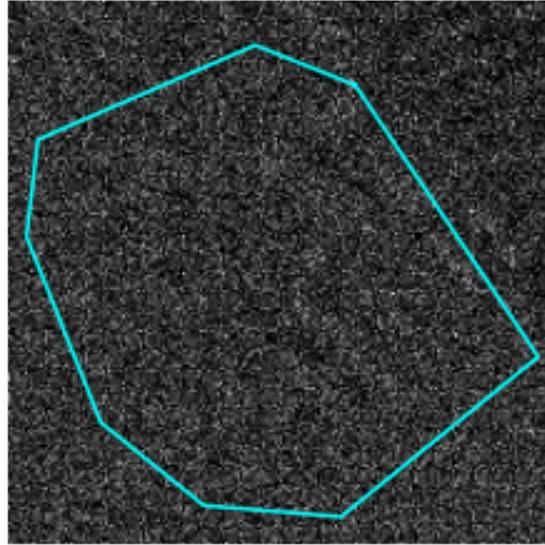


Figure S9: Barely significant at $\alpha = 0.05$ level with respect to persistent entropy and non-significant with respect to the ALPS statistic, one can make out some faint atomic features in $I_{5,761}$ for the top region. Boundaries of polygonal region R have been superimposed on the image.

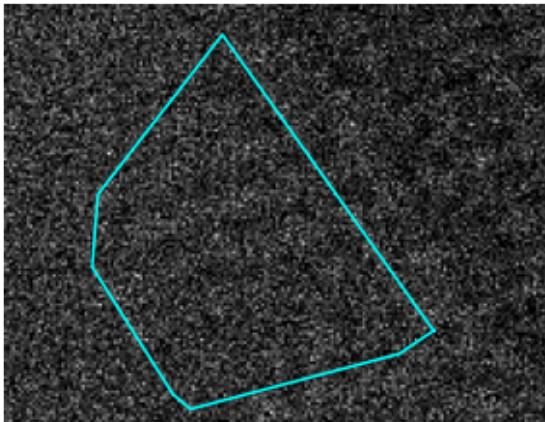


Figure S10: Barely significant at $\alpha = 0.05$ level with respect to persistent entropy and barely non-significant with respect to the ALPS statistic, one can make out some faint atomic features in $I_{5,501}$ for the bottom region. Boundaries of polygonal region R have been superimposed on the image.

References

Jean Babaud, Andrew P Witkin, Michel Baudin, and Richard O Duda. Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE transactions on pattern analysis and machine intelligence*, (1):26–33, 1986.

Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165 – 1188, 2001. doi: 10.1214/aos/1013699998. URL <https://doi.org/10.1214/aos/1013699998>.

Chao Chen and Herbert Edelsbrunner. Diffusion runs low on persistence fast. In *2011 International Conference on Computer Vision*, pages 423–430. IEEE, 2011.

Yu-Min Chung and Sarah Day. Topological fidelity and image thresholding: A persistent homology approach. *Journal of Mathematical Imaging and Vision*, 60(7):1167–1179, 2018.

David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have l_p -stable persistence. *Foundations of computational mathematics*, 10(2):127–139, 2010.

Noel Cressie. *Statistics for spatial data*. John Wiley & Sons, 1993.

Axel Gandy and Georg Hahn. MMCTest—a safe algorithm for implementing multiple monte carlo tests. *Scandinavian Journal of Statistics*, 41(4):1083–1101, 2014.

Pascal Getreuer. A survey of gaussian convolution algorithms. *Image Processing On Line*, 2013:286–310, 2013.

Yasuaki Hiraoka and Kenkichi Tsunoda. Limit Theorems for Random Cubical Homology. *Discrete & Computational Geometry*, 60(3):665–687, oct 2018. ISSN 0179-5376. doi: 10.1007/s00454-018-0007-z. URL <http://link.springer.com/10.1007/s00454-018-0007-z>.

Hui Kong, Hatice Cinar Akakin, and Sanjay E. Sarma. A generalized laplacian of gaussian filter for blob detection and its applications. *IEEE Transactions on Cybernetics*, 2013. ISSN 21682267. doi: 10.1109/TSMCB.2012.2228639.

Barnaby D.A. Levin, Ethan L. Lawrence, and Peter A. Crozier. Tracking the picoscale spatial motion of atomic columns during dynamic structural change. *Ultramicroscopy*, 213, 6 2020. ISSN 0304-3991. doi: 10.1016/j.ultramic.2020.112978.

Tony Lindeberg. Scale-space for discrete signals. *IEEE transactions on pattern analysis and machine intelligence*, 12(3):234–254, 1990.

Tony Lindeberg. Feature detection with automatic scale selection. *International journal of computer vision*, 30(2):79–116, 1998.

Bastian Rieck, Tristan Yates, Christian Bock, Karsten Borgwardt, Guy Wolf, Nicholas Turk-Browne, and Smita Krishnaswamy. Uncovering the topology of time-varying fMRI data using cubical persistence. *Advances in Neural Information Processing Systems (NeurIPS)*, 33(NeurIPS):6900–6912, 2020. ISSN 10495258.

Vanessa Robins, Peter John Wood, and Adrian P Sheppard. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *IEEE Transactions on pattern analysis and machine intelligence*, 33(8):1646–1658, 2011.