

Topological data analysis of black hole images

Pierre Christian¹,² Chi-kwan Chan,^{2,3,4} Anthony Hsu,⁵ Feryal Özel², Dimitrios Psaltis,² and Iniyan Natarajan^{6,7}

¹*Physics Department, Fairfield University, 1073 N Benson Road, Fairfield, Connecticut 06824, USA*

²*Department of Astronomy, University of Arizona, 933 North Cherry Avenue, Tucson, Arizona 85721, USA*

³*Data Science Institute, University of Arizona, 1230 N. Cherry Avenue, Tucson, Arizona 85721, USA*

⁴*Program in Applied Mathematics, University of Arizona, 617 N. Santa Rita, Tucson, Arizona 85721, USA*

⁵*Department of Computer Science, University of Arizona, 1040 4th Street, Tucson, Arizona 85721, USA*

⁶*Wits Centre for Astrophysics, University of the Witwatersrand,*

1 Jan Smuts Avenue, Braamfontein, Johannesburg 2050, South Africa

⁷*South African Radio Astronomy Observatory, Observatory 7925, Cape Town, South Africa*



(Received 1 February 2022; accepted 24 June 2022; published 15 July 2022)

Features such as photon rings, jets, or hot spots can leave particular topological signatures in a black hole image. As such, topological data analysis can be used to characterize images resulting from high-resolution observations (synthetic or real) of black holes in the electromagnetic sector. We demonstrate that persistent homology allows for this characterization to be made automatically by counting the number of connected components and one-dimensional holes. Further, persistent homology also allows for the distance between connected components or the diameter of holes to be extracted from the image. In order to apply persistent homology on synthetic black hole images, we also introduce metronization, a new algorithm to prepare black hole images in a form that is suitable for topological analysis.

DOI: [10.1103/PhysRevD.106.023017](https://doi.org/10.1103/PhysRevD.106.023017)

I. INTRODUCTION

The Event Horizon Telescope (EHT) recently showcased the feasibility of ground-based very-long-baseline interferometers (VLBIs) to produce images of supermassive black holes [1–6]. In performing the analysis for its data set, the EHT Collaboration utilized a suite of synthetic observations generated by ray-tracing general-relativistic magnetohydrodynamic (GRMHD) simulations assuming certain radiative transfer models [5].

As the EHT project expands, the variety of parameters included in the synthetic image calculations necessarily increases. Modifications to the synthetic image calculations can vary from simply increasing the range and resolution of previous calculations (e.g., computing more inclination angles), to changing the underlying astrophysical model (e.g., utilizing different radiative transfer or GRMHD models), to changing the fundamental physical theories (e.g., utilizing nongeneral-relativistic theories).

This bounty of synthetic images presents a problem. Given their quantity, it will be prohibitively expensive to compare each synthetic image with the observed data set. Thus, a method needs to be devised to quickly classify the synthetic images based on some broad features, which can then be quickly compared with either the observed data set or some prior knowledge of the black hole in question.

Further, interferometric imaging such as the one conducted by the EHT relies on reconstructing real-space images from incomplete observations of the target in

Fourier space [7–10]. Different reconstruction methods, as well as different choices of parameters within one reconstruction method, can produce widely differing images. Given the large amount of image-domain reconstructions that are produced in projects such as the EHT, it is also important to be able to quickly classify these images into bins based on some broad features.

Apparent in the first EHT observations [1,6] and in the synthetic observations of the M87 black hole [5] is a particular topological signature of the black hole image: a bright ring surrounding a dim region. We propose that topological signatures such as the existence of such a ring are a good set of features to classify black hole images.

Topological data analysis (TDA) is a set of techniques that allows for robust analysis of the topological characteristics of a data set to be made. In this work, we introduce a method to classify black hole images (synthetic or real) by the application of persistent homology, a TDA technique that measures certain topological invariants of a data set. In persistent homology, one first converts an image into a topological space associated with the data points of the image. Once in the form of a topological space, standard mathematical techniques can be used to compute the topological properties of the image. More details on TDA and persistent homology can be found in Refs. [11,12].

Persistent homology converts a high-dimensional data set (a black hole image) into more manageable low-dimensional information (its topological characteristics),

and is thus an example of a dimensionality reduction technique. In the study of black hole shadows, dimensionality reduction in the form of principal component analysis (PCA) has previously been utilized to decompose synthetic black hole images [13] and theoretical models of black hole shadows [14] into their eigencomponents. In contrast to PCA, where the transformation from high- to low-dimension spaces is linear, persistent homology applies a nonlinear transformation. This nonlinearity allows persistent homology to retain information that is lost by linear dimensionality reduction techniques.

In order to apply TDA to the large image set, we develop a two-step algorithm. The first step, which we call *metronization*, is a dimensional reduction procedure that condenses a black hole image into simple graphs that retain all of the topological information of the image. This step both transforms a black hole image into a form amenable to TDA and greatly accelerates the TDA calculations. The second step is to run persistent homology algorithms on the metronization output. Currently, there are many available open-source software packages to perform the later computational topology tasks. In this project, we utilize the open-source code `Dionysus 2`¹ [15]. Other examples of computational topology software include `Ripser++`² [16], which possesses graphics processing unit acceleration, and the `Topology Toolkit`³ [17].

This paper is organized as follows. In Sec. II, we give an introduction to persistent homology. In Sec. III, we describe metronization, our preprocessing technique to convert black hole images into a form suitable for persistent homology. In Sec. IV, we show examples of persistent homology applied to mock black hole images. In Sec. V, we supply our concluding remarks.

II. PERSISTENT HOMOLOGY

An abstract simplicial complex can be thought of as a collection of n -dimensional triangles, i.e., a collection of zero-dimensional points, one-dimensional lines, two-dimensional triangles, three-dimensional tetrahedrons, etc., (see, e.g., Ref. [18] for a rigorous definition). Given such a structure, one can compute the topological invariants associated with the homology groups—a set of mathematical objects that are closely related to the number of holes on a space. Considering our data points as points in \mathbb{R}^2 equipped with its standard metric,

$$d = \sqrt{x^2 + y^2}, \quad (1)$$

where (x, y) are coordinates on \mathbb{R}^2 and d is the distance on the image plane, we can transform our data points into a

simplicial complex using the Vietoris-Rips algorithm [19]. For a length scale l , the algorithm is as follows:

- (1) Consider each data point as a 0-simplex.
- (2) Two points within a distance l are connected with a 1-simplex.
- (3) Three points that are connected to each other by a 1-simplex form a filled triangle (a 2-simplex).

This algorithm can then be repeated for many l 's, and thus gives the topological invariants as a function of l .

For our purposes, the most relevant topological invariants are the Betti numbers, which can be computed from a simplicial complex in a straightforward manner. For a space X , they are defined as

$$b_n(X) = \text{Rank}[H_n(X)], \quad (2)$$

where $\text{Rank}[G]$ is the rank of a group G and $H_n(X)$ is the n th homology group of X , which is defined as

$$H_n = \text{Ker}(\partial_n) / \text{Im}(\partial_{n+1}), \quad (3)$$

where ∂_n is the n th-dimensional boundary operator of the simplicial complex and (as from now on) the dependence on the space X has been suppressed. Equivalently, the Betti number can be computed by

$$b_n = \dim[\text{Ker}(\partial_n)] - \dim[\text{Im}(\partial_{n+1})], \quad (4)$$

where $\dim[\]$ is the dimension operator.

As black hole images are two-dimensional and b_n is zero for n greater than the dimension of the space, we only need to consider b_0 , which counts the number of connected components, and b_1 , which counts the number of one-dimensional holes (i.e., circles). For example, a space with two holes will have

$$H_2 = \mathbb{Z} \oplus \mathbb{Z}, \quad (5)$$

which gives $b_2(U) = 2$.

A. Persistent homology in practice

A barcode diagram is a typical output of topological data analysis that encodes the Betti numbers found by the persistent homology algorithm as a function of l . A schematic example of a barcode diagram is shown in Fig. 1.

Suppose that the diagram in Fig. 1 is the b_1 barcode diagram (one barcode diagram can be plotted for each Betti number). The horizontal axis of a barcode diagram is l , the characteristic length scale. At a particular l , a manifold is formed out of the data points, the number of one-dimensional holes in this manifold is computed, and each hole is plotted as a point along the vertical axis at said l (at no particular order on the horizontal axis). Figure 1 then shows that when $l = 5$ there are three holes in the image,

¹<https://mrzv.org/software/dionysus2/>.

²<https://github.com/simonzhang00/ripser-plusplus>.

³<https://topology-tool-kit.github.io/>.




FIG. 1. Schematic example of a barcode diagram.

as there are three points plotted at $l = 5$. Similarly, when $l = 12$ there is only one hole in the image.

As l is increased, holes are *born* when l becomes large enough to connect the data points forming the boundary of the hole. However, these holes eventually *die* when l becomes large enough that they get filled (see Fig. 2). As such, for every detected hole, the barcode diagram for b_1 starts at the l where the hole is first detected and ends at the l where the hole gets filled in. The barcode diagram for b_0 can be analogously described, with the number of connected components taking the place of the number of holes.

Note that, unlike homology, persistent homology is sensitive to distances and sizes. The characteristic length where the barcode diagram ends, l_{death} , for the b_1 diagram is a measure of the size of the hole. In particular, for a data set with a circular hole of radius R ,

$$l_{\text{death}} = 2R \cos \pi/6 \approx 1.73R, \quad (6)$$

and for a data set consisting of two separate clumps of data points, l_{death} for the b_0 diagram is a measure of the distance between these two clumps.




FIG. 2. Birth and death of topological features: due to the gap in the data, l_1 is too short to connect all of the data points to form a circular object, and so $b_1(l_1) = 0$. In contrast, l_2 is long enough that the data points are connected together into a disk-like object with a hole in the middle, giving $b_1(l_2) = 1$. However, as l is increased, this hole eventually gets filled in, and is no longer detected when $l = l_3$, i.e., $b_1(l_3) = 0$.

Most topological signals are short lived, i.e., with a “lifespan”

$$\Delta l = l_{\text{death}} - l_{\text{birth}} \quad (7)$$

that is small. For our purposes, these short-lived topological signatures are not interesting, as they are typically the result of noise in the data spuriously forming topological objects such as small rings. As such, we define a cutoff L on the lifespan, where if

$$\Delta l > L, \quad (8)$$

we consider the topological signal to be robust or *persistent*. In the b_1 diagram, a cutoff of L means that only circles with radius at least $\approx 1.73L$ will pass our cutoff. For b_1 diagrams, we also impose a lower cutoff: any holes that are born at $l = 0$ must be due to the noise as at $l = 0$ the data points should all be disconnected.

Figure 3 shows a data set consisting of data points arranged in a noisy circle (by noisy we mean that the data set contain points that deviate from a perfect circle) and the barcode diagram corresponding to b_1 . As the barcode diagram attests, there is one component that is much longer than the rest. Figure 3 shows a dataset consisting of datapoints arranged in a noisy circle (by noisy we mean the dataset contain points that deviate from a perfect circle) and the barcode diagram corresponding to b_1 . As the barcode diagram attests, there is one component that is much longer than the rest. Figure 4 shows the b_0 diagram for a dataset consisting of three data clouds.

III. METRONIZATION

The persistent homology algorithm relies on building Vietoris-Rips complexes, which is computationally expensive when the data set is large enough, and exacerbated by the need to build one complex for every l . While a typical EHT reconstructed image has only $\sim 80 \times 80$ pixels, the latest EHT imaging surveys generate millions of them. Therefore, a special preprocessing technique is required to reduce the number of data points N in order to make this TDA algorithm practical for imaging surveys at the EHT scale.

Here we introduce “metronization” and demonstrate that it is an effective preprocessing technique for persistent homology. Metronization simplifies a VLBI image into a representation with less than 100 data points that still contain the same topological information of the original image. This leads to over a factor of a thousand increase in computational speed in the application of the TDA algorithm.

A secondary goal of metronization is to convert a black hole image—which is a three-dimensional data cloud in the space (x, y, I) , where x and y are the horizontal and vertical pixels (or angles) inscribed on the sky and I is the




FIG. 3. Data points in a ring (left) of diameter ~ 0.1 , along with its b_1 barcode diagram (right). The barcode diagrams show the detection of one persistent hole. The b_1 barcode terminates at $l = 2R \cos(\pi/6)$, as expected for an equilateral triangle inscribed in the circular hole.

brightness—into a series of two-dimensional data clouds each in the space (x, y) for every level of brightness. After metronization, emission rings will appear as circles in the two-dimensional pixel space. While two-dimensional holes in the three-dimensional space (x, y, I) or other high-dimensional topological characteristics can contain

important information on the image, we relegate its study to future work.

The metronization algorithm chains up different image processing methods to highlight the topological features in VLBI images. The steps involved in metronization preserve the topological characteristics of the original image, and




FIG. 4. Data points in three distinct connected components (top) and their associated b_0 barcode diagram (bottom). Note that the first connected component that trivially exists for any nonempty data set (with $l_{\text{birth}} = 0$ and $l_{\text{death}} \rightarrow \infty$) is not plotted. The barcode diagram detects three persistent connected components and displays the distances between each connected components in the l_{death} 's.

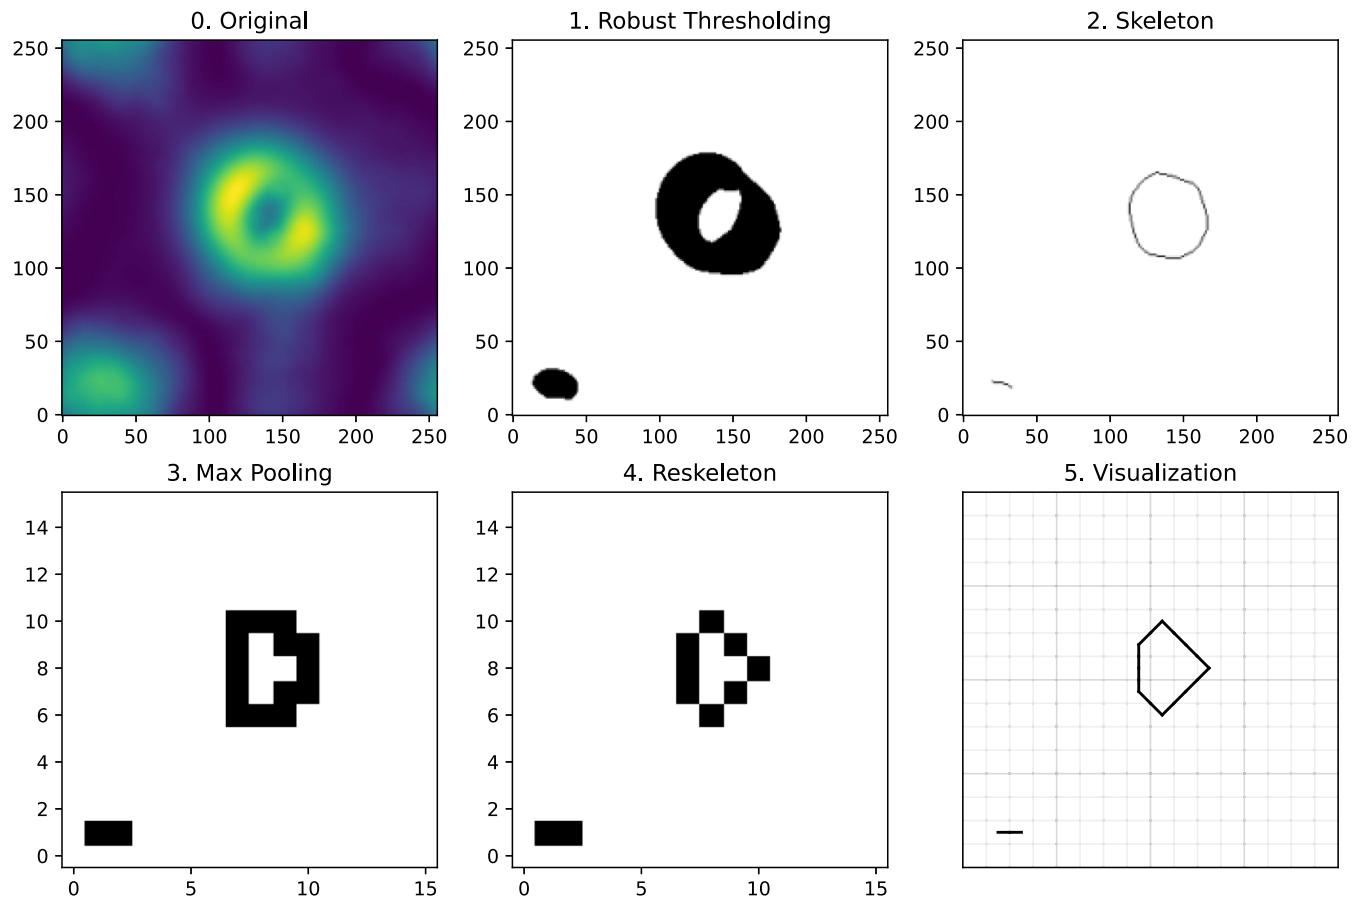


FIG. 5. Example of a synthetic image undergoing the metronization procedure. The robust thresholding parameter is $T = 0.7$. An emission ring can clearly be seen in the original image. While the metronization procedure wiped out most of the details in the original image, this topological feature is preserved.

thus the speedup obtained through metronization does not come at any cost in the accuracy of the TDA results. The different stages are shown in Fig. 5. The original image is shown in the top-left panel and the different steps are shown in the subsequent panels:

- (1) *Robust thresholding*. We first need to turn an image with scalar-valued pixels into one with Boolean-valued pixels. Instead of applying a user-supplied threshold based on the peak intensity, we use the following method. First, we sort the pixel values $I(x_i, y_j)$ into a monotonically increasing sequence $(I_1, I_2, \dots, I_{N^2})$. We then compute the cumulative sum of this sequence $(C_1, C_2, \dots, C_{N^2})$ such that $C_i = \sum_{j \leq i} I_j$. It is clear that C_i is also a monotonically increasing sequence. We then apply a user-supplied threshold $0 \leq T \leq 1$ by selecting the pixels corresponding to the indices i such that $T \leq C_i/C_{N^2}$. This step converts the black hole image from a three-dimensional data cloud in (x, y, I) into a series of two-dimensional data clouds in (x, y) for every T . Compared to thresholding based on the peak intensity, this method can be seen as an integration over

pixel values weighted by their appearance frequency, and is much more robust against the noise in the image.

- (2) *Skeleton*. We then apply a standard topological skeleton algorithm to the thresholded image at its native resolution. A large area of flagged pixels will be thinned in this step, and mid-sized “holes” in the image will be enlarged. Without this step, “holes” smaller than the max pooling resolution (see the next step) will be removed at the end.
- (3) *Max pooling*. Next, we downsample the image to, e.g., 16×16 pixels using max pooling. That is, if any of the pixels in the rebinned pixel is flagged, we flag the rebinned pixel.
- (4) *Reskeleton*. We then apply the skeleton algorithm again to the downsampled image. This step removes unnecessary corner pixels. The result is our metronized data points.
- (5) *Visualization*. Finally, we visually connect the flagged pixels like in a metro map. This step is for visualization only and is irrelevant to the application of persistent homology.

IV. CLASSIFYING SYNTHETIC IMAGES

We apply TDA to mock data by computing the persistent homology of the synthetic images. Our goal is to classify images that show emission ring features. We present two methods to convert the original image to a form suitable for TDA: the *grid-scan* method and the *bisection* method.

A. Grid-scan method

In the *grid-scan* method, we scan through a set of robust thresholding parameters of $T = (0.1, 0.2, \dots, 0.9)$, thus generating a grid in T space. Higher grid resolution will result in a finer classification scheme at the cost of more computational time. In our case, nine instances of b_1 barcode diagrams are generated per image. The persistence parameter is set to $L = 0.5$, but (as we will show in the next subsection) our result is not sensitive to changes in this parameter due to the large sizes of typical emission rings.

If a persistent hole is detected in any of the nine b_1 diagrams, we classify the image as having a hole. Whenever a hole is detected, the parameters l and T where the hole is detected is recorded. The parameter l gives the approximate size of the hole in the emission, while the parameter T encodes its brightness. The set (l, T) can then be used to further classify the images in terms of the size and brightness of their emission rings.

Figure 5 shows an example synthetic image being preprocessed by metronization with $T = 0.7$. At $T = 0.9$, the two bright blobs attached to the ring dominate, giving the image a zeroth Betti number of $b_0 = 2$ and a first Betti number of $b_1 = 0$. As T is lowered, the two blobs meld into a single connected component and a persistent hole can now be detected. Consequently, the Betti numbers change to $b_0 = 1$ and $b_1 = 1$ to reflect this qualitative change. As T is lowered further, the emission regions at the corners of the image start registering as separate connected components, and a corresponding increase in the zeroth Betti number b_0 is observed by the TDA algorithm. The changes in the number of connected components as T is varied can be used to detect qualitative features such as bright emission blobs or completely disconnected emission regions, as demonstrated in this example.

Further subclassifications can be made to characterize the images. For example, the T where a persistent hole is first detected (when scanning from high to low T) can be used to characterize the likelihood of the hole's detection, as a persistent hole detected at high T is less likely to be spontaneously created by noise. One can classify all images with holes detected at T greater than a cutoff value as images with "strongly detected" holes, and those with holes detected under the cutoff to be "weakly detected" holes. In Fig. 6, we show three images with "strongly detected" holes, which in this case means possessing a persistent hole that was detected at $T > 0.75$, as well as three "weakly detected" holes.

B. Bisection method

As exemplified by Fig. 5 in the previous subsection, changes in the Betti numbers carry elucidating information that can be used to classify the image. In the *bisection* method, we focus on locating the T where changes in the Betti number are detected.

Our bisection method takes inspiration from the bisection method from root finding, with a slight modification due to $b_i(T)$, the i th Betti number as a function of T , not being a monotonic function:

- (1) Start with a uniform grid in T space as in the grid method, each of size dT .
- (2) Assume that $b_i(T)$ is monotonic within a grid.
- (3) Suppose a Betti number changes between $T = T_{\text{high}}$ and $T = T_{\text{low}}$ (with $T_{\text{high}} - T_{\text{low}} < dT$).
- (4) Bisect the interval between T_{high} and T_{low} in half, and choose the interval where the Betti number changes for the next step.
- (5) Repeat the first step with the new interval.

The free parameter dT controls how fine the classification scheme is: the algorithm will not detect changes in $b_i(T)$ over intervals smaller than dT . Running the bisection method with $dT = 0.01$ allows us to determine borders in parameter space where different numbers of holes are detected in the image. When plotted in L - T space (see Fig. 7), this gives a "phase diagram," indicating regions where different numbers of holes are detected.

As seen in Fig. 7, for a large swath of parameter space, one hole is detected. When the brightness threshold is too large (above 0.8), dim pixels at the center of the emission ring start to contaminate the metronization result, filling the hole and thus resulting in zero holes being detected. As L increases, we lose sensitivities to holes with radii smaller than $\sim 1.73L$. Thus, when L is too large, the emission ring is no longer detected. In the low- L and $-T$ regime, noise dominates, resulting in multiple detected rings in the lower left region of Fig. 7.

C. Other parameters

While this work focuses on the detection of the emission rings themselves, we note that parameters other than L and T can be used to further classify black hole images or input *a priori* information into the classification system. Important examples are the characteristic lengths where the barcode diagram begins (l_{birth}) and ends (l_{death}), the analog of the lifespan in T space, as well as the zeroth Betti number, b_0 .

The parameter l_{birth} determines the largest gap in a possible ring that could be tolerated (see Fig. 2). Any persistent holes that are detected at large l_{birth} will have a large gap (or gaps) in its surrounding emission ring. If l_{birth} is set to be very high, crescents will be detected as holes in the b_1 diagram. A strategy to classify crescents in the images is therefore to find images with $b_1 = 1$ (a hole is detected) and $b_0 = 1$ (a single piece is present), but with large l_{birth} 's. Analogously, one can demand a maximum




FIG. 6. Examples of images classified as strong rings (top row) and weak rings (bottom row). From left to right, the strongly detected images possess persistent holes first detected at thresholds T of 0.82, 0.77, and 0.84, respectively. Similarly, the weakly detected images possess holes first detected at thresholds of 0.53, 0.52, and 0.64 respectively.

l_{birth} for an image to be counted as possessing a persistent hole in addition to the requirement that $\Delta l > L$ in order to remove crescent images from the population.

As mentioned in the previous sections, the parameter l_{death} contains information about the size of the detected hole. A data cloud with holes that are small compared to

l_{death} (i.e., with radii less than $\approx 1.73l_{\text{death}}$ for circular holes) will have the holes filled in when converted into simplicial complex. Thus, the parameter l_{death} can be used to classify holes based on their sizes. Setting a minimum l_{death} filters holes that are too small, while setting a maximum l_{death} filters holes that are too big. If there is *a priori* information on the possible sizes of the hole in the image (e.g., from an independent measurement of the black hole mass), then both a minimum and maximum l_{death} can be used to cull images possessing holes that are inconsistent with the *a priori* possible hole size.

A third parameter would be the analog of the lifespan in T space. For an image like that in Fig. 5, there will be no hole detected at high brightness thresholds, as the bright spots dominate the image. As T is lowered, the emission ring becomes dominant, and a hole is detected. However, as T is lowered further, the brightness depression at the center of the emission ring becomes incorporated into the simplicial complex, filling the hole. There is thus a threshold $T = T_{\text{max}}$ where the hole first is detected and another threshold T_{min} where the hole is no longer detected, and $\Delta T = T_{\text{max}} - T_{\text{min}}$. For emission rings, ΔT is directly related to the brightness contrast between the emission ring and the brightness depression of its center. Images with emission rings with a large contrast will possess high ΔT .




FIG. 7. “Phase diagram” of the number of holes in L - T space. Regions bordered in blue contain one hole, and regions filled in red contain four holes. Everywhere else, zero holes are detected.

Further, the zeroth Betti number b_0 , which counts the number of connected components in the image, can also be used as a proxy of how noisy the image is, as a noisy image will possess a large number of connected pieces that are disconnected from each other. The images can then be subclassified as either “noisy” or “clean” images based on the number of, for example, disconnected components larger than a certain size, detected at a certain threshold T . The first Betti number b_1 can be also used for this purpose: noisy images will have more spurious holes detected at low L and T than clean images.

V. CONCLUSION

In the near future, both image reconstructions from VLBI data sets (such as that obtained by the EHT) and synthetic observations generated from theoretical models will be poised to generate a very large library of black hole images. In this article, we proposed that persistent homology aided with metronization preprocessing is a useful tool to automatically classify these images based on their topological characteristics. This method is both computationally efficient and flexible, in the sense that it possesses free parameters that could be tweaked based on the needs of the problem under consideration.

In this article, we focused on the detection of emission rings based on the first persistent Betti number b_1 , and demonstrated how TDA can be used to not only separate images with and without rings, but also subclassify ringed images into “weak” or “strong” rings based on their detection threshold. We also discussed ways to further

leverage TDA to classify black hole images based on the sizes and brightness contrast of their emission rings, how crescent-like the image is, and the noise level of the image.

Black hole images can contain astrophysically significant features aside from rings which will also imprint topological information on the image, and our method has the potential to classify images with these features. Images with hot spots will have $b_0 > 1$ at high T 's, with the number of distinct connected components giving the number of hot spots and the analog of l_{death} for the b_0 barcode diagram giving a measure for the distance between hot spots in the image plane. Simulations of black hole accretion flows have also been shown to display a bright emission ring surrounding an “inner shadow” [20], giving a hole-within-a-ring structure to the image. Future high-resolution VLBI observations have the potential to observe this structure. Such a structure will generate images with two rings that can be detected with TDA, possibly at different brightness thresholds, and with different contrasts.

ACKNOWLEDGMENTS

The authors thank Shiro Ikeda for serving as the EHT Collaboration internal referee for this paper. C.-K. C. acknowledges the support of NSF PIRE Grant 1743747. I.N. acknowledges funding from the South African Research Chairs Initiative (Grant ID 77948) through the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation (NRF), an agency of the Department of Science and Innovation (DSI) of South Africa.

-
- [1] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L1 (2019).
 - [2] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L2 (2019).
 - [3] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L3 (2019).
 - [4] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L4 (2019).
 - [5] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L5 (2019).
 - [6] K. Akiyama, A. Alberdi *et al.* (Event Horizon Telescope Collaboration), *Astrophys. J. Lett.* **875**, L6 (2019).
 - [7] K. Akiyama, S. Ikeda, M. Pleau *et al.*, *Astron. J.* **153**, 159 (2017).
 - [8] K. L. Bouman, M. D. Johnson, D. Zoran *et al.*, in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Washington, DC, 2016), p. 913.
 - [9] A. A. Chael, M. D. Johnson, R. Narayan, S. S. Doeleman, J. F. C. Wardle, and K. L. Bouman, *Astrophys. J.* **829**, 11 (2016).
 - [10] M. Honma, K. Akiyama, M. Uemura, and S. Ikeda, *Publ. Astron. Soc. Jpn.* **66**, 95 (2014).
 - [11] G. Carlsson, *Bull. Am. Math. Soc.* **46**, 255 (2009).
 - [12] H. Edelsbrunner and J. L. Harer, *Computational Topology: An Introduction* (American Mathematical Society, Providence, Rhode Island, 2010).
 - [13] L. Medeiros, T. R. Lauer, D. Psaltis, and F. Özel, *Astrophys. J.* **864**, 7 (2018).
 - [14] L. Medeiros, D. Psaltis, and F. Özel, *Astrophys. J.* **896**, 7 (2020).
 - [15] D. Morozov, Dionysus 2, library for computing persistent homology, 2017, mrzv.org/software/dionysus2.
 - [16] S. Zhang, M. Xiao, and H. Wang, [arXiv:2003.07989](https://arxiv.org/abs/2003.07989).
 - [17] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, and M. Michaux, *IEEE Trans. Visual. Comput. Graph.* **24**, 832 (2018).
 - [18] J. M. Lee, *Introduction to Topological Manifolds* (Springer, New York, 2011).
 - [19] L. Vietoris, *Math. Ann.* **97**, 454 (1927).
 - [20] A. Chael, M. D. Johnson, and A. Lupasca, *Astrophys. J.* **918**, 6 (2021).