Citation: Brito-Pacheco, D., Giannopoulos, P. & Reyes-Aldasoro, C. C. (2026). Persistent Homology in Medical Image Processing: A Literature Review. ACM Transactions on Computing for Healthcare, doi: 10.1145/3820762

Persistent Homology in Medical Image Processing: A Literature Review

DANIEL BRITO-PACHECO, City St George's, University of London, United Kingdom
PANOS GIANNOPOULOS, City St George's, University of London, United Kingdom
CONSTANTINO CARLOS REYES-ALDASORO, City St George's, University of London, United Kingdom and The Institute of Cancer Research, United Kingdom

Medical image analysis has experienced significant advances with the integration of machine learning, deep learning, and other mathematical and computational methodologies into the pipelines of data analysis. One methodology that has received less attention is Persistent Homology (PH), which comes from the growing field of Topological Data Analysis and has the ability to extract features from data at different scales and build multi-scale summaries. In this work, we present a systematic review of PH applied in medical images. To illustrate the potential of PH, we introduce the main concepts of PH and demonstrate with an example of histopathology.

Thirty articles where PH was applied to medical image analysis tasks such as segmentation and classification were selected and reviewed. It was observed that PH is very versatile, as it can be applied in many different contexts and to different data types, whilst also showing great potential in increasing model accuracy in both classification and segmentation. It was also observed that image segmentation predominantly uses basic level-set filtration to calculate PH, while classification takes various approaches using filtration on more complex structures built from data. This review highlights PH as an important tool that can further advance medical image analysis.

CCS Concepts: • **Mathematics of computing → Algebraic topology**; • **Applied computing → Health informatics**; • **General and reference → Surveys and overviews**; • **Theory of computation** → *Design and analysis of algorithms*; • **Computing methodologies** → *Computer vision.*

Additional Key Words and Phrases: Persistent Homology, Biomedical Imaging, Medical Imaging, Image Processing, Image Analysis

1 Introduction

Medical image processing has seen steady growth since the 1970s when algorithmic approaches applied to analyse medical images started to show promising results [1–3]. Today, thousands of algorithms used to analyse medical images are published every year to much acclaim in areas such as histology [4] and radiology [5, 6]. Many of the algorithms described in these papers exploit techniques like Machine Learning (ML) [7], Mathematical Morphology [8], Fourier Analysis [9], and, more recently, Deep Learning (DL) [10–19], which is dominating the literature [20], for a variety of tasks such as nuclear detection, segmentation, classification and counting [21], predicting survival in cancer [22], out-of-distribution detection and localization [23].

Authors' Contact Information: Daniel Brito-Pacheco, City St George's, University of London, School of Science and Technology, London, United Kingdom, daniel.brito@citystgeorges.ac.uk; Panos Giannopoulos, City St George's, University of London, School of Science and Technology, London, United Kingdom, panos.giannopoulos@citystgeorges.ac.uk; Constantino Carlos Reyes-Aldasoro, City St George's, University of London, School of Science and Technology, London, United Kingdom and The Institute of Cancer Research, Integrated Pathology Unit, Division of Molecular Pathology, Sutton, United Kingdom, reyes@citystgeorges.ac.uk.

ACM 2637-8051/2026/6-ART
https://doi.org/10.1145/3820762

A series of mathematical techniques that have received less attention are those related to Persistent Homology (PH) [24–26]. PH considers data as a multi-scale representation of topological features and enables the analysis of higher dimensional representations where the data may have a predetermined *shape* [27]. PH allows for the identification of this shape from the data and has already been successfully used in areas as diverse as quantification of airway structures [28], analysis of functional brain activity [29], analysis of glass structure [30], the study of mitochondrial network connectivity [31], the segmentation of cell mitochondria [32], and the detection and segmentation of cell nuclei [33].

As such, PH has the potential to be a powerful instrument in the study of the immune system, especially when combined with other tools into image analysis algorithms and pipelines and may be particularly useful with higher dimensional data. In this review, a comparison and examination of the existing literature on PH in medical image processing is performed and the use of this technique in the analysis of medical datasets is explored.

2 Background

2.1 Persistent Homology

The Merriam-Webster dictionary presents four definitions of Homology: (1) similarity due to a common origin, (2) similarity of traits reflecting common descent and ancestry, (3) similarity of nucleotide or amino acid sequence, and (4) a branch of the theory of topology concerned with partitioning space into geometric components. Persistent Homology is closest to the fourth definition, which is the focus of this work, but the first two definitions are also relevant, as homology implies similarity of mathematical properties. It is interesting to notice that outside the topological context, there is discrepancy in the understanding of homology. Marabotti and Facchiano [34] insist that homology should only refer to *having a common evolutionary origin* and not as similarity that can be quantified, but only around 57% of articles published in 2007 used this and the other 43% used homology as a measure of similarity.

Persistent Homology is one of the main tools in the mathematical field of Topological Data Analysis (TDA), a field which uses concepts from the mathematical branch of Topology to create statistical models for analyzing data [35–37]. The essence of topology is to consider only some aspects of objects and completely disregard others. This was the approach taken by Euler to solve the problem of the seven bridges of Königsberg [38], which required to walk over each of the seven bridges of the city of Königsberg in Prussia exactly once and returning to the starting point. To solve this problem, the length, width or material of the bridges are irrelevant. The only important information is which islands of the city are connected by which bridges. The property of connectivity does not change if the bridges are deformed or change shape as long as they connect the same points in a map. Following this logic, a sphere can be deformed and maintain the same topological characteristics as an ellipsoid, a doughnut can be deformed and maintain the same characteristics as a cup with one handle or indeed a section of a blood vessel or the intestines. The key difference between the sphere and the doughnut relies on the continuity of the surface, which may have a number of *holes* or form *loops* that surround regions that are not part of the surface. It should be noted that both holes [39] and loops [40] are used in the literature. In the rest of this paper, we will use loops. A cup with two handles is different to the doughnut as it contains two loops, but it is similar or homologous to a double torus or a cross section of a heart as illustrated in Fig. 1.

2.2 Simplicial Complexes and Persistence

A *simplicial complex* is a mathematical structure that can be used to represent complex objects as topological spaces in a simplified manner. The complex is formed of points, edges, triangles, tetrahedra, etc., all of which are collectively known as *simplices*. Persistent Homology tracks the topological changes that occur to a topological space as it deforms according to a parameter, for instance, by slicing a 3D object *e.g.*, a plant (Fig. 2 (a)), or a bronchial tree (Fig. 2 (b)) at different levels. The parameter would be the height where the object is sliced and the deformation would follow to what is being sliced: roots, stem, branches and leaves, etc. (or larynx, trachea, bronchi, etc.). In the case of a greyscale image, an intensity threshold on the brightness of pixels could be lowered or raised progressively, creating a filtration from the pixels that are brighter than the threshold as illustrated by Fig. 2 and Fig. 3. As the threshold moves, regions above the threshold (Fig. 3(a)) become the components (Fig. 3(b)) and loops (Fig. 3(c)), which appear, or are *born* and represented by a bar in the Figure. The simplification of the components becomes the simplicial complex (Fig. 3(d)). When two components merge into a single component, one of them disappears, or *dies*. Typically, the oldest component survives (the one that was born earlier in the filtration). When a loop is filled in, it disappears and is also said to die at that time point. The birth and death of components and loops is encoded in a scatter plot known as the *Persistence Diagram* (Fig. 3(e)). The horizontal coordinates of circles (components) or triangles (loops) correspond to the time point of birth, and the vertical coordinate to the time point of death. The case may arise where two components (or loops) are born and die at the same times, thus leading to two different points having the same coordinates in the persistence diagram; there exist other visualisations, such as persistence barcodes, that do not have this problem [41]. The persistence diagram can be understood as a summary of the topological changes that occur through the sequence of deformations. It is a topological signature and can be used for comparing two shapes using well-known distance metrics [42, 43]. Finally, the number of components and loops in a complex are summarised metrics called *Betti Numbers*. The 0-th Betti Number (β_0) of a simplicial complex counts the amount of components, while the 1-st Betti Number (β_1) counts the amount of loops.

It is important to note that these definitions are informal and without mathematical notation. For a more formal introduction, the reader is referred to [24].

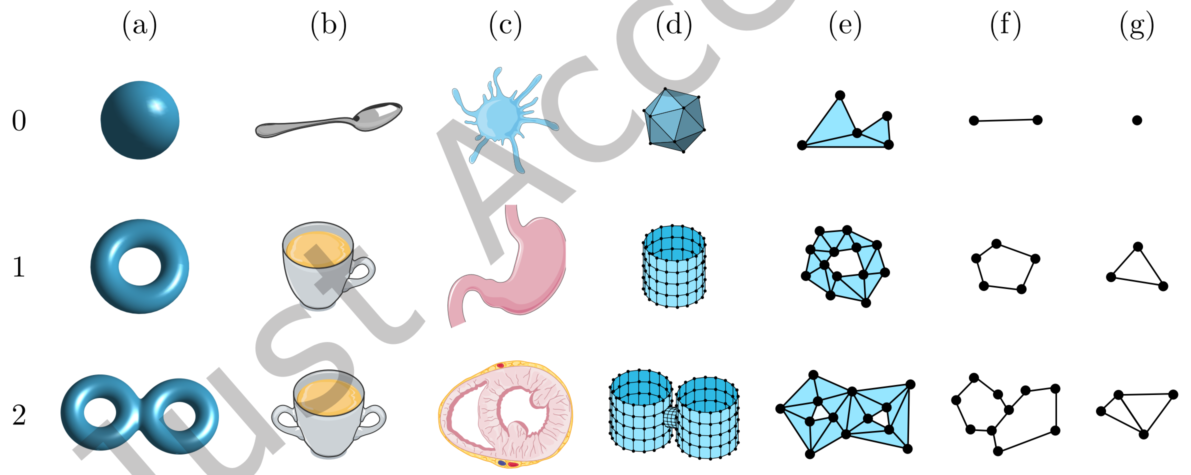

Fig. 1. Different objects grouped by the number of *loops* that they have. Row 0, no loops: (a) solid sphere, (b) spoon, (c) cartoon of a dendritic cell, (d) dodecahedron, (e) a union of triangles with no loops, (f) two points joined by an edge, (g) single point. Row 1, one loop: (a) solid 1-torus, (b) cup with a single handle, (c) cartoon of a stomach, (d) cylinder, (e) union of triangles surrounding a region of 2D space, (f) points joined by a cycle of 5 edges, (g) triangle not filled in. Row 2, two loops: (a) a solid double torus, (b) cup with two handles, (c) cartoon of a heart cross-section, (d) two cylinders joined together, (e) union of triangles surrounding two different regions, (f) points joined by two cycles of edges, (g) two triangles not filled in, joined by an edge.

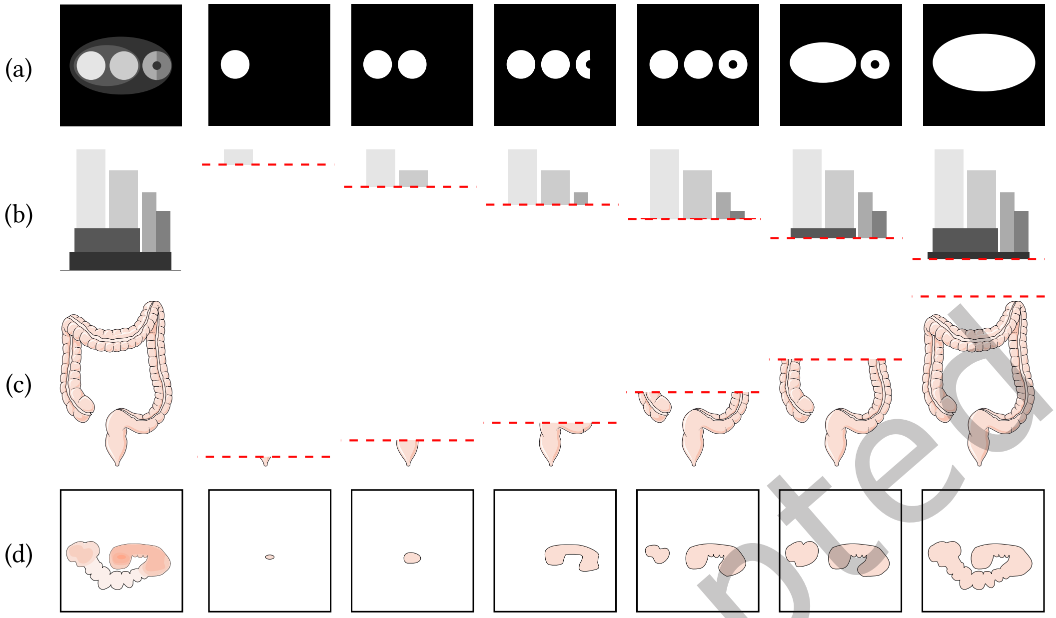

Fig. 2. Two independent examples of filtration by thresholding. (a) A varying intensity threshold is applied to a 2D greyscale image. As the threshold decreases six levels, regions with intensities above the threshold are coloured white and those pixels below the threshold are coloured black. (b) Illustration of the same intensities as in (a) shown from a sideway point of view. The threshold is indicated by a red dashed line that decreases its value and intersects the corresponding regions. (c) Illustration of a filtering with a threshold based on height over a three-dimensional structure, i.e., a human colon and rectum. In this case, the threshold (red dashed line) moves upwards and intersects at different levels of the colo-rectal structure. (d) Regions below the threshold are labelled in pink, whilst those regions above the threshold are labelled in white.

2.3 Level-Set Filtration

One of the most common techniques to calculate PH is through *level-set filtration*, which was briefly described in subsection 2.2. Consider a greyscale image of dimensions 8 × 8 where each pixel can have values between 0 and 255 (Fig. 4(a)). An intensity level threshold is used to segment the image into regions, starting at 255 and decreasing to 0 and as the threshold is lowered (Fig. 4(c)), new regions are born and then die as they merge (Fig. 4(d)). In a merger, the older component absorbs the one that dies. Next, the simplicial complexes are overlaid on the binary images (Fig. 4(e)). A vertex (0-simplex) is placed at each white pixel, neighbouring white pixels (including diagonal neighbours) are joined through an edge (1-simplex), and a triangle (2-simplex) is added when three edges create a cycle, which can be filled as in the first column of (Fig. 4(e)). Finally, the binary regions are removed to show only the simplicial complex of each step (Fig. 4(f)).

When the filtration process is finished, births and deaths are encoded into a Persistence Diagram (Fig. 4(b)). In this diagram, each point corresponds to a topological feature. The horizontal coordinate of the point corresponds to its time of birth while the vertical coordinate encodes the time of death.

It is important to highlight that there are other filtration strategies (like slicing a plant), however, level-set filtration is probably the most popular [44–46]. In fact, the way in which the filtration is built is where most applications of PH differ from one another and will be considered in the review below.

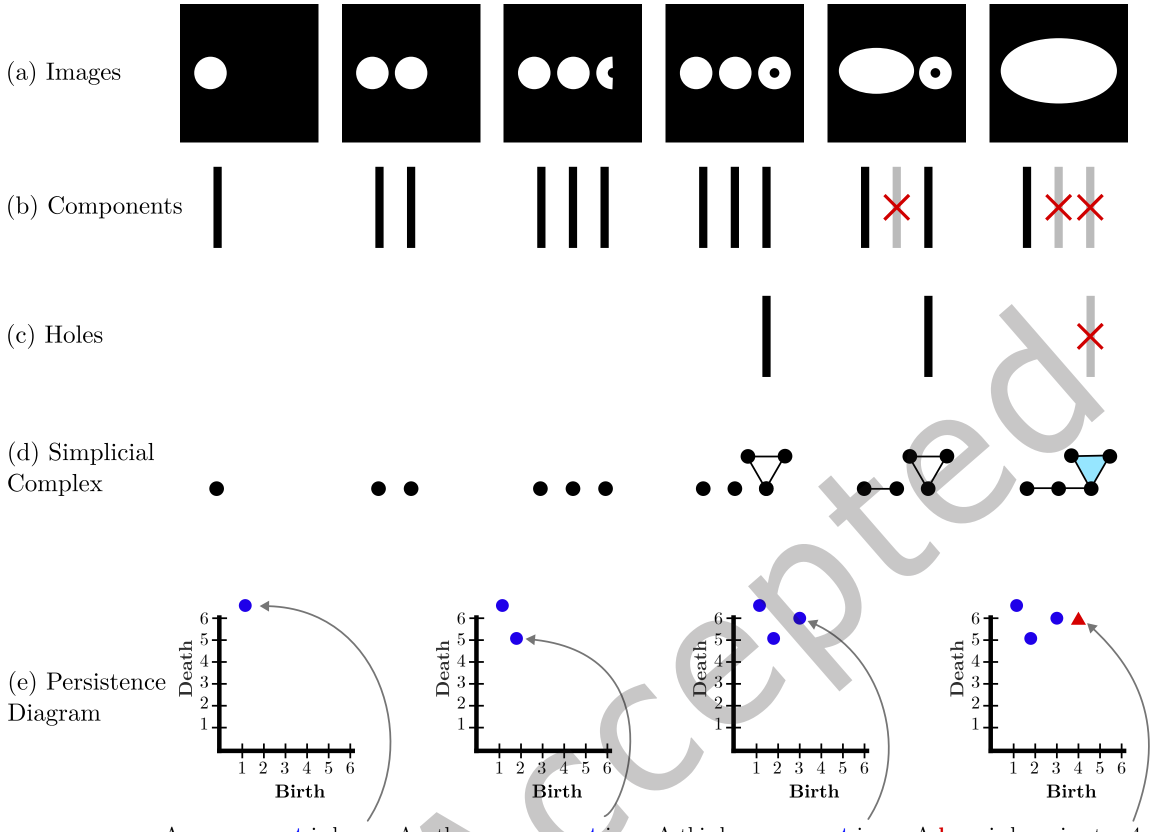

A **component** is born in step 1. This component lives past the last (6th) step.

Another **component** is born in step 2. This component dies in step 5.

A third **component** is born in step 3. This component dies in step 6.

A **loop** is born in step 4. It dies in step 6.

Fig. 3. Illustration of the main concepts of Persistence Homology. In this example, the birth and death of the loops and components are tracked within the six binary images created from the intensity threshold in Fig 2 (c) and noted in the persistence diagrams. (a) Six binary images with pixels that are above a threshold that is progressively decreased, from Fig 2 (c). (b) A black bar is added to illustrate the birth of a component and when the component dies, it is marked with a red cross and changed to grey. Following the elder rule, when two components merge, the eldest one survives, while the youngest dies. (c) The birth and death of loops is marked in the same way as for components. (d) Representation of the components as a simplicial complex. A point is added for each connected component, except if there is a loop, in which case a cycle of 3 points and 3 edges is added to represent said loop; the cycle is filled in when the loop dies. (e) Step by step formation of the persistence diagram. The first component that appears at step 1 will continue to live beyond the last step (6). The second component is born in step 2 and dies in step 5 when it merges with the first component. The oldest component always survives a merger. The third component appears in step 3 and dies in step 6. A loop is born in step 4 and dies in step 6.

It is now clear how level-set filtrations translate pixel-intensity structure into a sequence of topological summaries. The toy example in Fig. 4 uses a small greyscale image for clarity, but the same process extends naturally to real medical images. Before surveying the literature, it is instructive to demonstrate how PH behaves on real clinical data. The following subsection presents an example using histology images to show how level-set filtrations can extract topological features that meaningfully represent tissue architecture, serving as a typical and intuitive application of PH in medical imaging.

2.4 Persistent Homology in Histology: An Example

Persistent Homology can reveal meaningful differences in the architecture and texture of biological tissues, making it an effective feature-extraction tool for medical image analysis. To demonstrate this in a realistic setting, PH is applied to histology Whole Slide Images (WSI) of colorectal cancer (CRC) patients.

The dataset used for this example is NCT-CRC-HE-100K [47]. The set of images contains around 100,000 patches from H&E-stained human cancer tissue WSIs. All images in the set are 224 × 224 pixels, at a resolution of $5\mu m$ per pixel and normalized with the Macenko method [48]. Each patch belongs to one of nine possible tissue classes: adipose tissue (ADI), background (BACK), debris (DEB), lymphocytes (LYM), mucus (MUC), smooth muscle (MUS), normal colon mucosa (NORM), cancer-associated stroma (STR), and CRC epithelium (TUM). Fig. 5 shows five example images from each tissue class in the dataset.

Nine persistence diagrams, one for each type of tissue, are illustrated in Figure 6. Each patch is converted to greyscale and inverted (dark regions are made bright). Then, a 5 × 5 median filter is applied to the resulting image to reduce noise. Finally, the persistence diagram is computed by using the level-set filtration introduced in subsection 2.3, i.e., by applying a threshold that sweeps all the intensities of the greyscale patch to track the birth and death of components and loops. It is important to notice how the persistence diagrams capture the differences in textures and patterns visible in the different patches, in particular:

- Adipose tissue (ADI) forms bright bubbles (in light pink) with dark borders (in purple). In the greyscale inverted image, this is seen as dark regions surrounded by a light-grey / white membrane. This description is captured by the large group of blue circles concentrated towards the bottom of the persistence diagram: as the intensity threshold drops from 255, the components coalesce into a single component which persists for a long time (a blue circle with high persistence).
- The number of loops depends on the dark areas surrounded by brighter pixels: the amount of red triangles in the persistence diagram computed from the LYM patch is clearly greater than those calculated from the ADI patch.
- The greyscale version of the TUM patch has a few very dark spots (white in the colour version). This is captured by the persistence of red triangles (how high they are compared to the diagonal).
- The BACK patch has very little variation in pixel intensities, which is clearly identified in the tight clusters that both components and loops form in the persistence diagram.

The previous descriptions refer to characteristics of single patches. However, statistical features, such as the mean persistence of the components, can be calculated from each persistence diagram. In Figure 7, four statistical features are calculated from a sample of 500 patches of each class and plotted to illustrate clusters emerging from the topological features. The features that are calculated are: total number of components, mean persistence of components, standard deviation of the persistence of loops, median death of loops. Figure 7 (a) shows a scatterplot of the first two, while Figure 7 (b) shows a scatterplot of the third and last features. A representative patch from seven of the nine classes is shown connected to the point in the scatterplot corresponding to the features calculated from it. In this figure, it is noticeable that the cluster for the LYM class is very distinctive on both plots. This is likely due to the characteristic size and distribution of lymphocytes, which, in patches of consistent size, result in a relatively uniform tissue texture. Additionally, in Figure 7 (a) two distinct groups of the ADI class are distinguishable. The representative patches show that this may be due to two different subclasses of adipose tissue, one that forms with a darker histogram than the other.

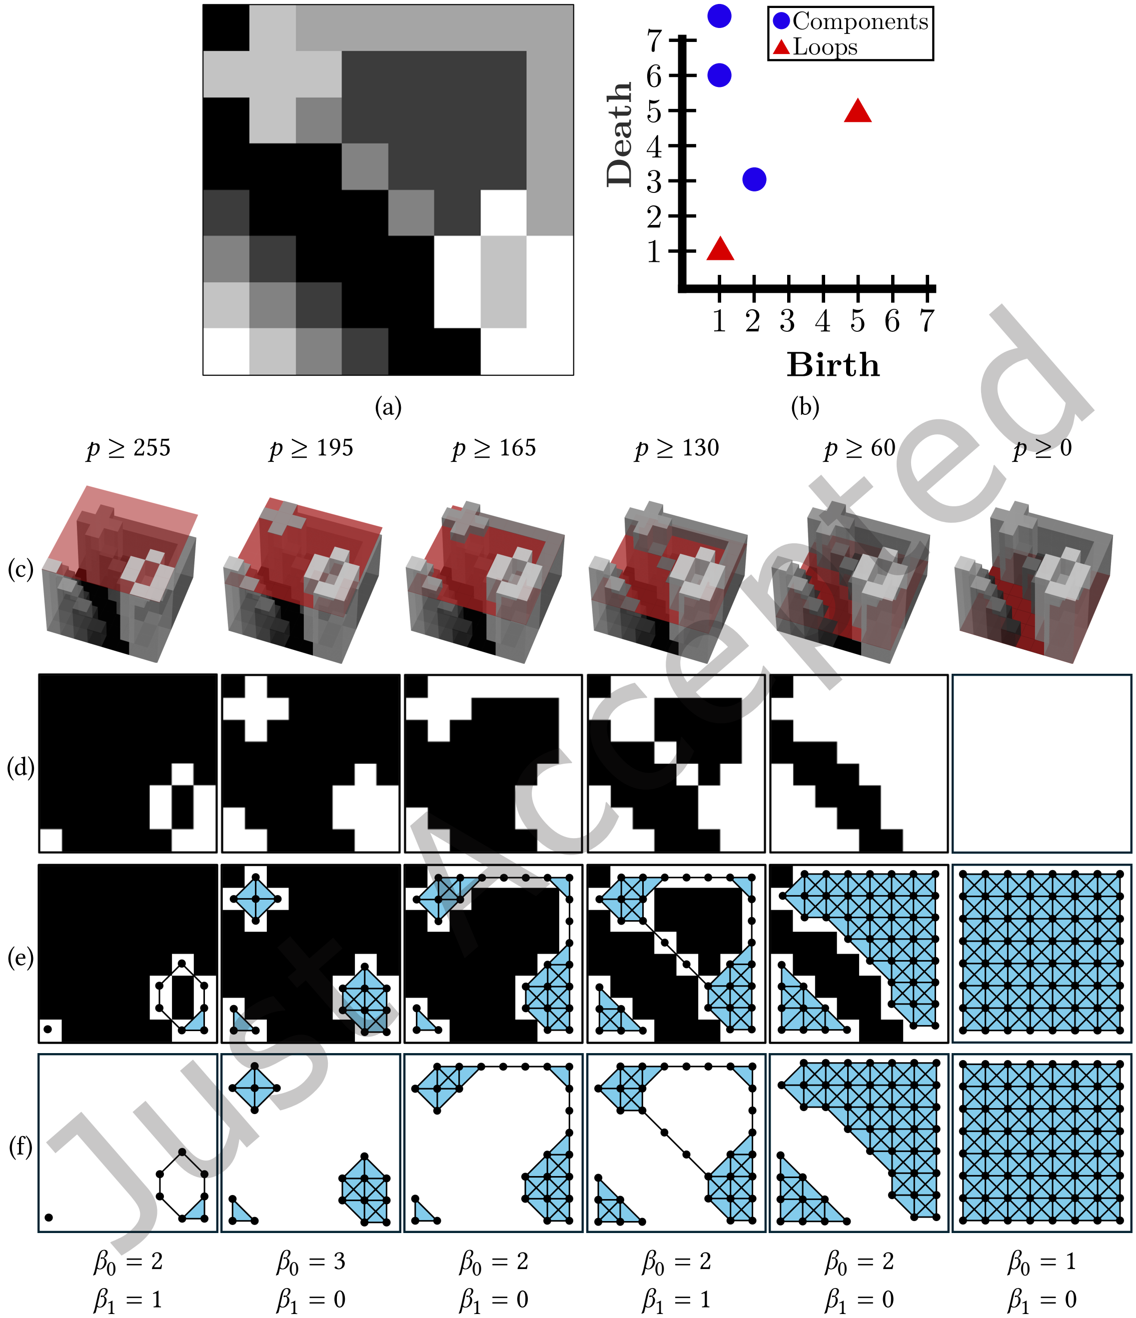

Fig. 4. Illustration of the filtration and PH calculation of a greyscale image with level-set thresholds. (a) Original greyscale image with pixel intensities in the range $[0, 255]$. (b) Persistence diagram. (c) 3D representation of the greyscale image with thresholds illustrated by a red transparent surface. (d) Binary images of pixels above the threshold. (e) Filtration overlaid on the binary images. A vertex (point) is placed at each white pixel. An edge is added between neighbouring pixels. A triangle is added when cycles of three edges are formed. (f) Filtration with the pixels removed. Betti numbers for each case are added at the bottom of the figure.

2.5 Stability of Persistent Homology

An important aspect of PH is its robustness to noise, in the sense that small deviations in the input data produce small deviations in the resulting persistence diagram. This is formally known as the *Stability Theorem*. The formal statement of this theorem is beyond the scope of this review, but this result guarantees that persistence diagrams behave predictably under realistic imaging variations. In the case of medical images, noise and artefacts from acquisition are common. However, the stability theorem ensures that the topological descriptors extracted by PH remain meaningful across samples. The reader is directed to work by Cohen-Steiner *et al.* [43] and generalisations by Chazal *et al.* for the theoretical details [49].

3 Materials and Methods

A systematic review to identify the use of PH in Medical Image Processing literature was performed following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) [50, 51] and the guidelines proposed in [52]. The steps of the review are illustrated in Fig. 8. The following research questions were formulated:

1. What filtration methodologies are applied to medical images to calculate Persistent Homology?
2. Can Persistent Homology improve performance of image processing algorithms?

Three sources were considered:

- Medline, queried through PubMed (https://pubmed.ncbi.nlm.nih.gov/) using the following search string: *("Topological Data Analysis" OR "Computational Topology" OR "Persistent Homology") AND ("microscop*" OR "medical imag*" OR "biomedical imag*") NOT Review*
- The proceedings for the annual conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) were accessed through Springer Link (https://link.springer.com/conference/miccai). Within each volume for each year, the search string *persisten* homolog** was used to query the papers in the proceedings.
- The proceedings for the annual conference on Neural Information Processing Systems (NIPS) were queried through Google and Google Scholar using the search string *"Topological Data Analysis" OR "persistent homology" OR "computational topology") AND ("microscopy" OR "microscope" OR "medical imaging" OR "medical image" OR "biomedical image" OR "biomedical imaging").*

Only primary studies published between 01/01/2012 and 31/08/2025 were considered. Additionally, to exclude entries, the abstract was carefully analysed. Articles were excluded if the answer to any of the following questions was no:

- Does the given paper focus its research on images (as opposed to other data types, such as point clouds)?
- Are the images of a medical or biomedical nature?
- Does the paper use Persistent Homology as a key step in the image processing algorithm?

Then, the full text of the article was read. This closer inspection revealed that some papers were not considered to address the research questions. In addition, papers that lacked details about how PH was used were discarded.

4 Results

The initial search results returned 140 entries. Some of these results (seven), were duplicates from PubMed and MICCAI, so they were removed. After the initial screening, 82 entries were excluded, and in the final selection step, another 21 entries were discarded according to the steps described in Section 3. This process returned 30 records to be included in the systematic review (Table 1).

Two main types of tasks using PH were identified, namely, image classification and image segmentation. Segmentation was performed in 12 studies (40%), classification in 11 studies (36.7%) , both segmentation and

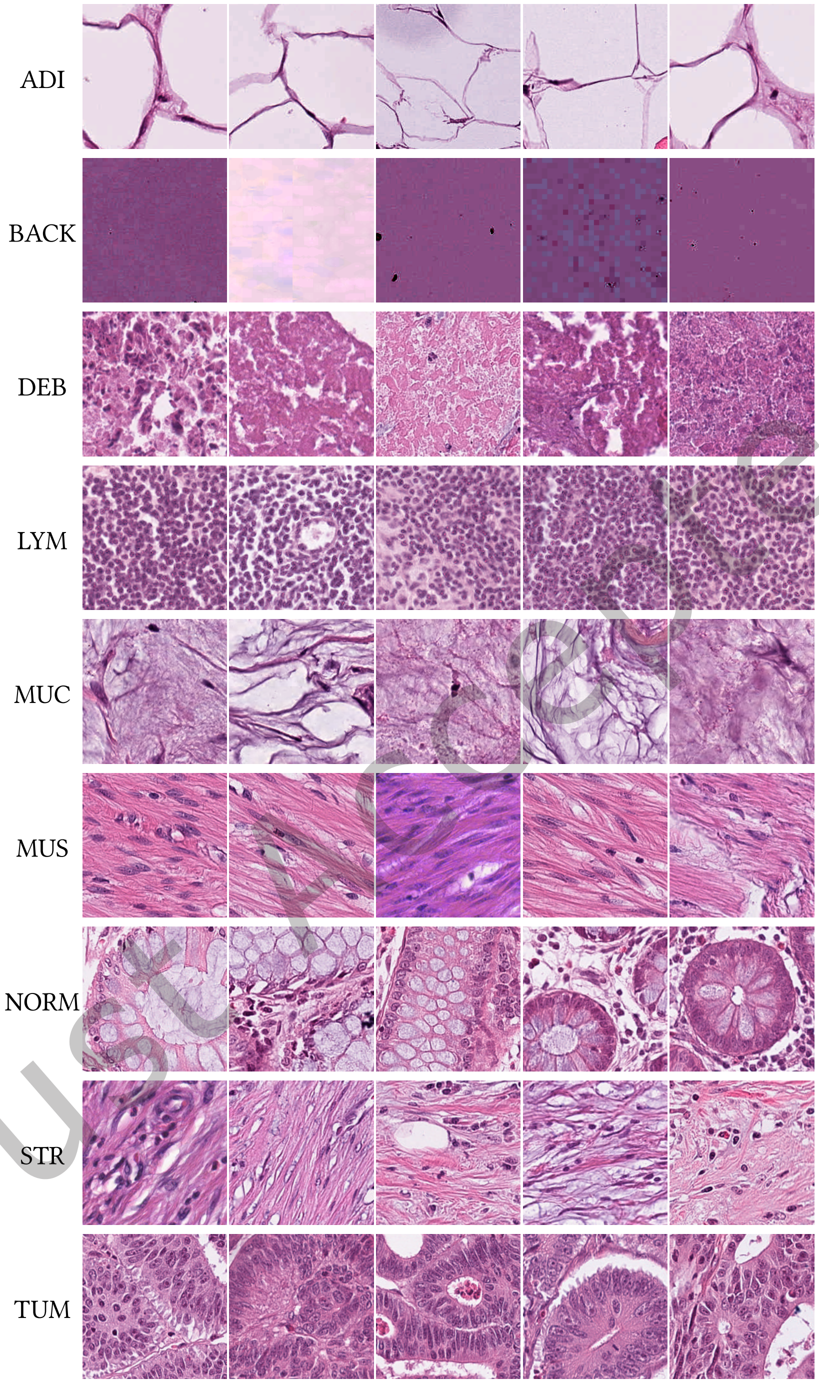

Fig. 5. Illustration of histological tissues of 9 classes. Each row contains five representative examples from each of the nine classes in the NCT-CRC-HE-100K dataset [47]. ADI: adipose tissue; BACK: background; CRC: colorectal cancer; DEB: debris; LYM: lymphocytes; MUC: mucus; MUS: smooth muscle; NORM: normal colon mucosa; STR: cancer-associated stroma; TUM: colorectal adenocarcinoma epithelium.

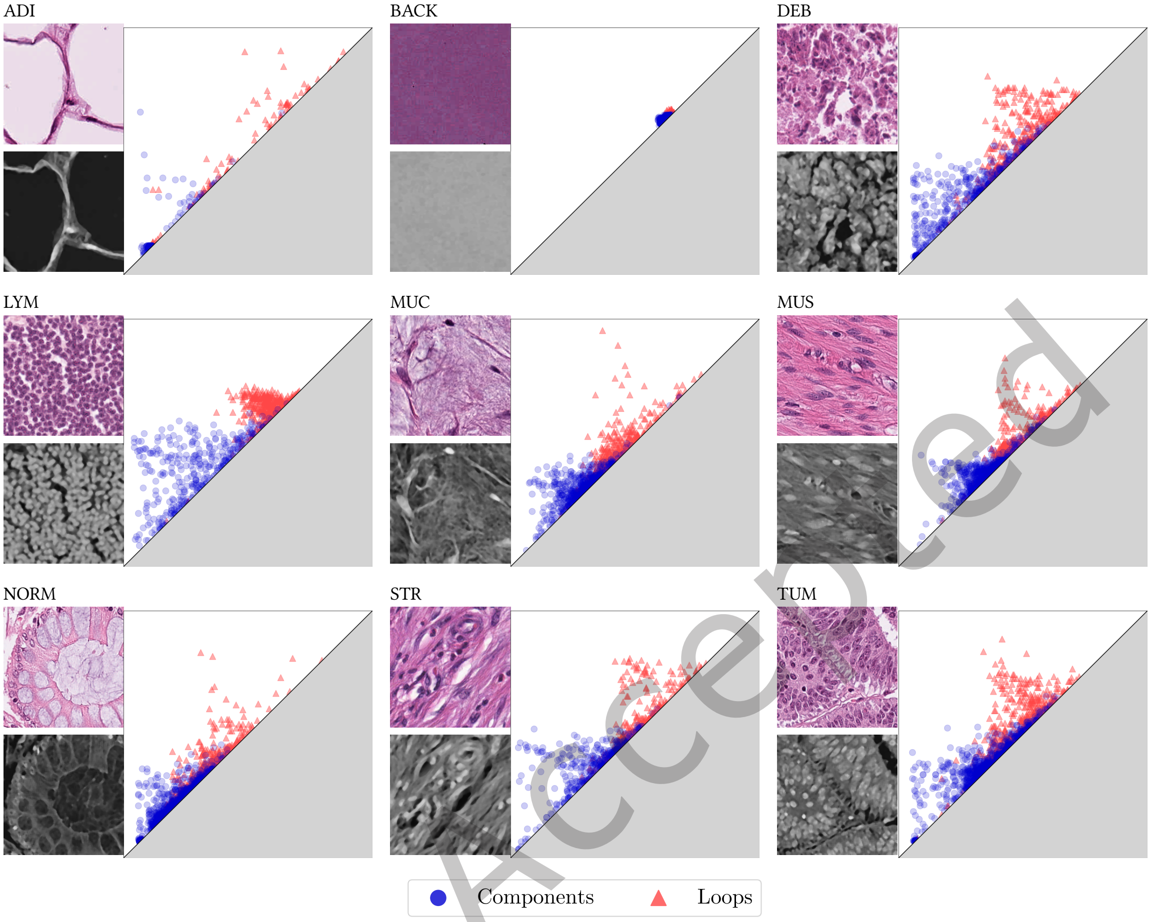

Fig. 6. Illustration of the persistence diagrams of the histological tissues presented in Figure 5. One representative patch from each of the classes (ADI, BACK, DEB, LYM, MUC, MUS, NORM, STR, TUM) from the NCT-CRC-HE-100K dataset [47] is converted to grayscale and inverted. Noise is removed on the grayscale image by applying a 5 × 5 median filter. A persistence diagram is calculated on the smoothed grayscale image where blue circles are components and red triangles are loops. The distribution of the scatterplots in the persistence diagram capture differences in the textures of different tissues; for instance, the persistence diagram of adipose tissue sees most components born and die early since bright areas in the grayscale join relatively quickly as the threshold lowers; this contrasts with Lymphocytes where the components take a longer time to join into one single component. ADI: adipose tissue; BACK: background; CRC: colorectal cancer; DEB: debris; LYM: lymphocytes; MUC: mucus; MUS: smooth muscle; NORM: normal colon mucosa; STR: cancer-associated stroma; TUM: colorectal adenocarcinoma epithelium.

classification in one study and six (20%) other papers performed tasks that cannot be considered segmentation or classification.

4.1 Segmentation by Level-Set Thresholding

Segmentation by level-set thresholding as previously described was used in 6 papers. This is by far the most common and straightforward way to make a filtration from a greyscale image. In brief, this process sequentially

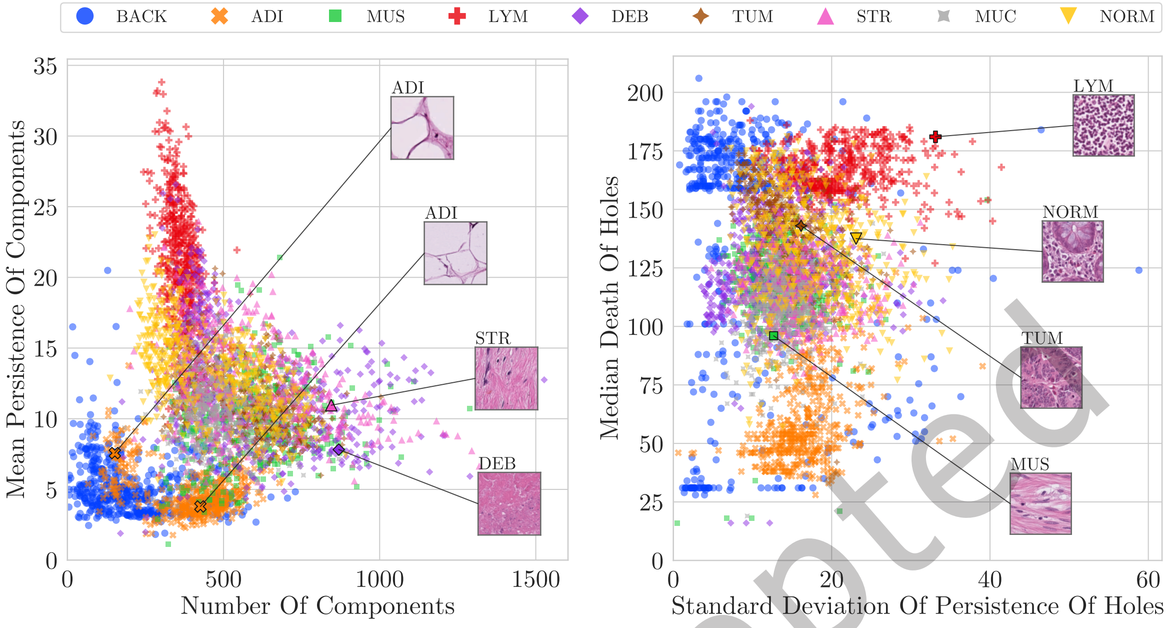

(a) (b)

Fig. 7. Illustration of metrics extracted from the persistence diagrams. Different statistical attributes were calculated from the persistence diagrams of 500 representative patches of each tissue. Namely, the total number of components, mean persistence of components, standard deviation of persistence of loops and median death of loops were calculated. (a) Number of components against mean persistence of components from each of the nine classes of tissue. (b) Standard deviation of the persistence of loops against median death of loops. In both cases it is possible to observe how PH features such as these could be used to distinguish between tissue types, as, for example, LYM patches (red crosses) show a tight cluster easily distinguishable in both plots. The background (blue circles) is concentrated in the lower left in (a), but spread over the vertical axis in (b), this may be due to variations on the intensity (dark/bright) that are visible in Figure 5. It can also be observed that ADI seems to be split in two regions in (a). ADI: adipose tissue; BACK: background; DEB: debris; LYM: lymphocytes; MUC: mucus; MUS: smooth muscle; NORM: normal colon mucosa; STR: cancer-associated stroma; TUM: colorectal adenocarcinoma epithelium.

creates a sequence of complexes that form a filtration by lowering (or raising) a threshold directly on the intensity of the pixels. In some cases, a filtration is used to train a CNN through a topological loss function [44, 55].

Level-set thresholding where *n* levels were used to create *n* binary images was used by Rojas-Moraleda [33]. The connected components of these images are then tracked as the threshold varies in order to calculate their persistence, this helps identify individual objects, like cells, which may be very close and appear as merged into one with traditional segmentation methods. Haft-Javaherian *et al.* applied a threshold directly on 3D volumes to build a filtration based on the connected components of the resulting binary volumes [44]. Using the resulting filtration, a topological loss-function encoding the target Betti Numbers was combined with a customised cross-entropy function from the DeepVess neural network architecture to train a neural network [78].

Qaiser *et al.* introduced a novel concept known as Persistent Homology Profiles (PHPs) used as features built by level-set thresholding patches of greyscale 2D images [46]. The PHPs encode the 0 and 1-dimensional Betti Numbers at each threshold, as well as their ratio, as a real-valued function of the threshold. These features

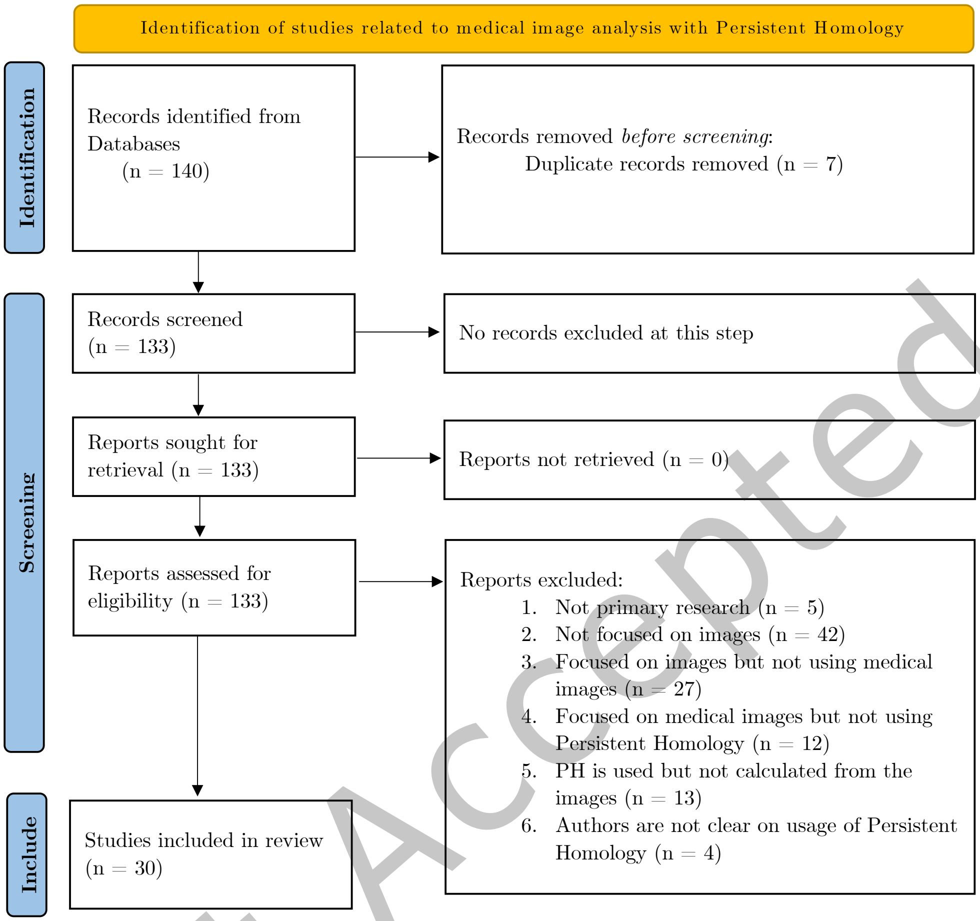

Fig. 8. The PRISMA statement flowchart shows details of the steps taken to select studies which were included in the literature review.

were then used to train a deep Convolutional Neural Network (CNN) to perform classification of tumour vs non-tumour patches. The final segmentation was given by the pixels in those patches that were classified as being part of a tumour. This description forms the base for the two algorithms built by Qaiser *et al.*, the first is known as the *Fast* algorithm while the *Accurate* algorithm added deep convolutional features into the feature extraction and segmentation process.

Banerjee *et al.* combined PH with Discrete Morse Theory (DMT) [79] [80] to perform segmentation on Bright-Field (BFI), Fluorescent Whole Slide (WSI), and Fluorescent Serial Two-Photon Tomography (STPT) images of the brain with tracers [53]. Initially, images were smoothed with a Gaussian filter. Then, a level-set filtration was built from the smoothed images. DMT was used to track the formation and merging of components and loops,

Paper	Task	Filtration	Type of Imaging Data	Dimension
[53]	Segmentation	Level-Set	BFI, STPT, WSI	2D
[44]			Microscopy	3D
[45]			Microscopy	2D
[54]			Microscopy	3D
[46]			WSI	2D
[33]			WSI	2D
[55]		Probability Map	CT scan	2D
[56]			Echocardiography	2D
[57]			MRI	2D, 3D
[58]			Many	2D
[59]			Microscopy	2D
[60]		SEDT	Angiography	3D
[61]	Classification	Level-Set	CT Scan	3D
[62]			MRI	3D
[63]			MRI	3D
[64]			Microscopy	2D
[65]		Height	CT Scan	3D
[66]		SEDT	Microscopy	2D
[67]		Delaunay	Microscopy	2D
[68]		Vietoris-Rips	Microscopy	2D
[69]		Vietoris-Rips	MRI	3D
[70]		Vietoris-Rips, Level-Set	CT Scan	2D
[71]		Multiple Filtration Functions	Ultrasound	2D
[72]	Both	Cech	Microscopy	2D
[73]	Other	Level-Set	Microscopy	2D
[74]		Level-Set	Microscopy	2D
[75]		Level-Set	MRI	3D
[76]		Multiple Filtration Functions	Microscopy	2D
[28]		Distance	CT Scan	3D
[77]		Probability Map	MRI	3D

Table 1. Full list of articles categorised by image processing task, filtration, acquisition technique: bright-field images (BFI), serial two-photon tomography (STPT), Whole slide imaging (WSI), Computed Tomography (CT), Magnetic resonance imaging (MRI).

which were used to calculate persistence. Less persistent features were discarded and the final segmentation was obtained by combining the topological features with a Deep CNN architecture known as ALBU [81].

Panconi *et al.* employed a human-in-the-loop approach for their segmentation model [45]. First, a level-set filtration was obtained from greyscale images to produce an initial segmentation. Since objects of interest tend to be brighter than the background (in this case) and merge with more persistent areas later in the filtration than noise does, the segmentation was simply made by thresholding for connected components that have a high birth value and a high persistence. These thresholds were manually raised or lowered to match the expected number of connected components in the greyscale image. This initial segmentation was used to isolate inner and outer sections of cells and structures with a variation of the Swinging Arm method [82]. The algorithm was tested on synthetic data and microscopy images of HEK293 cells and the R265 strain of *C. gattii* cells. This approach is later expanded upon in [54] to perform segmentations in 3D volumes and spatiotemporal images (video). The computation of PH is adapted to work directly on these domains as opposed to working with individual 2D slices, achieving high specificity and sensitivity scores.

4.2 Probability Map Filtrations and Topological Loss Functions for Segmentation

Another common way to do segmentation using topological information is to build on top of a first "guess" of the segmentation, called the probability map. This map usually comes from a CNN, which looks at the image and produces a map of probabilities where each pixel in the image is given a value that reflects how likely it is to belong to a certain class. The topological information is extracted via a filtration from this map and used to train a loss function to update the CNN's weights. A straightforward example of this is seen in Shin *et al.*'s work, where a filtration on the probability map is used to compute persistence and fine-tune a U-net to segment CT Scans of the small bowel[55].

Sun *et al.* segmented the myocardium in echocardiography images [56]. This article tackles two problems with existing CNN segmentation algorithms. First, a smooth boundary is sought, in a segmentation but this can lead to many inaccuracies. The second problem is that due to poor image quality, an automatically segmented myocardium can be broken up into two or more pieces, which is anatomically incorrect. Thus, the model used PH in two ways:

- PH was used to build a loss function that assigns weights given to certain pixels on the boundary of the myocardium to update the weights of a pre-existing segmentation model.
- PH was used to maintain the integrity of the myocardium. This means that the final segmentation must consist of a single connected component with no loops, while the background is expected to be one connected component with one loop.

In both cases, PH was calculated, separately, from the ground truth images and a 2D probability map. These 2D structures were first transformed into a 3D point cloud, where the x, y values of a point are given simply by its position on the image, while the height z of the point is given by its assigned pixel/probability value. On the point cloud, a filtration of Alpha Complexes is calculated [83]. This filtration does not require thresholding pixel intensities.

Also working on myocardium images, were Byrne *et al.*, who expanded these concepts to perform multi-class segmentation on Cardiovascular MRIs (CMR) [57]. U-Net is prone to making topological errors such as spurious connected components or broken loops in the myocardium. To solve this, a filtration is built from the class-specific probability maps, then the persistence of topological features is calculated. The loss function is designed to maximise the persistence of expected topological features and minimise the persistence of spurious ones and used as to fine-tune a pre-trained network. Experiments on 2D short-axis and 3D whole-heart CMR datasets show significant improvements in topological correctness with little effect on F_1/Dice score.

Similarly, Berger *et al.* expanded an already existing topological loss function to work with multi-class 2D segmentation [58]. For multi-class segmentation, a filtration was built from each class, based on the predicted probability map, and PH was used to extract topological features, which were weighted by their persistence. The Betti-Matching Loss works by matching predicted these weighted topological features to features of the binary ground truth segmentation and penalising discrepancies in matched features and the presence of unmatched features [84]. Summing across all classes yielded the topological multi-class loss, which was used to train a CNN to reduce topological errors in segmentation.

Yang *et al.* leveraged PH to make an attention-based Long Short-Term Memory Convolutional Neural Network (ConvLSTM) for segmentation [59, 85]. First, the ConvLSTM processed sequences of consecutive 2D slices, producing probability maps. A topological attention module then used PH to extract critical topological features from these maps, which were aggregated into query and key vectors that were used to find the similarity maps, guiding the ConvLSTM's attention using a topological loss function. On the CREMI dataset, the method improved Dice scores over competing CNNs, while on ISBI12/13, Dice scores were not better. However, across all datasets the model achieved the lowest Betti error, demonstrating superior topological fidelity compared to the baselines.

Molina-Abril and Frangi sought to improve segmentation of blood vessels through the use of PH [60]. Segmentations of blood vessels often yield an incorrect result with separate blood vessels that touch each other. The proposed method takes an initial segmentation, skeletonises it, and uses a measure known as the Kissingness Measure (KS) to identify critical points. This measure takes into account the average differences in radius and average angle formed between branches at intersection points. The centrelines with the higher KS form what the authors called the suspicious kissing set. To create a filtration, a spring force was applied to the initial segmentation towards the intersection point until there is a topological change around the point. The smaller the forces required to make this change, the higher the probability that the point will be selected as critical. An algorithm based on active contours is then used to remove the critical points [86].

In the article by Edwards *et al.* classification and segmentation were both performed [72]. PH was used to classify between cells that have been treated with different perturbations that affect the actin cytoskeleton. The treatments used were small-molecule inhibitors and genetic modifications. Initially, small circular patches were taken from the original image and high-intensity pixels were selected inside the patches, along with evenly spaced pixels on the boundary of the patch. The pixels form a point cloud and a filtration of Cech complexes was created by varying a distance parameter [37]. Persistence landscapes were calculated from the filtration and used to train Support Vector Regression (SVR) classifiers [87]. The SVR classifiers classified patches individually. The classification of an image was given by the average label assigned to its patches by the SVR classifiers. The segmentation was created by mapping values assigned to patches by the SVR classifiers back to pixels on the image. When visualised on top of the original image, with negative-valued pixels as one colour and positive-valued pixels as another, the different colours overlap with areas where treatment has affected the cells.

4.3 Performance of Segmentation Methods

Table 2 shows performance metrics reported by each author that dealt with a segmentation problem. All metrics were calculated at the pixel-level, unless stated otherwise (like the algorithms tested by [46]). It is important to note that in the case of [33], the authors did not explicitly state the F_1 score, but it was possible to calculate through the confusion matrix provided. Additionally, [33] and [56] tested the corresponding methodologies on two different datasets and obtained one score for each, these scores were averaged to obtain a single F_1 score for each algorithm.

From Table 2 it is seen that all algorithms except [56] and [58] reached an F_1 score of 0.85 or higher, which indicates a sufficiently good segmentation. In the case of [53], the segmentation was compared against other state-of-the-art methods (U-Net [88], SegNet [89], ALBU [81]), outperforming all of them. As with the previous article, the authors in [44] compared their method against state-of-the-art CNN segmentation methods ([78, 88, 90–92]) outperforming all in Dice and Jaccard index. The method in [55] was compared with a decoder network without any topological information incorporated into the training process: it was found that the loss function with topological prior improves the segmentation performance. The method in [33] was not directly compared to other segmentation models, however, the obtained F_1 score was above 0.9, which indicates an extremely accurate segmentation. The model presented in [56] outperformed the state-of-the-art methods it was compared to (clDice [93], TopoNet [94], and U-Net [88]). Berger *et al.*'s loss function was also compared to other loss functions, generally outperforming all of them [58].

Model	Evaluation Metric
[53]	F_1: 0.86
[55]	F_1: 0.85
[44]	F_1: 0.87
[57]	F_1: $0.93^{(*m)}$
[56]	F_1: 0.82 $^{(*)}$
[58]	F_1: 0.77 $^{(*)}$
[59]	F_1: 0.95 $^{(*)}$
[33]	F_1: 0.91 $^{(*)(**)}$
[45]	Sensitivity: 0.88, Specificity: 0.88 $^{(***)}$
[54] 3D	Sensitivity: 0.96, Specificity: 0.95 $^{(***)}$
[46] Fast	Patchwise F_1: 0.83
[46] Accurate	Patchwise F_1: 0.84

Table 2. Segmentation models together with the performance metrics reported by the authors. $^{(*m)}$: multi-class average. $^{(*)}$: averaged from testing on multiple different datasets. $^{(**)}$ calculated from confusion matrix. $^{(***)}$: tested on synthetic data.

For both of Panconi *et al.*'s articles, [45, 54], specificity and sensitivity were reported. From these metrics, it is impossible to calculate the F_1 score or the Jaccard index without knowing the proportion of pixels of the object to the pixels of the background.

Yang *et al.* compared did not significantly outperform others in F_1-score, but by measuring Betti Error, it was shown that the model makes less topological mistakes.

In Qaiser *et al.* [46], two different models were tested. In their methodology, a segmentation was produced by taking small patches from the complete image and classifying each patch into one of two classes. In this case, patch-level F_1 was reported.

It is important to note that in Edwards *et al.* [72] and Molina *et al.* [60] no segmentation performance metrics were reported. In [60] PH is used to detect errors in a pre-existing volumetric segmentation, therefore detection and correction rates are provided: 81.80% and 68.77%, respectively.

As with segmentation, the most common way to build a filtration is by Level-Set Thresholding. Of the articles where classification is done, 4 built a filtration through this technique. In classification tasks, with this filtration it is possible to find topological changes at differing brightness levels, which are then typically used to train ML models.

Rammal *et al.*'s goal was to estimate the Gleason Score from Spatial Inference Light Microscopy (SLIM) images of cancerous prostate glands [64]. Overlapping patches of the greyscale images were processed individually to obtain the following eight features in each: mean pixel intensity, standard deviation of pixel intensity, mean lifespan of components and loops, standard deviation of the lifespan of components and loops, the persistent entropy for components and loops. The latter six are all topological features. Details on persistent entropy can be found in [95]. The features were used to train five different ML methods for 5-class classification (one class per possible Gleason Score). A decision tree classifier was the highest performer.

Huang *et al.* and Wang *et al.* both leveraged PH to guide the attention of CNN classifiers in [61, 62], respectively. Differing in the implementation, and imaging modalities (esophageal CT scans vs breast parenchyma MRIs), both identified salient features through their persistence which were combined with the CNN's feature maps to guide the training to "look at" more topologically meaningful features in the image.

Du *et al.* distilled topological information to aid in training a CNN to perform classification on MRI breast images, predicting whether a breast cancer patient achieves pathological complete response (pCR) after neoadjuvant therapy [63]. The topological information was obtained through Betti Curves similar to [46]: a level-set filtration was built directly from the 3D MRI volumes, the corresponding Betti number was mapped to each step in the filtration was computed. These maps were distilled using DenseNet into probability distributions for pCR and non-pCR patients. A student CNN was trained with a soft cross-entropy loss to match those distributions.

4.4 The Vietoris-Rips Filtration for Classification

The Vietoris-Rips filtration, builds simplicial complexes through a distance parameter d in the range 0-∞. Specifically, it considers points in n-dimensional space for the initial step of the filtration, made up entirely of vertices. An edge is placed between two vertices only when d has grown to the equivalent distance between the vertices; a triangle is placed once three edges make a cycle; a tetrahedron is placed when three triangles make a cycle; etc [37]. When applied to image classification, the authors extract locations of important objects (like cells) from the image, which when viewed as a point cloud, a Vietoris-Rips filtration can reveal meaningful topological information about their distribution and used for classification.

Xin *et al.* use PH *recursively* to diagnose and provide prognosis of Multiple Sclerosis (MS) through lesions appearing in MRI scans of the human brain [69]. Lesions visible on brain MRI scans were converted to a point cloud in 3D space by matching the centroid of each region to a point. A Vietoris-Rips filtration was built from the point cloud and named the G-Net. Taking the G-Net, a Community-Level network (C-Net) was obtained by matching k-simplices in the G-Net to vertices in a graph G. Edges between vertices in G are drawn if the corresponding k-simplices in G-Net share a $(k-1)$-simplex. This yields a simpler representation of the local topology of the lesions. A filtration in C-Net naturally arises as simplices appear and fuse in the G-Net - the lifespan of a feature is given by the steps in the filtration where the members of a community (the vertices) do not change. Using the C-Net, statistical features from each community were obtained and used to train two different classifiers to distinguish between MS and healthy patients.

Two other articles used the Vietoris-Rips filtration. Wang *et al.* [68] made use of PH (persistent homology) and graph neural networks (GNN) to classify tumours in H& E-stained images or immunofluorescence microscopy images. The method was described in three steps: (1) cell locations were identified computationally and their phenotypes were determined by clustering or using prior biological knowledge; this process yielded point cloud data of the cells and their phenotypes. (2) Cell Clusters (CCs) were identified in the point clouds using DBSCAN and the Vietoris-Rips filtration was built, a topological descriptor coined as a Cell Community Forest (CCF) was built based on the hierarchical structures of CCs. (3) Appearance features, spatial microenvironment, and the CCF topology were then integrated within a GNN to classify pathology images. Iqbal *et al.* [70] classified lung CT scans into health and COVID-19 cases. To do this, the authors created a two-branch pipeline. The first captured structures associated with lung shape, such as components, loops, and loops by grouping centroids of connected components of points at height bins (using the pixel intensity as the height). The new point cloud of the centroids had less points and made PH computation less expensive. It was used to compute a Vietoris-Rips filtration and 0-, 1-, 2-dimensional persistence diagrams. The second branch identified clinically important features like ground-glass opacity through a Level-Set Filtration. Both branches produced persistence diagrams, which were approximated via a 2D probability distribution and the dimensionality reduced to classify through SVMs.

4.5 Other Types of Filtering for Classification

Belchi *et al.* studied 3D CT scans of patients' lungs with the objective to characterise lung scans of patients with Chronic obstructive pulmonary disease (COPD) versus healthy patients, and smokers versus non-smokers [65]. A centreline was fitted to the 3D scans and metrics from this line were extracted via PH. A filtration was made by sequentially thresholding the centreline at different heights, creating a different graph at each height. The number of times the centreline changes vertical direction and the number of branching points were tracked. These metrics together with the length of the bronchial tree were enough to distinguish Healthy Non-Smokers from both COPD groups; and healthy smokers from patients with mild COPD through Kolmogorov-Smirnov tests.

Pritchard *et al.* used PH to study the internal microstructure of bone samples in [66]. The authors wished to identify differences in male and female groups of mice both with and without osteocalcin-expressing cells in TPaF and SHG images. First, greyscale images were binarised using Otsu's method, and patches of the binary image were processed individually in the next steps [96]. A function called Signed Euclidean Distance Transform (SEDT) was applied to each pixel; it assigns the distance to the region of the opposite colour to each pixel, with black pixels assigned a negative sign. The transformed patch was thresholded sequentially to create a sequence of binary images. Similar to the level-set filtration, a vertex is added when a new region appears in the sequence of images and an edge is added between vertices if the corresponding pixels are neighbours (not including diagonals). The resulting persistence diagram was used to calculate statistical features. For each persistent statistic, hypothesis tests were done to estimate if two distinct groups give statistically different results. It was found that the statistics differ significantly for most permutations of the groups. Finally, a support vector machine (SVM) was trained on the statistical features to classify a brand-new patch [97].

Hartsock *et al.* study the organization of stem cell colonies under varying Doxycycline concentrations using 2D multichannel fluorescence microscopy [67]. Some cultures were stained using NANOG and GATA6 and others using NANOG and HA. Images were first segmented via histogram thresholding to locate cells and quantify fluorescence signals. Thresholds were chosen to determine cell types according to the signal expressions and for each type, a Delaunay filtration was used to compute persistence landscapes. These landscapes were used for statistical analysis and machine learning classification. The method revealed differences in spatial organization that was correlated with loss of pluripotency and enabled more accurate classification of treatment groups than cell counts alone.

Notably, only one article dealt with ultrasound imaging, this was the article produced by Wei *et al.* [71]. In it the authors built a novel pipeline for tumour classification (benign vs malignant). This pipeline, called WT-TDA, incorporated wavelet transforms and PH to extract the features for classification. The pipeline was tested on three datasets of cancer on different organs: breast, thyroid, and kidney. The datasets were previously annotated on the location and malignity of the tumours. WT-TDA works in the following steps: segment the tumour, apply a wavelet transform to the segmented image, make many filtrations from the transformed images, compute Betti numbers and persistence landscapes from the persistence diagrams of the filtration, input into an ML algorithm, and use SHAP analysis to select features. The filtrations that were built were: level-set filtration, radial filtration (distance to a fixed pixel), density filtration, dilation filtration (minimum distance to pixels valued at 1), erosion filtration (minimum distance to pixels valued at 0), SEDT. Tests were run using different wavelet transformations, different filtration techniques and different ML algorithms.

As mentioned earlier, in the article by Edwards *et al.* classification was also performed [72]. The details can be found in 4.1.

Algorithm	ROC-AUC	F_1
[72]	$0.97^{(*)}$	$0.94^{(*)}$
[66]	-	$0.78^{(*)}$
[64]	0.99	-
[61]	-	$0.76^{(*)}$
[62]	0.92	-
[63]	0.87	-
[70]	-	$0.99^{(*)}$
[68]	0.91	0.75
[71]	$0.87^{(*)}$	-
[69]	0.88	0.87

Table 3. Classification models and ROC-AUC, F_1, metrics reported by the authors. $^{(*)}$: averaged from testing on multiple different datasets.

4.6 Performance of Classification Methods

In image classification tasks, reported metrics are more standardised than in segmentation. Even though some authors may report other metrics as well, this section focuses solely on the ROC-AUC (Receiver Operating Characteristic - Area Under the Curve) and F_1 score of the predictions. Table 3 shows the results reported by the authors that dealt with an image classification task.

Notably, the scores reported for the model in Edwards *et al.* [72] are very close to 1. This indicates an outstanding performance in classification. It should be noted that the results were averaged from scores obtained on twelve different datasets. However, the lowest F_1 obtained for an individual dataset was 0.816, on which it also reached a very high ROC-AUC of 0.908. This was not compared against other models. Similarly, Iqbal *et al.* [70] obtained a very high F_1 score that was the average of experiments on two datasets.

An extremely high ROC-AUC of 0.99 was achieved in Rammal *et al.* [64] when the features were classified using a Decision Tree Classifier (DTC). The highest accuracy obtained by the DTC was 95%.

The method in [66] used an SVM to test whether topological features are enough to distinguish between control and test populations. Two different datasets were used (from two different image acquisition techniques). An F_1 score of 0.78 was obtained when averaged over both datasets. This is a comparatively low score, as is the one obtained by [61]. For the latter, it is due to a significant difference in precision (0.64) and Recall scores (0.92). The reported accuracy was extremely high (0.98) so these scores could be explained by a stark class imbalance, as the proposed method outperforms other existing methods. Wang *et al.*'s method outperformed DL models in tumour type classification (TCGA dataset) and pre-class classification (BRACS dataset) [68]. Likewise, the proposed methodology in [62] outperformed conventional CNNs and other state-of-the-art topology-based methods in accuracy and AUC.

The model presented in [69] was tested against other models based on PH and an ROC-AUC of 0.88 and an F_1 score of 0.87 were obtained, which means that this model outperformed all the others it was compared to. A similar ROC-AUC score was obtained by the model in [71], however, it was only compared against two traditional DL models.

Hartsock *et al.* [67] did not report AUC nor F_1-score. However accuracy was reported for an SVM that correctly classified 71.8% of patches in the pan-GATA6 group and 73.5% in the HA group, which were higher to just using cell counts for classification.

Notably, no performance metric was reported in [65]. However, Kolmogorov-Smirnov tests revealed that topological features were statistically significantly different between healthy smokers and moderate COPD patients, as well as between healthy non-smokers and moderate COPD patients.

4.7 Other Types of Analyses

Another class of articles did not perform any segmentation or classification of images but instead used PH to aid in other tasks. In this subsection those articles are presented.

Herron *et al.* aimed to find the structure and organization of phagocytic podosomes [73]. This was done by first culturing Macrophages and introducing Immunoglobulin G (IgG) to induce the formation of podosomes with well-defined geometry. PH was used to localise the podosome sites in images of the cultures. Images were acquired using Structured Illumintation Microscopy (SIM) and Interferometric Photoactivated Localization Microscopy (iPALM). To find the podosome sites, the image was first smoothed using a Gaussian filter and small Gaussian noise was also added so that each pixel became distinct from its neighbours. A level-set filtration was built and a persistence diagram was obtained from the filtration, then k-means clustering was used to distinguish between relevant features and noise. The steps described previously were done twice, the first application uses a small Gaussian kernel, and was used to calculate loops to locate podosomes; while the second application used a larger Gaussian kernel and was used to calculate 1-dimensional loops to locate phagocytic sites. After this, some extra refining steps were done to improve the analysis. Podosomes that were far away from a phagocytic site were ignored, phagocytic sites that were only associated to 2 podosomes or less were also ignored, then the centre-point of the site was determined by the barycentre of all the podosomes associated to it.

Berger *et al.* also employed a level-set filtration on microscopy images [74]. The authors spiked pig and human blood samples with argatroban (an anticoagulant). The samples were then either diluted or coagulation was induced. Viscoelasticity was tested and confocal microscopy was used to take images of the fibrin networks in the plasma. The pre-processed images were used to compute persistence diagrams using the level-filtration; thresholds were then applied to distinguish true features from noise. Wasserstein distances were computed between baseline and treated samples, as well as between each persistence diagram and the diagonal-only (empty) diagram, to quantify structural differences. Haussdorf dimensions and weighted silhouettes were also calculated [98, 99]. The analysis revealed that both dilution and thrombin inhibition (with argatroban) led to detectable changes in the topology of fibrin networks.

Time-varying data from fMRI images of adult and children patients, who were watching a movie, were studied by Rieck *et al.* in [75]. At each time step, the persistence diagram of the 3D volume was computed from the level-set filtration and three analyses performed: (1) the diagrams were reduced to scalar time series via summary statistics and used for age prediction; (2) the diagrams were vectorised into persistence images to analyse brain state trajectories; (3) differences at event boundaries of the movie were compared.

Kaji *et al.*, describe changes to the structure of airways in COPD patients' lungs through the use of PH, in CT scans [28]. Two metrics were created for this purpose. The first, radialH_0, was defined through a filtration on the airway's radius to the centreline. The filtration was built by sequentially including voxels based on a moving threshold on their distance to the centreline. A persistence diagram of dimension 0 was computed using this filtration, and relevant points (within a meaningful range) were kept. The second metric, treeH_0, was computed by the 0-dimensional persistence diagram of a filtration built by voxels' distance to the trachea carina along the centreline tree.

Koseki*et al.* studied skin function looking to make a regression model that predicts Transepidermal Water Loss (TEWL) [76]. Greyscale images of the skin were binarised using Otsu's method, then downsampled. Three different filtration functions were tested: k-Nearest Neighbour (k-NN) function, signed distance transform, and level-set thresholding. The k-NN function assigns to each pixel a value based on the distance to the k-th nearest white pixel. The signed distance transform assigns to each white pixel, the Manhattan distance to the nearest black region and, conversely, to each black pixel, the negative Manhattan distance to the nearest white region. In each case, the values assigned to the pixels by the functions were thresholded sequentially, creating binarised images and a filtration from the sequence. PH features were stored as a vector known as a Persistence Image [100]. Combined with other features such as age, sex, temperature, and humidity, several ML methods were tested to predict TEWL.

In [77], PH is used to count lesions in brain MRI images. This novel method called P-Count is proposed as a more accurate way of counting lesions than the current standard. In clinical practice, lesion counts are typically obtained by applying a direct threshold to soft probability maps (from software such as SAMSEG-lesion). P-Count leverages PH by calculating the persistence of connected components as the threshold on the soft probability map varies. Connected components with low persistence are regarded as noise and ones with higher persistence are counted as lesions.

4.8 Performance of Other Types of Analyses

Due to the differences in the tasks performed by the models that did not deal with Classification or Segmentation, standard metrics were not compared. In [65], regression was performed using topological features obtained from images as independent variables to predict TEWL. The best model obtained an R^2 of 0.524 and a Root Mean-Squared Error of 3.0 was obtained, indicating a sufficiently good fit. In the case of the model for localization of podosomes in [73] was compared against manually identified phagocytic podosomes, obtaining a false discovery rate of 0.01.

In [74], the differences in populations were significant for certain treatments and closely matched visual and viscoelastic observations, demonstrating that PH captures subtle structural alterations in clot architecture.

Rieck *et al.* performed three analyses (1) age prediction using topological information was accurate; (2) they found that young children exhibit simple, linear trajectories, while older children and adults display more complex dynamics; (3) topological differences were significant across cohorts at event boundaries of the movie, especially in occipital-temporal regions.

In [28], Kaji *et al.* built two metrics to compare samples of patients: radialH_0 was found to be significantly higher in samples of COPD patients with moderate to large decline in FEV1, whereas treeH_0 was significantly higher in COPD patients with moderate FEV1 decline.

The experiments in [77] showed that P-Count is more accurate than the current practice. This was shown by comparing the graph of time vs lesions, noting that the proposed method is closer in shape and space to the ground truth.

5 Discussion

5.1 Impact in Clinical Settings

In this study a literature review of PH-based methods for performing image processing and analysis was presented. To complement academic research, a brief investigation into the impact that persistent homology is having in real-world clinical scenarios was also performed. A search on Google Patents (https://patents.google.com/) using the following string *("Topological Data Analysis" OR "Persistent Homology" OR "Persistence Homology") AND ("medicine" OR "clinical") AND ("images" OR "imaging")* yielded 147 results. While it is out of the scope of this literature review, it is possible to draw a conclusion that the techniques and ideas explored in the academic research are being translated into real-world patents for medical devices. However, to the best knowledge of the authors these have not yet been incorporated into AI-enabled medical devices approved by the U.S. Food and Drug Administration (FDA). Searching on the official FDA list of AI-enabled medical devices (https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices) for topological-adjacent terminology returned no results. This is supporting evidence for a gap that exists between academia and industry in using PH-based methods for medical image analysis.

5.2 Patterns Across Reviewed Studies

A number of broad observations related to the performance of PH methodologies applied in medical image processing reviewed in this paper can be drawn. Below, differences observed between segmentation and classification studies are explored, followed by an examination of failure modes: conditions under which PH-based methods underperformed.

Segmentation vs. Classification. It is evident that in segmentation tasks, models often use Level-Set filtrations on pixel intensities or probability maps, thus suggesting that there is a higher degree of consensus in PH-based segmentation methods. In comparison, classification models show more variation in the type of filtration strategies used, including some that were based on distances instead of values assigned to specific pixels/voxels (i.e. brightness or otherwise). For example, [67] located cells and then used their locations to build a Delaunay filtration. Across both tasks, PH was rarely used alone. In segmentation, PH was often used to enforce topological correctness in an existing CNN pipeline, through a topological loss function [55–59, 78]. Improvements on pre-existing pipelines can be measured using topological metrics, instead of simply the F_1-score. An example of this was seen in Yang *et al.* [59] where F_1-improvement was modest but the Betti error indicated better topological fidelity. This suggests that the primary contribution of PH in segmentation is topological correctness rather than raw accuracy, a distinction that standard metrics like F_1 or Dice do not fully capture. In classification, a similar pattern emerged: PH features were commonly combined with conventional ML or DL classifiers. For example, Huang *et al.* [61] and Wang *et al.* [62] both used PH to guide CNN attention, while Edwards *et al.* [72] combined persistence landscapes with SVR classifiers. The strong performance of these hybrid approaches suggests that PH is most effective when it complements existing feature extraction methods.

Failure Modes. There are cases where the application of Persistence Homology as a methodology may lead to failures or underperformance. The following cases have been identified.

First, low scores in classification could be obtained, such as those reported by Pritchard *et al.* [66] and Huang *et al.* [61], where the topological features extracted do not capture sufficient structure to discriminate between classes. Second, methods relying on distance-based filtrations such as Vietoris-Rips face a high computational cost. This cost can be reduced, as demonstrated by Iqbal *et al.*, who replaced raw pixels with centroids of connected components and made the computation much simpler [70]. However, it must be noted that, if not done carefully, this approach runs a risk of discarding relevant spatial information. Third, methods that depend on a preliminary segmentation or probability map before computing PH may inherit any errors from that initial step, as demonstrated in Yang *et al.*, where Dice score did not improve on the ISBI12/13 datasets, which could possibly be explained by errors from the initial probability maps produced by ConvLSTM before using the topological attention module.

5.3 Computational Complexity

For a filtration with n simplices (total), the theoretical worst-case complexity of computing persistence via the standard matrix-reduction algorithm is $O(n^3)$ [25]. This worst case arises because the column reduction step on the boundary matrix can be cubic in the number of cells. For an image of k pixels, the usual triangulation from the grid of pixels yields a number of simplices of $n = O(k)$, so the worst-case complexity based on number of pixels is $O(k^3)$. It must be emphasised that this is an extremely pessimistic bound and, in reality, computation of persistence is much better behaved, thanks to sparsity, compression optimisations and cubical methods. In fact, in [102, 103] an optimised algorithm is presented, which, when averaged over four 3D datasets,

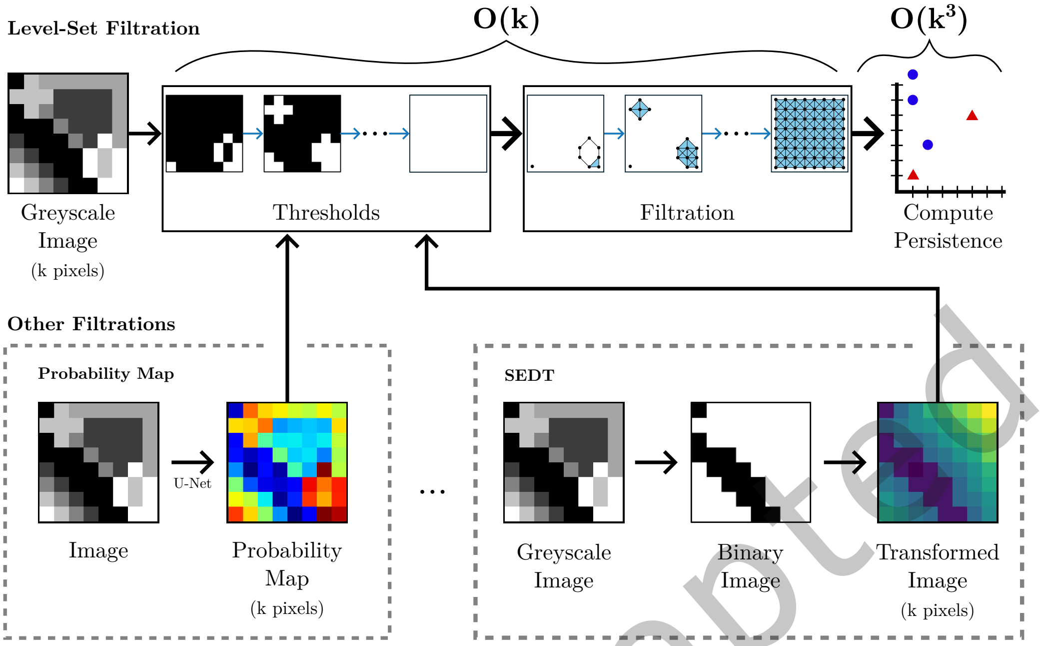

Fig. 9. Many filtration strategies are based on a Level-Set filtration. That is, data will be preprocessed with a transformation (U-Net, SEDT, etc., denoted by dotted boxes). The transformed image data, will become the input to the a level set filtration. Then, data will follow the basic Level-Set pipeline, where threshold moves over the intensities of the image. Building the filtration is efficient (linear on the number of pixels) [101], but a theoretical worst-case bound of $O(k^3)$ for an image with k pixels arises from column reduction on the boundary-matrix. The exact complexity of computing persistence with a filtration strategy where the image is first transformed, depends on the chosen transformation on the image.

reduces the execution time over 600% (from to 7.64 to 1.18 minutes), and the number of add operations is also significantly reduced. Moreover, for level-set filtrations on images, Wagner, Chen, and Vuçini show that the filtration can be built in linear time $O(n)$ [101], but the reduction step still governs the theoretical worst-case cost of $O(n^3)$. Some of the filtration strategies found in the review essentially boil down to a level-set filtration. Take the work by Byrne *et al.* [57], for example, a DL model is used to output a probability map and a moving threshold parameter over the image of the probability map is used to build the filtration and compute persistence (equivalent to the intensity threshold of pixel values in the Level-Set Filtration). Similarly, papers where the image is first transformed using the SEDT, or any other transformation, can be viewed as a level-set filtration on the transformed image (Fig. 9). These types of filtrations will then also have an additional complexity term depending on the type of transform done before thresholding.

For distance-based filtrations such as Vietoris-Rips or Cech, building the filtration grows exponentially with the number of points m, since the number of simplices grows with order of $O(2^m)$. This growth in the number of simplices means the complexity in time taken to compute persistence also grows exponentially as the number of points grows. This, however, proves how valuable the work done by Iqbal *et al.* was [70]. As mentioned in Subsection 4.4, in this article, the authors managed to reduce the number of points needed to perform classification by identifying connected components at different intensity bins and taking their centroids as the point cloud to compute persistence.

5.4 Scalability and Dimensionality

All filtration strategies shown in this review can be scaled up to any dimension of data: the pixel intensity thresholds can easily be extended to work with voxel intensities, any transformation that is done on 2D images (such as the SEDT, or Probability Map) can be extended to transform 3D volumes instead, and distance-based filtrations working with points in 2D can be defined in any n-dimensional Euclidean space. Height thresholds such as the ones used by Belchi *et al.* [65] can be defined in any dimension by choosing a vector v and thresholding along the direction defined by v.

5.5 Software Tools

There exist many options for researchers and amateurs to begin working with PH easily, without the need to develop bespoke implementations themselves. Specifically: the packages for C++ and Python bindings Ripser and Ripser++ allow fast computation of Vietoris–Rips complexes [104], Gudhi provides a comprehensive TDA framework [105], and Dionysus provides a simple and flexible way to build custom filtrations [106]. These libraries also offer compatibility with common ML pipelines, for example, by exporting persistence diagrams as feature vectors suitable for scikit [107], or by incorporating PH-based loss functions into PyTorch [108]. Another library that is worth noting is PHAT, a C++ library targeted towards software developers that improves the efficiency on the boundary matrix reduction [103]. More graphical interfaces exist for users without a programming background, such as the Topology Toolkit (TTK) which is integrated with ParaView [109].

5.6 Limitations of the Review

While this review offers plenty of insight into the current state of the art of PH in the context of medical image analysis and processing, it must be noted that there are limitations to the approach through which the selection of articles was done. The PubMed database offers a selection of most relevant medical research, MICCAI is a venue famous for harbouring a lot of the research done specifically on applications of computing in medical image analysis, and NIPS is known for high-quality machine learning research and has been the venue for several seminal contributions to the field of computer vision and deep learning [110–112]. Nonetheless, these three selections do not make up the entirety of research in this field. Particularly missing from this approach are contributions from other databases and venues such as IEEE Xplore or specialised conferences in applied mathematics and biomedical engineering. However, given the strong representation of medical imaging research within PubMed, MICCAI, and NIPS, the selected sources still provide a representative overview of the state of the art.

A further constraint arose from the reviewed studies themselves. Since they do not share common datasets, use the same evaluation protocols, nor have the same experimental settings, the comparison between techniques and methods was inherently limited.

6 Conclusion

In this review, an in-depth analysis of 30 different articles was performed. These articles all use PH in some key step of a model used to analyse medical images, for different tasks such as segmentation, classification, or other applications. It is worth highlighting that, in those papers where a performance metric was reported, the models obtained a sufficiently high score, independently of the task it was used for. In most cases, PH is used to extract multi-scale topological features from images that perhaps other algorithms are not able to. The high-scoring performance metrics obtained by the models underscore the fact that these topological features provide information which is beneficial to existing models used to segment or classify images.

Additionally, this review showcases the diversity of applications of PH. This is mostly shown by the five papers presented in subsection 4.7. For example, while segmentation and classification are classic image analysis tasks, in [73] the structure and organization of podosomes is described using PH, while in [76], topological features obtained with PH are inputted into a linear regression model to predict TEWL from images. The first of these two papers showcases the power of topology to describe the overall shape of data. The second, demonstrates that PH can capture useful features from an image of skin that adds information to regression models.

In subsections 4.1 to 4.6 it is seen that PH is applied to both classification and segmentation equally, with sufficiently good performances. The main way that PH captures features from images to perform segmentation is through building a level-set filtration from the pixel intensities in images or patches of a whole image. Through thresholding, this filtration captures the most persistent objects in greyscale images while discarding regions that are essentially noise. Adding to this, filtrations are also commonly built from the probability maps predicted by an initial CNN segmentation model, PH then helps to update the weights according to topologically-meaningful results [55–59]. In classification, the reported performance metrics are also high compared to other state-of-the-art methodologies. For this task, the level-set filtration is also the most prevalent, but, a larger variety of ways to build a filtration is found. In the case of [66], SEDT is first applied to previously binarised images, this function is then thresholded to build a filtration. Xin *et al.* in [69] filter the Vietoris-Rips complex to build a filtration [37]. Then, this same filtration is used to build a higher-level filtration that captures the dynamics of local clusters (communities) of simplices in the filtration.

It is worth mentioning that PH does present some downsides/challenges. For example, PH requires high-level understanding of topological concepts in order to build on existing ideas. However, this high-level understanding is not necessary to perform analysis with PH. From a practical perspective, articles within this review have differed in their implementation for building filtrations and computing persistence, from using dedicated libraries, to implementing their own algorithms. But, dedicated toolkits are becoming evermore available, as discussed in subsection 5.5, thus facilitating the application of PH in medical imaging. However, challenges still remain regarding scalability, standardization, and user-friendly documentation. Addressing these issues will be essential for PH to transition from an emerging research tool to a widely adopted component of clinical imaging workflows.

Considering the above, the benefits of integrating PH into image processing pipelines seem to outweigh the challenges. For this we present the following arguments:

(1) As mentioned previously, PH can extract features from data that may not be initially attainable by other methods. These types of features can easily be used as priors to improve existing image processing pipelines.
(2) The most popular methods for image processing recently have been those that involve CNNs. PH can be easily integrated into CNNs by adding a topological term to the loss function, as seen in [44] and [55].
(3) The calculation of Betti numbers (and thus PH) becomes a relatively simple task once a filtration is built. This can easily be done by representing the simplicial complex as a Boundary Matrix. Betti numbers can be found by using simple matrix multiplication (which is very fast) [24].
(4) In this review, it has been shown that PH can help improve the performance of image processing models. Whether it is compared with SegNet, U-Net, and ALBU, as shown by Banerjee *et al.* [53], or demonstrates clear improvements over a CNN without a topological prior, as shown by Shin *et al.* [55]. It is evident that PH provides benefits when integrated into image processing models.

6.1 Future Directions

Although significant progress has been made in the application of PH to medical image analysis, there are still gaps in the field. Based on the findings of this literature review, it is possible to draw investigation paths that still require further research:

- **Classifying histology patches:** the example in Section 2.4 can be expanded into a complete classification model. As seen in Figure 7, statistical attributes of the persistence diagrams show distinct distributions across samples of tissue types. These attributes are similar to the ones computed by Pritchard *et al.* [66]. This, with other features can be used to reveal important information about histopathological datasets, seen in [113]. In [68], a classifier for histopathology achieves high ROC-AUC but fails when it comes to multi-class classification.

- **Analysis of topological features in disease progression over time:** many diseases present with structural changes as time progresses. PH presents a way to track the topological changes over different time points. Potential applications could include quantifying gradual structural changes in brain MRIs due to neurodegenerative diseases such as Parkinson's or Alzheimer's, analysis of how tumors evolve over time in response to treatment, or tracking structural changes in heart tissue over time. In this review, two attempts to work with spatio-temporal data have been presented [54, 75]. These are, however, centred on video, instead of a longitudinal study to reveal changes in topological structures.
- **Application of PH in 3D and 4D medical imaging:** as shown in this review, most PH-based studies focus on 2D medical images, but many imaging modalities produce 3D volumes. PH can be extended to higher-dimensional datasets by considering higher-dimensional loops [54, 63]. However, the effectiveness of using PH features in these datasets is still under-researched. A direct application where this idea might be useful is in volumes of Electron Microscopy images [32]. Dark spots such as mitochondria in a volume of HeLa cells could potentially be captured by the level-set filtration as they resemble loops inside a cell.
- **Sensitivity to parameters:** when building a filtration, some parameters are arbitrarily chosen. For example, in the level-set filtration, usually all possible pixel intensity values are chosen as the level sets (*e.g.*, all values in the rage $[0, 255]$). However, reducing the number of levels (*e.g.*, six different pixel intensities as in Figure 4) could have an effect on the persistence diagrams and classification accuracy. Additionally, different filtration functions (*e.g.*, SEDT) transform the image in different ways, highlighting different parts and features of the images. Yet, to the best knowledge of the authors, no studies systematically compare their effects on the same classification or segmentation task. For example, an intensity-based filtration may highlight cell density variations, while distance-based filtrations (*e.g.*, SEDT) may emphasize spatial organization. Comparing these effects could refine topological biomarkers for cancer diagnosis. Additionally, in diseases like COPD, tissue degradation leads to structural changes in the bronchial tree, as shown in [65]. Combining these findings with descriptors derived from other filtrations could capture previously unknown aspects of lung texture and airways, helping in severity assessment.
- **Training classifying models using persistence image:** the Persistence Image (PI) encodes PH features into vectors [100]. Koseki *et al.*, and Rieck *et al.* used PI for regression [75, 76], but a similar approach could be applied to classification and segmentation. For example, in the case of patches of histology WSI (Figure 6), the PI can be computed for each tissue type, which can be used to train a classification ML model (*e.g.*, Support Vector Machines, Logistic Regression, Deep Learning Classifiers).
- **Distance metrics between distributions:** as mentioned in the previous point, and seen in Figure 7, the statistical attributes of persistence diagrams vary according to the tissue classes. A simple way to compare the distributions is by the difference in their means. However, this approach does not take into account the standard deviations and overall shapes of the distributions. More complex distance functions such as the Bayes Discriminability Index [114], the Bhattacharyya Distance [115], Bhattacharyya Space [116] and the Kullback-Leibler Divergence [117] better capture the dissimilarities between distributions. Even more sophisticated measures, such as the Total Variation Distance [118] and the Wasserstein Metric [118], permit a rigorous mathematical quantification of differences between distributions. All these measures can be used to enhance statistical analysis and classification of tissue types by the scatterplots generated from the statistical attributes that each distribution yields.

- **Distance metrics between persistence diagrams:** distances between Persistence Diagrams can be mathematically defined in a similar way as the metrics between distributions. These distances include the Wasserstein Distance for Persistence Diagrams [119] and the Bottleneck Distance [43]. These metrics can be used to compare individual persistence diagrams but could also be extended to compare groups of persistence diagrams quantitavely. This has been tested by [58], but, the results that these comparisons have remain generally unknown. It is possible that these can help to quantify structural changes as an early biomarker to the progression of diseases such as cancer, or cardiovascular disease.

7 Conflict of Interests

The authors declare no conflict of interests.

References

[1] J. M. S. Prewitt, Parametric and nonparametric recognition by computer: An application to leukocyte image processing, Advances in Computers 12 (1972) 285–414. doi:10.1016/S0065-2458(08)60511-2.
[2] E. B. Carstens, R. G. Marusyk, An image analysis study of adenovirus type 5-induced crystalline inclusions, Journal of General Virology 29 (3) (1975) 249–256. doi:10.1099/0022-1317-29-3-249.
[3] T. Caspersson, R. Sennerstam, Low level absorption microscope image analysis, Experimental Cell Research 92 (2) (1975) 333–338. doi:10.1016/0014-4827(75)90387-0.
[4] A. Echle, N. G. Laleh, P. L. Schrammen, N. P. West, C. Trautwein, T. J. Brinker, S. B. Gruber, R. D. Buelow, P. Boor, H. I. Grabsch, P. Quirke, J. N. Kather, Deep learning for the detection of microsatellite instability from histology images in colorectal cancer: A systematic literature review, ImmunoInformatics 3–4 (2021) 100008. doi:10.1016/j.immuno.2021.100008.
[5] R. Ortiz-Ramón, A. Larroza, S. Ruiz-España, E. Arana, D. Moratal, Classifying brain metastases by their primary site of origin using a radiomics approach based on texture analysis: a feasibility study, European Radiology 28 (11) (2018) 4514–4523. doi:10.1007/s00330-018-5463-6.
[6] L. Misera, G. Müller-Franzes, D. Truhn, J. N. Kather, Weakly supervised deep learning in radiology, Radiology 312 (1) (2024) e232085. doi:10.1148/radiol.232085.
[7] N. A. Valous, F. Popp, I. Zörnig, D. Jäger, P. Charoentong, Graph machine learning for integrated multi-omics analysis, Br. J. Cancer 131 (2) (2024) 205–211. doi:10.1038/s41416-024-02706-7.
[8] R. Berenguer-Vidal, R. Verdú-Monedero, J. Morales-Sánchez, I. Sellés-Navarro, R. Del Amor, G. García, V. Naranjo, Automatic segmentation of the retinal nerve fiber layer by means of mathematical morphology and deformable models in 2D optical coherence tomography imaging, Sensors (Basel) 21 (23) (2021) 8027. doi:10.3390/s21238027.
[9] J. Ning, L. Zhang, Fourier analysis of corneal irregular astigmatism after small-incision lenticule extraction and transepithelial photorefractive keratectomy: A comparative study, Medicine (Baltimore) 103 (9) (2024) e37340. doi:10.1097/MD.0000000000037340.
[10] R. Wang, S. A. Khurram, H. Walsh, L. S. Young, N. Rajpoot, A novel deep learning algorithm for human papillomavirus infection prediction in head and neck cancers using routine histology images, Modern Pathology: An Official Journal of the United States and Canadian Academy of Pathology, Inc 36 (12) (2023) 100320. doi:10.1016/j.modpat.2023.100320.
[11] N. Zamanitajeddin, M. Jahanifar, M. Bilal, M. Eastwood, N. Rajpoot, Social network analysis of cell networks improves deep learning for prediction of molecular pathways and key mutations in colorectal cancer, Medical Image Analysis 93 (2024) 103071. doi: 10.1016/j.media.2023.103071.
[12] M. Bilal, S. E. A. Raza, A. Azam, S. Graham, M. Ilyas, I. A. Cree, D. Snead, F. Minhas, N. M. Rajpoot, Development and validation of a weakly supervised deep learning framework to predict the status of molecular pathways and key mutations in colorectal cancer from routine histology images: a retrospective study, The Lancet. Digital Health 3 (12) (2021) e763–e772. doi:10.1016/S2589-7500(21)00180-1.
[13] K. Silina, F. Ciompi, Cancer-associated lymphoid aggregates in histology images: Manual and deep learning-based quantification approaches, Methods Mol. Biol. 2864 (2025) 231–246. doi:10.1007/978-1-0716-4184-2_12.
[14] J. Wang, R. Gong, S. Heidari, M. Rogers, T. Tani, H. Abe, N. Ichinohe, A. Woodward, P. J. Delmas, A deep learning-based pipeline for segmenting the cerebral cortex laminar structure in histology images, Neuroinformatics 22 (4) (2024) 745–761. doi:10.1007/s12021-024-09688-0.
[15] C. Yoon, E. Park, S. Misra, J. Y. Kim, J. W. Baik, K. G. Kim, C. K. Jung, C. Kim, Deep learning-based virtual staining, segmentation, and classification in label-free photoacoustic histology of human specimens, Light: Science & Applications 13 (1) (2024) 226. doi: 10.1038/s41377-024-01554-7.
[16] D. Reinecke, N. Maarouf, A. Smith, D. Alber, J. Markert, N. K. Goff, T. C. Hollon, A. Chowdury, C. Jiang, X. Hou, A.-K. Meissner, G. Fürtjes, M. I. Ruge, D. Ruess, T. Stehle, A. Al-Shughri, L. I. Körner, G. Widhalm, T. Roetzer-Pejrimovsky, J. G. Golfinos, M. Snuderl,

V. Neuschmelting, D. A. Orringer, Fast intraoperative detection of primary central nervous system lymphoma and differentiation from common central nervous system tumors using stimulated raman histology and deep learning, Neuro-Oncology 27 (5) (2024) 1297–1310. doi:10.1093/neuonc/noae270.

[17] K. Yasaka, C. Sato, H. Hirakawa, N. Fujita, M. Kurokawa, Y. Watanabe, T. Kubo, O. Abe, Impact of deep learning on radiologists and radiology residents in detecting breast cancer on CT: a cross-vendor test study, Clinical Radiology 79 (1) (2024) e41–e47. doi:10.1016/j.crad.2023.09.022.

[18] S. B. U. Haque, A. Zafar, Robust medical diagnosis: A novel two-phase deep learning framework for adversarial proof disease detection in radiology images, Journal of Imaging Informatics in Medicine 37 (1) (2024) 308–338. doi:10.1007/s10278-023-00916-8.

[19] K. Doi, H. Takegawa, M. Yui, Y. Anetai, Y. Koike, S. Nakamura, N. Tanigawa, M. Koziumi, T. Nishio, Deep learning-based detection of patients with bone metastasis from japanese radiology reports, Japanese Journal of Radiology 41 (8) (2023) 900–908. doi:10.1007/s11604-023-01413-2.

[20] Y. Arafat, C. C. Reyes-Aldasoro, Computational image analysis techniques, programming languages and software platforms used in cancer research: A scoping review, in: G. Yang, A. Aviles-Rivero, M. Roberts, C.-B. Schönlieb (Eds.), Medical Image Understanding and Analysis, Vol. 3413 of Lecture Notes in Computer Science, Springer International Publishing, Cham, 2022, p. 833–847. doi:10.1007/978-3-031-12053-4_61.

[21] S. Graham, Q. D. Vu, M. Jahanifar, M. Weigert, U. Schmidt, W. Zhang, J. Zhang, S. Yang, J. Xiang, X. Wang, J. L. Rumberger, E. Baumann, P. Hirsch, L. Liu, C. Hong, A. I. Aviles-Rivero, A. Jain, H. Ahn, Y. Hong, H. Azzuni, M. Xu, M. Yaqub, M.-C. Blache, B. Piégu, B. Vernay, T. Scherr, M. Böhland, K. Löffler, J. Li, W. Ying, C. Wang, D. Snead, S. E. A. Raza, F. Minhas, N. M. Rajpoot, C. C. Consortium, Conic challenge: Pushing the frontiers of nuclear detection, segmentation, classification and counting, Medical Image Analysis 92 (2024) 103047. doi:10.1016/j.media.2023.103047.

[22] J. N. Kather, J. Krisam, P. Charoentong, T. Luedde, E. Herpel, C.-A. Weis, T. Gaiser, A. Marx, N. A. Valous, D. Ferber, L. Jansen, C. C. Reyes-Aldasoro, I. Zörnig, D. Jäger, H. Brenner, J. Chang-Claude, M. Hoffmeister, N. Halama, Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study, PLOS Medicine 16 (1) (2019) e1002730. doi:10.1371/journal.pmed.1002730.

[23] D. Zimmerer, P. M. Full, F. Isensee, P. Jager, T. Adler, J. Petersen, G. Kohler, T. Ross, A. Reinke, A. Kascenas, B. S. Jensen, A. Q. O'Neil, J. Tan, B. Hou, J. Batten, H. Qiu, B. Kainz, N. Shvetsova, I. Fedulova, D. V. Dylov, B. Yu, J. Zhai, J. Hu, R. Si, S. Zhou, S. Wang, X. Li, X. Chen, Y. Zhao, S. N. Marimont, G. Tarroni, V. Saase, L. Maier-Hein, K. Maier-Hein, Mood 2020: A public benchmark for out-of-distribution detection and localization on medical images, IEEE transactions on medical imaging 41 (10) (2022) 2728–2738. doi:10.1109/TMI.2022.3170077.

[24] H. Edelsbrunner, J. Harer, Persistent homology—a survey, in: J. E. Goodman, J. Pach, R. Pollack (Eds.), Contemporary Mathematics, Vol. 453, American Mathematical Society, Providence, Rhode Island, 2008, p. 257–282. doi:10.1090/conm/453/08802.

[25] Edelsbrunner, Letscher, Zomorodian, Topological persistence and simplification, Discrete & Computational Geometry 28 (4) (2002) 511–533. doi:10.1007/s00454-002-2885-2.

[26] A. Zomorodian, G. Carlsson, Computing persistent homology, Discrete & Computational Geometry 33 (2) (2005) 249–274. doi:10.1007/s00454-004-1146-y.

[27] C. S. Pun, S. X. Lee, K. Xia, Persistent-homology-based machine learning: a survey and a comparative study, Artificial Intelligence Review 55 (7) (2022) 5169–5213. doi:10.1007/s10462-022-10146-z.

[28] S. Kaji, N. Tanabe, T. Maetani, Y. Shiraishi, R. Sakamoto, T. Oguma, K. Suzuki, K. Terada, M. Fukui, S. Muro, S. Sato, T. Hirai, Quantification of airway structures by persistent homology, IEEE transactions on medical imaging 43 (8) (2024) 2758–2768. doi:10.1109/TMI.2024.3376683.

[29] M. Mijangos, L. Pacheco, A. Bravetti, N. González-García, P. Padilla, R. Velasco-Segura, Persistent homology reveals robustness loss in inhaled substance abuse rs-fmri networks, PLOS ONE 19 (9) (Sep. 2024). doi:10.1371/journal.pone.0310165.

[30] S. S. Sørensen, T. Du, C. A. N. Biscio, L. Fajstrup, M. M. Smedskjaer, Persistent homology: A tool to understand medium-range order glass structure, Journal of Non-Crystalline Solids: X 16 (2022) 100123. doi:10.1016/j.nocx.2022.100123.

[31] Y.-M. Chung, C.-S. Hu, E. Sun, H. C. Tseng, Morphological multiparameter filtration and persistent homology in mitochondrial image analysis, PLOS ONE 19 (9) (2024) e0310157. doi:10.1371/journal.pone.0310157.

[32] D. Brito-Pacheco, C. Karabağ, C. Brito-Loeza, P. Giannopoulos, C. C. Reyes-Aldasoro, Relationship between irregularities of the nuclear envelope and mitochondria in hela cells observed with electron microscopy, in: 2024 IEEE International Symposium on Biomedical Imaging (ISBI), 2024, p. 1–5. doi:10.1109/ISBI56570.2024.10635639.

[33] R. Rojas-Moraleda, W. Xiong, N. Halama, K. Breitkopf-Heinlein, S. Dooley, L. Salinas, D. W. Heermann, N. A. Valous, Robust detection and segmentation of cell nuclei in biomedical images based on a computational topology framework, Medical Image Analysis 38 (2017) 90–103. doi:10.1016/j.media.2017.02.009.

[34] A. Marabotti, A. Facchiano, When it comes to homology, bad habits die hard, Trends in Biochemical Sciences 34 (3) (2009) 98–99. doi:10.1016/j.tibs.2008.12.001.

[35] G. Carlsson, M. Vejdemo-Johansson, Topological Data Analysis with Applications, Cambridge University Press, Cambridge, 2021. doi:10.1017/9781108975704.

[36] F. Chazal, B. Michel, An introduction to topological data analysis: Fundamental and practical aspects for data scientists, Frontiers in Artificial Intelligence 4 (2021) 667963. doi:10.3389/frai.2021.667963.

[37] L. Wasserman, Topological data analysis, Annual Review of Statistics and Its Application 5 (Volume 5, 2018) (2018) 501–532. doi:10.1146/annurev-statistics-031017-100045.

[38] L. Euler, Solutio problematis ad geometriam situs pertinentis, Commentarii academiae scientiarum Petropolitanae (1741) 128–140.

[39] J. Beuria, Persistent homology of collider observations: When (w)hole matters, Physics Letters B 846 (2023) 138188. doi:10.1016/j.physletb.2023.138188.

[40] D. Bhaskar, W. Y. Zhang, I. Y. Wong, Topological data analysis of collective and individual epithelial cells using persistent homology of loops, Soft Matter 17 (17) (2021) 4653–4664. doi:10.1039/d1sm00072a.

[41] R. Ghrist, Barcodes: the persistent topology of data, Bulletin of the American Mathematical Society 45 (1) (2008) 61–75. doi:10.1090/S0273-0979-07-01191-3.

[42] M. Kerber, D. Morozov, A. Nigmetov, Geometry helps to compare persistence diagrams, ACM J. Exp. Algorithmics 22 (2017) 1.4:1–1.4:20. doi:10.1145/3064175.

[43] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, Discrete & Computational Geometry 37 (1) (2007) 103–120. doi:10.1007/s00454-006-1276-5.

[44] M. Haft-Javaherian, M. Villiger, C. B. Schaffer, N. Nishimura, P. Golland, B. E. Bouma, A topological encoding convolutional neural network for segmentation of 3d multiphoton images of brain vasculature using persistent homology, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2020, p. 4262–4271. doi:10.1109/CVPRW50498.2020.00503.

[45] L. Panconi, M. Makarova, E. R. Lambert, R. C. May, D. M. Owen, Topology-based fluorescence image analysis for automated cell identification and segmentation, Journal of Biophotonics 16 (3) (2023) e202200199. doi:10.1002/jbio.202200199.

[46] T. Qaiser, Y.-W. Tsang, D. Taniyama, N. Sakamoto, K. Nakane, D. Epstein, N. Rajpoot, Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features, Medical Image Analysis 55 (2019) 1–14. doi:10.1016/j.media.2019.03.014.

[47] J. N. Kather, N. Halama, A. Marx, 100,000 histological images of human colorectal cancer and healthy tissue (Apr. 2018). doi:10.5281/ZENODO.1214455.

[48] M. Macenko, M. Niethammer, J. S. Marron, D. Borland, J. T. Woosley, X. Guan, C. Schmitt, N. E. Thomas, A method for normalizing histology slides for quantitative analysis, in: 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, 2009, p. 1107–1110. doi:10.1109/ISBI.2009.5193250.

[49] F. Chazal, S. Y. Oudot, M. Glisse, V. de Silva, The Structure and Stability of Persistence Modules, SpringerBriefs in Mathematics, Springer Verlag, 2016. doi:10.1007/978-3-319-42545-0.

[50] A. Liberati, D. G. Altman, J. Tetzlaff, C. Mulrow, P. C. Gøtzsche, J. P. A. Ioannidis, M. Clarke, P. J. Devereaux, J. Kleijnen, D. Moher, The prisma statement for reporting systematic reviews and meta-analyses of studies that evaluate healthcare interventions: explanation and elaboration, BMJ 339 (2009) b2700. doi:10.1136/bmj.b2700.

[51] D. Moher, L. Shamseer, M. Clarke, D. Ghersi, A. Liberati, M. Petticrew, P. Shekelle, L. A. Stewart, P.-P. Group, Preferred reporting items for systematic review and meta-analysis protocols (prisma-p) 2015 statement, Systematic Reviews 4 (1) (2015) 1. doi:10.1186/2046-4053-4-1.

[52] B. Kitchenham, S. Charters, Guidelines for performing systematic literature reviews in software engineering, Tech. rep., Keele University (2007).
URL https://legacyfileshare.elsevier.compromis_misc/525444systematicreviewsguide.pdf

[53] S. Banerjee, L. Magee, D. Wang, X. Li, B.-X. Huo, J. Jayakumar, K. Matho, M.-K. Lin, K. Ram, M. Sivaprakasam, J. Huang, Y. Wang, P. P. Mitra, Semantic segmentation of microscopic neuroanatomical data by combining topological priors with encoder–decoder deep networks, Nature Machine Intelligence 2 (10) (2020) 585–594. doi:10.1038/s42256-020-0227-9.

[54] L. Panconi, A. Tansell, A. J. Collins, M. Makarova, D. M. Owen, Three-dimensional topology-based analysis segments volumetric and spatiotemporal fluorescence microscopy, Biological Imaging 4 (2024) e1. doi:10.1017/S2633903X23000260.

[55] S. Y. Shin, S. Lee, D. Elton, J. L. Gulley, R. M. Summers, Deep small bowel segmentation with cylindrical topological constraints, in: A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M. A. Zuluaga, S. K. Zhou, D. Racoceanu, L. Joskowicz (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2020, Springer International Publishing, Cham, 2020, p. 207–215. doi:10.1007/978-3-030-59719-1_21.

[56] S. Sun, Y. Wang, J. Yang, Y. Feng, L. Tang, S. Liu, H. Ning, Topology-sensitive weighting model for myocardial segmentation, Computers in Biology and Medicine 165 (2023) 107286. doi:10.1016/j.compbiomed.2023.107286.

[57] N. Byrne, J. R. Clough, I. Valverde, G. Montana, A. P. King, A persistent homology-based topological loss for cnn-based multiclass segmentation of cmr, IEEE Transactions on Medical Imaging 42 (1) (2023) 3–14. doi:10.1109/TMI.2022.3203309.

[58] A. H. Berger, L. Lux, N. Stucki, V. Bürgin, S. Shit, A. Banaszak, D. Rueckert, U. Bauer, J. C. Paetzold, Topologically faithful multi-class segmentation in medical images, in: M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, J. A. Schnabel (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, Springer Nature Switzerland, Cham, 2024, p. 721–731. `doi:10.1007/978-3-031-72111-3_68`.

[59] J. Yang, X. Hu, C. Chen, C. Tsai, A topological-attention convlstm network and its application to em images, in: M. de Bruijne, P. C. Cattin, S. Cotin, N. Padoy, S. Speidel, Y. Zheng, C. Essert (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2021, Springer International Publishing, Cham, 2021, p. 217–228. `doi:10.1007/978-3-030-87193-2_21`.

[60] H. Molina-Abril, A. F. Frangi, Topo-geometric filtration scheme for geometric active contours and level sets: application to cerebrovascular segmentation, Medical image computing and computer-assisted intervention: MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention 17 (Pt 1) (2014) 755–762. `doi:10.1007/978-3-319-10404-1_94`.

[61] C. Huang, F. Guo, S. Lin, Y. Dai, Q. Chen, S. Zhang, X. Xu, Attentive feature interaction based persistent homology-augmented network for esophageal cancer lesion detection, Medical Physics 52 (6) (2025) 3927–3939. `doi:10.1002/mp.17707`.

[62] F. Wang, Z. Zou, N. Sakla, L. Partyka, N. Rawal, G. Singh, W. Zhao, H. Ling, C. Huang, P. Prasanna, C. Chen, Topotxr: A topology-guided deep convolutional network for breast parenchyma learning on dce-mris, Medical Image Analysis 99 (2025) 103373. `doi:10.1016/j.media.2024.103373`.

[63] S. Du, Q. Lao, Q. Kang, Y. Li, Z. Jiang, Y. Zhao, K. Li, Distilling knowledge from topological representations for pathological complete response prediction, in: L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, S. Li (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2022, Springer Nature Switzerland, Cham, 2022, p. 56–65. `doi:10.1007/978-3-031-16434-7_6`.

[64] A. Rammal, R. Assaf, A. Goupil, M. Kacim, V. Vrabie, Machine learning techniques on homological persistence features for prostate cancer diagnosis, BMC Bioinformatics 23 (1) (2022) 476. `doi:10.1186/s12859-022-04992-5`.

[65] F. Belchi, M. Pirashvili, J. Conway, M. Bennett, R. Djukanovic, J. Brodzki, Lung topology characteristics in patients with chronic obstructive pulmonary disease, Scientific Reports 8 (1) (2018) 5341. `doi:10.1038/s41598-018-23424-0`.

[66] Y. Pritchard, A. Sharma, C. Clarkin, H. Ogden, S. Mahajan, R. J. Sánchez-García, Persistent homology analysis distinguishes pathological bone microstructure in non-linear microscopy images, Scientific Reports 13 (1) (2023) 2522. `doi:10.1038/s41598-023-28985-3`.

[67] I. Hartsock, E. Park, J. Toppen, P. Bubenik, E. S. Dimitrova, M. L. Kemp, D. A. Cruz, Topological data analysis of pattern formation of human induced pluripotent stem cell colonies, Scientific Reports 15 (1) (2025) 11544. `doi:10.1038/s41598-025-90592-1`.

[68] H. Wang, G. Huang, Z. Zhao, L. Cheng, A. Juncker-Jensen, M. L. Nagy, X. Lu, X. Zhang, D. Z. Chen, CCF-GNN: A unified model aggregating appearance, microenvironment, and topology for pathology image classification, IEEE Transactions on Medical Imaging 42 (11) (2023) 3179–3193. `doi:10.1109/TMI.2023.3249343`.

[69] B. Xin, J. Huang, L. Zhang, C. Zheng, Y. Zhou, J. Lu, X. Wang, Dynamic topology analysis for spatial patterns of multifocal lesions on MRI, Medical Image Analysis 76 (2022) 102267. `doi:10.1016/j.media.2021.102267`.

[70] S. Iqbal, H. F. Ahmed, T. Qaiser, M. I. Qureshi, N. Rajpoot, Classification of covid-19 via homology of CT-SCAN, Computers in Biology and Medicine 193 (2025) 110226. `doi:10.1016/j.compbiomed.2025.110226`.

[71] Q. Wei, S. Zeng, Z. Jiang, F. Xu, S. Shi, W. Wen, Z. Qin, Z. Lou, K. Li, Topological representation based on wavelet transform as a novel imaging biomarker for tumor diagnosis in ultrasound images: A comprehensive study, Computer Methods and Programs in Biomedicine 269 (2025) 108859. `doi:10.1016/j.cmpb.2025.108859`.

[72] P. Edwards, K. Skruber, N. Milićević, J. B. Heidings, T.-A. Read, P. Bubenik, E. A. Vitriol, Tdaexplore: Quantitative analysis of fluorescence microscopy images through topology-based machine learning, Patterns 2 (11) (2021) 100367. `doi:10.1016/j.patter.2021.100367`.

[73] J. C. Herron, S. Hu, T. Watanabe, A. T. Nogueira, B. Liu, M. E. Kern, J. Aaron, A. Taylor, M. Pablo, T.-L. Chew, T. C. Elston, K. M. Hahn, Actin nano-architecture of phagocytic podosomes, Nature Communications 13 (1) (2022) 4363. `doi:10.1038/s41467-022-32038-0`.

[74] M. Berger, T. Hell, A. Tobiasch, J. Martini, A. Lindner, H. Tauber, M. Bachler, M. Hermann, Analysis of fibrin networks using topological data analysis – a feasibility study, Scientific Reports 14 (1) (2024) 13123. `doi:10.1038/s41598-024-63935-7`.

[75] B. Rieck, T. Yates, C. Bock, K. Borgwardt, G. Wolf, N. Turk-Browne, S. Krishnaswamy, Uncovering the topology of time-varying fMRI data using cubical persistence, in: Advances in Neural Information Processing Systems, Vol. 33, Curran Associates, Inc., 2020, p. 6900–6912. `doi:10.5555/3495724.3496303`.

[76] K. Koseki, H. Kawasaki, T. Atsugi, M. Nakanishi, M. Mizuno, E. Naru, T. Ebihara, M. Amagai, E. Kawakami, Assessment of skin barrier function using skin images with topological data analysis, npj Systems Biology and Applications 6 (1) (2020) 1–9. `doi:10.1038/s41540-020-00160-8`.

[77] X. Hu, A. Sorby-Adams, F. Barkhof, W. Taylor Kimberly, O. Puonti, J. E. Iglesias, P-count: Persistence-based counting of white matter hyperintensities in brain MRI, in: C. Chen, Y. Singh, X. Hu (Eds.), Topology- and Graph-Informed Imaging Informatics, Springer Nature Switzerland, Cham, 2025, p. 100–110. `doi:10.1007/978-3-031-73967-5_10`.

[78] M. Haft-Javaherian, L. Fang, V. Muse, C. B. Schaffer, N. Nishimura, M. R. Sabuncu, Deep convolutional neural networks for segmenting 3d in vivo multiphoton images of vasculature in alzheimer disease mouse models, PloS One 14 (3) (2019) e0213539. `doi:10.1371/journal.pone.0213539`.

[79] R. Forman, Morse theory for cell complexes, Advances in Mathematics 134 (1) (1998) 90–145. `doi:10.1006/aima.1997.1650`.

[80] R. Forman, A user's guide to discrete Morse theory, Séminaire Lotharingien de Combinatoire 48 (2002) B48c, 35 p., electronic only. URL http://eudml.org/doc/123837

[81] A. Buslaev, S. Seferbekov, V. Iglovikov, A. Shvets, Fully convolutional network for automatic road extraction from satellite imagery, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2018, p. 197–1973. `doi:10.1109/CVPRW.2018.00035`.

[82] A. Galton, M. Duckham, What is the region occupied by a set of points?, in: M. Raubal, H. J. Miller, A. U. Frank, M. F. Goodchild (Eds.), Geographic Information Science, Springer, Berlin, Heidelberg, 2006, p. 81–98. `doi:10.1007/11863939_6`.

[83] H. Edelsbrunner, E. P. Mücke, Three-dimensional alpha shapes, ACM Trans. Graph. 13 (1) (1994) 43–72. `doi:10.1145/174462.156635`.

[84] N. Stucki, J. C. Paetzold, S. Shit, B. Menze, U. Bauer, Topologically faithful image segmentation via induced matching of persistence barcodes, in: Proceedings of the 40th International Conference on Machine Learning, PMLR, 2023, p. 32698–32727. `doi:10.5555/3618408.3619764`.

[85] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-k. Wong, W.-c. Woo, Convolutional lstm network: a machine learning approach for precipitation nowcasting, in: Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1, Vol. 1 of NIPS'15, MIT Press, Cambridge, MA, USA, 2015, p. 802–810. `doi:10.5555/2969239.2969329`.

[86] V. Caselles, Geometric models for active contours, in: Proceedings., International Conference on Image Processing, Vol. 3, 1995, p. 9–12 vol.3. `doi:10.1109/ICIP.1995.537567`.

[87] B. E. Boser, I. M. Guyon, V. N. Vapnik, A training algorithm for optimal margin classifiers, in: Proceedings of the fifth annual workshop on Computational learning theory, COLT '92, Association for Computing Machinery, New York, NY, USA, 1992, p. 144–152. `doi:10.1145/130385.130401`.

[88] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: N. Navab, J. Hornegger, W. M. Wells, A. F. Frangi (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Springer International Publishing, Cham, 2015, p. 234–241. `doi:10.1007/978-3-319-24574-4_28`.

[89] V. Badrinarayanan, A. Kendall, R. Cipolla, Segnet: A deep convolutional encoder-decoder architecture for image segmentation, IEEE Transactions on Pattern Analysis and Machine Intelligence 39 (12) (2017) 2481–2495. `doi:10.1109/TPAMI.2016.2644615`.

[90] P. S. Tsai, J. P. Kaufhold, P. Blinder, B. Friedman, P. J. Drew, H. J. Karten, P. D. Lyden, D. Kleinfeld, Correlations of neuronal and microvascular densities in murine cortex revealed by direct counting and colocalization of nuclei and vessels, Journal of Neuroscience 29 (46) (2009) 14553–14570. `doi:10.1523/JNEUROSCI.3287-09.2009`.

[91] P. Teikari, M. Santos, C. Poon, K. Hynynen, Deep learning convolutional networks for multiphoton microscopy vasculature segmentation, CoRR abs/1606.02382 (2016). `arXiv:1606.02382`.
URL http://arxiv.org/abs/1606.02382

[92] S. Gur, L. Wolf, L. Golgher, P. Blinder, Unsupervised microvascular image segmentation using an active contours mimicking neural network, in: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), 2019, p. 10721–10730. `doi:10.1109/ICCV.2019.01082`.

[93] S. Shit, J. C. Paetzold, A. Sekuboyina, I. Ezhov, A. Unger, A. Zhylka, J. P. W. Pluim, U. Bauer, B. H. Menze, clDice - a novel topology-preserving loss function for tubular structure segmentation, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2021, p. 16555–16564. `doi:10.1109/CVPR46437.2021.01629`.

[94] X. Hu, L. Fuxin, D. Samaras, C. Chen, Topology-preserving deep image segmentation, in: Proceedings of the 33rd International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA, 2019, p. 5657–5668. `doi:10.5555/3454287.3454795`.

[95] N. Atienza, R. Gonzalez-Diaz, M. Rucco, Persistent entropy for separating topological features from noise in Vietoris-Rips complexes, Journal of Intelligent Information Systems 52 (3) (2019) 637–655. `doi:10.1007/s10844-017-0473-4`.

[96] T. Kurita, N. Otsu, N. Abdelmalek, Maximum likelihood thresholding based on population mixture models, Pattern Recognition 25 (10) (1992) 1231–1240. `doi:10.1016/0031-3203(92)90024-D`.

[97] C. Cortes, V. Vapnik, Support-vector networks, Machine Learning 20 (3) (1995) 273–297. `doi:10.1007/BF00994018`.

[98] D. Schleicher, Hausdorff dimension, its properties, and its surprises, The American Mathematical Monthly 114 (6) (2007) 509–528.
URL https://www.jstor.org/stable/27642249

[99] F. Chazal, B. T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, Stochastic convergence of persistence landscapes and silhouettes, in: Proceedings of the thirtieth annual symposium on Computational geometry, SOCG'14, Association for Computing Machinery, New York, NY, USA, 2014, p. 474–483. `doi:10.1145/2582112.2582128`.

[100] H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta, L. Ziegelmeier, Persistence images: a stable vector representation of persistent homology, The Journal of Machine Learning Research 18 (1) (2017) 218–252. `doi:10.5555/3122009.3122017`.

[101] H. Wagner, C. Chen, E. Vuçini, Efficient computation of persistent homology for cubical data, in: Topological Methods in Data Analysis and Visualization II, Mathematics and Visualization, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, p. 91–106. `doi:10.1007/978-3-642-23175-9_7`.

[102] C. Chen, M. Kerber, Persistent homology computation with a twist, in: European Workshop on Computational Geometry, 2011, p. 197–200.
URL https://research-explorer.ista.ac.at/record/3270

[103] U. Bauer, M. Kerber, J. Reininghaus, H. Wagner, Phat – persistent homology algorithms toolbox, Journal of Symbolic Computation 78 (2017) 76–90. `doi:10.1016/j.jsc.2016.03.008`.

[104] C. Tralie, N. Saul, R. Bar-On, Ripser.py: A lean persistent homology library for python, Journal of Open Source Software 3 (29) (2018) 925. `doi:10.21105/joss.00925`.

[105] The GUDHI Project, GUDHI User and Reference Manual, 3rd Edition, GUDHI Editorial Board, 2025.
URL https://gudhi.inria.fr/doc/3.11.0/

[106] D. Morozov, Dionysus 2: library for computing persistent homology (2025).
URL https://mrzv.org/software/dionysus2/

[107] S. v. d. Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, scikit-image: image processing in python, PeerJ 2 (2014) e453. `doi:10.7717/peerj.453`.

[108] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems, Vol. 32, Curran Associates, Inc., 2019, p. 8026–8037. `doi:10.5555/3454287.3455008`.

[109] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, M. Michaux, The topology toolkit, IEEE Transactions on Visualization and Computer Graphics 24 (1) (2018) 832–842. `doi:10.1109/TVCG.2017.2743938`.

[110] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18, Curran Associates Inc., Red Hook, NY, USA, 2018, p. 6572–6583. `doi:10.5555/3327757.3327764`.

[111] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, Curran Associates Inc., Red Hook, NY, USA, 2017, p. 6000–6010. `doi:10.5555/3295222.3295349`.

[112] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20, Curran Associates Inc., Red Hook, NY, USA, 2020, p. 6840–6851. `doi:10.5555/3495724.3496298`.

[113] D. Brito-Pacheco, R. Ibadulla, X. Fernández, P. Giannopoulos, C. C. Reyes-Aldasoro, Persistent homology and gabor features reveal inconsistencies between widely used colorectal cancer training and testing datasets, in: S. Ali, D. C. Hogg, M. Peckham (Eds.), Medical Image Understanding and Analysis, Springer Nature Switzerland, Cham, 2025, p. 87–101. `doi:10.1007/978-3-031-98688-8_7`.

[114] A. Das, W. S. Geisler, A method to integrate and classify normal distributions, Journal of Vision 21 (10) (2021) 1. `doi:10.1167/jov.21.10.1`.

[115] K. Fukunaga, Introduction to Statistical Pattern Recognition, 2nd Edition, Elsevier Science, San Francisco, 2014. `doi:10.1016/C2009-0-27872-X`.

[116] C. C. Reyes-Aldasoro, A. Bhalerao, The bhattacharyya space for feature selection and its application to texture segmentation, Pattern Recognition 39 (5) (2006) 812–826. `doi:10.1016/j.patcog.2005.12.003`.

[117] S. Kullback, R. A. Leibler, On information and sufficiency, The Annals of Mathematical Statistics 22 (1) (1951) 79–86.
URL https://www.jstor.org/stable/2236703

[118] C. Villani, Optimal Transport, Springer, Berlin, Heidelberg, 2009, Ch. The Wasserstein distances, p. 93–111. `doi:10.1007/978-3-540-71050-9\_6`.

[119] Y. Mileyko, S. Mukherjee, J. Harer, Probability measures on the space of persistence diagrams, Inverse Problems 27 (12) (2011) 124007. `doi:10.1088/0266-5611/27/12/124007`.