



Train-Free Segmentation in MRI with Cubical Persistent Homology

Anton François, Raphaël Tinarrage

► To cite this version:

Anton François, Raphaël Tinarrage. Train-Free Segmentation in MRI with Cubical Persistent Homology. 2024. [⟨hal-04369459⟩](#)

HAL Id: hal-04369459

<https://hal.science/hal-04369459v1>

Preprint submitted on 2 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License

Train-Free Segmentation in MRI with Cubical Persistent Homology

Anton François*, Raphaël Tinarrage†

* MAP5, Université de Paris-Cité, France. <https://antonfrancois.github.io/>

† FGV EMAP, Rio de Janeiro, Brazil. <https://raphaeltinarrage.github.io/>

Abstract. We describe a new general method for segmentation in MRI scans using Topological Data Analysis (TDA), offering several advantages over traditional machine learning approaches. It works in three steps, first identifying the whole object to segment via automatic thresholding, then detecting a distinctive subset whose topology is known in advance, and finally deducing the various components of the segmentation. Although convoking classical ideas of TDA, such an algorithm has never been proposed separately from deep learning methods. To achieve this, our approach takes into account, in addition to the homology of the image, the localization of representative cycles, a piece of information that seems never to have been exploited in this context. In particular, it offers the ability to perform segmentation without the need for large annotated data sets. TDA also provides a more interpretable and stable framework for segmentation by explicitly mapping topological features to segmentation components. By adapting the geometric object to be detected, the algorithm can be adjusted to a wide range of data segmentation challenges. We carefully study the examples of glioblastoma segmentation in brain MRI, where a sphere is to be detected, as well as myocardium in cardiac MRI, involving a cylinder, and cortical plate detection in fetal brain MRI, whose 2D slices are circles. We compare our method to state-of-the-art algorithms.

Keywords. Topological Data Analysis · Segmentation · Computational methods in biology · Image processing

MSC2020 codes. 55N31 · 68-04 · 92-08 · 68U10

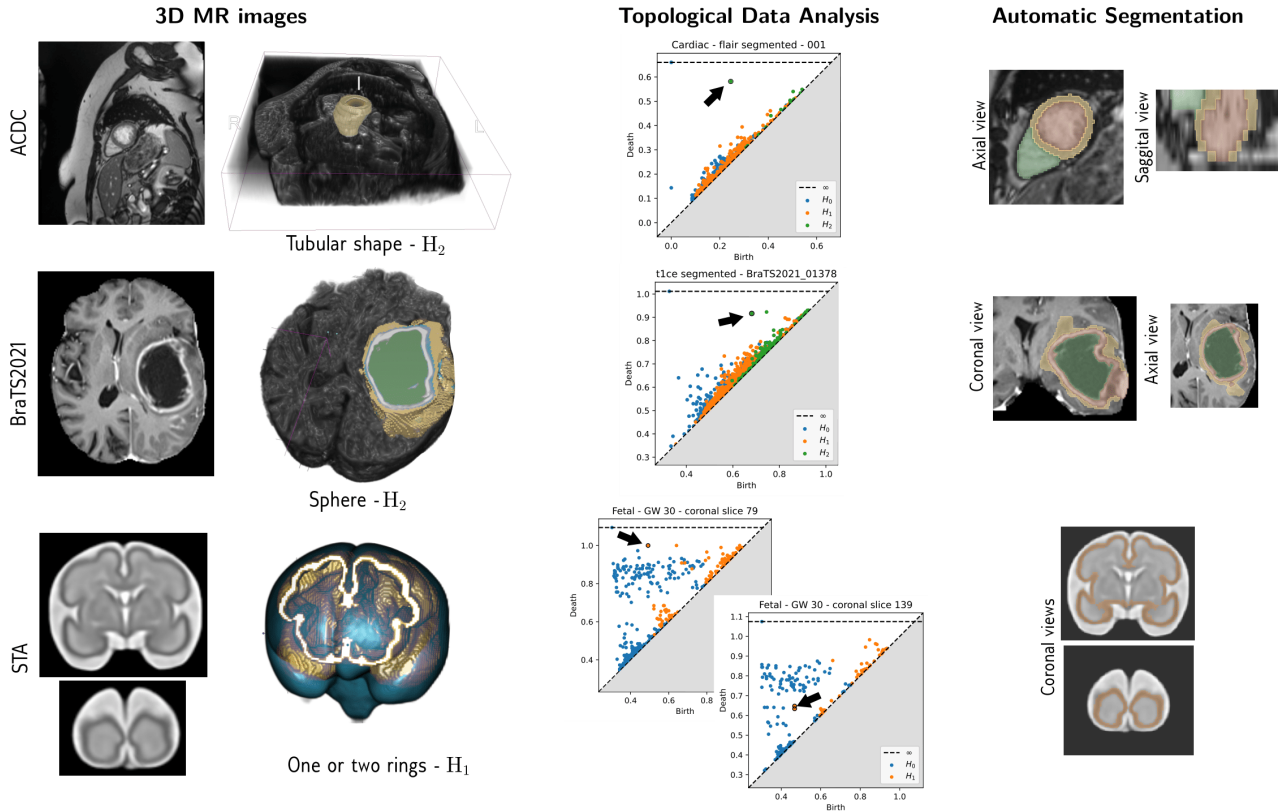


Fig. 1: TDA segmentation overview: We demonstrate a simple concept for segmentation using Topological Data Analysis (TDA) on three data sets. Each row focuses on a particular organ to segment, from top to bottom: the myocardium in ACDC, glioblastoma in BraTS 2021 and cortical plates in STA. In each case, we are provided with a 3D MRI and aim at detecting a connected component of a given topology. We automatically select components, as illustrated by the persistence diagram into the TDA column. Eventually, through the strategies detailed in this article, we deduce a segmentation.

1 Introduction

Motivation. Anatomical segmentation in Magnetic Resonance Imaging (MRI) refers to the process of identifying and separating different structures within an MRI scan of the body. It can be performed by a computer algorithm, or a human operator using specialised software. The algorithm or operator segments the scan into different regions of interest, based on differences in image intensity, shape and size.

A particularly well-studied case is that of *glioblastoma*, the most common brain tumour, whose segmentation is commonly divided into three classes: the primary tumour mass, its necrotic contour, and the infected tissues. Other important examples include cardiac segmentation, where the *myocardium* is to be separated from the ventricles, as well as *cortical plate* segmentation, the embryonic precursor of the cortex. Having access to such segmentations is essential for patient management, disease diagnosis, risk evaluation, as well as treatment decisions [1–3].

To this end, several methods have been proposed, including the U-Net architectures, that recently achieved excellent scores, quantified by the Sørensen-Dice coefficient [4–6]. These methods, however, come with known limitations: their training requires large annotated datasets, expansive to build, is prone to over-fitting, and is oblivious to certain geometric properties of the tissues, sometimes yielding anatomically impossible results.

Our contributions. In this paper, we convolve tools from Topological Data Analysis (TDA), namely persistent homology (PH), to devise a novel algorithm for segmentation in MRI. It is divided into three simple modules: (1) identification of the whole object to segment, (2) detection of a particular subset via TDA, and (3) deduction of the other classes of the segmentation. The second step involves prior morphological knowledge, presenting a topology that TDA can detect. We demonstrate our algorithm on three segmentation problems: glioblastoma from brain MRI, for which we take advantage of the spherical shape of the Enhancing Tumor; cardiac segmentation from Cardiovascular Magnetic Resonance Imaging (CMR), based on the cylindrical shape of the myocardium; and cortical plate from fetal brain MRI, using the fact that slices of cortical plate are composed of one or two circles.

Our algorithm addresses the common pitfalls of previously reported methods. First of all, state-of-the-art deep learning algorithms—namely CNNs and U-Nets—, sometimes yield anatomically unrealistic segmentations. This has been observed in glioblastoma [7] or cardiac segmentation [5, 8]. By leveraging tools from TDA, our algorithm is backed by theoretical guarantees, ensuring the topological quality of the output segmentations. Namely, the whole object is provably connected, and geometric regions exhibit the correct topology (spherical, circular).

We note that incorporating topological constraints into neural networks has already been considered, for instance in TopoNet and TopoCP [9, 10]. These methods, however, share the second issue faced by deep learning models: their training involves an annotated dataset, requiring a highly trained specialist and hours of process. Our method, on the other hand, needs no training set. It merely consists in fitting an explicit morphological model.

We draw the reader’s attention to the fact that, differently from other geometric methods, our algorithm is intrinsically topological. As a consequence, it adapts to many different morphologies and is less sensitive to certain data acquisition errors. To our knowledge, this is the first purely PH-based technique of segmentation, i.e., not involving deep learning. Moreover, since the algorithm consists of three simple steps, its outputs are highly understandable. This addresses the issue of interpretability, a crucial property of data-driven healthcare, allowing clinicians to make informed judgements [11].

In addition, our method has the potential to be generalized to different datasets and segmentation needs. It can be adapted as long as one of the regions of the segmentation presents a non-trivial homology, that this region appears hyper-intense or hypo-intense in the image, and that the other components are located either inside or outside the region.

In summary, our main contributions are:

- (i) The description of a novel TDA-based algorithm for segmentation from MRI scans, that does not require any training data and produces directly interpretable results.
- (ii) A study of the three modules composing the algorithm, and their generalization to different biomedical contexts.
- (iii) Its validation on three datasets: BraTS 2021 (glioblastoma), ACDC (cardiac) and STA (fetal), along with a comparison to state-of-the-art methods.

We stress that our method is not exempt from limitations. In particular, its accuracy relies on the assumption that the input MRI satisfies a topological model. For images that do not fall under the hypotheses of the given model, our method is bound to fail, and our results are significantly worse than other techniques. By formulating a precise definition of the model, we identified that a third of the BraTS database satisfies the hypotheses, with scores approaching state-of-the-art. Besides, the images of cardiac and fetal datasets consistently exhibit common characteristic features, allowing us to seamlessly identify an appropriate topological model. This results in good segmentation scores as is.

Outline. The remainder of this paper is organized as follows. Related works are described Section 2, both in the fields of AI and TDA. We give in Section 3 a concise introduction to persistent homology for images, allowing to describe our methods in Section 4. Experimental results are given in Section 5, where glioblastoma, cardiac and fetal segmentations are respectively studied in Sections 5.1, 5.2 and 5.3. We gather additional comments in Section 6 and conclude in Section 7.

Code and data availability. The code for this project is fully available on GitHub at https://github.com/antonfrancois/gliomaSegmentation_TDA for anyone to use and contribute to. We experiment on the dataset BraTS 2021, available under request from <http://braintumorsegmentation.org/>, as well as ACDC, available from <https://www.creatis.insa-lyon.fr/Challenge/acdc/>, and STA from <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/WE9JVR>. In addition, some video animations, illustrating PH on MRI, are gathered at https://www.youtube.com/playlist?list=PL_FkltNTklAxtaMrQB5UCbobxRsiw8sE.

2 Related works

Glioblastomas. In the domain of biomedical segmentation, a significant effort has been made regarding *glioblastomas*. It is the most common brain tumour, diffuse, of variable degree of aggressiveness, and whose medical prognosis is difficult to establish. Its segmentation involves three regions: the peritumoral Edema (ED), composed of invaded tissue; the Tumorous Core (TC), representing the primary tumour mass; and the Enhancing Tumor (ET), usually the surgical target, along with the necrotic portions of the tumour (see Fig. 2). The union of these classes is referred to as Whole Tumour (WT), the entirety of the disease.

Accurate segmentation of glioblastomas is important for several reasons. Firstly, it enables medical professionals to make informed treatment decisions, such as the choice of surgical intervention or radiation therapy, by providing a clear understanding of the size and location of the tumour. Secondly, it is a valuable tool for monitoring disease progression and evaluating the effectiveness of treatment over time. Finally, glioblastoma segmentation is essential for the development of computer-aided diagnosis systems, which have the potential to significantly improve the accuracy and efficiency of many medical imaging algorithms [1].

In this context, the *Center for Biomedical Image Computing & Analytics* of the Perelman School of Medicine run the Brain

Tumour Segmentation (BraTS) challenge for ten years. In this article, we exploit data from BraTS 2021 [12–14]. The dataset consists of 1251 MRIs, coming in four modalities: Native (T1), contrasted and enhanced (T1ce), T2-weighted (T2) and T2-fluid Attenuated Inversion Recovery (FLAIR), along with a “groundtruth” segmentation for each patient. All imaging volumes were segmented using the STAPLE fusion [15] of top-ranked BraTS 2020 algorithms [16–18], then refined manually by neuroradiology experts and approved by experienced neuroradiologists.

Although not systematically, the sub-regions of the segmentations present distinctive characteristics. ET exhibits increased signal intensity in T1ce modality, compared to healthy white matter; TC typically appears hypo-intense in T1ce relative to T1; and WT is typically distinguished by a hyper-intense signal in FLAIR. We note the rare exception of the astrocytomas (IDH-mutant, 1p19q non-codeleted) that present a T2/FLAIR mismatch [19]. As we will present in Section 3, a hyper-intense region (resp. hypo-intense) appears early in the sublevel sets of the image (resp. superlevel sets), hence WT and TC can be identified via threshold selection. This observation is at the basis of our method.




Fig. 2: Glioblastoma segmentations from BraTS 2021. Each row contains horizontal MRI slices of the same patient, in modalities T1, T1ce, T2 and FLAIR. TC in red, ET in orange and ED in blue.

Myocardium. In clinical cardiology, it is essential to measure cardiac functions, such as strokes volumes, mass of the ventricles, or thickness of the myocardium [2, 20]. These measurements are commonly obtained from segmented Cardiac Magnetic Resonance Images (CMR), which are multi-slice 2D MRIs. The segmentation task consists in three classes—myocardium (Myo), left ventricular cavity (LV) and right ventricle (RV) (see Fig. 3).

In this paper, we will use the Automated Cardiac Diagnosis Challenge (ACDC) dataset [5], part of 2017 MICCAI-ACDC challenge.¹ It gathered 150 patients with diverse cardiac profiles (normal subjects, systolic heart failure, dilated or hypertrophic cardiomyopathy). The dataset contains, for each patient, two CMR scans: one at end diastolic and one at end systolic phase. The ground truth segmentations, encompassing myocardium, LV and RV, were drawn and double-checked by two experts.




Fig. 3: Examples of coronal segmentation in the ACDC dataset. The first image is a slice of a CMR at end diastolic phase, and the second one the corresponding segmentation. Myocardium is represented in orange, LV in red and RV in blue.

It must be noted that CMR segmentation involves specific challenges. The images are inherently noisy, due to the motion of the heart, and present large variability from one patient to another. Moreover, the myocardium is surrounded by tissues of similar intensity, making its automatic segmentation delicate. This is visualized on Fig. 3. On the good side, the left and right ventricular cavities, through which blood flows, appear significantly luminous, allowing for a relatively direct detection in superlevel sets.

Cortical plate. Several neurological deficits—such as ventriculomegaly, responsible for schizophrenia, autism and epilepsy—can be detected during fetal development [3]. In particular, the correct maturation of the fetal brain can be observed through the *gyrification* of the Cortical Plate (CP), the embryonic precursor of cerebral cortex. From the tenth to the thirty-fifth gestational week, the CP changes from a smooth surface to a highly convoluted one, making its segmentation a difficult task (see Fig. 4).

Among the publicly available data, the Fetal Tissue Annotation and Segmentation Dataset (FeTA) [6], led by the University Children’s Hospital Zürich and the University of Zürich, gathered 50 manually segmented pathological and non-pathological fetal magnetic resonance brain volume reconstructions², across a range of gestational ages (20 to 33 weeks) into 7 different tissue categories (external cerebrospinal fluid, grey matter, white matter, ventricles, cerebellum, deep grey matter, and brainstem/spinal cord). A version of this dataset is part of MICCAI 2021 challenge.³

Another dataset is the Spatiotemporal Atlas (STA) of [21].⁴ It contains representations of the *average* fetal brain, at one-week intervals between 21 and 38 weeks gestational age. These images are obtained via diffeomorphic deformable registration of 81 T2-weighted MRI scans of healthy fetuses. As opposed to the clinical dataset FeTA, atlas data and segmentations are smoother and easier to work with, and will be employed in this paper.




Fig. 4: Examples of cortical plate segmentations in the STA dataset. The rows contain sagittal slices of the same MRI, for gestational week 21 (top), 30 (middle) and 38 (bottom).

AI-methods for segmentation. Early efforts in automatic segmentation employed hand-crafted features engineering, along with traditional machine learning methods. Among these, one can find *atlas-based* segmentations, which utilize a representative brain to propagate a segmentation on an unknown one [22]. Besides, *decision forests* allow to classifying each pixel in an image as belonging to one of the tissue types. It is employed by Zikic et al. [23] for the segmentation of high-grade gliomas from multi-channel MR images, by extracting tissue characteristic features in the different channels. In addition, *Conditional Random Fields* (CRFs) are a type of probabilistic graphical model that can be

²FeTA data <https://www.synapse.org/#!Synapse:syn23747212/wiki/608434>

³FeTA MICCAI data <https://www.synapse.org/#!Synapse:syn25649159/files/>

⁴STA data <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/WE9JVR>

¹ACDC data <https://www.creatis.insa-lyon.fr/Challenge/acdc/>.

used to model inter-pixel spatial relationships. It has been applied to brain tumours by Wu et al. [24].

The rise of GPU processing capabilities has led to a shift in focus towards deep learning for brain tumour segmentation. In practice, to address such challenges, a multi-level segmentation approach is often used. This involves first using a simpler method to segment the image and then using a deep learning method to refine the segmentation. For example, Islam et al. [25] propose a four-step multilevel segmentation method. First, they perform a common preprocessing/filtering phase. Then, they cluster the MRIs using k-means clustering. Next, they use an optimal threshold and watershed segmentation technique to segment the tumour. Finally, they use a morphological operation to refine the segmentation. Another approach consists in using a fusion of different modalities. For example, Zhang et al. [26] propose to fuse FLAIR and T2 images to improve the segmentation of the tumour boundaries. They then train a dense 2D-CNN using novel architectures and loss functions. Finally, the Liu et al. [27] introduce a method to solve the inter-class ambiguity problem in segmentation. This problem occurs when the tumor and the surrounding tissue have similar intensities. The authors propose to use a convolutional attention network (CANet) to address this problem.

Driven by the BraTS competition, a growing number of deep learning algorithms for tumour segmentation have been developed since 2014. The winners of the most recent competition have all employed such networks, which demonstrates the superior performance of this approach when more data is available. In fact, the eight selected teams of the BraTS 2021 challenge used deep-learning techniques for segmentation, the most efficient ones being the ones using nnU-Nets architectures. The winning team had Dice coefficients of 0.8835, 0.8878 and 0.9319 for the enhancing tumour (ET), the tumour core (TC), and the whole tumour (WT) respectively. The other teams got very similar scores (inter-study variance being of order 10^{-2}) [28–35].

Coronal and fetal segmentations are no exception: in the ACDC challenge, nine out of ten participants implemented a CNN, most of which U-Nets. At end diastolic phase, the best reported Dice scores for myocardium, LV and RV are respectively 0.902, 0.946 and 0.968, obtained with an ensemble of 2D and 3D U-Nets [36]. The results for end systole are comparable. Similarly, most of the twenty participants of the FeTA 2021 challenge proposed variants of the U-Net architecture [37]. The best average Dice score reported, over the seven classes, is 0.786. It is worth mentioning another dataset, presented in Dou et al. [38], consisting of MRI of 57 fetuses, scanned in the 16–39 weeks gestational week range, along with the manual segmentation of their cortical plate. The authors attain, via a CNN-based architecture, a score of 0.87.

If deep learning is becoming the norm, they still need big annotated data sets, computationally costly to train, and produce non-interpretable results—which is a problem in case of failure. Among the attempts to solve the latter problem, Topological Data Analysis (TDA) has been invested recently, as a mean to constrain the topology of the outputs, as reviewed below. However, such methods still rely on annotated data. The aim of this paper is to propose a purely TDA-based algorithm, free from neural networks, opening the door to solving both problems.

TDA for MRI analysis. Topological Data Analysis (TDA) is a field at the intersection of computational geometry, algebraic topology and data analysis. It aims at capturing relevant geometric and topological information from datasets. Since its emergence in the 2000s, it has been applied to a wide range of problems, from medicine, physics, computer vision and machine learning, among others [39–41]. We provide a short introduction to Persistent Homology (PH), TDA’s most popular tool, in Section 3, and refer the reader to [42, 43] for an extended presentation. As far as this section is concerned, it is enough to know that PH yields *persistence diagrams*, which are summaries of the homology of the image at different scales. They contain information relative to the topological features present in the image: connected components, 1-dimensional cycles, or sphere-shaped objects.

One can distinguish three types of application of TDA to MRI

analysis. A first application consists in designing segmentation by *constraining their topology*. This idea has been proposed first in the context of CMR data by Clough et al. [44, 45]. In addition to the usual loss used to train a U-Net, the authors add a *topological loss*, calculated using TDA. Based on a prior topological knowledge, the segmentation is constrained to be close to a pre-defined shape. For instance, in the problem of segmenting the myocardium and the left ventricle, the authors use the knowledge that the myocardium is ring-shaped, achieving a mean Dice scores of 0.899 and 0.906 in the ACDC dataset for myocardium at end-diastole and end-systole respectively. The same idea has been applied to the problem of placenta segmentation, knowing that it forms one connected component with no holes [7, 46].

Another sort of topological loss has been proposed by Hu et al. [9], under the name of TopoNet: instead of constraining the persistence diagram to an explicit prior knowledge, the neural network is trained to ensure that the persistence diagrams of the images and their groundtruths are close. In the task of neuronal membrane segmentation, where the first homology group H_1 (1-dimensional cycles) is relevant, they attain Dice scores comparable to those achieved by U-Net, while exhibiting significantly superior topological quality. This loss has been generalized to arbitrary homology dimension, and applied to cortical plate segmentation by de Dumast et al. [10, 47]. Their algorithm, called TopoCP, is trained on the FeTA dataset. Applied on the atlas images of STA, it yields a mean Dice score of 0.79 ± 0.05 , while a simple U-Net shows 0.77 ± 0.05 . Recent developments of this method include the use of Wasserstein distance to compare the persistence diagrams [48], yielding a score of 0.7383, 0.8005, 0.8793 for ET, TC and WT on BraTS 2019; the construction of a dictionary of typical persistence diagrams in the latent space [49]; the addition of topology-based weighting schemes [50]; and the joint regularization of the neural network through anatomical priors [51].

As a second application, TDA can be used to *identify the topology* of the components of the image. An example is given by Qaiser et al. [52, 53] in the context of colorectal cancer tumour segmentation, based on Hematoxylin and Eosin stained images. The authors determine which patches of the images exhibit a tumour via their persistent homology: infected patches correspond to those that contain more holes and connected components. This can be attributed to the fact that, in infected tissues, nuclei tend to have atypical characteristics, irregular shape and size.

It is interesting to observe that in all the articles cited so far, TDA is always used in conjunction with deep learning. This is easily understood: the persistence diagrams are only used as machine learning features, compared between each others, or used as is. This overlooks the fact that other information can be extracted: points in persistence diagrams can be matched with subsets of the image, through the notion of *representative cycles*. As already exploited in image analysis, mapping pixels with elements of the persistence diagram allows for purely TDA-based segmentation methods [54, 55], also applicable to point cloud data [56, 57]. Our algorithm leverages this idea, with the novelty that we do not only consider PH in degree 0, as it is the case in these works, but also the homology groups H_1 and H_2 (1- and 2-dimensional cycles).

As a last application of TDA, it can be a *feature for other regression of classification tasks*. In this context, no topological prior is known, and TDA is seen as an exploratory tool. When provided as an input to statistical models, topological features of brain MRI (such as Smooth Euler Characteristic Transform, persistent entropy, or Betti Curves) enable personalized diagnosis and prediction of clinical outcomes in glioblastoma [58–60] or Alzheimer’s disease [61]. In this context, the most common features are the Persistence Images, used in hepatic tumor classification [62–64] and lung cancer survival prediction [65, 66].

For completeness, we stress that TDA is not only applicable to images, but also to different structures, such as graphs, point clouds, or real-valued functions. For instance, in [67] is built a graph from FDG-PET brain scans, whose PH allows to identify patients with attention-deficit hyperactivity or autism disorder. Similar constructions have been used in the context of electroencephalography brains signals [68, 69] or brain arteries net-

work [70]. Last, it must be mentioned that TDA is a growing field, with new methods still under development. The Decorated Merge Trees, for instance, have been proposed recently in [71], including a preliminary experiment with glioblastoma segmentations.

3 Background on TDA

This section serves as a self-contained introduction to persistent homology (PH), the most popular technique of TDA. It is a theoretical framework that allows inferring the homology groups of a dataset [72–74]. We will focus on *cubical* persistent homology, the incarnation of PH for image data, and use brain MRIs as illustrations. We refer the reader to [75] for a thorough presentation of homology, and [42, 43] for an extended introduction to TDA.

Homology. Many theories of homology exist, and we will consider here a particular one: the singular homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$, the finite field with two elements (0 and 1). It associates to any topological space X a sequence of $\mathbb{Z}/2\mathbb{Z}$ -vector spaces $H_0(X)$, $H_1(X)$, $H_2(X)$, \dots , called *singular homology groups*. Instead of presenting their construction, which is based on a combination of topology and algebra, we will explain roughly what they represent. Given $i \in \mathbb{N}$, and if the i^{th} homology group $H_i(X)$ is finitely generated, which we will suppose, then it admits a finite dimension, denoted $\dim H_i(X)$, called the i^{th} *Betti number* of X . It carries topological information about X :

- $\dim H_0(X)$ is the number of connected components of X .
- $\dim H_1(X)$ is the number of “independent loops” in X . It is equal to 1 for the circle, and 0 for the sphere.
- $\dim H_2(X)$ is the number of “independent voids” in X . It is equal to 0 for the circle, and 1 for the sphere.

Except for dimension 0, this list should not be treated as a formal mathematical result, but only as heuristic interpretations.

In order to use homology in the context of images, one has to transform them into topological spaces. This is easily done when the image is *binary*, that is when it only has black and white pixels. Let $I : \Omega \rightarrow \{0, 1\}$ be a binary image with domain Ω , and consider the collection $I^{-1}(\{1\}) \subset \Omega$ of its black pixels. Seen as a subset of Ω , this collection can be seen as a topological space, hence its homology is well-defined. As an example, we represent in Fig. 5 a 2-dimensional image, whose homology groups satisfy $\dim H_0(X) = 3$ and $\dim H_1(X) = 2$.




Fig. 5: A 2d binary image, consisting of three connected components and two “holes”.

Filtrations. The aim of this work is to analyse MRIs. These are 3-dimensional greyscale images, not binary, and there is no obvious way to convert them canonically into topological spaces. Therefore, homology theory cannot be used directly. In order to circumvent this issue, we can follow the pipeline of persistent homology. The idea consists in building not one but a collection of topological spaces, called a *filtration*. It is an increasing family of binary images, denoted $\{I^t \mid t \in [0, 1]\}$, indexed by a parameter $t \in [0, 1]$. Given a greyscale image $I : \Omega \rightarrow [0, 1]$, various popular filtrations can be defined, such as the height, radial and density filtrations, implemented in *giotto-tda*, or the sublevel and superlevel sets filtrations, in *cubical ripser* and *gudhi* [76–78]. Our work is concerned with the last two filtrations.

For any $t \in [0, 1]$, let us define I^t as the set of pixels with intensity at most t . It is a topological space—a union of cubes—and we have the relation $I^s \subset I^t$ for any $s, t \in [0, 1]$ such that $s \leq t$. The family $\{I^t \mid t \in [0, 1]\}$ is called the *sublevel sets filtration*. Note that, since the image has values in $[0, 1]$, the parameter t parses all the possible values a pixel can take, and the last image, I^1 , is

equal to the whole domain Ω . Similarly, if we define I^t as the set of pixels with intensity at least $1 - t$, then we obtain the *superlevel sets filtration*. Both filtrations are illustrated in Fig. 6, where I is the T2 modality of a MRI of a healthy brain.

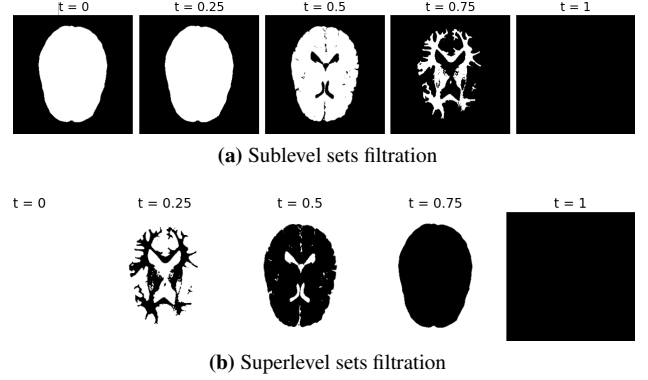


Fig. 6: Examples of filtrations on a healthy brain MRI.

Within a filtration, one can understand the parameter t as a temporal value. The more t increases, the more pixels are added to the image. In the case of the sublevel sets filtration, pixels of low intensity are added first. In Fig. 6a, these are the pixels of the background (filtration value $t = 0$). On the contrary, in a superlevel sets filtration, pixels of high intensity are added first. In Fig. 6b, we see that this corresponds to the ventricles and the grey matter. In other words, the most luminous parts of the image appear the earliest in filtration. This idea will be used in Section 4 when devising a segmentation method for glioblastoma. Indeed, on a FLAIR modality MRI, glioblastoma tends to be represented by the pixels of the highest intensity, hence we expect the tumour to be the first element to appear in the filtration. An illustration of this phenomenon is given in Fig. 7 for a brain presenting a tumour.




Fig. 7: Superlevel sets filtration on a FLAIR MRI with tumour. One observes that, at the value $t = 0.3$, only the tumour is present.

Persistent homology. The main objects of TDA, the *persistence modules*, are obtained by applying the i^{th} homology functor to a filtration $\{I^t \mid t \in [0, 1]\}$. This yields family of vector spaces $\{H_i(I^t) \mid t \in [0, 1]\}$. It is a purely algebraic object, which gathers the homology groups of I at various scales. Actually, in a persistence module, we can extract more information than just the homology groups. Using the inclusion maps $I^s \hookrightarrow I^t$ for $s \leq t$, we can “track” the evolution of the homological features. That is, we are able to tell whether a cycle of $H_i(I^s)$ is still alive in $H_i(I^t)$. The interval on which a feature exists is called its *persistence*. In practice, one interprets cycles of large persistence as relevant features of the dataset, and cycles of low persistence as noise.

The persistence of all the cycles is recorded in the *persistence diagram* of the persistence module. It is a set of points of the form $p = (t_b, t_d)$, with $t_b \leq t_d$, interpreted as a “homological feature” born at time t_b and dead at time t_d (see Fig. 8). To each point of the persistence diagram corresponds a *birth* pixel p_b , which gives birth to a cycle (a new connected component, a new H_1 cycle, etc.), and a *death* pixel p_d , which kills the cycle (merge the component to another one, fill the H_1 cycle, etc.). In particular, in the H_0 persistence diagram, there always is a point that “dies at infinity”: it represents the connected component of all pixels.

There exists a duality between the persistence diagrams of the sublevel and superlevel sets filtration of an image, allowing to transform one into the other, via an explicit process [79]. Hence

both contain the same information. In what follows, we choose to work with the superlevel sets filtration, which is easier to interpret.

Representative cycles. Each point $p = (t_b, t_d)$ of the diagram represents a cycle, born at t_b and killed at t_d . In all the images I' , with $t_b \leq t < t_d$, one is able to define a *representative cycle*, that is, a subset that “identifies” the cycle. Formally, it is a singular chain whose image in $H_i(I')$ is the cycle. In the case of H_0 , the representative cycle is unique, and simply is the connected component of the pixel x_b that gave birth to p . We will denote it $CC^{t_b}(x_b)$.

However, in higher degrees H_i , $i \geq 1$, a particular difficulty arises: the cycles are not uniquely localised. That is, several sets of pixels may represent the same cycle. Moreover, their identification is a challenging computational task [80]. In this paper, instead of computing a representative cycle, we will consider, as a proxy, the connected component $CC^{t_b}(x_b)$ of the birth pixel. We defer to Section 6 a discussion regarding potential extensions of this strategy.

As an illustration, we provide in Fig. 8 the persistence diagram of the superlevel sets filtrations of a 2D slice of a healthy brain. We employ the SRI24 Atlas as a normal human brain anatomy standard reference [81]. The H_0 -cycles are represented as red dots, and H_1 -cycles as blue dots. We stress that, in the following section, we will consider the persistent homology of the full 3D image, and not only of its slices. We choose here to analyse a 2-dimensional slice for visualisation ease.




Fig. 8: Persistence diagram of the superlevel sets filtration of a 2D slice of the SRI template (middle), identification of the top H_0 -cycles (left) and representatives of the top H_1 -cycles (right).

One identifies on the figure three red dots significantly distant from the diagonal. They correspond to connected components, that evolve independently, without merging into each other. For each of these persistent cycles, we consider the corresponding point $p = (t_b, t_d)$ of the diagram, extract the pixel of birth, and plot its connected component at time t_d . It represents the component just before it merges with another one. We see that these components are part of the grey matter, disconnected in this slice.

Similarly, four blue dots stand away from the diagonal, with one point particularly off. In order to represent the corresponding holes, we circled them on the figure, choosing arbitrary representatives. We see that they correspond to the lateral and third ventricles (again, disconnected in this slice).

4 Methods

We now describe our algorithm for segmentation in MRI. We start with a brief overview, showing its ability to adapt to different segmentation challenges. The following sections present the three modules that make up the algorithm in greater detail. For the sake of simplicity, we detail them on the glioblastoma problem only.

4.1 Overview

Model. Since our algorithm is based on TDA, it is worth giving the topological insights which support the method. It makes use of two sorts of information: the intensity of the components to segment, and their shape. The first piece of information is exploited through the fact that hyper-intense (resp. hypo-intense) regions can be found in the superlevel sets (resp. sublevel sets) of the image. As for the shape, it is used to detect the object on the persistence diagram, via its homology.

Our method can be applied as long as there exists one component of the segmentation, referred to as a *geometric object*, whose homology is non-trivial, and which separates the components from each other. Additionally, the union of all the components called whole object, must appear hypo- or hyper-intense in the image. More precisely, we assume that the ground truth segmentation satisfies the following hypotheses:

- (H1) The whole object is connected, has no holes, and appears luminous in (a certain modality of) the image.
- (H2) The geometric object has non-trivial homology, and appears significantly luminous in (a certain modality of) the image.
- (H3) Other components lie inside or outside the geometric object.

We deliberately formulate the model in a rather non-rigorous way, since it only serves, for the moment, to guide intuition. The hypotheses will be made more precise in Section 5 for each dataset. In the case of glioblastoma, we will quantify how many images fall under them. As far as cardiac segmentation is concerned, we will see that (H1) is not satisfied by the whole object but only some of its components. We will adjust the algorithm accordingly. Last, for the cortical plate, only (H2) will be used.

The three datasets considered in this paper have been presented in Section 2. The glioblastoma (collection BraTS 2021), cardiac (collection ACDC) and fetal segmentations (atlas SAT) respectively consist of 3, 3 and 1 components. We stress that BraTS 2021 is the only dataset for which more than one modality is provided (we make use of FLAIR and T1ce). Although not systematically satisfied, we found that the components regularly exhibit common characteristics, gathered in Table 1. We therefore choose, for the role of the geometric object, the component TC in glioblastoma, Myo in cardiac, and CP in fetal brains segmentation.

	TC	ET	ED
BraTS	hyper (FLAIR)	hyper (FLAIR)	hyper (FLAIR)
	hypo (T1ce)	hyper (T1ce) sphere-shaped	hypo (T1ce)
ACDC	Myo	LV	RV
	hypo-intense tubular-shaped	hyper-intense	hyper-intense
SAT	CP		
	hypo-intense sphere-shaped in 3D circular in slices		

Table 1: Characteristics shared by most of the images, for each dataset and each component of the segmentation. Columns indicate the component label, its typical intensity (hypo- or hyper-intense), and its expected topology (in blue), if any.

Structure of the algorithm. It is naturally split into three steps, which will be further explained in the next sections. Let Ω be the domain of the MRI, $I: \Omega \rightarrow [0, 1]$ the image, and denote by X_{whole} and $X_{\text{geometric}} \subset \Omega$ the whole object and the geometric object.

Module 1: Identification of the whole object (Section 4.2) According to (H1), the image shows a hyper-intense area, corresponding to the union of the segmentation components. Consequently, X_{whole} can be found in a superlevel set I^t of the image. The parameter is identified through an analysis of the curve $t \mapsto \#I^t$ of the number of active voxels at time t .

Module 2: Detection of the geometric object (Section 4.3) Following (H2), the component $X_{\text{geometric}}$ of the segmentation has a distinctive homology (that of a sphere, cylinder or circle), hence corresponds to a persistent H_2 -cycle in the persistence diagram. Thus, we compute the persistent homology of the superlevel sets filtration of I restricted to the whole object X_{whole} obtained above and select the most persistent cycle (in H_2 or H_1). By denoting p_b and t_b as the voxel and time of birth of this cycle, we define $X_{\text{geometric}}$ as its connected component $\text{CC}^{t_b}(p_b)$. Illustrations are given in Fig. 1.

Module 3: Deduction of the other components (Section 4.4) To obtain final segmentation, we consider the whole object X_{whole} , from which we remove $X_{\text{geometric}}$. This binary image can be partitioned into connected components. We identify the other components according to whether they belong inside or outside $X_{\text{geometric}}$, as suggested by the hypothesis (H3).

Notations. In order to detail precisely these modules, we choose, in the next sections, to focus on the glioblastomas only. We remind the reader that, in this case, two modalities of MRI are employed, and the corresponding images will be denoted I_{FLAIR} and I_{T1ce} : $\Omega \rightarrow [0, 1]$. As suggested in Table 1, I_{FLAIR} will be used in Module 1, and I_{T1ce} in Module 2. The geometric object is, in this case, the component ET. The three components of the segmentation are denoted X_{ET} , X_{TC} and X_{ED} . They are subsets of Ω . By *whole tumour* or *global segmentation*, we refer to the subset $X_{\text{WT}} = X_{\text{ET}} \cup X_{\text{TC}} \cup X_{\text{ED}}$. Throughout these sections, we will use as running examples a few MRIs from the collection BraTS 2021.

4.2 Module 1: Identification of whole object

Selection of the superlevel set. In this first step, we wish to select the largest hyper-intense region present in the FLAIR image, supposedly corresponding to X_{WT} . To achieve this, we analyse the number of voxels in the filtration over time, starting from $t = 1$ and moving to $t = 0$. While the number of voxels increases steadily, we anticipate a sharp increase, precisely when the voxels corresponding to the white and grey matter get included. By selecting a t just before this event, we obtain a reliable estimation of X_{WT} .

More precisely, let $t \mapsto \#I_{\text{FLAIR}}^t$ be the map that indicates the number of voxels of intensity not lower than t , and $t \mapsto d\#I_{\text{FLAIR}}^t$ is its derivative, computed by finite difference. Moreover, we fix a positive real number `dt_threshold`, treated as a parameter. We then identify t as the first value, starting from 1, for which $d\#I_{\text{FLAIR}}^t$ exceeds `dt_threshold`. Once the optimal value of t has been selected, we define X_{WT} as the largest connected component of I_{FLAIR}^t (i.e., that containing the most voxels). We draw the reader's attention to the fact that this procedure is reminiscent of classical binarization algorithms, such as Otsu's method.

The procedure is illustrated in Fig. 9, where are represented the map $t \mapsto \#I_{\text{FLAIR}}^t$ and its derivative, alongside the corresponding brightest components. On the right, one sees that the selected component is close to the ground truth, with a Dice score of 0.9523. In this experiment, we set the value of `dt_threshold` to be the area under the curve. This choice, although arbitrary, yields good results, and will be used throughout the rest of this article. Potential improvements will be discussed in Section 6.

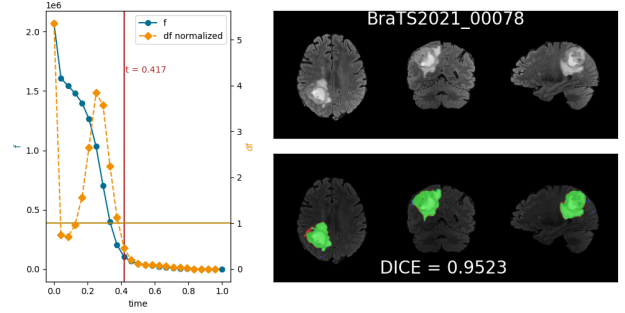


Fig. 9: Identification of X_{WT} via Module 1. Left: the map $t \mapsto \#I_{\text{FLAIR}}^t$ in blue and its derivative in orange, with axes located respectively on the left and right side. Right: Superposition of the estimated segmentation with the ground truth. Green areas correspond to a correct match, red areas to false positives (voxels identified but not in the ground truth segmentation), and blue to false negatives (ground truth voxels not selected).

Refinement: filling the holes. We perform a post-processing step on the segmentation X_{WT} , which consists of filling its holes, as suggested by (H1) of our model. More precisely, this is achieved by computing the connected components of the binary image $\Omega \setminus X_{\text{WT}}$, discarding the background, and adding them to X . In practice, we use the function `ndimage.binary_fill_holes` of `scipy` [82]. From a biomedical point of view, we observed that the presence of holes in X_{WT} is often caused by the necrosis, which occasionally appears darker on FLAIR images.

4.3 Module 2: Detection of geometric object

For the second step, we will use the image I_{T1ce} , as well as X_{WT} computed previously, to obtain the enhancing tumour X_{ET} . According to (H2), this component is the boundary of the tumorous core and is highly intense in I_{T1ce} . Hence, using a superlevel sets filtration, we expect to see a sphere, formed by the boundary of the tumour, and represented by a cycle in the H_2 -persistence diagram.

More precisely, this procedure is automatized as follows. First, we compute the persistent homology of the superlevel sets filtration of the image I_{T1ce} restricted to X_{WT} . Then, we select the H_2 -feature of highest persistence, that is, the point (t_b, t_d) of the diagram that maximizes $|t_d - t_b|$. Let $x_b \in \Omega$ be the voxel that gave birth to it. Following our strategy outlined in Section 3, we identify X_{ET} by taking the connected component of x_b in the binary image $I_{\text{T1ce}}^{t_b}$. We draw the reader's attention to the fact that this component may not be a representative cycle of the homology class—it only contains one. We have found that this strategy, which is particularly easy to compute, gives correct results.

An example is given in Fig. 10, still using the same data as in Fig. 2. On the diagram, one green point appears particularly far away from the diagonal: it is the persistent cycle we are looking for. The resulting segmentation can be seen on the left, presenting a Dice score of 0.9047. In practice, we compute persistent homology with `cubical_ripser` [77], for it appeared to run the fastest among TDA libraries for 3D images.




Fig. 10: Persistence diagram of I_{T1ce} restricted to X_{WT} , with top H_2 -cycle circled (right), and the resulting segmentation X_{ET} (left).

4.4 Module 3: Deduction of components

We now aim to identify the components X_{TC} and X_{ED} . This last step does not depend on the initial MRI, but only on the subsets X_{WT} and X_{ET} estimated previously.

Following the hypothesis (H3), TC corresponds to the part of the tumour that lies *inside* ET, and ED to the part that lies *outside* ET (and still within X_{WT}). In order to identify these parts, we consider the subset $X_{WT} \setminus X_{ET}$ —the complementary of ET in the whole segmentation—, and compute its connected components. Note that we may have more than two connected components. The outer component is identified as that containing the background, and its restriction to X_{WT} is saved in X_{TC} . The others are considered inner and are added to X_{ED} .

The resulting segmentation of our running example is visualized in Fig. 11. This is a successful case, presenting the Dice scores of 0.94, 0.90 and 0.89 for EC, ET and ED respectively.

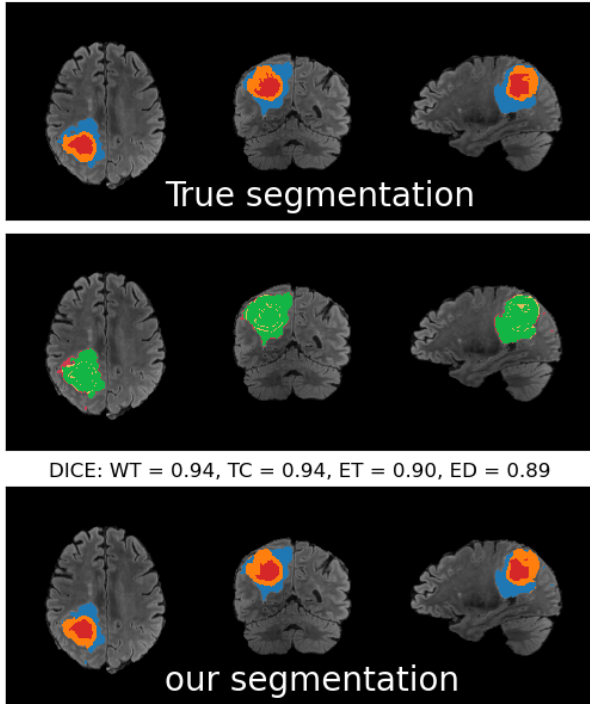


Fig. 11: BraTS2021 full segmentation. On the middle image, green means that the segmentation is correct, orange means that the tissue was mislabeled but is part of the tumour, and in red the tissues that should not have been segmented.

5 Results

This section is concerned with the application of our algorithm, described in Section 4, to three segmentation challenges. We start with glioblastoma from the BraTS 2021 database (Section 5.1), for which the algorithm can be used as is. We then turn to cardiac segmentation in the ACDC dataset (Section 5.2), which calls for a slight adaptation of Module 1. Two versions are proposed: either segmenting the myocardium based on the whole 3D CMR, or working slice by slice. Last, we segment the fetal plate in the STA atlas (Section 5.3), considering only a 2D approach.

5.1 Glioblastoma segmentation

Dataset. For our first segmentation task, we focus on glioblastomas, employing the BraTS 2021 dataset presented in Section 2. The segmentation is made of three classes (TC, ED and ET), whose union is referred to as WT. To perform the segmentation, we use our algorithm described in Section 4.

As a pre-processing step, we applied to all the images a 0-1 normalization. In addition, we found that best results were obtained with a Gaussian blur of standard deviation $\sigma = 1$, and a binary dilation with a 3D ball of radius 3. We used the implementations `gaussian_filter` and `grey_erosion` of `scipy.ndimage`.

The state-of-the-art methods and the winners of BraTS2021 segmentation challenges, are all U-Net [28–35]. We trained a conventional U-Net model on all four modalities available in the BraTS2021 dataset and achieved consistent scores compared to those reported in the literature when restricted to the test set. We recognize that evaluating a neural network on the same dataset it was trained on may not provide an accurate representation of its true performance. Nevertheless, we consider this a suitable benchmark for its peak performance. While this approach may not fully favour our Persistence Homology solution, we consider the U-Net scores as an upper bound on the accuracy achievable by these types of methods. Moreover, it’s crucial to note that we will be comparing our traditional algorithm, which utilizes only two modalities (FLAIR and T1ce), against the U-Net, which leverages four. The aforementioned U-net consists of four Encoder and Decoder layers, inputting and outputting data from a bottleneck layer. Each encoder layer consists of two 3D convolution layers, followed by instance normalization. A Leaky ReLU activation is applied, and 3D max-pooling is utilized to decrease spatial resolution. Each decoder layer begins with upsampling, followed by two 3D convolution layers, instance normalization, and Leaky ReLU activation. If a residual connection is feasible, features from the corresponding encoder are concatenated before convolution.

Scores. We applied our algorithm to the whole collection of MRIs of BraTS 2021, consisting of 1251 images. For each image, we computed the Dice coefficient for each segmentation class (TW, TC, ED and ET), between our result and the segmentation provided by the specialists. The results are gathered in Fig. 12. We recall that the Dice coefficient between two binary images $X, Y : \Omega \rightarrow \{0, 1\}$ is defined as

$$\text{Dice}(X, Y) = \frac{2\#(X \cap Y)}{\#X + \#Y}.$$

By inspecting the boxplots, we observe that the results seem correct for the whole segmentation (mean value approximately 0.711 ± 0.283) but relatively poor for TC, ED and ET (respectively 0.369 ± 0.358 , 0.477 ± 0.306 and 0.457 ± 0.305). In comparison, the scores of U-Net are respectively 0.868 ± 0.123 , 0.73 ± 0.282 , 0.821 ± 0.154 and 0.815 ± 0.202 .

One can take a closer look at the images where the algorithm performs poorly in Fig. 13b, which reveals low scores are mainly attained by images that do not satisfy the model we introduced in Section 4. We observed that either the enhancing tumour does not surround the necrosis, or only partially, forming a perforated sphere. In both cases, the algorithm cannot partition the domain into the interior and exterior of ET, leading subsequently to an incorrect estimation of the other components of the segmentation. We examine this question further below.

Validation of the model. In order to evaluate the importance of the morphological model, we restricted the analysis to only a subset of images, those that satisfy the Hypotheses (H1) to (H3) described in Section 4. Formally, we consider images such that:

- (H1') X_{TW} consists of one connected component, or potentially more, the other ones being 10 times smaller. Moreover, X_{TW} is at most 50 times larger than X_{TC} . Also, the most intense voxel of the whole tumour in FLAIR belongs to X_{TC} or X_{ET} .
- (H2') After 3 binary dilations, X_{ET} divides the space into two connected components. Moreover, the most intense voxel of the whole tumour in T1ce belongs to X_{ET} .
- (H3') Applying a binary dilatation to X_{TC} (resp. X_{ED}) yields new pixels of which at least (resp. at most) half belongs to X_{ET} .

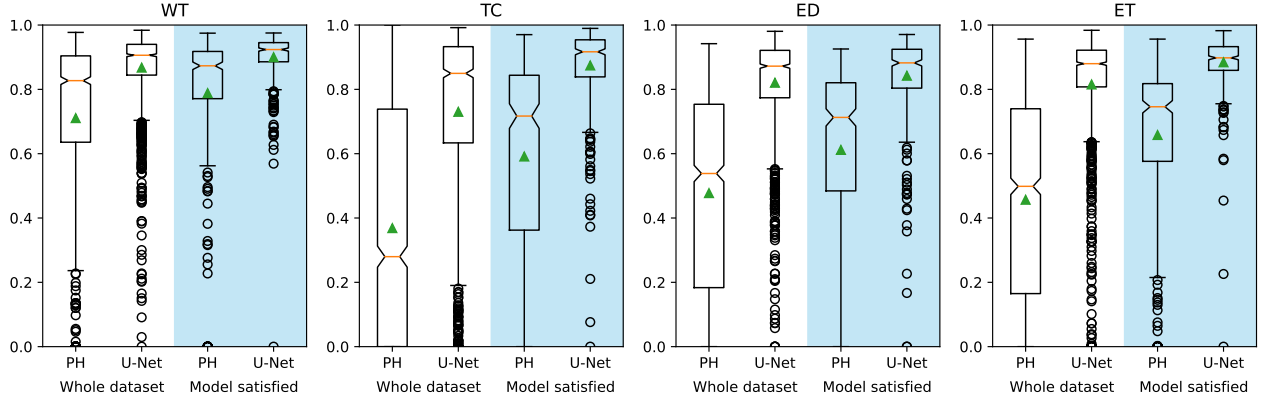


Fig. 12: Boxplots of the Dice coefficients on segmentations of the BraTS 2021 dataset, for the four regions WT, TC, ET and ED, in four scenarios: respectively for the whole dataset (1251 MRIs) and segmentations of our method (PH) and U-net, and on the restricted dataset satisfying the model (390 MRIs, 31% of the dataset) with our method and U-net.

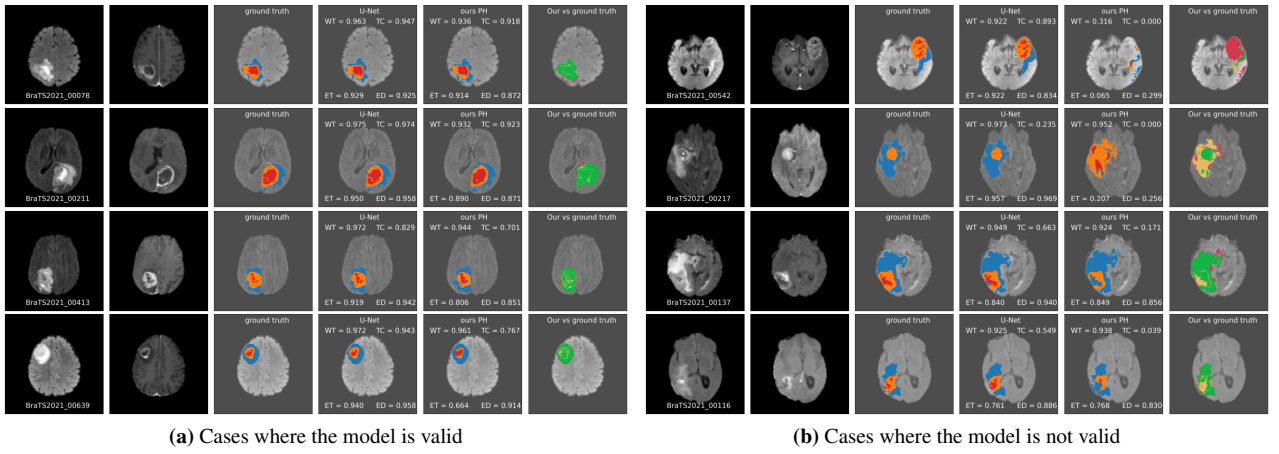


Fig. 13: Results of our algorithm, in cases where the model is valid or not. On each row, from left to right: an image of modality FLAIR, of modality T1ce, the segmentation provided by the specialists and the segmentation we obtained with both U-Net and the PH method along with a comparison (green indicating well segmented, orange for mislabeled and red otherwise).

We point out that the values of 10 and 50 in (H1') are arbitrary, chosen to represent a non-significant component. Besides, the 3 binary dilations in (H2') have been chosen so as to avoid the effects of "thin-edged" tumors, that form an open sphere. This point will be discussed further in Section 6 (see Fig. 19). Last, the hypothesis (H3') is to be understood as weakening of the properties of being "inside" or "outside" of X_{ET} . Indeed, a connected component within ET is expected to have all its boundary voxels in contact with ET, while a connected component outside of ET is expected to have at least half its boundary voxels out of it.

We computed that 390 out of 1251 images satisfy the assumptions (31.2%). The Dice coefficients of the segmentations restricted to these images are presented in Fig. 12. One sees that, in this case, segmentations are of higher quality: the mean scores for the WT, TC, ED and ET are respectively 0.788 ± 0.229 , 0.591 ± 0.317 , 0.612 ± 0.267 and 0.658 ± 0.241 . In comparison, U-Net attains 0.9 ± 0.08 , 0.874 ± 0.13 , 0.842 ± 0.133 and 0.884 ± 0.087 . In conclusion, our method, while not reaching U-net scores, comes close to a subset. Furthermore, the segmentations obtained are, qualitatively, of good quality, as we study below. It is worth noting a significant discrepancy between the mean and median values, equal to 0.873, 0.717, 0.713 and 0.746 respectively.

Qualitative evaluation. Overall, most of the output produced by our method yields biologically plausible segmentations, as exemplified in Fig. 13a. Notably, in the second row, our method accurately captured fine details of the tumour outline, and in all

four examples, all labels were correctly identified.

Several cases can cause our algorithm to fail, all falling outside the scope of our hypotheses. Examples of these cases are shown in Fig. 13b. In the first row, the tumour in the FLAIR image is not hyper-intense. Therefore, the selection of the entire tumour failed, resulting in a very poor final segmentation. In the three other rows, the first step was successful, but the labelling was incorrect. In the second row, the hyper-intense component is not a sphere but a ball (with no hole), so Module 2 selected the wrong component. In the third row, the ET was correctly selected, but the produced segmentation was open, lacking a topological interior or exterior, causing Module 3 to fail. Finally, in the last row, the T1ce image is hypo-intense at the tumour location, making it difficult to select ET. As can be seen, our method and U-Net produced similar predictions for ET. This error, like in the third row, caused Module 3 to fail.

While U-Net frequently attains higher Dice scores, our method still delivers very comparable results. When comparing the mean Dice scores of both approaches, we observe that our method outperforms or matches U-Net in 47.56% (595 out of 1251 cases) across the entire dataset, with the whole-tumour (WT) segmentation performing better than individual labelling. (Our method surpasses U-Net for WT 410 times, ET 147 times, TC 207 times, and ED 96 times out of 1251 images. The mean of the mean Dice score for these images is 0.83 ± 0.14).

On the data subset where images satisfy our hypotheses, the frequency at which our method outperforms or is equivalent to U-Net on at least one label rises to 37.79% (147 out of 390). No-

tably, our method outperforms or is equivalent to U-Net 100 times for WT, 61 times for ET, 63 times for TC, and 38 times for ED. These promising results suggest that our method is approaching the precision level of a U-Net model.

We can also report a couple of occurrences where a part of the brain was occulted for some reason. In these cases, while the shape of the inferred segment was plausible, the outlines were leaking, while our method gave a satisfactory response. Moreover, similarly to what was mentioned in [12, §2.1.5], both U-Net and PH commit several common errors. For example, as the vessels within the edematous area are hyperintense, they end up labelled ET or ED when they are close to the tumoural region. Also, white matter hyperintensities are sometimes identified as tumours, as can be seen in Fig. 13b(first row).

5.2 Cardiac segmentation

Dataset. We now turn to the task of segmentation of myocardium (Myo) and left/right ventricles (LV and RV) from CMR data, based on the ACDC dataset [5]. It is a collection of 300 CMR scans, two for each of the 150 patients, at end-systolic (ES) and end-diastolic (ED) phases. In the context of the MICCAI-ACDC 2017 challenge, 100 patients were reserved for training the models, and 50 for testing. We point out that U-net architectures show very good scores in this context. The best team in the challenge obtained 0.899 as the lowest Dice score, over all three classes and the two end phases [36]. We aim to investigate whether our persistent homology-based methodology also performs well. Since our algorithm requires no training, we will only consider the 100 testing images (two for each patient).

As a pre-processing step, we applied to all the images a 0-1 normalization and a Gaussian blur ($\sigma = 1.5$ in 2D, and 2.5 in 3D).

Adaptation of the method in 2D. In order to adapt our algorithm in Section 4 to cardiac segmentation, two specific issues must be addressed. First, the whole object—union of the myocardium and the two ventricles—presents heterogeneous intensity and hence cannot be selected as a whole via Module 1. Besides, the myocardium forms a cylinder, which is topologically equivalent to a circle. Therefore, detecting the homology of a cylinder in the persistence diagram will not yield global information. In order to solve these issues, we will start by devising a 2-dimensional version of our algorithm, i.e., working slice by slice. In slices, the myocardium forms a circle, and only the first issue remains. We shall study a 3D solution in the next paragraph.

Given a slice $I: \Omega \rightarrow [0, 1]$, and as already visualized in Fig. 3, both ventricles appear hyper-intense, and the myocardium hypointense. Moreover, the latter is surrounded by structures that are equally weak in intensity, making it impossible to estimate the union myocardium/ventricles by a simple selection of superlevel sets. In order to circumvent this issue, we propose to first detect LV, then RV, and eventually deduce the myocardium.

To start, the left ventricle is to be found as a connected component of a superlevel set of the slice. However, a number of other objects are hyper-intense, hence are also present in these level sets. We thus propose a localized form of Module 1: given a pixel $x \in \Omega$ with intensity $I(x)$, we consider $t \in [0, I(x)] \mapsto \#CC^t(x)$, the map that computes the number of pixels in the connected component of x in the binary image I^t . Just as in Section 4.2, the derivative $t \mapsto d\#CC^t(x)$ is expected to present a peak precisely when the whole object has been formed, and starts to connect to others. Thus we define t as the last value in $[0, I(x)]$ such that $d\#CC^t(x)$ is greater to the parameter `dt_threshold`. Last, we associate to x the binary image $CC^t(x)$.

In order to apply this localized variation of Module 1, we are left with detecting a pixel x that belongs to LV. This is done by computing the H_0 -persistence diagram of the superlevel sets of I , and selecting the most persistent N points, where N is chosen equal to 10 in our experiments. Next, for each of these N points we consider the pixel that gives birth to the topological feature, and compute the corresponding connected component $CC^t(x)$.

We eventually select the the component that is the most “disk-shaped”. Indeed, as visualized in Fig. 14, slices of the left ventricle tend to be spherical. In practice, we measure sphericity as follows: we consider the minimal enclosing disk D of $CC^t(x)$, and compute the Dice similarity between these binary images. Our segmentation of LV is chosen as the component $CC^t(x_{LV})$ that minimizes this quantity.

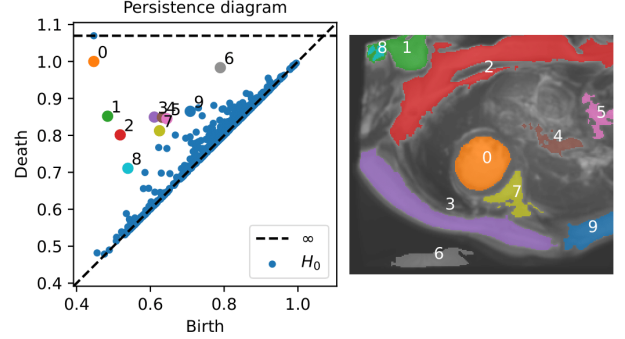


Fig. 14: Visualization of segmentation of LV. For each of the ten most persistent H_0 -features of the slice, we estimate the corresponding connected component via Localized Module 1, and select the one that is the most disk-shaped (label 0 here). The closest component, labelled 7, is used as a segmentation of RV.

Once LV has been segmented, we find RV as the connected component, obtained by the localized Module 1, generated by the birth pixel x_{RV} whose connected component $CC^t(x_{RV})$ is the closest to that of LV. As a last step, and in order to also cover the myocardium, we dilate the component of LV until it reaches RV. We will use the union of these components as our segmentation of the whole object.

Based on the segmentation of the union, one applies Module 2 as is, selecting the most persistent cycle in the H_1 -persistence diagram, whose connected component is associated to the myocardium. Module 3 is employed to conclude the segmentation, defining the interior as LC and the exterior as RV. Before presenting the scores we obtained this way, we propose another strategy, taking into account the 3-dimensional image directly.

Adaptation of the method in 3D. The first step of the method described above directly translates in the 3-dimensional setting: by defining I as the whole CRM, and not a slice, one considers the H_0 -persistence diagram of its sublevel sets, select a few top points of the diagram, compute the corresponding connected components of their birth pixels, and identify LV as the component the most “cylindrical”. For the sake of simplicity, we measure the cylindricity of a binary image via the mean value of the sphericity of its slices, where sphericity has been defined in the previous paragraph. Just as it is the case in dimension 2, the most cylindrical component corresponds, most of the time, to the groundtruth location of LV. As before, RV is segmented as the component whose birth pixel is the closest to that of LV, and the whole object is obtained by dilating LV until reaching RV.

Next, when we come to apply Module 2, we face an issue. The myocardium, that we wish to extract through persistent homology, has the homotopy type of an (open) cylinder. However, from a homological point of view, a cylinder have trivial group H_2 . Capturing the cycle in H_1 , which has dimension 1, would not help to identify the cylinder, since a representative of this cycle could only be a 1-dimensional subset of it. This issue is solved by artificially adding a slice of pixels of value 1 at the top and the bottom of the CMR, thus transforming the cylinder into a sphere. Module 2 is applied to select the most persisting feature in the sublevel sets of I , restricted to the union segmentation, whose representative connected component is our segmentation of the myocardium. The ventricles are then identified via Module 3.




Fig. 15: Boxplots of the Dice coefficients on segmentations of the ACDC dataset, for the four regions: whole object, myocardium, left ventricle (LV) and right ventricle (RV), using our 2- our 3-dimensional algorithm, and restricted to end diastolic (ED) or end systolic (ES) phases. Furthermore, “oracle step 1” denotes the same methods, but fed as an input the exact segmentation of the whole object.

Scores. As one sees from the Dice scores gathered in Fig. 15, the 2D version of the algorithm performs better than 3D. For instance, at end-diastolic phase, we obtain, with 2D, the scores 0.701 ± 0.156 , 0.486 ± 0.156 , 0.565 ± 0.221 and 0.752 ± 0.157 for the classes whole object, Myo, LV and RV respectively. In contrast, the 3D version yields 0.484 ± 0.201 , 0.332 ± 0.175 , 0.324 ± 0.229 and 0.478 ± 0.235 . This discrepancy could be explained by the fact that CMR images are not exactly imaging methods favourable to 3D analysis. Indeed, resolution along the vertical axis is low (10 axial slices). We will clarify this point in the next paragraph.

We note that, for both methods, mean scores in ES are lower than ED. For instance, with the 2D algorithm, the scores on end-systolic images are 0.528 ± 0.297 , 0.29 ± 0.221 , 0.331 ± 0.329 and 0.62 ± 0.324 . This phenomenon is also observed in the literature, for U-Net methods. As a matter of fact, the most difficult class to segment, RV at ES, is considered as a delicate task, even for experimented observers [5].

In addition to the scores of our algorithm, Fig. 15 contains, under the name “oracle step 1”, the result of our method when the whole object is given as an input. In other words, these results are not representative of what we can achieve right now, but the best results that can be expected, if we improved Module 1, or use an other efficient method. With the 2D version, remarkable scores are obtained. At ED, one reads 0.805 ± 0.069 , 0.898 ± 0.103 and 0.918 ± 0.041 for Myo, LV and RV; at ES, it is 0.82 ± 0.053 , 0.852 ± 0.127 and 0.832 ± 0.095 . These prospective results, which, it should be remembered, were only obtained using topological ideas and without training, highlight the potential of TDA for biomedical image segmentation.

Validation of the model. In the context of the ACDC dataset, we propose to reformulate the Hypotheses (H1) to (H3) as follows:

- (H1') The whole object consists of one connected component.
- (H2') X_{myo} divides the space into two connected components.
- (H3') The inner component of X_{myo} contains X_{LV} .

These hypotheses can be either understood axial slice by axial slice for the 2D version of the algorithm, or globally, for the 3D version. As it turns out, among the 300 CMR images of the whole collection (both phases ED and ET considered), 81.33% verify the hypotheses in 2D, and only 11% in 3D. As a matter of fact, most of the images do not satisfy (H2') in 3D. Fig. 16 illustrates the issue: the vertical resolution is so low that “gaps” appear when moving from one axial slice to another. As a correction, one can apply a binary dilation to X_{myo} . For instance, if chosen of radius 3, then the proportion of images satisfying (H2') raises to 69%. However, applying numerous dilations reduces segmentation precision. In our experiments, we chose not to apply dilations.




Fig. 16: **Left:** Superposition of the groundtruth segmentation of the myocardium in two consecutive axial slices (red and pink) at ES phase. **Right:** Several coronal slices, with myocardium in red.

5.3 Fetal cortical plate segmentation

Dataset. Our third segmentation challenge is that of cortical plate in fetal brain MRI. As discussed in Section 2, popular datasets include FeTA and STA [6, 21]. The latter does not contain actual MRI but average images, obtained by diffeomorphic deformable registration. They are consequently smoother than the originals, and we choose to work with those for simplicity. The STA collection consists of a 3D image for each gestational week between 21 and 38. In this context, the task consists in identifying the cortical plate only—it is a segmentation with a single class.

Up to our knowledge, two articles used STA as a test dataset for segmentation of cortical plate: TopoCP [10], announcing a mean Dice score of 0.70, followed by an improvement of 0.79 in [47]. In both cases, authors train a neural network on the FeTA dataset, and use STA as a way of measuring its generalization on new data. We will use the latter score as a reference.

In our experiments, we applied to each image a simple preparation, consisting of a 0-1 normalization, a Gaussian blur of standard deviation $\sigma = 0.5$, as well as a dilation of radius 1.

Adaptation of the method in 2D. The cortical plate being the only class to identify, we expect to only use Module 2 of our method. However, as visualized on Fig. 1, it forms a perforated sphere, opened at the level of the cerebellum. Consequently, the homology groups of the cortical plate are all trivial, and it cannot be detected with our method. In order to circumvent this issue, one could modify the image to produce the apparition of homology—for instance by “closing” the sphere, as we did with the myocardium in 3D. Instead, we suggest a different strategy: we study the image slice-by-slice, in the coronal plane.

As it can be observed on Fig. 4, slices are of three types: the

cortical plate forms either (i) one circle, (ii) two disjoint circles or (iii) an open circle. These cases are identified by inspection of the persistence diagram, as follows.

- We consider the sublevel sets filtration of the slice, and compute its H_1 -persistent homology.
- Next, we discard points of the diagram that correspond to implausible segmentations. This is done by identifying the corresponding connected component, and checking that the hole it encloses has a number of pixels included in $[N_{\min}, N_{\max}]$. In practice, we found that the values 25% and 75% of the slice are suitable.
- Last, we select the earliest born from the remaining points. We identify that the slice belongs to type (ii) if there exists another point of the diagram at distance at most ϵ , where $\epsilon = 0.03$ is a fixed threshold, and type (i) otherwise. Slices of type (iii) are treated as if they were of first type.

The segmentation is eventually defined as the connected component corresponding to the selected point(s) of the diagram: one point for type (i) and (iii), and two points in type (ii). As described in Module 2, the components are defined as that of the birth pixel at birth time.

Scores. Over the whole STA collection of 18 images, and compared with the ground-truth segmentation, we obtained a mean Dice score of 0.69. Fig. 17 contains the scores as a function of the gestational age (in weeks). One observes that scores decline over time, corresponding to the fetal plate becoming increasingly coiled. This phenomenon is observed in [47, Fig. 11]. Compared with the score of 0.79 obtained in the aforementioned article, one could consider our scores as correct.




Fig. 17: Evolution of the Dice scores for segmentation of cortical plate on the STA dataset, as a function of the gestational week.

Validation of the model. Since our task of fetal brain segmentation only involves one class, the cortical plate, only the second hypothesis of our model in Section 4.1 is relevant. Remembering that, in coronal slices, the cortical plate is expected to form one or two circles, we propose the following reformulation:

- (H2') In a cortical slice, the fetal plate divides the plane into several connected components. After removing the background and the components of cardinal lower than a hundredth of the slice, exactly one or two components remain.

We computed that, averaging over the whole collection of 18 images, 56.95% of the non-empty cortical slices verify the hypothesis. As a matter of fact, the faulty slices can be divided into two types, visualized in Fig. 18. Either the cortical plate consists of several convex components, as it is the case in the top and bottom slices, or it forms an open circle, as observed in the middle slices. Although not studied further in our work, special attention could be paid to these slices, in order to improve the scores.

6 Discussion and potential improvements

Representative cycle identification. Let $I : \Omega \rightarrow [0, 1]$ be an image, and $X \subset \Omega$ a geometric object. In our experiments on




Fig. 18: Cortical slices of a fetal brain from the STA atlas, with segmentation of cortical plate indicated in orange. One sees that in the first, middle and last slices, the CP has trivial H_1 -homology: it forms convex components, or it is an open circle.

glioblastoma or myocardium 3D, X is a sphere, and in the case of myocardium 2D or fetal plate 2D, it is a circle. At the core of our approach, in Module 2, we suppose that X corresponds to a point in the i^{th} persistent diagram of I , where $i = 1$ or 2 . Say $p = (t_b, t_d)$ is the point of the persistence diagram, and p_b is its birth pixel. In order to go back from the diagram to the image, we identify X as $CC^{t_b}(p_b)$, the connected component of p_b in the superlevel set I^{t_b} . However, this method does not guarantee that $CC^{t_b}(p_b)$ has the correct homology, i.e., that it represents a singular chain whose homology class in $H_i(I^{t_b})$ is the selected homology class. In particular, it may not be a sphere or a circle. Indeed, $CC^{t_b}(p_b)$ only contains the singular chain as a subset. In this article, we've settled for this method, for the sake of simplicity, and because it produced convincing results.

As a matter of fact, identifying a *representative cycle* is a current problem in TDA, that faces two difficulties, as reviewed in [80]. First, representative cycles are not uniquely defined. It is thus common to add an extra condition, such as minimizing the number of 1-simplices [83, 84], the volume (number of "enclosed" 2-simplices) [85], the diameter of the chain [86], or minimal with respect to a fixed lexicographic order on the simplices [87]. Secondly, finding optimal representative cycles is a computationally expansive task. From a formal point of view, it has been shown that certain formulations of the problem are NP-complete [86, 88]. In practice, it is time-consuming, on the order of several hours for reasonably large data [85, Table 1]. Up to our knowledge, only two available libraries implement this task: OptiPersLP⁵ for edge-minimal cycles [84], and homcloud⁶ for volume-minimal cycles [89]. Including these computations will be the object of future works.

Preprocessing. Presently, in our three experiments, we opted for a simple pre-processing strategy, including: an affine 0-1 normalization, a slight Gaussian blur and a dilation of the image. In our research efforts, however, a comprehensive exploration and testing were conducted to identify the optimal preprocessing pipeline. Various methods, including simple, local, and patch-wise normalization, pixel-wise equalization and noise reduction techniques (such as SUSAN and Non-Local Means), as well as intensity enhancement approaches, were rigorously examined. Surprisingly, these methods not only failed to enhance the results but, in some cases, even worsened them. We believe that the issue is twofold, the first reason lying in the significant variability of the images. Indeed, these preprocessing steps enhanced segmentation scores only for a certain number of images. Further research is needed to identify an algorithmic criterion to decide on the relevance of applying a specific preprocessing step. For example, increasing the contrast of an image may help better select the edges of a connected component corresponding to a tumour. However, in some cases, the tumour exhibits an intensity gradient that results in the absence of well-defined edges, leading to the improper selection of the threshold t in Module 1, and erroneous segmentation. The second problem, when it comes to pre-processing data, is that TDA can react very differently to certain transformations, as discussed below.

First, it is worth noting that monotone bijective transformations of the pixels have little effect on TDA. Indeed, given such a map $f : [0, 1] \rightarrow [0, 1]$, and an image $I : \Omega \rightarrow [0, 1]$, then the persistence

⁵OptiPersLP <https://bitbucket.org/remere/optiperslp/>

⁶homcloud <https://homcloud.dev/>

diagram $\text{Diag}(f \circ I)$ of the transformed image $f \circ I$, for sublevel or superlevel sets filtrations, is equal to the transformed diagram

$$f \circ \text{Diag}(I) = \{(f(t_b), f(t_d)) \mid (t_b, t_d) \in \text{Diag}(I)\}.$$

In particular, points of large persistence remain so, as long as f does not disturb too much the values. Clearly, it is the case for 0-1 normalization. This remark also holds for our Module 1, since it uses, as unique information, the number of pixels in level sets.

On the opposite, TDA is highly sensitive to local transformations, such as blurring or enhancement. This is because topological features—such as connected components, or higher dimensional cycles—can be formed or destroyed with the addition or deletion of a pixel only. This can be seen from the stability theorem of persistent homology [90]: given two images $I, J : \Omega \rightarrow [0, 1]$, the bottleneck distance between their persistence diagrams is upper bounded by the *sup norm* between the images:

$$d_{\text{bottleneck}}(\text{Diag}(I), \text{Diag}(J)) \leq \|I - J\|_{\infty}.$$

Unlike the Euclidean norm between images, the sup norm can be significantly altered by a single pixel.

On the bright side, this also means that these local transformations can “help” TDA to identify geometric objects. We exemplify this idea in Fig. 19, using a slice of an MRI from the BraTS 2021 dataset. We consider the T1ce modality and restrict the image to the segmentation of the whole tumour. As provided in the dataset, the image is noisy, thus the corresponding persistence diagram contains many points. Applying a light Gaussian blur (standard deviation $\sigma = 0.5$) allows for the number of features: a few connected components (blue points), provoked by isolated bright pixels, disappear, and H_1 -cycles (orange points) get filled earlier. However, the circle drawn by the tumour’s contour (ET) is not seen clearly on the diagram. This is because the tumour is open—one sees dark pixels at the bottom left—, thus the whole circle only forms late in the filtration. This is remedied via a dilatation (of radius 3), allowing to “close” the tumour, and resulting in the detachment of a point from the noise, in H_1 .




Fig. 19: Effect of local transformations on the superlevel sets persistent homology of a MRI. **Left:** A slice of a T1ce-modality MRI from BraTS 2021 is restricted to the whole tumour, and its persistence is computed. **Middle:** Applying a Gaussian blur allows to reduce the number of points in the diagram, that were artifacts of the noise. **Right:** After dilation, the tumour contour is closed, and a H_1 -persistent cycle appears far away from the diagonal.

Improvement of Module 1. Our first module really is a simple automatic thresholding method, reminiscent of Otsu’s. The difference lies in the fact that Otsu’s method binarizes the image by thresholding it at a single intensity value, and outputs the binary image as is. This can result in many small-intensity islands, which are not desirable in our applications. Our method, on the opposite, selects in this binary image the largest connected component.

We point out that Module 1 has the disadvantage of setting the threshold by using global information, while what we seek is a precisely localized object. Indeed, as described in Section 4.2, the threshold is defined as the first peak of the curve of number of pixels in the level sets. In the case of glioblastoma, which are consistently the brightest component in FLAIR modality MRI, this approach was sufficient (see Fig. 9). Cardiac MRIs, however, contain several components of similar brightness (see Fig. 14). We circumvented this issue in Section 5.2 by defining a Localized Module 1, fixing a pixel $x \in \Omega$ as an additional input, and counting the number of pixels in the component of x only. In order to find the correct input pixel, we chose the birth pixels of the most persistent H_0 cycles, and selected the one that yielded, via Localized Module 1, the most “spherical” connected component.

Among the improvements that could be implemented in Module 1, one thinks of using information other than just the number of pixels. In the BraTS collection in particular, for which our method yields a Dice score of 0.711 (see Fig. 12), we have observed that several tumours display a heterogeneous intensity, with border blending with the surrounding tissue. In this context, an improvement could consist in first estimating the boundary of the tumour, via edge detection techniques, and secondly selecting a threshold consistent with the edges. Initial attempts have shown, however, that tumour boundaries are not regular, and can have large gradient differences from one side to the other.

Refining the models. Module 1, when applied on glioblastoma or myocardium segmentation, is purely homological: it selects in the persistence diagram the most persistent H_2 -feature, and returns its connected component. However, in the context of fetal plate, our experiments have shown that this procedure is insufficient; we finally opted for the feature that encloses the most pixels. Arguably, this sort of geometric information could be incorporated in the first two problems.

On brain MRI, for instance, one could think of selecting H_2 -features based on their adequacy with a geometric model of the tumour, in the manner of a template registration. Such a registrations techniques have been applied successfully to other structures of the brain [91]. However, one would soon come up against the problem of variability of the tumors: they come in different shapes and sizes. As a matter of fact, only 25% of the BraTS 2021 database satisfies the simple model we proposed in Section 5.1. This calls for the creation of an atlas of tumour topologies, enabling homological profiling, and adaptation of the model for each case.

Combination with U-Net. Although this paper argued for a segmentation algorithm free of neural networks, both technologies may be combined. We have reviewed in Section 2 how TDA has already been incorporated in neural networks, through “topological losses”. They either work by matching the obtained persistence diagram with a reference diagram [9, 45], or by computing statistics from the diagram [53]. We remind the reader that, in addition to the diagram itself, our approach employs the localization of the topological features inside the image. We envision introducing this spatial information in a neural network in two ways.

First, we remind the reader that our algorithm relies, through Module 2, on the selection of a persistent cycle, from which we deduce the connected component of the geometric object. This cycle is selected as that maximizing the persistence, in the case of glioblastoma and cardiac, or that maximizing the inner volume in fetal brains. However, and as it is visualized in Figs. 14 and 19, the obtained persistence diagrams often consists of many points, with several away from the diagonal, making the selection of a point a delicate task.

For the second idea, one could think of introducing our estimation of components in the architecture of a U-Net itself. This would have the benefit of a warm-start for the training. In addition, it would as well constrain the segmentations to prescribed topologies, an important feature, already attempted in the above-mentioned works.

7 Conclusion

This study explored further the potential of Topological Data Analysis for medical image segmentation, by utilizing not only persistence diagrams information but also the cycle localisations. Capitalizing on its train-free nature and interpretable results, our modular algorithm design facilitates integration into existing pipelines and adapts to diverse applications.

Through detailed examples on different datasets, we used that the segmentation problem can be decomposed into two subtasks: organ localization and labellisation. We demonstrated that our segmentations, with a relatively simple topological model, can in some cases, achieve the precision of a U-Net.

For data with relatively low variability, such as myocardium segmentation, a single TDA-based module can achieve comparable performance to state-of-the-art methods. However, for more complex and variable data, such as cancer segmentation, a hybrid approach combining TDA with statistical analysis may be necessary.

Overall, our work highlights the versatility and potential impact of TDA in medical image segmentation, paving the way for further development and clinical applications.

References

- [1] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- [2] Perry M Elliott, Aris Anastasakis, Michael A Borger, Martin Borggrefe, Franco Cecchi, Phillippe Charron, Albert Alain Hagege, Antoine Lafont, Giuseppe Limongelli, Heiko Mahrholdt, et al. 2014 esc guidelines on diagnosis and management of hypertrophic cardiomyopathy. *Kardiologia Polska (Polish Heart Journal)*, 72(11):1054–1126, 2014.
- [3] Antonios Makropoulos, Serena J Counsell, and Daniel Rueckert. A review on automatic fetal and neonatal brain mri segmentation. *NeuroImage*, 170:231–248, 2018.
- [4] Dmitry Lachinov, Evgeny Vasiliev, and Vadim Turlapov. Glioma segmentation with cascaded UNet. In *International MICCAI Brainlesion Workshop*, pages 189–198. Springer, 2018.
- [5] Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging*, 37(11):2514–2525, 2018.
- [6] Kelly Payette, Priscille de Dumast, Hamza Kebiri, Ivan Ezhov, Johannes C Paetzold, Suprosanna Shit, Asim Iqbal, Romesa Khan, Raimund Kottke, Patrice Grethen, et al. An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific data*, 8(1):167, 2021.
- [7] Nick Byrne, James R Clough, Giovanni Montana, and Andrew P King. A persistent homology-based topological loss function for multi-class cnn segmentation of cardiac mri. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 3–13. Springer, 2021.
- [8] Nathan Painchaud, Youssef Skandarani, Thierry Judge, Olivier Bernard, Alain Lalande, and Pierre-Marc Jodoin. Cardiac segmentation with strong anatomical guarantees. *IEEE transactions on medical imaging*, 39(11):3703–3713, 2020.
- [9] Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. *Advances in neural information processing systems*, 32, 2019.
- [10] Priscille de Dumast, Hamza Kebiri, Chirine Atat, Vincent Dunet, Mériam Koob, and Meritxell Bach Cuadra. Segmentation of the cortical plate in fetal brain mri with a topological loss. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis: 3rd International Workshop, UNSURE 2021, and 6th International Workshop, PIPPI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, October 1, 2021, Proceedings 3*, pages 200–209. Springer, 2021.
- [11] Colleen E Charlton, Michael TC Poon, Paul M Brennan, and Jacques D Fleuriot. Development of prediction models for one-year brain tumour survival using machine learning: a comparison of accuracy and interpretability. *Computer methods and programs in biomedicine*, 233:107482, 2023.
- [12] Ujjwal Baid and et al. The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv:2107.02314*, 2021.
- [13] Bjoern H. Menze and et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015.
- [14] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific Data*, 4(1), 2017.
- [15] Simon K Warfield, Kelly H Zou, and William M Wells. Simultaneous truth and performance level estimation (staple): an algorithm for the validation of image segmentation. *IEEE transactions on medical imaging*, 23(7):903–921, 2004.
- [16] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.
- [17] Richard McKinley, Raphael Meier, and Roland Wiest. Ensembles of densely-connected cnns with label-uncertainty for brain tumor segmentation. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part II 4*, pages 456–465. Springer, 2019.
- [18] Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis*, 36:61–78, 2017.
- [19] Shoichi Deguchi, Takuma Oishi, Koichi Mitsuya, Yuko Kakuda, Masahiro Endo, Takashi Sugino, and Nakamasa Hayashi. Clinicopathological analysis of t2-flair mismatch sign in lower-grade gliomas. *Scientific reports*, 10(1):1–6, 2020.
- [20] Christopher A Miller, Peter Jordan, Alex Borg, Rachel Argyle, David Clark, Keith Pearce, and Matthias Schmitt. Quantification of left ventricular indices from ssfp cine

- imaging: impact of real-world variability in analysis methodology and utility of geometric modeling. *Journal of Magnetic Resonance Imaging*, 37(5):1213–1222, 2013.
- [21] Ali Gholipour, Caitlin K Rollins, Clemente Velasco-Annis, Abdelhakim Ouaham, Alireza Akhondi-Asl, Onur Afacan, Cynthia M Ortinau, Sean Clancy, Catherine Limperopoulos, Edward Yang, et al. A normative spatiotemporal mri atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific reports*, 7(1):476, 2017.
- [22] Stefan Bauer, Christof Seiler, Thibaut Bardyn, Philippe Buechler, and Mauricio Reyes. Atlas-based segmentation of brain tumor images using a markov random field-based tumor growth model and non-rigid registration. In *2010 annual international conference of the IEEE engineering in medicine and biology*, pages 4080–4083. IEEE, 2010.
- [23] Darko Zikic, Ben Glocker, Ender Konukoglu, Antonio Criminisi, Cagatay Demiralp, Jamie Shotton, Owen M Thomas, Tilak Das, Raj Jena, and Stephen J Price. Decision forests for tissue-specific segmentation of high-grade gliomas in multi-channel mr. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2012: 15th International Conference, Nice, France, October 1-5, 2012, Proceedings, Part III 15*, pages 369–376. Springer, 2012.
- [24] Wei Wu, Albert YC Chen, Liang Zhao, and Jason J Corso. Brain tumor detection and segmentation in a crf (conditional random fields) framework with pixel-pairwise affinity and superpixel-level features. *International journal of computer assisted radiology and surgery*, 9:241–253, 2014.
- [25] Rafiqul Islam, Shah Imran, Md Ashikuzzaman, and Md Munim Ali Khan. Detection and classification of brain tumor based on multilevel segmentation with convolutional neural network. *Journal of Biomedical Science and Engineering*, 13(4):45–53, 2020.
- [26] Xiaobing Zhang, Yin Hu, Wen Chen, Gang Huang, and Shengdong Nie. 3d brain glioma segmentation in mri through integrating multiple densely connected 2d convolutional neural networks. *Journal of Zhejiang University-SCIENCE B*, 22(6):462–475, 2021.
- [27] Zhihua Liu, Lei Tong, Long Chen, Feixiang Zhou, Zheheng Jiang, Qianni Zhang, Yinhai Wang, Caifeng Shan, Ling Li, and Huiyu Zhou. Canet: Context aware network for brain glioma segmentation. *IEEE Transactions on Medical Imaging*, 40(7):1763–1777, 2021.
- [28] Huan Minh Luu and Sung-Hong Park. Extending nn-unet for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 173–186. Springer, 2022.
- [29] Yading Yuan. Evaluating scale attention network for automatic brain tumor segmentation with large multi-parametric mri database. In *International MICCAI Brainlesion Workshop*, pages 42–53. Springer, 2022.
- [30] Michał Futrega, Alexandre Milesi, Michał Marcinkiewicz, and Pablo Ribalta. Optimized u-net for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 15–29. Springer, 2022.
- [31] Md Mahfuzur Rahman Siddiquee and Andriy Myronenko. Redundancy reduction in semantic segmentation of 3d brain tumor mris. *arXiv preprint arXiv:2111.00742*, 2021.
- [32] Jun Ma and Jianan Chen. Nnnet with region-based training and loss ensembles for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 421–430. Springer, 2022.
- [33] Krzysztof Kotowski, Szymon Adamski, Bartosz Machura, Lukasz Zarudzki, and Jakub Nalepa. Coupling nnu-nets with expert knowledge for accurate brain tumor segmentation from mri. In *International MICCAI Brainlesion Workshop*, pages 197–209. Springer, 2022.
- [34] Jianxun Ren, Wei Zhang, Ning An, Qingyu Hu, Youjia Zhang, and Ying Zhou. Ensemble outperforms single models in brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 451–462. Springer, 2022.
- [35] Haozhe Jia, Chao Bai, Weidong Cai, Heng Huang, and Yong Xia. Hnf-netv2 for brain tumor segmentation using multimodal mr imaging. *arXiv preprint arXiv:2202.05268*, 2022.
- [36] Fabian Isensee, Paul F Jaeger, Peter M Full, Ivo Wolf, Sandy Engelhardt, and Klaus H Maier-Hein. Automatic cardiac disease assessment on cine-mri via time-series segmentation and domain specific features. In *Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges: 8th International Workshop, STACOM 2017, Held in Conjunction with MICCAI 2017, Quebec City, Canada, September 10-14, 2017, Revised Selected Papers 8*, pages 120–129. Springer, 2018.
- [37] Kelly Payette, Hongwei Bran Li, Priscille de Dumast, Roxane Licandro, Hui Ji, Md Mahfuzur Rahman Siddiquee, Daguang Xu, Andriy Myronenko, Hao Liu, Yuchen Pei, et al. Fetal brain tissue annotation and segmentation challenge results. *Medical Image Analysis*, 88:102833, 2023.
- [38] Haoran Dou, Davood Karimi, Caitlin K Rollins, Cynthia M Ortinau, Lana Vasung, Clemente Velasco-Annis, Abdelhakim Ouaham, Xin Yang, Dong Ni, and Ali Gholipour. A deep attentive convolutional neural network for automatic cortical plate segmentation in fetal mri. *IEEE transactions on medical imaging*, 40(4):1123–1133, 2020.
- [39] Steve Y Oudot. Persistence theory: from quiver representations to data analysis. *Mathematical Surveys and Monographs*, 209:218, 2015.
- [40] Yara Skaf and Reinhard Laubenbacher. Topological data analysis in biomedicine: A review. *Journal of Biomedical Informatics*, 130:104082, 2022.
- [41] Yashbir Singh, Colleen M Farrelly, Quincy A Hathaway, Tim Leiner, Jaidip Jagtap, Gunnar E Carlsson, and Bradley J Erickson. Topological data analysis in medical imaging: current state of the art. *Insights into Imaging*, 14(1):1–10, 2023.
- [42] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
- [43] Frédéric Chazal and Bertrand Michel. An introduction to Topological Data Analysis: fundamental and practical aspects for data scientists. *Frontiers in Artificial Intelligence*, 4, 2021.
- [44] James R Clough, Ilkay Oksuz, Nicholas Byrne, Julia A Schnabel, and Andrew P King. Explicit topological priors for deep-learning based image segmentation using persistent homology. In *International Conference on Information Processing in Medical Imaging*, pages 16–28. Springer, 2019.
- [45] James Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A. Zimmer, Julia A. Schnabel, and Andrew King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE TPAMI*, 2020.
- [46] James Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A Zimmer, Julia A Schnabel, and Andrew King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

- [47] Priscille de Dumast, Hamza Kebiri, Vincent Dunet, Mériam Koob, and Meritxell Bach Cuadra. Multi-dimensional topological loss for cortical plate segmentation in fetal brain mri. *arXiv preprint arXiv:2208.07566*, 2022.
- [48] Andac Demir, Elie Massaad, and Bulent Kiziltan. Topology-aware focal loss for 3d image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 580–589, 2023.
- [49] Ainkaran Santhirasekaram, Mathias Winkler, Andrea Rockall, and Ben Glocker. Topology preserving compositionality for robust medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 543–552, 2023.
- [50] Song Sun, Yonghuai Wang, Jinzhu Yang, Yong Feng, Lingzhi Tang, Shuo Liu, and Hongxia Ning. Topology-sensitive weighting model for myocardial segmentation. *Computers in Biology and Medicine*, 165:107286, 2023.
- [51] Liu Li, Qiang Ma, Zeju Li, Cheng Ouyang, Weitong Zhang, Anthony Price, Vanessa Kyriakopoulou, Lucilio C Grande, Antonis Makropoulos, Joseph Hajnal, et al. Fetal cortex segmentation with topology and thickness loss constraints. In *Workshop on the Ethical and Philosophical Issues in Medical Imaging*, pages 123–133. Springer, 2022.
- [52] Talha Qaiser, Korsuk Sirinukunwattana, Kazuaki Nakane, Yee-Wah Tsang, David Epstein, and Nasir Rajpoot. Persistent homology for fast tumor segmentation in whole slide histology images. *Procedia Computer Science*, 90:119–124, 2016.
- [53] Talha Qaiser, Yee-Wah Tsang, Daiki Taniyama, Naoya Sakamoto, Kazuaki Nakane, David Epstein, and Nasir Rajpoot. Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features. *Medical image analysis*, 55:1–14, 2019.
- [54] William J Beks and Nikolaos Papanikolopoulos. 3d region segmentation using topological persistence. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1079–1084. IEEE, 2016.
- [55] Rabih Assaf, Alban Goupil, Valeriu Vrabie, Thomas Boudier, and Mohammad Kacim. Persistent homology for object segmentation in multidimensional grayscale images. *Pattern Recognition Letters*, 112:277–284, 2018.
- [56] Primoz Skraba, Maks Ovsjanikov, Frederic Chazal, and Leonidas Guibas. Persistence-based segmentation of deformable shapes. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 45–52. IEEE, 2010.
- [57] Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Persistence-based clustering in riemannian manifolds. *Journal of the ACM (JACM)*, 60(6):1–38, 2013.
- [58] Matteo Rucco and Giovanna Viticchi. Fast glioblastoma detection in fluid-attenuated inversion recovery (flair) images by topological explainable automatic machine learning. *arXiv preprint arXiv:1912.08167*, 2019.
- [59] Matteo Rucco, Giovanna Viticchi, and Lorenzo Falsetti. Towards personalized diagnosis of glioblastoma in fluid-attenuated inversion recovery (flair) by topological interpretable machine learning. *Mathematics*, 8(5):770, 2020.
- [60] Lorin Crawford, Anthea Monod, Andrew X Chen, Sayan Mukherjee, and Raúl Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.
- [61] Ameer Saadat-Yazdi, Rayna Andreeva, and Rik Sarkar. Topological detection of alzheimer’s disease using betti curves. In *Interpretability of Machine Intelligence in Medical Image Computing, and Topological Data Analysis and Its Applications for Medical Data*, pages 119–128. Springer, 2021.
- [62] Asuka Oyama, Yasuaki Hiraoka, Ippei Obayashi, Yusuke Saikawa, Shigeru Furui, Kenshiro Shiraishi, Shinobu Kumagai, Tatsuya Hayashi, and Jun’ichi Kotoku. Hepatic tumor classification using texture and topology analysis of non-contrast-enhanced three-dimensional t1-weighted mr images with a radiomics approach. *Scientific reports*, 9(1):8764, 2019.
- [63] Takashi Teramoto, Toshiya Shinohara, and Akihiro Takiyama. Computer-aided classification of hepatocellular ballooning in liver biopsies from patients with nash using persistent homology. *Computer Methods and Programs in Biomedicine*, 195:105614, 2020.
- [64] Yashbir Singh, William Jons, Gian Marco Conte, Jaidip Jagtap, Kuan Zhang, Joseph D Sobek, Pouria Rouzrokh, John E Eaton, and Bradley J Erickson. Persistent homology approach distinguishes potential pattern between “early” and “not early” hepatic decompensation groups using mri modalities. *Current Directions in Biomedical Engineering*, 7(2):488–491, 2021.
- [65] Eashwar Somasundaram, Adam Litzler, Raoul Wadhwa, Steph Owen, and Jacob Scott. Persistent homology of tumor ct scans is associated with survival in lung cancer. *Medical physics*, 48(11):7043–7051, 2021.
- [66] Chul Moon, Qiwei Li, and Guanghua Xiao. Using persistent homology topological features to characterize medical images: Case studies on lung and brain cancers. *The Annals of Applied Statistics*, 17(3):2192–2211, 2023.
- [67] Hyeekyoung Lee, Hyejin Kang, Moo K Chung, Bung-Nyun Kim, and Dong Soo Lee. Persistent brain network homology from the perspective of dendrogram. *IEEE transactions on medical imaging*, 31(12):2267–2277, 2012.
- [68] Arshi Khalid, Byung Sun Kim, Moo K Chung, Jong Chul Ye, and Daejong Jeon. Tracing the evolution of multi-scale functional networks in a mouse model of depression using persistent brain network homology. *NeuroImage*, 101:351–363, 2014.
- [69] Aina Ferrà, Gloria Cecchini, Fritz-Pere Nobbe Fisas, Carles Casacuberta, and Ignasi Cos. A topological classifier to characterize brain states: When shape matters more than variance. *Plos one*, 18(10):e0292049, 2023.
- [70] Paul Bendich, James S Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *The annals of applied statistics*, 10(1):198, 2016.
- [71] Justin Curry, Haibin Hang, Washington Mio, Tom Needham, and Osman Berat Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, pages 1–58, 2022.
- [72] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. In *Proceedings 41st annual symposium on foundations of computer science*, pages 454–463. IEEE, 2000.
- [73] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

- [74] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.
- [75] Allen Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2000.
- [76] Guillaume Tauzin, Umberto Lupo, Lewis Tunstall, Julian Burella Pérez, Matteo Caorsi, Anibal M Medina-Mardones, Alberto Dassatti, and Kathryn Hess. giotto-tda: A topological data analysis toolkit for machine learning and data exploration. *J. Mach. Learn. Res.*, 22(39):1–6, 2021.
- [77] Shizuo Kaji, Takeki Sudo, and Kazushi Ahara. Cubical ripser: Software for computing persistent homology of image and volume data. *arXiv preprint arXiv:2005.12692*, 2020.
- [78] Clément Maria, Jean-Daniel Boissonnat, Marc Glisse, and Mariette Yvinec. The gudhi library: Simplicial complexes and persistent homology. In *International congress on mathematical software*, pages 167–174. Springer, 2014.
- [79] Adélie Garin, Teresa Heiss, Kelly Maggs, Bea Bleile, and Vanessa Robins. Duality in persistent homology of images. *arXiv preprint arXiv:2005.04597*, 2020.
- [80] Lu Li, Connor Thompson, Gregory Henselman-Petrusek, Chad Giusti, and Lori Ziegelmeier. Minimal cycle representatives in persistent homology using linear programming: an empirical study with user’s guide. *Frontiers in artificial intelligence*, 4:681117, 2021.
- [81] Rohlfing Torsten, Natalie M. Zahr, Edith V. Sullivan, and Adolf Pfefferbaum. The SRI24 multichannel atlas of normal adult human brain structure. *Human brain mapping*, 31(5):798–819, 2010.
- [82] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [83] Tamal K Dey, Anil N Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 221–230, 2010.
- [84] Emerson G Escolar and Yasuaki Hiraoka. Optimal cycles for persistent homology via linear programming. In *Optimization in the Real World: Toward Solving Real-World Optimization Problems*, pages 79–96. Springer, 2016.
- [85] Ippei Obayashi. Volume-optimal cycle: Tightest representative cycle of a generator in persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 2(4):508–534, 2018.
- [86] Chao Chen and Daniel Freedman. Hardness results for homology localization. *Discrete & Computational Geometry*, 45(3):425–448, 2011.
- [87] David Cohen-Steiner, André Lieutier, and Julien Vuillamy. Lexicographic optimal homologous chains and applications to point cloud triangulations. *Discrete & Computational Geometry*, 68(4):1155–1174, 2022.
- [88] Tamal K Dey, Tao Hou, and Sayan Mandal. Persistent 1-cycles: Definition, computation, and its application. In *Computational Topology in Image Context: 7th International Workshop, CTIC 2019, Málaga, Spain, January 24-25, 2019, Proceedings 7*, pages 123–136. Springer, 2019.
- [89] Ippei Obayashi, Takenobu Nakamura, and Yasuaki Hiraoka. Persistent homology analysis for materials research and persistent homology software: Homcloud. *Journal of the physical society of japan*, 91(9):091013, 2022.
- [90] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*, volume 10. Springer, 2016.
- [91] Jaeil Kim, Maria del C Valdés Hernández, Natalie A Royle, Susana Munoz Maniega, Benjamin S Aribisala, Alan J Gow, Mark E Bastin, Ian J Deary, Joanna M Wardlaw, and Jinah Park. 3d shape analysis of the brain’s third ventricle using a midplane encoded symmetric template model. *Computer methods and programs in biomedicine*, 129:51–62, 2016.