

Short Paper

Persistent Homology With Improved Locality Information for More Effective Delineation

Doruk Oner ¹, Adélie Garin, Mateusz Koziniński ², Kathryn Hess ³, and Pascal Fua ¹, *Fellow, IEEE*

Abstract—Persistent Homology (PH) has been successfully used to train networks to detect curvilinear structures and to improve the topological quality of their results. However, existing methods are very global and ignore the location of topological features. In this paper, we remedy this by introducing a new filtration function that fuses two earlier approaches: thresholding-based filtration, previously used to train deep networks to segment medical images, and filtration with height functions, typically used to compare 2D and 3D shapes. We experimentally demonstrate that deep networks trained using our PH-based loss function yield reconstructions of road networks and neuronal processes that reflect ground-truth connectivity better than networks trained with existing loss functions based on PH.

Index Terms—Aerial images, connectivity, map reconstruction, road network reconstruction.

I. INTRODUCTION

In many image segmentation tasks, the topology of the resulting mask is as important as, if not more than, its pixel-wise accuracy. For example, a model of an aortic valve that does not form a ring is biologically implausible. Similarly, networks of curvilinear structures—be they roads in aerial images, blood vessels in Computer Tomography (CT) scans, or dendrites and axons in Light Microscopy (LM) image stacks—should not feature breaks that disrupt connectivity or false connections between disjoint structures. Unfortunately, deep networks trained by minimizing pixel-wise loss functions, such as the cross-entropy or the mean square error, are subject to such mistakes. This is in part because it often takes very few mislabeled pixels to alter the topology significantly with little impact on the pixel-wise accuracy. In other words, it is possible for a network trained in this manner to deliver both a good pixel classification accuracy and an incorrect topology.

Specialized solutions to this problem have been proposed in the form of loss functions that compare the topology of the prediction to that of the annotation. They are effective for specific applications but do not naturally generalize. For example, the perceptual loss of [1] penalizes topological differences between the prediction and the ground truth,

Manuscript received 3 March 2022; revised 23 December 2022; accepted 2 February 2023. Date of publication 1 March 2023; date of current version 30 June 2023. This work was supported in part by Swiss National Science Foundation Sinergia under Grant 177237 and in part by FWF Austrian Science Fund Lise Meitner under Grant M3374. Recommended for acceptance by K. Schindler. (Corresponding author: Mateusz Koziniński.)

Doruk Oner and Pascal Fua are with the Computer Vision Laboratory, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland (e-mail: doruk.oner@epfl.ch; pascal.fua@epfl.ch).

Adélie Garin and Kathryn Hess are with the Laboratory for Topology and Neuroscience, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland (e-mail: adélie.garin@epfl.ch; kathryn.hess@epfl.ch).

Mateusz Koziniński is with the Institute of Computer Graphics and Vision, TU Graz, 8010 Graz, Austria (e-mail: mateusz.kozinski@gmail.com).

Digital Object Identifier 10.1109/TPAMI.2023.3246921

but cannot be guaranteed to detect them all. Similarly, minimizing the MALIS loss for segmenting electron microscopy scans [2], [3] yields better region boundaries but does not penalize interruptions in loopy linear structures. This has been addressed by [4] for delineation of 2D road networks but the proposed solution is not applicable to 3D image stacks.

Persistent Homology (PH) [5], an elegant approach to describing and comparing topological structure of data, offers the promise to address the connectivity problem in a generic way, both for 2D and 3D images. Homology is the study of topological features in an object, such as its connected components (0-homology classes), loops (1-homology classes), and closed surfaces (2-homology classes). Persistent homology detects homology classes in objects *filtered* at different *scales*. A homology class that appears at a particular scale and disappears at a larger one is represented by a scale interval called the *persistence interval*. The set of persistence intervals for all the homology classes characterizes the overall topology of the structure. It can be represented by a *persistence diagram*. The similarity of these diagrams across two different structures can then be used to quantify their topological similarity. This has been successfully exploited to train deep networks for delineation [6], image segmentation [6], [7], [8] and crowd counting [9].

We show that these methods fail to unleash the full power of persistent homology, because they discard too much information about the structure of the prediction and the annotation when encoding them in the form of persistence diagrams. As shown in Fig. 1, this can result in networks that still fail to enforce the proper topology. To remedy this, we introduce a new approach to computing persistence diagrams that increases their descriptive power, as shown in Fig. 3. Our main contribution is a novel filtration technique that combines two approaches to filtration commonly used in topological data analysis (TDA): thresholding-based-filtration [6], [7], [8] and the height function [10]. It yields a loss function that can be used for both 2D and 3D images and significantly improves performance compared to state-of-the-art topological methods, as we will demonstrate in our experiments.

II. RELATED WORK

Training a deep network that produces topologically correct segmentations has typically been done by designing loss functions that, when minimized, favor plausible topology. In this section, we briefly review first those that do not rely on Persistent Homology, and then those that do.

A. Losses Designed to Enforce Topological Correctness

Several such losses have been proposed already to go beyond pixel-wise classification accuracy by encoding more global properties. In [11], the connectivity between neighboring pixel pairs is used as




Fig. 1. *2D and 3D delineation.* (a) Aerial image and slice of a microscopy stack. (b) A network trained using a standard homology-based loss yields road and neurite interruptions. (c) One trained using our localized loss is more topologically accurate and produces predictions that closely resemble the ground truth (d).

an additional source of supervision. This approach has been shown to improve connectivity, but since disconnections or false connections are not penalized explicitly, there is no guarantee it captures all such errors. The perceptual loss of [1] is based on the assumption that a pre-trained neural network can capture differences of connectivity between the prediction and the ground-truth. However, even though it has been shown experimentally to improve the topology of masks produced by a deep net, there is no guarantee that this assumption holds in general. Making the Rand index of segmentations produced by the network similar to that of ground truth ones [2], [3] helps when modeling tree-like structures, both in 2D and in 3D, but cannot prevent disconnections in loopy structures. This shortcoming has been addressed by [4] by detecting disconnections of 2D loopy structures as interconnections of background regions, but the proposed solution does not generalize to 3D.

B. Losses That Rely on Persistent Homology

Persistent Homology (PH) [12], [13] is a well-established topological data descriptor. One of its important applications is comparing the topological structures of binary images, for example by enforcing the correct Betti number on binary masks resulting from inference in Markov Random Fields [14]. More recently, it has been shown that persistence diagrams can also be computed for grayscale images and differentiated with respect to the pixel values [6], [7], [15], [16], [17]. Hence, they can be incorporated into loss terms and used to train deep networks. In this vein, a loss term that enforces a sequence of desired Betti numbers on the predicted segmentation was introduced in [7]. This approach was further extended to a loss function that tends to equalize the Betti number of the prediction and the ground truth [8]. In a similar vein, the loss term of [6] relies on comparing persistence diagrams of the prediction and the ground truth, where the persistence diagrams are obtained by thresholding. As discussed in the next section, for binary ground truth this results in degenerate persistence diagrams that only encode the Betti number. Thus, this approach can be interpreted as equalizing the Betti numbers of the prediction and the ground truth, as in [8]. This was improved upon in [18] by applying PH to distance maps instead of binary annotations or class affinity maps. We show in the next section that this makes the loss function better at detecting and penalizing topological errors. Unfortunately, even this improved technique is susceptible to incorrectly matching the persistence diagrams

of the prediction and the ground truth because it throws away location information. By incorporating such information into our diagrams, our method makes them more informative and alleviates this problem.

It has also been proposed to detect disconnections in predicted 2D and 3D structures using Discrete Morse Theory [19]. Topological features that are inconsistent with the ground truth are then penalized in the loss function. However, when the annotations lack spatial precision, which is often the case for neurite and road centerline annotation like the ones studied here, ground-truth inaccuracies may confuse the network. By contrast, our technique allows for considerable misalignment between the prediction and the ground truth.

III. METHOD

We first introduce Persistent Homology and its application to characterizing two-dimensional images and three-dimensional image stacks. As PH provides global descriptors that ignore the location of topological features, we then introduce our approach to accounting for it.

A. Persistent Homology

In the interest of simplicity, we introduce PH for binary images and image stacks, where homology classes are limited to connected components, loops, and closed surfaces. We refer the interested reader to the review [12] for a more general treatment, applicable to non-image and higher-dimensional data.

At the heart of PH is detecting homology classes—connected components, loops, closed surfaces—at many different scales. The ones that exist over a wide scale range are called persistent and deemed more likely to represent true features, as opposed to sampling artifacts or noise. Here, scale has a very specific meaning. It refers to the parameter of a filtration function F that is applied to an image \mathbf{Y} to produce topological objects called *cubical complexes*. They arise when filtering images and their properties are described in [20], [21] for instance. A reader not familiar with algebraic topology can think of them as binary masks. The masks obtained for different scales form a sequence of inclusions, that is, for a pair of scale parameters $s_1 < s_2$, the mask $F(\mathbf{Y}, s_1)$ is entirely contained within the mask $F(\mathbf{Y}, s_2)$. The simplest example of a function for filtering grayscale images is thresholding, where the threshold acts as the scale, as shown in Fig. 2.




Fig. 2. *Filtration*. When the distance map shown on the left is filtered by thresholding, the loop h emerges at scale b^h and is filled at scale d^h . This gives rise to the point (b^h, d^h) in the persistence diagram shown on the right. Here, thresholding means retaining all pixels whose value is lower than the threshold.

As the scale changes, homology classes in the filtered cubical complex emerge and disappear. To capture this, the scale range is sampled from small to large, the image is filtered at the selected scale values, homology classes in the resulting binary masks are detected algebraically [12], and correspondence is established between the homology classes found at consecutive scales. For each class, this yields a pair (b, d) , where b is the scale at which the homology class appears and d the scale at which it disappears. We will refer to them as *birth* and *death* times and to the interval $[b, d]$ as the *persistence interval* of the homology class. The set $P_{\mathbf{Y}} = \{(b^h, d^h)\}_{h \in H_{\mathbf{Y}}}$, where $H_{\mathbf{Y}}$ is the set of all homology classes found in the filtered image \mathbf{Y} , is called the *persistence diagram* of \mathbf{Y} , and was first introduced by [22]. In practice, we use the Gudhi library [23] to compute persistence diagrams from images. Fig. 2 depicts the birth and death of a specific homology class.

To compare images \mathbf{Y}_1 and \mathbf{Y}_2 , one-to-one matching is performed between their persistence diagrams, $P_{\mathbf{Y}_1}$ and $P_{\mathbf{Y}_2}$, with the cost of matching a homology $g \in H_{\mathbf{Y}_1}$ to a homology $h \in H_{\mathbf{Y}_2}$ set to $c_{g,h} = (b^g - b^h)^2 + (d^g - d^h)^2$ and the cost of leaving an interval $[b, d]$ unmatched is set to the distance between the point (b, d) and the diagonal in \mathbb{R}^2 . The optimal matching can be found using the Hungarian algorithm. Its cost that we denote as $C(\mathbf{Y}_1, \mathbf{Y}_2)$ quantifies the topological discrepancy between \mathbf{Y}_1 and \mathbf{Y}_2 by penalizing differences between corresponding homology classes and ones that only appear in either \mathbf{Y}_1 or \mathbf{Y}_2 .

B. Training Deep Networks Using PH

Let f be a network that associates to an input image \mathbf{X} a segmentation mask $\mathbf{Y} = f(\mathbf{X})$ such that for all pixels or voxels $p \in \mathbf{Y}$, $0 \leq \mathbf{Y}[p] \leq 1$ and let $\hat{\mathbf{Y}}$ be the corresponding ground-truth mask. A natural idea then is to train f by minimizing

$$L_{\text{tot}}(\mathbf{Y}, \hat{\mathbf{Y}}) = L(\mathbf{Y}, \hat{\mathbf{Y}}) + \alpha C(\mathbf{Y}, \hat{\mathbf{Y}}), \quad (1)$$


where L is the standard loss function, either the Mean Square Error, or the Cross Entropy, and α is a hyper-parameter, which we set to 0.01 in practice. This is possible because C is sub-differentiable with respect to its inputs when filtration is achieved by thresholding, as shown in [6], [7], [16]. However, when the ground truth $\hat{\mathbf{Y}}$ is binary, as it often is, all structures emerge at scale zero and disappear at scale one. Hence, as shown in Fig. 3(a) the persistence intervals all are $[0, 1]$. In other words, all points in a ground truth persistence diagram are located in its upper-left corner, and the only difference between diagrams obtained for annotations of different images is the number of points they contain. Unlike in classical applications of PH [12], where persistence diagrams serve as rich topological descriptors, such diagrams only encode the Betti number of the annotation. An approach

to handling this difficulty is to replace the binary ground truth by its distance transform that can be thresholded over a wide range of threshold values to create different binary masks [18]. Unfortunately, computing the persistence diagram of a ground truth distance transform still yields persistence diagrams in which the topological features of the original, binary ground truth are spread along the ‘death’ axis but not along the ‘birth’ one: The distance value at the structures themselves is zero and, as a result, all the loops of the ground truth mask appear as soon as the scale value becomes positive. As shown in Fig. 3(b), this may lead to erroneous matches between persistence diagrams, which encourages the deep network to produce wrong segmentations.


C. accounting for the Location of Topological Features During Filtration

Our goal is therefore to prevent erroneous matches between topological features of the prediction and of the ground truth. To this end, we want to use the features’ image location to characterize them. However, re-defining the matching cost to include a position-dependent term would be difficult, because topological features extend across the scale-space, and because there is no natural notion of distance between them. Hence, instead of modifying the matching cost, we propose a new filtration function that distinguishes features at different positions. We draw our inspiration from a filtration technique called the height function [10]. It was originally designed for three-dimensional meshes and can be applied to binary images by assigning to each pixel a *height* value that is the coordinate of its projection along a selected straight line. Filtration is carried out by forming binary masks made of pixels whose height is smaller than the scale parameter [24]. As the scale is increased, the binary image is revealed in scan-lines perpendicular to the height axis, one scan-line at a time. The birth and death times are the heights of pixels responsible for the emergence and disappearing of homology classes. As a result, the persistence diagram contains partial information about the location of topological features. Moreover, both birth and death times of different homology classes are distributed across scales. Additionally, it has been shown that a binary image can be reconstructed from as few as four persistence diagrams obtained with height functions with well-chosen directions [25]. A height function is only defined for binary images, but the abovementioned result inspired us to extend its definition by combining it with thresholding distance maps. Given a scale s , the value of the filtered binary mask at coordinates \mathbf{p} is taken to be

$$F(\mathbf{Y}, s)[\mathbf{p}] = \mathbb{1}(\mathbf{Y}[\mathbf{p}] + \rho(\mathbf{p}) < s), \quad (2)$$



(a) **Filtration by thresholding** binary ground truth and predicted class affinity maps. Here, filtration involves decreasing the threshold from 1 to 0, and retaining the pixels greater than the threshold. Note, that the binary masks resulting from filtering the ground truth at different scales are all the same and that all points in the ground truth persistence diagram (*top-right*) coincide. This results in erroneous matches between the predicted and ground truth homology classes. Minimizing a loss function based on such a filtration can magnify the errors.



(b) **Filtration by thresholding distance maps** distributes the topological features of the ground truth along the vertical but not the horizontal axis. This still results in erroneous matching between the predicted and ground truth homology classes: Loop D' in the prediction emerges when the threshold is high enough to make the road break disappear. Hence, it remains unmatched and the E' loop created by the false positive road is matched to the ground truth loop D .



(c) **Our localized filtration of distance maps** distributes the persistence diagram of the ground truth across the plane, promoting correct matches between predicted and ground truth homology classes. The gray arrow represents the direction of the height axis used by the filtration function.

Fig. 3. **Comparing filtration functions on synthetic data.** The binary ground truth road annotation (*top-left* in each table part) contains four loops, marked with cyan dashed lines. We synthesized a predicted class affinity map (*bottom-left* in each part) by extending one road to the left and interrupting another. In consequence, loop B and D from the ground truth are joined into B' in the prediction, and A is split into A' and E' . For each filtration method, we show binary masks resulting from filtration at different scales, pairs of persistence diagrams, and their optimal matches.

where $\mathbb{1}(\cdot)$ evaluates to one if the condition in the bracket is satisfied and to zero otherwise. In essence, this amounts to thresholding the sum of the height function ρ and the pixel values. From the perspective of TDA, such combination of two filtration functions can be seen as a line in the fibered barcode defined by [26].

In its simplest form, ρ is a linear function of pixel coordinates, and the region highlighted for any s extends along a line perpendicular to

the height axis, as shown in Fig. 3(c). But other forms of ρ are also possible. We tested

- linear functions $\rho(\mathbf{p}) = \mathbf{w}^T \mathbf{p}$, where \mathbf{w} is a two-vector hyper-parameter encoding the orientation of the height axis and the slope of the height function;
- a scaled distance to a point \mathbf{q} in the image, $\rho(\mathbf{p}) = a \|\mathbf{p} - \mathbf{q}\|_2$, where \mathbf{q} and a are hyper-parameters;

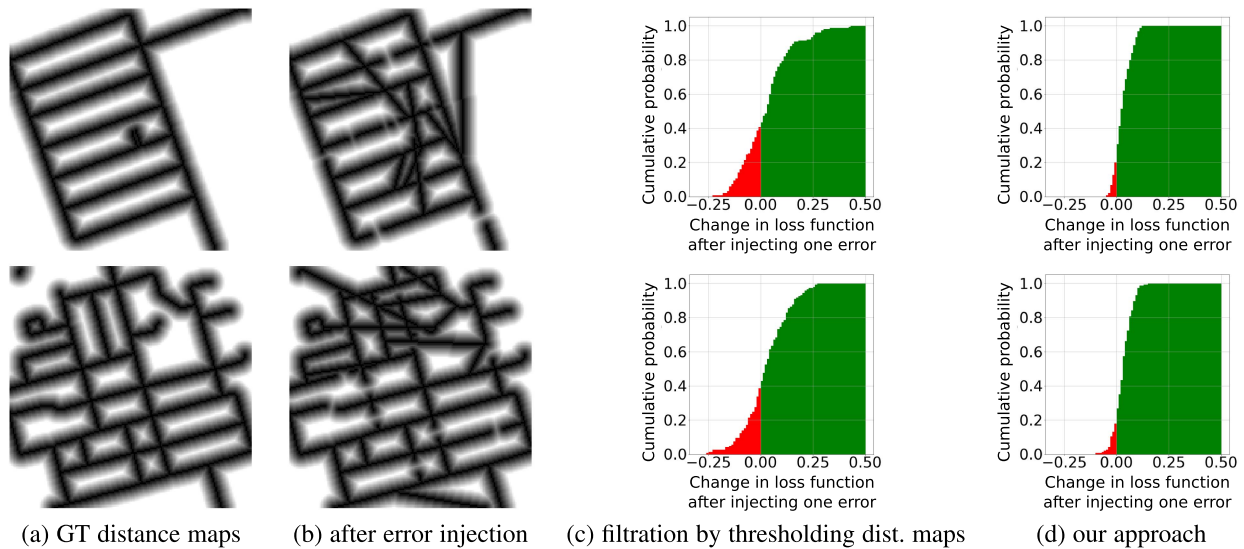


Fig. 4. *Sensitivity of the topological loss term C to the number of injected errors* (a) Ground truth distance maps of road networks. (b) Distance maps corrupted by introducing false roads and interruptions. We randomly injected one error at a time, obtaining corrupt distance maps with 30 errors. We repeated this simulation 10 times. (c,d) The cumulative distribution function of change in the loss term in response to injecting one error. In (c), C is evaluated using the filtration by thresholding distance maps, whereas in (d) we use our filtration. The probability of decreasing the existing loss term by injecting additional errors is around 0.4, whereas for our loss term it drops to 0.2. We conclude that our loss term is more monotonic with respect to the error number.

- the square of the height function $\rho(\mathbf{p}) = \mathbf{p}^T \mathbf{W} \mathbf{p}$, where $\mathbf{W} = \mathbf{w}^T \mathbf{w}$, and \mathbf{w} is the hyper parameter encoding the slope of the function and the orientation of the height axis;

The function ρ introduces localization information of the topological features into the persistence diagram. This is illustrated by Fig. 3 where different values of the scale parameter make homology classes appear in different parts of the image. But, because the scale parameter must be a scalar, it can only pinpoint location of topological features in 2D or 3D images along one direction. This could be addressed by evaluating the loss function many times for many different orientations of the height axis, or more generally, for many different hyper-parameters of ρ . This approach is legitimized by the theoretical result of [25] that states that four well chosen filtration directions suffice to completely represent a binary image. The problem of combining a number of different filtration functions is known in topological literature as multipersistence [27]. But current multipersistence techniques are not easily plugged into a deep learning framework for lack of results on their differentiability. Moreover, filtering the data along multiple directions would considerably slow down the training. Instead, we randomly draw the hyper-parameters of the height function at each training iteration. We show in the supplementary material, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2023.3246921> that, in practice, the simple linear function performs best.

IV. EXPERIMENTS

We first demonstrate that our loss function correlates with the number of topological errors better than standard PH-based losses. We then evaluate its performance in training deep networks to delineate road networks and neuronal arborizations.

A. Correlation to the Number of Topological Errors

We motivated our filtration technique by the fact that it introduces partial localization of topological features into the persistence diagrams

and better spreads the diagrams across the plane. Here, we validate this on synthetic data to show that it correlates better with the number of errors injected into a distance map than the baseline loss, which is based solely on thresholding distance maps. To this end, we took two crops of ground truth road graphs of the *RTracer* dataset [28] and generated faulty synthetic distance maps by injecting thirty errors one at a time. They were selected randomly and with equal probability between a road disconnection and a false interconnection. After each error injection, we evaluated the topological loss term C of (1) using either filtration by thresholding distance maps or our combined filtration. Ideally, we would expect C to increase every time an error is added. Hence, we repeated the experiment ten times. For each crop, we plot the distribution of the increment in C resulting from adding one error, when using the baseline loss in Fig. 4(c) and ours in Fig. 4(d). The parts of the distributions shown in green correspond to positive increments, which are what we expect, and those in red denote the negative ones, which are essentially erroneous. Note that the red parts are far smaller when using our loss than the baseline one.

B. Performance in Training Deep Networks

Having shown that our loss function captures topological correctness better than existing PH-based methods, we now compare the performance of deep nets trained with our and existing losses.

1) Datasets: We experimented on three datasets.

- *RTracer*. A dataset of high-resolution satellite images covering urban areas of forty cities in six countries [28]. The ground truth was obtained from OpenStreetMap. Like [28], [29], [30], [31], we used twenty five cities as the training set and the remaining fifteen as the test set.
- *Massachusetts*. The Massachusetts dataset [32] features both urban and rural neighborhoods, with many different kinds of roads ranging from small paths to highways. For a fair comparison to [6], we split the data into three equal folds and performed a three-way cross validation.
- *Neurons*. The dataset is a part of a proprietary 3D, 2-photon microscopy scan of a whole mouse brain. It contains 14 stacks

TABLE I

VALIDATION RESULTS ON THE MASSACHUSETTS DATASET. OUR LOSS FUNCTION OUTPERFORMS ALL PH-BASED LOSS FUNCTIONS. WE REPORT MEANS AND STANDARD DEVIATIONS OVER THREE INDEPENDENT TRAINING RUNS

Method	Connectivity-oriented				pixel-based
	<i>APLS</i> \uparrow	<i>TLTS</i> \uparrow	<i>JCT</i> \uparrow	<i>Betti</i> \downarrow	<i>CCQ</i> \uparrow
<i>UNet-CE</i>	60.9 \pm 3.9	41.6 \pm 4.1	72.0 \pm 2.7	3.12 \pm 0.6	66.9 \pm 2.6
<i>UNet-MSE</i>	61.3 \pm 3.7	41.9 \pm 4.2	71.9 \pm 2.9	3.09 \pm 0.7	67.3 \pm 2.3
<i>DMT</i>	64.7 \pm 2.9	45.8 \pm 2.8	80.6 \pm 2.4	0.99 \pm 0.4	74.9 \pm 1.9
<i>ConnLoss</i>	73.4 \pm 3.6	53.2 \pm 4.4	81.4 \pm 1.9	1.29 \pm 0.5	75.8 \pm 2.2
<i>Homo-Pre</i>	62.5 \pm 1.9	42.1 \pm 1.9	74.2 \pm 1.7	1.28 \pm 0.3	69.3 \pm 1.9
<i>Homo-Reg</i>	65.0 \pm 2.2	45.6 \pm 1.8	76.9 \pm 1.9	1.09 \pm 0.2	71.8 \pm 2.1
<i>Homo-Ours</i>	68.7 \pm 1.2	50.6 \pm 2.3	79.2 \pm 2.6	0.90 \pm 0.3	74.9 \pm 1.8

TABLE II

OUR LOSS FUNCTION OUTPERFORMS ALL PH-BASED LOSS FUNCTIONS ON THE RTRACER DATASET. WE REPORT MEANS AND STANDARD DEVIATIONS OVER CITIES FROM THE TEST SET

Method	Connectivity-oriented				pixel-based
	<i>APLS</i> \uparrow	<i>TLTS</i> \uparrow	<i>JCT</i> \uparrow	<i>Betti</i> \downarrow	<i>CCQ</i> \uparrow
<i>UNet-CE</i>	63.4 \pm 1.6	37.5 \pm 1.9	78.0 \pm 1.0	3.08 \pm 0.6	59.7 \pm 2.2
<i>UNet-MSE</i>	66.3 \pm 1.9	40.0 \pm 2.0	77.5 \pm 1.3	2.99 \pm 0.5	59.5 \pm 1.9
<i>Segmentation</i>	62.5 \pm 1.5	33.0 \pm 1.6	78.2 \pm 1.5	3.04 \pm 0.6	54.4 \pm 1.0
<i>RoadTracer</i>	59.1 \pm 0.8	40.6 \pm 1.5	81.2 \pm 1.6	2.85 \pm 0.7	47.8 \pm 1.6
<i>Seg-Path</i>	68.1 \pm 1.4	46.5 \pm 1.7	75.4 \pm 1.3	2.31 \pm 0.4	54.0 \pm 1.4
<i>RCNNU-Net</i>	48.2 \pm 1.6	18.4 \pm 1.9	75.9 \pm 1.4	3.25 \pm 0.7	62.8 \pm 1.5
<i>DeepRoad</i>	24.6 \pm 2.2	6.4 \pm 0.9	51.4 \pm 1.5	4.95 \pm 1.0	43.6 \pm 2.0
<i>PolyMapper</i>	61.3 \pm 2.3	31.5 \pm 1.9	80.0 \pm 1.2	2.90 \pm 0.4	35.7 \pm 1.4
<i>ConnLoss</i>	75.4 \pm 1.6	49.6 \pm 1.4	82.6 \pm 0.6	1.30 \pm 0.4	68.4 \pm 0.9
<i>Homo-Pre</i>	67.3 \pm 1.7	42.3 \pm 1.1	78.7 \pm 0.9	1.32 \pm 0.3	61.9 \pm 1.9
<i>Homo-Reg</i>	69.9 \pm 1.6	45.1 \pm 1.4	79.6 \pm 1.3	1.07 \pm 0.3	63.2 \pm 1.6
<i>Homo-Ours</i>	73.8 \pm 1.8	47.8 \pm 0.9	81.3 \pm 1.6	0.89 \pm 0.2	66.3 \pm 1.7

of size $250 \times 250 \times 200$ voxels and a spatial resolution of $1.0 \times 0.3 \times 0.3 \mu m$. We used ten stacks for training and the remaining four for testing.

- *Brain*. The dataset contains two 3D images of neurons in a mouse brain. The axons and dendrites have been outlined manually while viewing the sample under a microscope and the image has been captured later. The sample deformed in the meantime, resulting in a misalignment between the annotation and the image. To ensure that the test and training data comes from the same distribution, we split the two scans into stacks of $150 \times 200 \times 200$ voxels and a spatial resolution of $1 \mu m$, and randomly divided the resulting data set into a training set of twelve stacks and a test set of ten scans.

2) *Methods Tested*: To test the impact of our proposed filtration functions, we used the standard U-Net architecture [33], with four blocks, each with two sequences of convolution-ReLU-batch normalization. Max-pooling in 2×2 windows followed each of the blocks. The initial feature size was set to 32 and grew to 512 in the smallest feature map in the network. We augmented the training data with vertical and horizontal flips and random rotations, and used the ADAM algorithm [34] with the learning rate set to $1e - 4$. We then used different version of the L_{tot} of Eq. 1 we minimized to train the network. We tested the following as baselines:

- *UNet-CE*. L is the Cross Entropy loss for pixel classification and there is no topological discrepancy loss, that is, $\alpha = 0.0$. Binary masks are used as ground truth.

- *UNet-MSE*. L is the mean squared error of the truncated distance to the closest foreground pixel, with no topological discrepancy loss.

- *Homo-Pre*. L is the cross Entropy loss and we compute C by thresholding pixel classification maps., as in [6], [7], [8].

- *Homo-Reg*. L is the mean squared error and we compute C by thresholding the truncated distance maps, as in [18].

- *Homo-Ours*. L is the mean squared error and we compute C using our proposed filtration function.

Based on the results of the ablation study on the *Massachusetts* data set, presented in the supplementary material, available online, in all our experiments with *Homo-Pre*, *Homo-Reg*, and *Homo-Ours*, we set $\alpha = 0.01$ and compute the loss in windows of size 64×64 pixels. Like [6], we limit the method to homology classes order 1, that is, loops. This has two advantages. First, by convention, loops are created by the borders of the window, making disconnections in dead-ending roads or neurites detected as broken loops. Second, detection of homology classes is computationally expensive, and the time grows cubically with the number of pixels. In our current setup, computing the loss for a single window takes 0.5 seconds. Similarly to [6], we did not observe any performance gain due to using homology classes of order 0—connected components—in addition to loops.

For completeness, we also compared our approach to recent techniques *not* relying on persistent homology: *Segmentation* [28], *RoadTracer* [28], *Seg-Path* [31], *RCNNU-Net* [30], *DeepRoad* [35], *PolyMapper* [29], *DMT* [19], and *ConnLoss* [4]. *Segmentation*,

TABLE III
COMPARATIVE RESULTS ON THE *NEURONS* DATASET. OUR LOSS OUTPERFORMS ALL THE BASELINES. WE REPORT MEANS AND STANDARD DEVIATIONS OVER THREE INDEPENDENT TRAINING RUNS

Method	Connectivity-oriented			pixel-based
	<i>APLS</i> \uparrow	<i>TLTS</i> \uparrow	<i>Betti</i> \downarrow	<i>CCQ</i> \uparrow
<i>UNet-CE</i>	79.9 \pm 1.5	80.8 \pm 2.2	2.33 \pm 0.6	90.6 \pm 2.0
<i>UNet-MSE</i>	80.2 \pm 1.6	80.9 \pm 2.0	2.31 \pm 0.7	90.4 \pm 1.9
<i>Homo-Pre</i>	83.5 \pm 1.0	82.1 \pm 1.7	1.06 \pm 0.2	91.2 \pm 1.8
<i>Homo-Reg</i>	85.4 \pm 1.2	83.4 \pm 1.5	0.91 \pm 0.2	92.5 \pm 1.6
<i>Homo-Ours</i>	86.9 \pm 1.1	85.2 \pm 1.9	0.80 \pm 0.2	93.3 \pm 1.9

TABLE IV
COMPARATIVE RESULTS ON THE *BRAIN* DATASET. OUR LOSS OUTPERFORMS ALL PH-BASED LOSSES. MEANS AND STANDARD DEVIATIONS OVER THREE INDEPENDENT TRAINING RUNS AS PRESENTED

Method	Connectivity-oriented			pixel-based
	<i>APLS</i> \uparrow	<i>TLTS</i> \uparrow	<i>Betti</i> \downarrow	<i>CCQ</i> \uparrow
<i>UNet-CE</i>	65.8 \pm 1.8	63.6 \pm 1.3	2.89 \pm 0.4	70.4 \pm 1.9
<i>UNet-MSE</i>	66.0 \pm 1.6	63.9 \pm 1.4	2.92 \pm 0.5	70.6 \pm 1.8
<i>Homo-Pre</i>	67.6 \pm 1.5	65.3 \pm 1.0	1.39 \pm 0.2	71.5 \pm 1.4
<i>Homo-Reg</i>	70.5 \pm 1.5	68.8 \pm 0.9	1.22 \pm 0.3	72.6 \pm 1.3
<i>Homo-Ours</i>	73.4 \pm 1.4	70.1 \pm 1.1	1.06 \pm 0.2	73.2 \pm 1.2

RoadTracer, *RCNNU-Net*, and *PolyMapper* do not explicitly enforce topology constraints, while the others do and are discussed in the related work section. The outputs of these methods were shared by the authors directly with us or on the Internet, and we computed all the performance metrics.

3) *Performance Metrics*: Comparing connectivity of segmentation masks is difficult, because the reconstructions rarely overlap with the ground truth, and often deviate from it significantly. There seems to be no consensus concerning the best evaluation technique; we found five connectivity-oriented metrics in concurrently published recent work. To provide an exhaustive evaluation, we used all of them.

- *APLS*. Average Path Length Similarity aggregates relative length differences of shortest paths between pairs of corresponding points in the ground truth and predicted maps [36].
- *TLTS* is a statistics of lengths of shortest paths between corresponding pairs of end points randomly selected in the predicted and ground-truth networks [37]. We report the fraction of paths with relative length difference within 5%.
- *JCT*. It is a junction score that considers the number of roads intersecting at each junction [28]. It consists of road recall, averaged over the intersections of the ground-truth and road precision, averaged over the intersections of the prediction. We report the corresponding F1 score.
- *Betti*. The Betti error [6] is an average absolute difference between the number of topological structures seen in the ground truth and predicted delineations. We take random patches sized 64×64 from predictions, compute the number of 1-homology classes (loops) and compare the numbers computed for the prediction and the ground truth. We average this difference over 10 trials. In practice, to compute the error we use the code made publicly available by the authors.
- *CCQ*. We complement the connectivity-oriented metrics with the most popular metric that measures spatial co-occurrence of annotated and predicted road pixels. The Correctness, Completeness and Quality are equivalent to precision, recall and

intersection-over-union, with the definition of a true positive relaxed from spatial coincidence of prediction and annotation to co-occurrence within a distance of 5 pixels [38]. We report the Quality as our single-number metric.

4) *Comparative Results*: We present validation results for the *Massachusetts* data set in Table I, and test results for the *RTracer* data set in Table II. Our method outperforms the other methods based on Persistent Homology, which demonstrates that our approach to filtering is truly effective. It also outperforms the other 2D tracing algorithms targeted at handling aerial images, *RoadTracer*, *Seg-Path*, *DeepRoad*, and *PolyMapper*, with the exception of *ConnLoss* that does marginally better. This is presumably because *ConnLoss* explicitly penalizes each disconnection of the prediction, whereas a persistence diagram is a lossy topological descriptor that may fail to penalize some errors. However, *ConnLoss* does not naturally extend to 3D data, whereas our method does. On the 3D *Neurons* data set, it outperforms the competing algorithms, as evidenced by the test results shown in Tables III and IV. We provide qualitative results in the supplementary material, available online.

V. CONCLUSION

We demonstrated a fault in the design of existing methods to employ Persistent Homology to train deep networks in delineating curvilinear structures: by using inadequate filtration functions, they severely reduce the information content of the persistence diagrams, hampering performance of the trained network. We proposed an improved approach, based on combining filtration by thresholding with the height function, that increases the descriptive power of the diagrams, and gives PH a place among the best-performing methods to train topologically accurate deep networks.

The proposed approach is limited by the need to randomly select the parameters of the height function at each training iteration, because some orientations of the height axis might result in a failure to detect topological errors, or provoke erroneous matches between

the persistence diagrams of the prediction and the ground truth. We therefore plan to investigate the use of multipersistence [27] for less random and more effective supervision. Another limitation stems from the fact that our loss function has sparse gradients that only depend on values at pixels that are critical for emergence and disappearance of topological features. This limits robustness and our future work will focus on developing topological descriptors with more smooth gradients. While our loss function improves the topological correctness of the segmentation masks, some bio-medical applications require full confidence of correctness of anatomy models, which current methods cannot guarantee. This also motivates us to investigate the use of topological methods to highlight the regions of the segmentation masks that require manual correction, thereby facilitating proof-reading of segmentation results.

REFERENCES

- [1] A. Mosińska, P. Marquez-Neila, M. Kozinski, and P. Fua, "Beyond the pixel-wise loss for topology-aware delineation," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 3136–3145.
- [2] K. Briggman, W. Denk, S. Seung, M. Helmstaedter, and S. Turaga, "Maximin affinity learning of image segmentation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1865–1873.
- [3] J. Funke et al., "Large scale image segmentation with structured loss based deep learning for connectome reconstruction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 7, pp. 1669–1680, Jul. 2019.
- [4] D. Oner, M. Kozifski, L. Citraro, N. C. Dadap, A. G. Konings, and P. Fua, "Promoting connectivity of network-like structures by enforcing region separation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, pp. 5401–5413, 2022.
- [5] H. Edelsbrunner and J. Harer, "Persistent homology—A survey," *Contemporary Math.*, vol. 453, pp. 257–282, 2008.
- [6] X. Hu, F. Li, D. Samaras, and C. Chen, "Topology-preserving deep image segmentation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 5658–5669.
- [7] J. Clough, I. Oksuz, N. Byrne, J. Schnabel, and A. King, "Explicit topological priors for deep-learning based image segmentation using persistent homology," in *Proc. Inf. Process. Med. Imag.*, 2019, pp. 16–28.
- [8] J. Clough, N. Byrne, I. Oksuz, V. Zimmer, J. Schnabel, and A. King, "A topological loss function for deep-learning based image segmentation using persistent homology," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 12, pp. 8766–8778, Dec. 2022.
- [9] S. Abousamra, M. Hoai, D. Samaras, and C. Chen, "Localization in the crowd with topological constraints," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, pp. 872–881.
- [10] K. Turner, S. Mukherjee, and D. Boyer, "Persistent homology transform for modeling shapes and surfaces," *Inf. Inference*, vol. 3, no. 4, pp. 310–344, Dec. 2014.
- [11] X. Li, Y. Wang, L. Zhang, S. Liu, J. Mei, and Y. Li, "Topology-enhanced urban road extraction via a geographic feature-enhanced network," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 12, pp. 8819–8830, Dec. 2020.
- [12] H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," in *Proc. IEEE 41st Annu. Symp. Foundations Comput. Sci.*, 2000, pp. 454–463.
- [13] A. Zomorodian and G. Carlsson, "Computing persistent homology," in *Proc. ACM Symp. Comput. Geometry*, 2004, pp. 347–356.
- [14] C. Chen, D. Freedman, and C. Lampert, "Enforcing topological constraints in random field image segmentation," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2011, pp. 2089–2096.
- [15] R. Gabriellsson, B. Nelson, A. Dwaraknath, and P. Skraba, "A topology layer for machine learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 1553–1563.
- [16] J. Leygonie, S. Oudot, and U. Tillmann, "A framework for differential calculus on persistence barcodes," *Foundations Comput. Math.*, vol. 22, pp. 1069–1131, 2021.
- [17] M. Carriere, F. Chazal, M. Glisse, Y. Ike, H. Kannan, and Y. Umeda, "Optimizing persistent homology based functions," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 1294–1303.
- [18] F. Wang, H. Liu, D. Samaras, and C. Chen, "TopoGAN: A topology-aware generative adversarial network," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 118–136.
- [19] X. Hu, Y. Wang, L. Fuxin, D. Samaras, and C. Chen, "Topology-aware segmentation using discrete morse theory," in *Proc. Int. Conf. Learn. Representations*, vol. abs/2103.09992, 2021.
- [20] A. Garin, T. Heiss, K. A. R. Maggs, B. Bleile, and V. Robins, "Duality in persistent homology of images," 2020, *arXiv:2005.04597*.
- [21] B. Bleile, A. Garin, T. Heiss, K. Maggs, and V. Robins, "The persistent homology of dual digital image constructions," 2021, *arXiv:2102.11397*.
- [22] S. Barannikov, "The framed morse complex and its invariants," *Adv. Sov. Math.*, vol. 21, pp. 93–115, 1994.
- [23] C. Maria, J. D. Boissonnat, M. Glisse, and M. Yvinec, "The Gudhi library: Simplicial complexes and persistent homology," in *Proc. Int. Congr. on Math. Softw.*, 2014, pp. 167–174.
- [24] A. Garin and G. Tauzin, "A Topological 'Reading' lesson: Classification of MNIST using TDA," in *Proc. Int. Conf. Mach. Learn. Appl.*, 2019, pp. 1551–1556.
- [25] L. Bethausser, "Topological reconstruction of grayscale images," Ph.D. dissertation, Dept. Math., Univ. Florida, Gainesville, Florida, 2018.
- [26] M. Carrière and A. Blumberg, "Multiparameter persistence image for topological machine learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 22 432–22 444.
- [27] G. Carlsson and A. Zomorodian, "The theory of multidimensional persistence," *Discret. Comput. Geom.*, vol. 42, no. 1, pp. 71–93, 2009.
- [28] F. Bastani et al., "RoadTracer: Automatic extraction of road networks from aerial images," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Salt Lake City, UT, USA, 2018, pp. 4720–4728.
- [29] Z. Li, J. Wegner, and A. Lucchi, "Topological map extraction from overhead images," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 1715–1724.
- [30] X. Yang, X. Li, Y. Ye, R. Y. K. Lau, X. Zhang, and X. Huang, "Road detection and centerline extraction via deep recurrent convolutional neural network U-net," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 9, pp. 7209–7220, Sep. 2019.
- [31] A. Mosińska, M. Kozinski, and P. Fua, "Joint segmentation and path classification of curvilinear structures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 6, pp. 1515–1521, Jun. 2020.
- [32] V. Mnih, "Machine learning for aerial image labeling," Ph.D. dissertation, Dept. Comput. Sci., Univ. Toronto, Toronto, Canada, 2013.
- [33] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Conf. Med. Image Comput. Comput. Assist. Intervention*, 2015, pp. 234–241.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimisation," in *Proc. Int. Conf. Learn. Representations*, vol. abs/1412.6980, 2015.
- [35] G. Mátyus, W. Luo, and R. Urtasun, "Deeproadmapper: Extracting road topology from aerial images," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 3458–3466.
- [36] A. V. Etten, D. Lindenbaum, and T. Bacastow, "SpaceNet: A remote sensing dataset and challenge series," 2018, *arXiv:1807.01232*.
- [37] J. Wegner, J. Montoya-Zegarra, and K. Schindler, "A higher-order CRF model for road network extraction," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 1698–1705.
- [38] C. Wiedemann, C. Heipke, H. Mayer, and O. Jamet, "Empirical evaluation of automatically extracted road axes," in *Empirical Evaluation Techniques in Computer Vision*, Hoboken, NJ, USA: Wiley, 1998, pp. 172–187.