

1 We thank the reviewers for constructive feedback. We will improve our manuscript accordingly. Some concerns were
2 raised regarding the reproducibility. We will address them below. If accepted, we will release the code immediately.

3 **Q1 (R1,R3): How is the optimal matching computed? How is $Dgm(f)$ computed? Small noisy components?**

4 **A:** The matching algorithm is as follows. A total of $k = (\text{Betti number})$ dots from ground truth ($Dgm(g)$) are at the
5 upper-left corner $p = (0, 1)$ (in Fig.4 (b) of the paper). We find the k dots closest to the corner p in $Dgm(f)$ and
6 match them to the ground truth dots. The remaining dots in $Dgm(f)$ are matched to the diagonal line. The algorithm
7 computes and sorts the squared distances from all dots in $Dgm(f)$ to p . The complexity is $O(n \log n)$, $n = \#$ of dots in
8 $Dgm(f)$. In a general setting, the state-of-the-art diagram matching algorithm [Kerber et al. 2017] takes $O(n^{3/2})$.

9 The dots matched to the diagonal line correspond to small noisy components or noisy loops. These dots will be
10 pushed to the diagonal. And their corresponding components/loops will be removed or merged with others.


11 To compute the persistence diagram $Dgm(f)$, we use the classic algorithm of [13,14]: We first discretize an image
12 patch into vertices (pixels), edges and squares. The adjacency relationship between these elements and their likelihood
13 function values are encoded in a boundary matrix, whose rows and columns correspond to vertices/edges/squares. The
14 matrix is reduced using a modified Gaussian elimination. The pivoting entries of the reduced matrix correspond to all
15 the dots in $Dgm(f)$. Efficient software packages are publicly available (GUDHI, RIPSER, PHAT, etc).

16 **Q2 (R1,R3): How is the prediction across patches aggregated? Handle the boundaries between patches?**

17 **A:** We do not partition the image into patches. Instead, we randomly and densely sample patches which can overlap.
18 Our loss enforces correct topology within each sampled patch. These overlaps between patches propagate correct
19 topology everywhere. Correct topology within a patch means the segmentation can be a deformation of the ground
20 truth. But the deformation is constrained within the patch. The patch size controls the tolerable geometric deformation.

21 **Q3 (R1,R3): Ablation study of lambda. lambda=0 should also be evaluated as a baseline. Report standard deviation.**

22 **A:** Figures (a)(b) show ablation studies of lambda on CREMI
23 w.r.t. accuracy, Betti error and convergence rate. As we increase
24 lambda, per-pixel accuracy is slightly compromised. The Betti
25 error decreases first but increases later. One important observation
26 is that a certain amount of topological loss improves the conver-
27 gence rate significantly. These observations suggest choosing
28 lambda = 1×10^{-4} . We will also add ablation study w.r.t. patch size and report standard deviation in the final version.
29 Since our method is based on the DIVE architecture, the baseline DIVE is equivalent to the case when lambda = 0.




30 **Q4 (R1): Training with a dilated ground truth mask?**

31 **A:** A dilated mask cannot help. The topological loss identifies difficult locations (e.g., blurred membrane locations in
32 Figure 1 of the paper) as critical pixels and forces the network to learn to classify them correctly. Dilated ground truth
33 masks cannot help with this. For CREMI dataset, using a topology-agnostic model with dilated masks (1 and 2 pixel
34 dilation), we have Betti Error 4.126 and 4.431 respectively, still significantly worse than TopoNet (Betti Error = 1.113).

35 **Q5 (R1): Usefulness of the proof. What is the benefit of the loss or the proof?**

36 **A:** The correctness proof ensures that the topological loss enforces topological constraints. Topo-
37 logical constraints can significantly reduce the size of the solution space, making the optimization
38 much easier. We prove this for a simple connectivity constraint and leave general cases for future
39 work. For an $m \times m$ image patch, the space of all binary segmentations has $\text{DoF} = m \times m$ Degrees
40 of Freedom. But if we force the segmentation to be connected, the solution space is much smaller.



41 See the figure on the right. Within each column (or row) of pixels, the segmentation is an interval with $\text{DoF} = 2$
42 (locations of two ends). Over all rows and columns, the DoF of the connected segmentation space is only $4m$.

43 Furthermore, in practice, the topological loss combats sample bias. As noted in Q4, a network learns to predict most
44 pixels except for a few difficult locations. There are never sufficient samples at such difficult locations. A topological
45 loss mitigates such sample bias by identifying these locations as critical pixels and increasing their weights in the
46 training. Thus, training converges faster (Figure (b)) and topology-relevant metrics improve.

47 **Q6 (R2): What if cross entropy loss is turned off or downweighted? Relative importance of the two losses.**

48 **A:** Per Q5, the topological loss complements the cross-entropy loss by improving training efficiency and mitigating
49 sampling bias. However, without cross-entropy loss, inferring topology from a completely random likelihood map is
50 meaningless. Cross-entropy loss finds a reasonable likelihood map so that the topological loss can improve its topology.

51 **Q7 (R2): Robust to geometry errors: help in misalignment between the ground truth and training images?**

52 **A:** Absolutely. This is an additional benefit of the topological loss that increases robustness to geometry errors.

53 **Q8 (R3): Can both types of metrics be maximized? If not, can post-processing help?**

54 With sufficient samples for the difficult locations (blurred locations), we can maximize both types of metrics. However,
55 this is not the case in reality. We use topological loss to force the model to memorize patterns at the difficult locations,
56 at the expense of overfitting and consequently slightly compromised per-pixel accuracy. Post-processing at test time
57 should help in general, but does not identify these difficulty locations for specific handling.