Stony Brook University

Learning Topological Representations for Deep Image Understanding

Xiaoling Hu

Committee: Prof. Dimitris Samaras, Prof. Haibin Ling
Prof. Chao Chen (Advisor), Prof. Fuxin Li

Topology is everywhere

Image

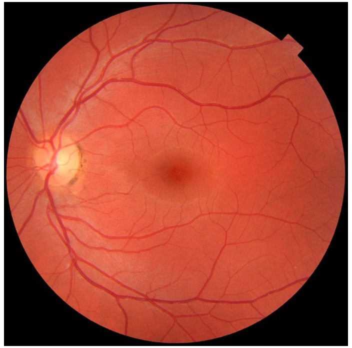

Point Cloud

Graph

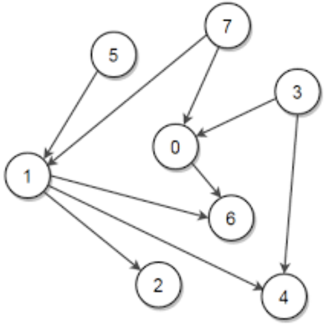

Network

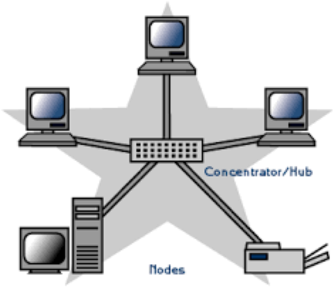

...

Road network reconstruction

Topological correctness is one of the main challenges in road network reconstruction!

Zhu, Lingli, et al. "Automated 3D scene reconstruction from open geospatial data sources: Airborne laser scanning and a 2D topographic database." ***Remote Sensing*** (2015).

Neuron reconstruction

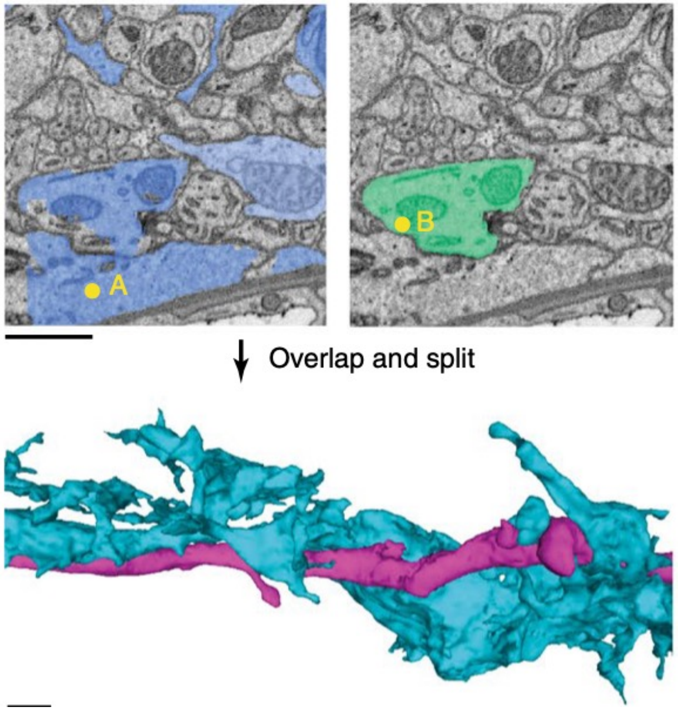

Incorrect topology results in incorrect merge or split!

Januszewski, Michał, et al. "High-precision automated reconstruction of neurons with flood-filling networks." ***Nature methods*** (2018).

Vasculature morphology measurement

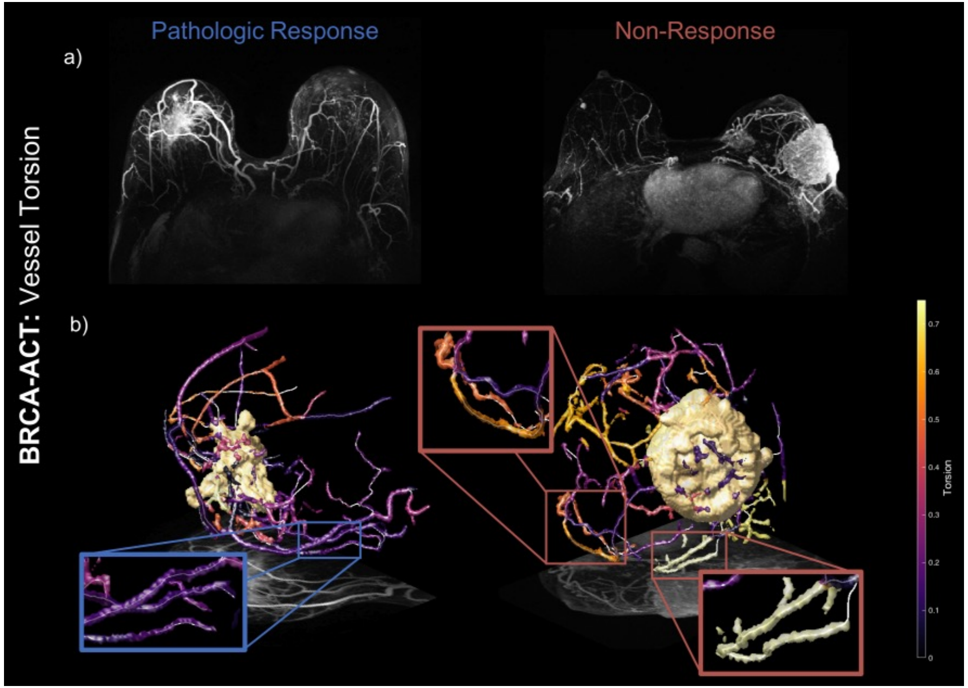

Correct measurement of topology leads to powerful biomarkers!

Braman, Nathaniel, et al. "Novel Radiomic Measurements of Tumor-Associated Vasculature Morphology on Clinical Imaging as a Biomarker of Treatment Response in Multiple Cancers." ***Clinical Cancer Research*** (2022).

Quantification of mouse brain

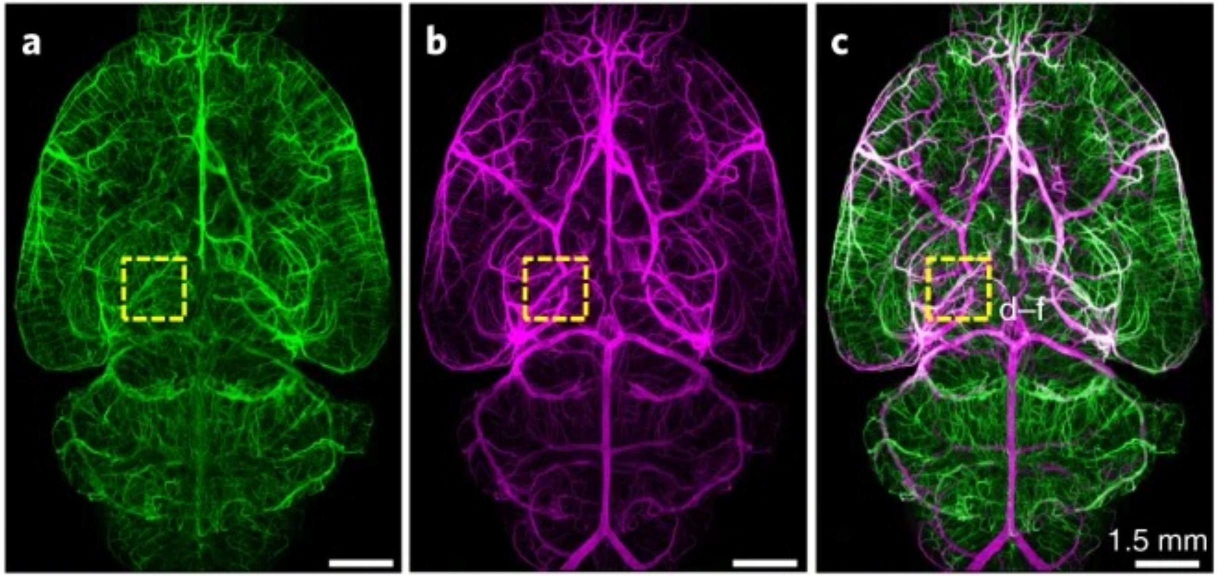

Structure yields biological insights into the vascular function of the brain!

Todorov, Mihail Ivilinov, et al. "Machine learning analysis of whole mouse brain vasculature." ***Nature methods*** (2020).

Data with rich structural information

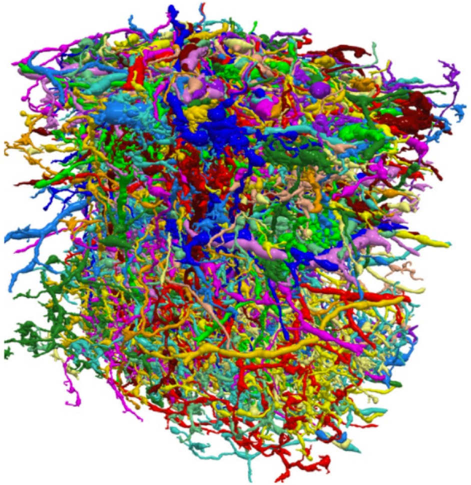

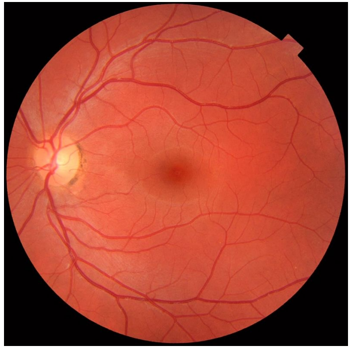

- **Challenges**
  - ➢ Complex topology/structure
  - ➢ Noisy data
  - ➢ Limited labels

Difficult to analyze!

My research: Structure-informed image analysis

- **Problems**
  - ➢ Analyze structures with complex topology and geometry
- **Applications**
  - ➢ Vessels, neurons, cells, etc.
- **Challenges**
  - ➢ Incorrect topology (extraction)
  - ➢ Purely data-driven methods are topology-agnostic
  - ➢ Annotations are hard to obtain

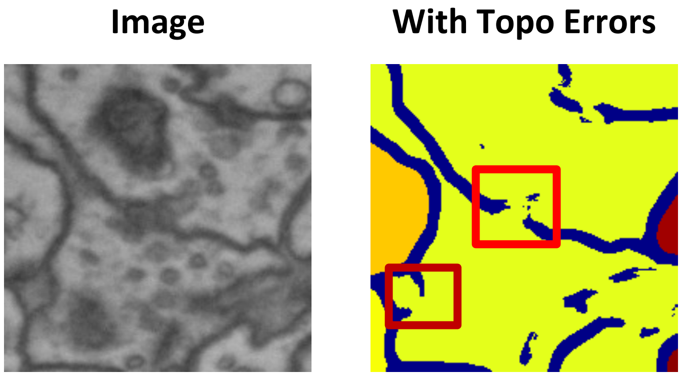

Hu et al. *NeurIPS'19*.

Key methodology: Topology

- **Explicit modeling of complex structures from data**
  - Math theory: algebraic topology
    - Persistent homology
    - Discrete Morse theory
    - Homotopy warping
- **Combine topology with deep learning organically**
  - Differentiability

DL methods incorporating classical topological theory

- **Persistent homology**

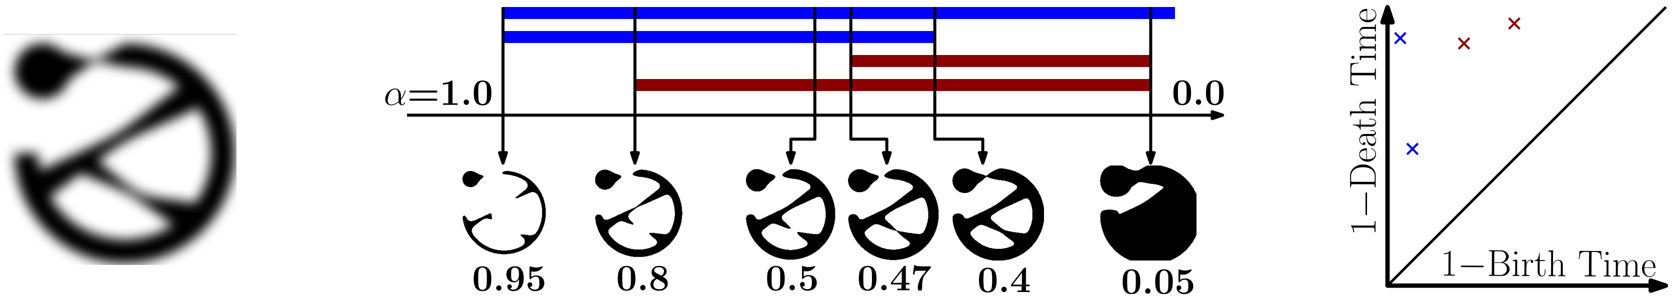

- **Discrete Morse theory**

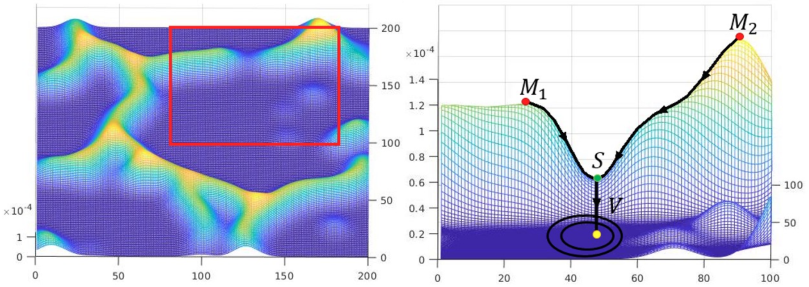

- **Homotopy warping**

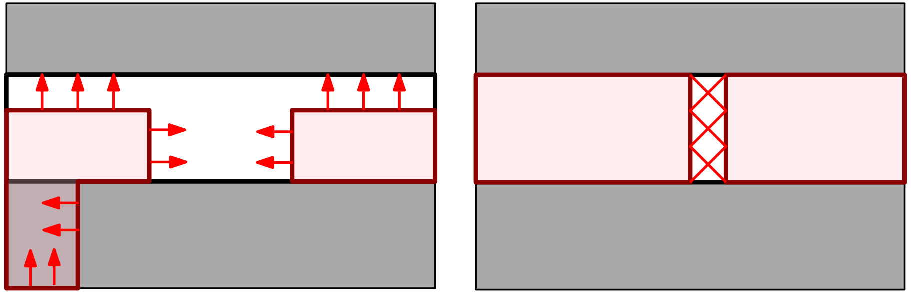

Hu et al. *NeurIPS'19*, Hu et al. *ICLR'21*, Hu et al. *NeurIPS'22.*

Research overview

Methodology

- *Teach neural network to think in* ***topology/structure***
- *Topological/structural* ***prior/regularization/feature***

Toolbox

- Persistent homology
- Discrete Morse theory
- Digital topology

Applications

- Image segmentation
- Uncertainty estimation
- ……

Future Directions

- Imperfect data
- Topological quantification
- ……

Outline

- **Topological loss for image segmentation**
  - Persistent homology based loss (NeurIPS'19, MICCAI'21)
  - Homotopy warping loss (NeurIPS'22)
- **Beyond pixel-wise representation**
  - Discrete Morse theory loss (ICLR'21, Spotlight)
  - Topological/Structural representation of images (ICLR'23, Spotlight)
- **Extensions**
  - Trojan detection with topological prior (ICLR'22)
- **Future work**

CNN are prone to topological errors

- Existing methods optimize w.r.t. per-pixel accuracy
- Topological errors
  - ➢ broken connection, missing components
- Broken membranes
  - ➢ small per-pixel error
  - ➢ topological error – large error in downstream analysis

Image | GT | Classic CNN | Our method

Solution: Topological loss [NeurIPS'19]

- Loss function – train the model to be topology-preserving
  - Repeatedly evaluate the accuracy in topology and force the model to correct mistakes
- Betti number
  - 0-dim betti number b_0: number of connected components
  - 1-dim betti number b_1: number of cycles
  - 2-dim betti number b_2: number of voids
  - ……

$$b_0 = 1, b_1 = 11, b_{n\geq 2} = 0$$

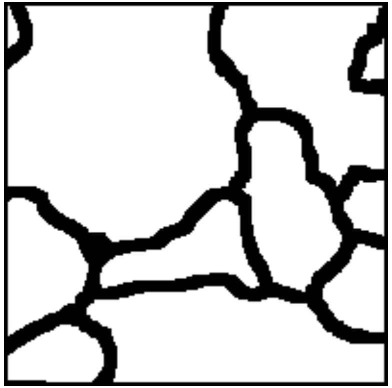

Solution: Topological loss [NeurIPS'19]

- Betti number error – topological difference between seg. and GT

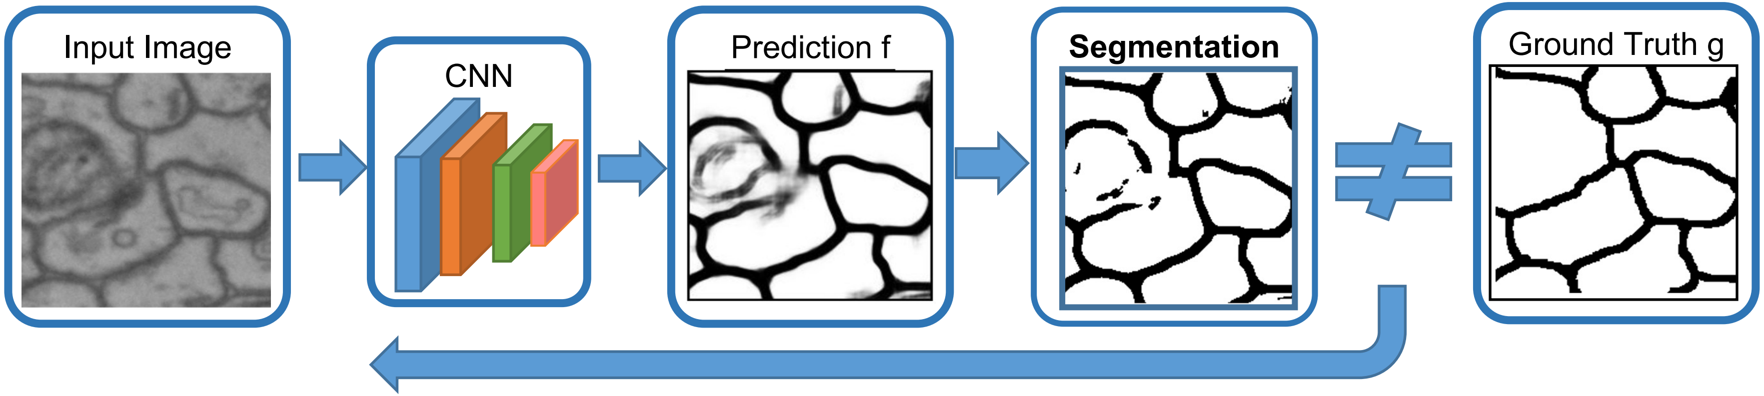

- Challenge: **not differentiable**;
- Solution: persistent homology

Hu, Xiaoling, et al. "Topology-Preserving Deep Image Segmentation." ***NeurIPS*** (2019).

Intuition: focus on likelihood f

- How far is f from generating a segmentation X with correct topology?
- What is the best way to fix X by changing f?
- Compare with a worse likelihood f' (same segmentation X)

Likelihood f'

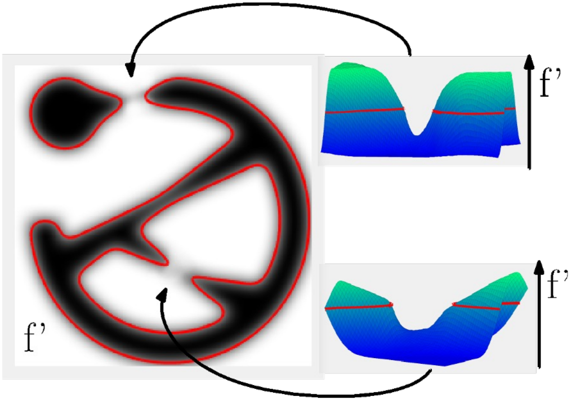

Ground Truth

$\neq$

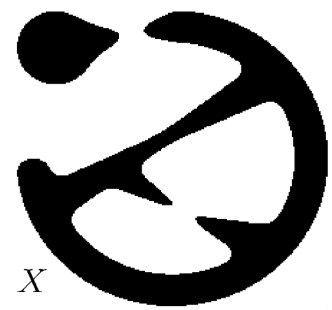

Segmentation

Likelihood

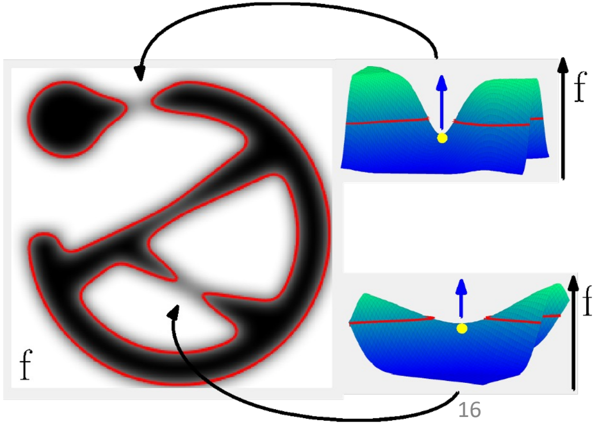

Persistent homology

- Superlevel set: $f^{\alpha} = \{x | \, f(x) \geq \ \alpha\}$
- Dgm(f) = {death(p), birth(p)}, 0-dim (blue), 1-dim (burgundy)

Likelihood f

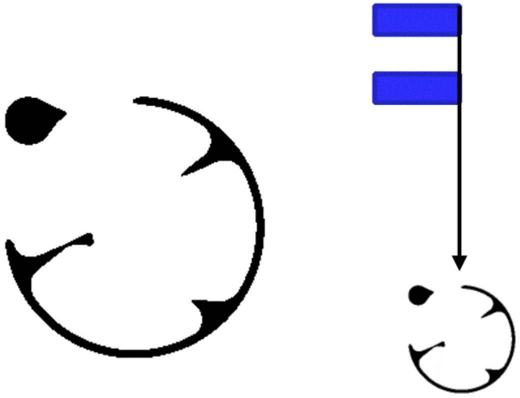

Persistence barcodes/diagrams

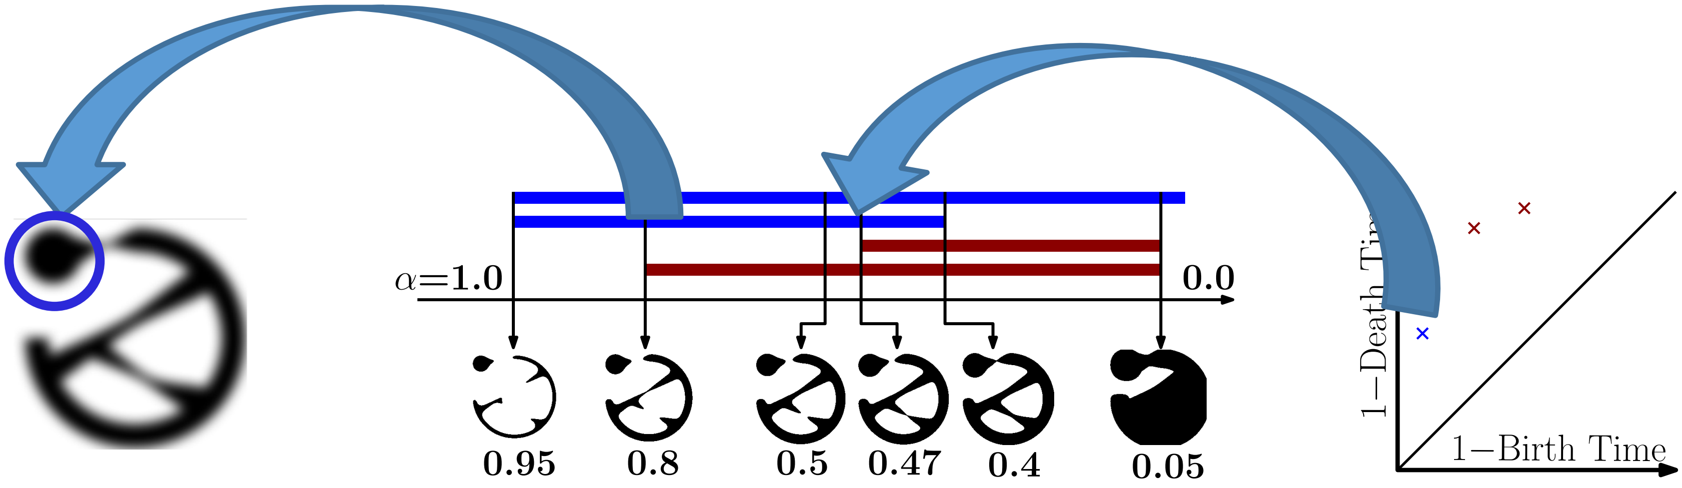

Dgm(g)

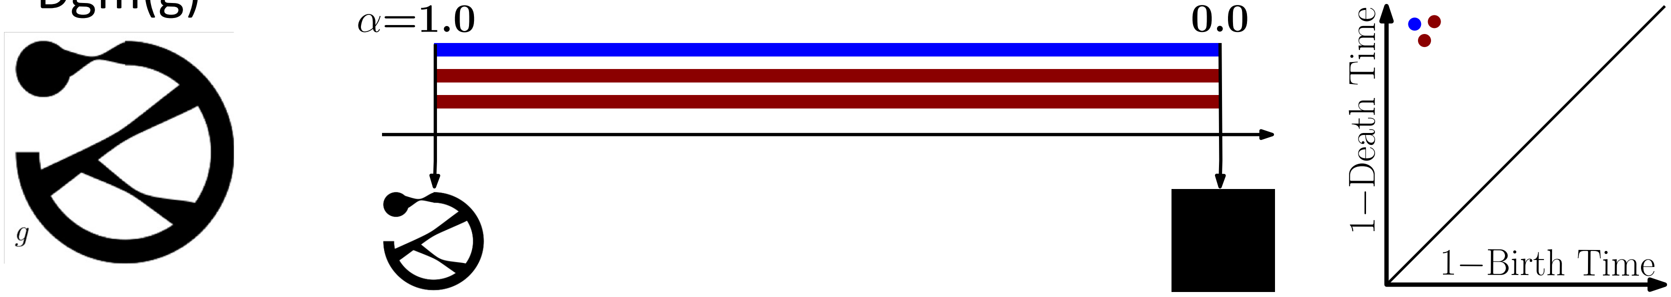

Topological loss (L_{topo}) = Distance between diagrams

$$\min_{\gamma\in\Gamma} \sum_{p\in \mathrm{Dgm}(f)} ||p-\gamma(p)||^2 = \sum_{p\in \mathrm{Dgm}(f)} [\mathrm{birth}(p) - \mathrm{birth}(\gamma^*(p))]^2 + [\mathrm{death}(p) - \mathrm{death}(\gamma^*(p))]^2$$

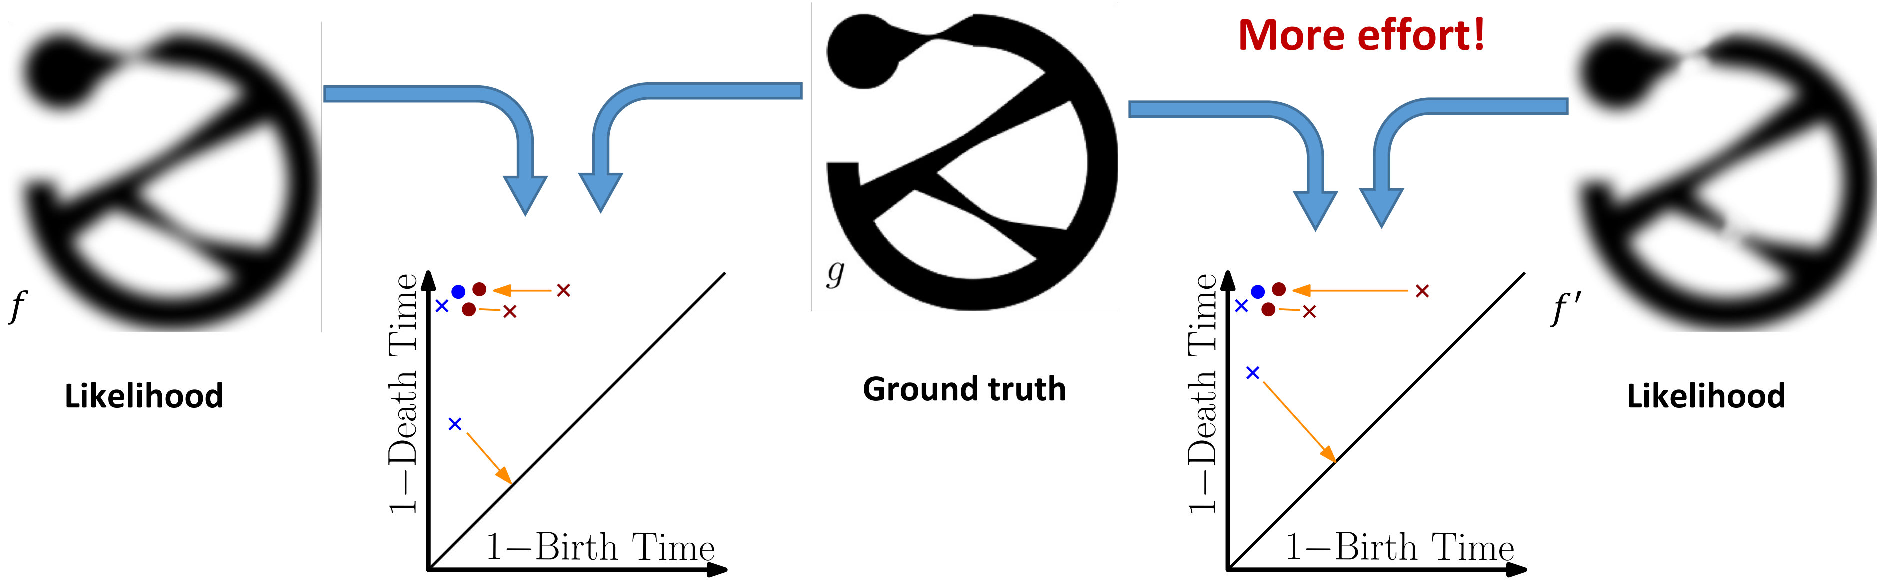

Differentiable?

$$\min_{\gamma\in\Gamma} \sum_{p\in \mathrm{Dgm}(f)} ||p-\gamma(p)||^2 = \sum_{p\in \mathrm{Dgm}(f)} [\mathrm{birth}(p) - \mathrm{birth}(\gamma^*(p))]^2 + [\mathrm{death}(p) - \mathrm{death}(\gamma^*(p))]^2$$

- Depending on critical thresholds: topological changes happen, e.g., birth/death times
- Uniquely determined by the critical locations/pixels (including maxima, minima and saddle points)

Topological loss is actually defined on specific pixels!

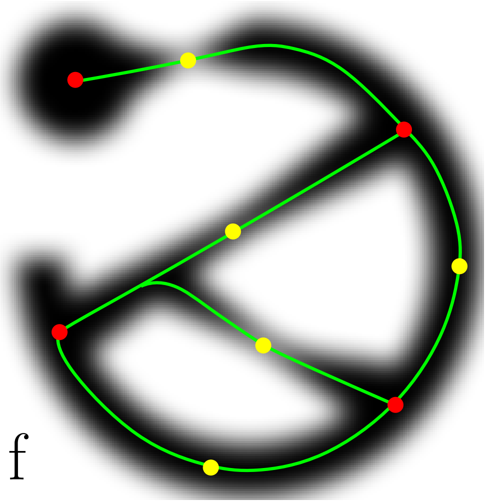

Gradient

- $c_b(p)$ – birth critical point of p; $c_d(p)$ – death critical point of p
- Assume γ^* fixed, $c_b(p)$ and $c_d(p)$ are fixed

$$\sum_{p \in \mathrm{Dgm}(f)} 2[f(c_b(p)) - \mathrm{birth}(\gamma^*(p))] \boxed{\frac{\partial f(c_b(p))}{\partial \omega}} + 2[f(c_d(p)) - \mathrm{death}(\gamma^*(p))] \boxed{\frac{\partial f(c_d(p))}{\partial \omega}}$$

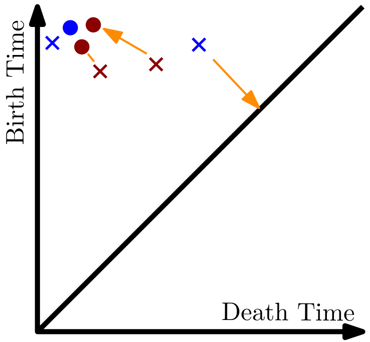

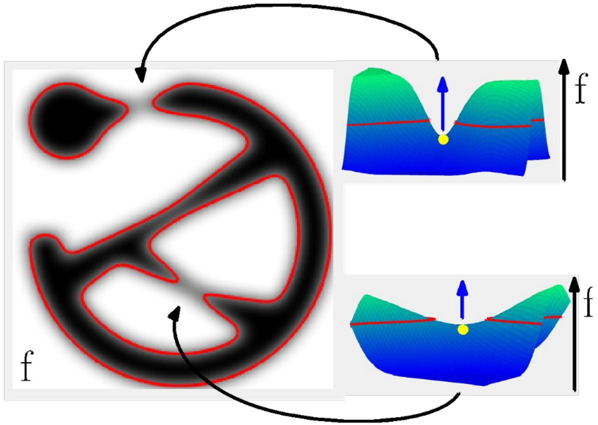

How it looks like in training?

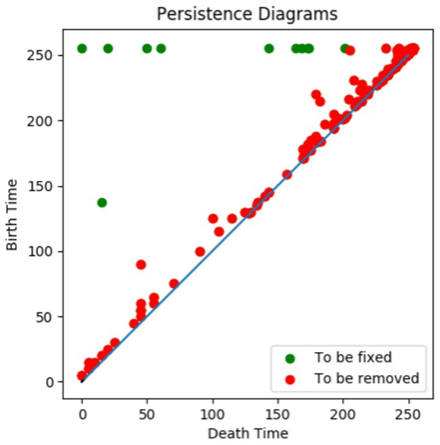

Epoch 10

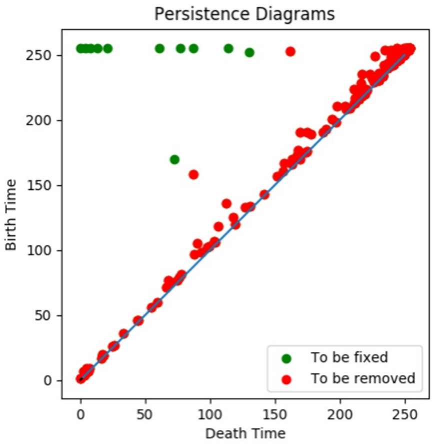

Epoch 30

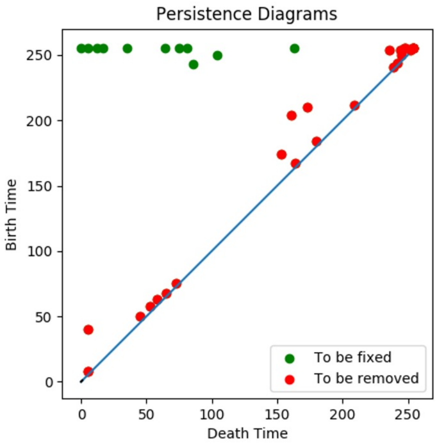

Epoch 50

- Green dots: to be fixed; Red dots: to be removed
- As training continues
  - ➢ fewer red dots as noises are removed
  - ➢ green dots move closer to upper left corner (getting fixed)

Localized topological loss

- Evaluate topology in whole image
  - ➢ Expensive
  - ➢ Too many dots: matching is difficult
- Localized topological measure: relative homology
- Sample patches over images and evaluate topology loss

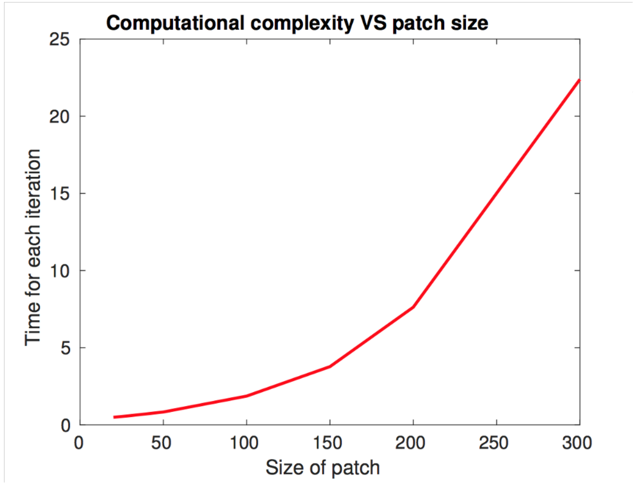

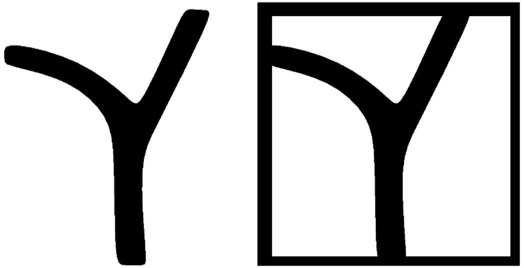

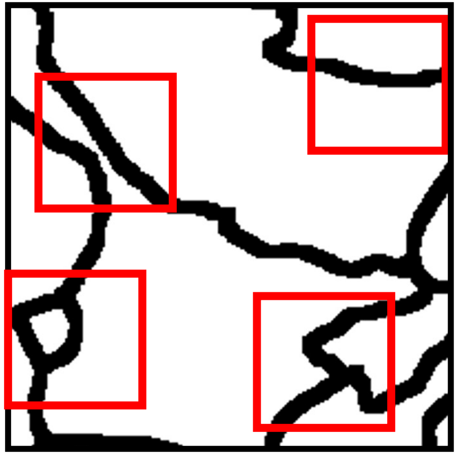

Overview

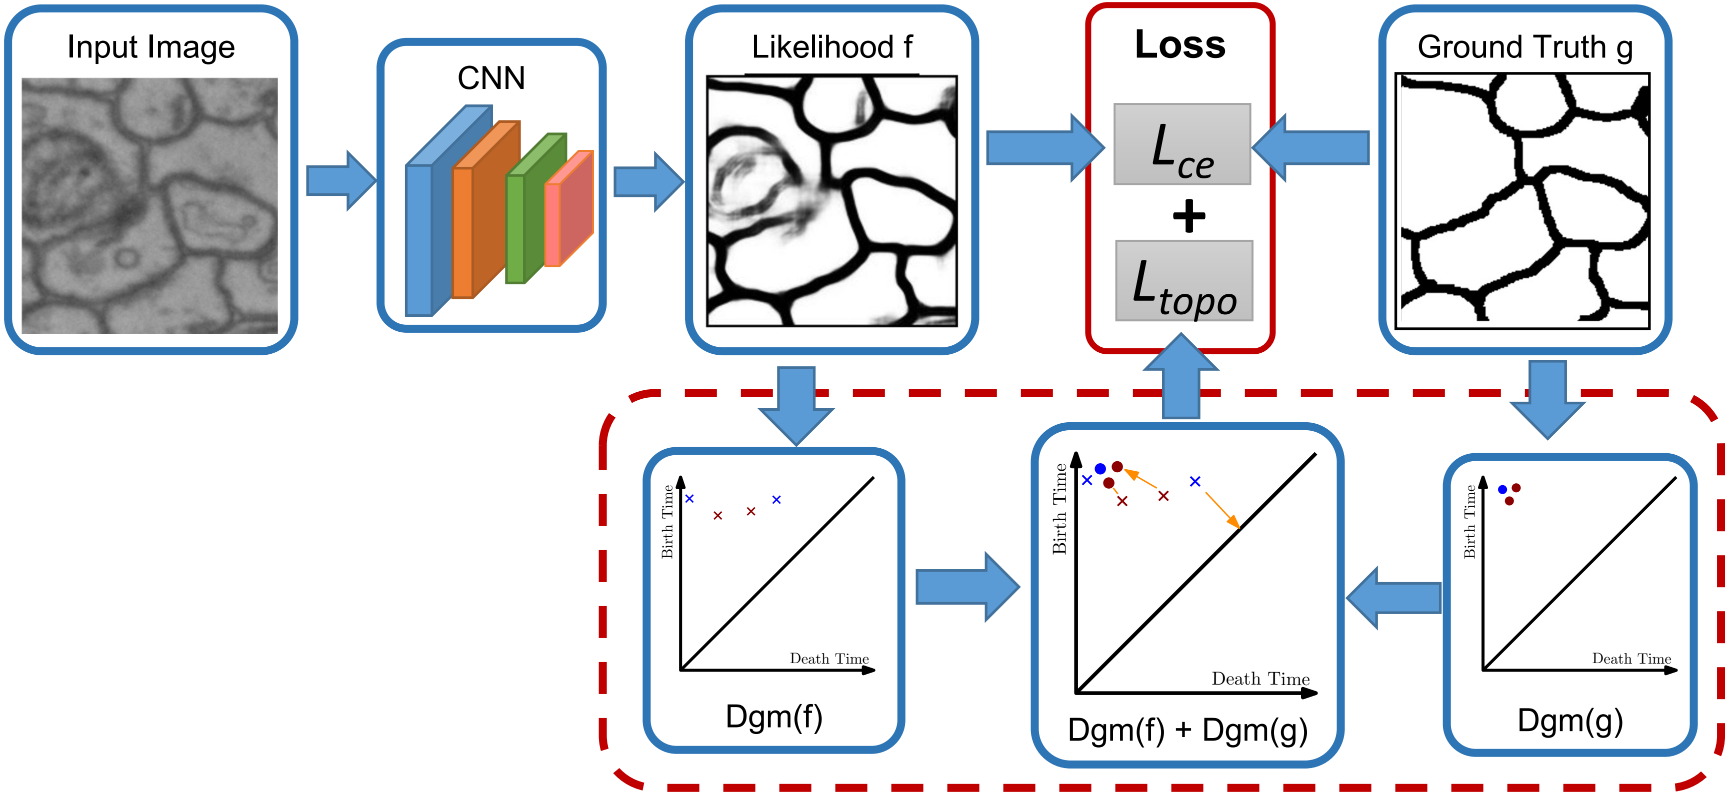

Input Image
CNN
Likelihood f
Loss
L_{ce}
+
L_{topo}
Ground Truth g
Birth Time
Death Time
Dgm(f)
Dgm(f) + Dgm(g)
Dgm(g)

Qualitative results

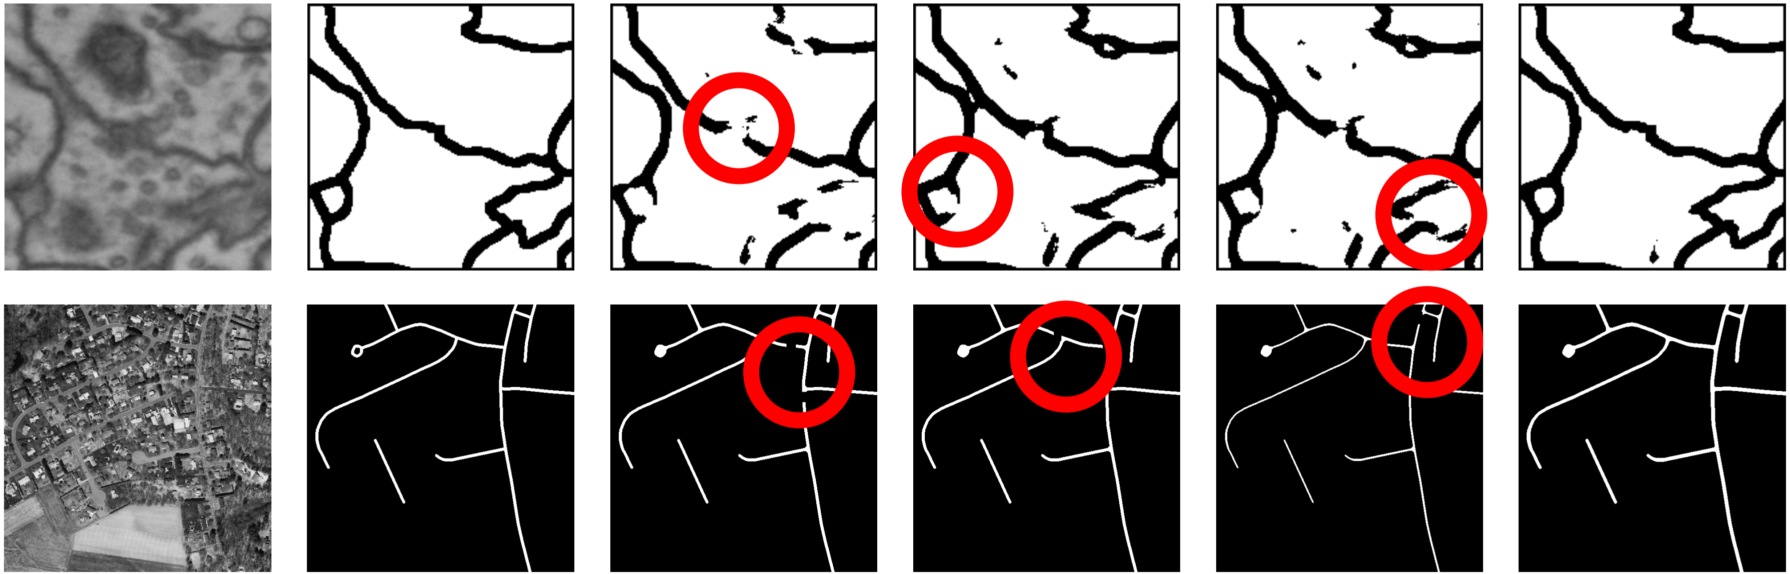

Input	Ground Truth	DIVE [Bioinformatics'16]	U-Net [MICCAI'16]	VGG-UNet [CVPR'16]	TopoNet

Quantitative results

- Per-pixel accuracy, Adapted Rand Index, Variation of Information
- Betti number error
  - Measures topological differences/errors directly
- EM (neuron) image segmentation datasets

Dataset	Method	Accuracy	ARI	VOI	Betti Error
ISBI12	**DIVE**	0.9640 ± 0.0042	0.9434 ± 0.0087	1.235 ± 0.025	3.187 ± 0.307
	U-Net	$\mathbf{0.9678 \pm 0.0021}$	0.9338 ± 0.0072	1.367 ± 0.031	2.785 ± 0.269
	Mosin.	0.9532 ± 0.0063	0.9312 ± 0.0052	0.983 ± 0.035	1.238 ± 0.251
	TopoLoss	0.9626 ± 0.0038	$\mathbf{0.9444 \pm 0.0076}$	$\mathbf{0.782 \pm 0.019}$	$\mathbf{0.429 \pm 0.104}$
ISBI13	**DIVE**	$\mathbf{0.9642 \pm 0.0018}$	0.6923 ± 0.0134	2.790 ± 0.025	3.875 ± 0.326
	U-Net	0.9631 ± 0.0024	0.7031 ± 0.0256	2.583 ± 0.078	3.463 ± 0.435
	Mosin.	0.9578 ± 0.0029	0.7483 ± 0.0367	1.534 ± 0.063	2.952 ± 0.379
	TopoLoss	0.9569 ± 0.0031	$\mathbf{0.8064 \pm 0.0112}$	$\mathbf{1.436 \pm 0.008}$	$\mathbf{1.253 \pm 0.172}$
CREMI	**DIVE**	$\mathbf{0.9498 \pm 0.0029}$	0.6532 ± 0.0247	2.513 ± 0.047	4.378 ± 0.152
	U-Net	0.9468 ± 0.0048	0.6723 ± 0.0312	2.346 ± 0.105	3.016 ± 0.253
	Mosin.	0.9467 ± 0.0058	0.7853 ± 0.0281	1.623 ± 0.083	1.973 ± 0.310
	TopoLoss	0.9456 ± 0.0053	$\mathbf{0.8083 \pm 0.0104}$	$\mathbf{1.462 \pm 0.028}$	$\mathbf{1.113 \pm 0.224}$

Segmentation during training

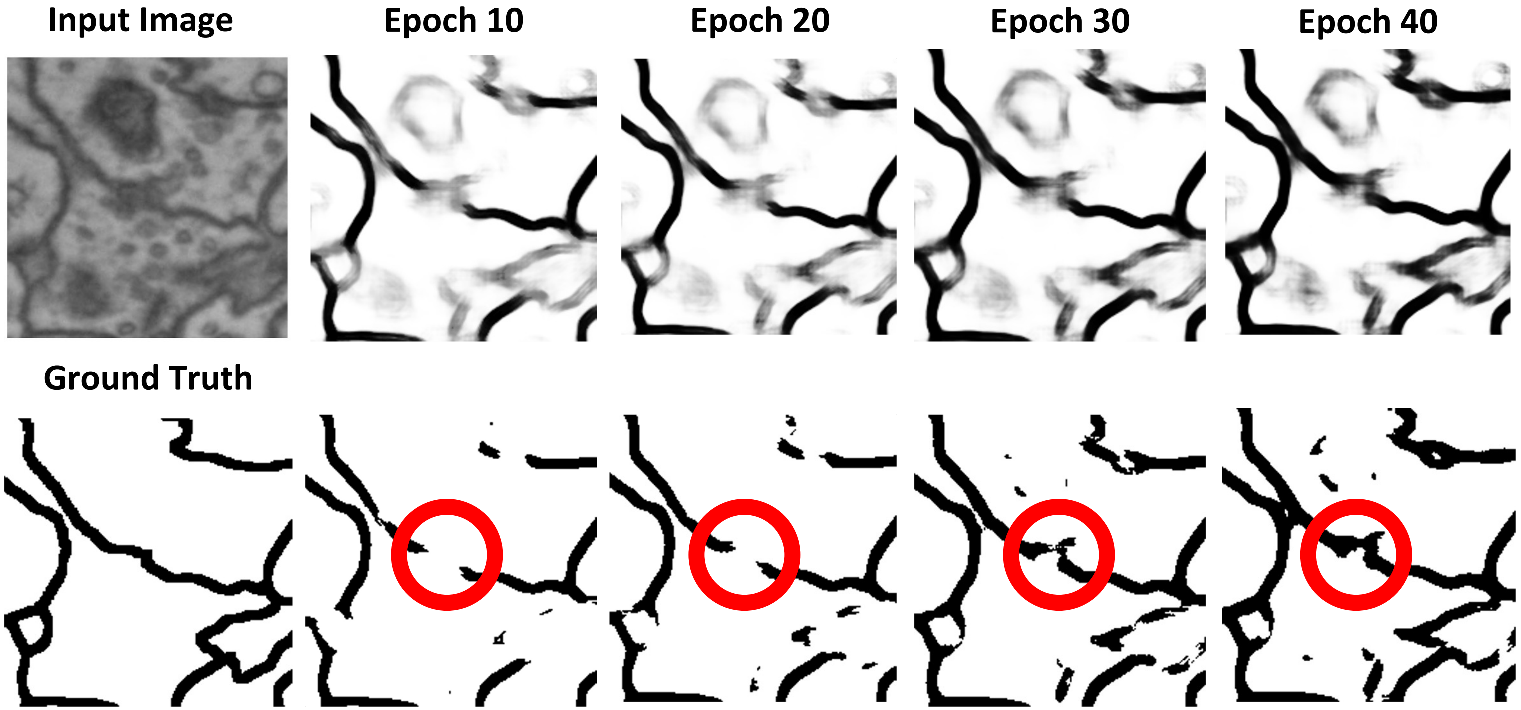

Rationale of topological loss

- Likelihood map/segmentation are stabilized globally
  - Mostly because of cross-entropy loss
- Topological errors are gradually fixed by the topological loss
  - Likelihood map only change at topology-relevant locations
- Topological loss compliments cross-entropy
  - Combating sampling bias
  - Identifying difficult locations and increase their weights in training
- Without cross-entropy loss, inferring topology from a completely random likelihood map is meaningless

3D topology-preserving segmentation (MICCAI'21)

- Topological loss for 3D anisotropic cases
  - *Spatial topological-attention*: Spatial topological information across adjacent slices
  - *Iterative topological-attention:* Improve the stability of the topologically critical maps

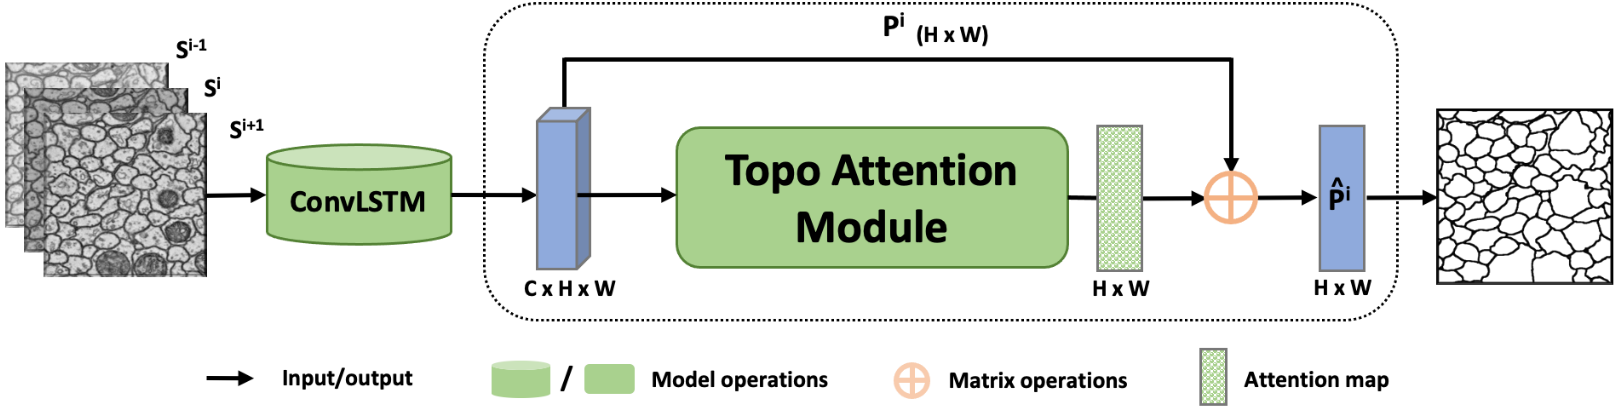

Yang, Jiaqi*, **Hu, Xiaoling***, et al. "A Topological-Attention ConvLSTM Network and Its Application to EM Images." ***MICCAI*** (2021).

Homotopy warping loss (NeurIPS'22)

- **Efficiently identify topological critical points**
  - [Hu et al. NeurIPS'19] – Topological loss by matching persistence diagram

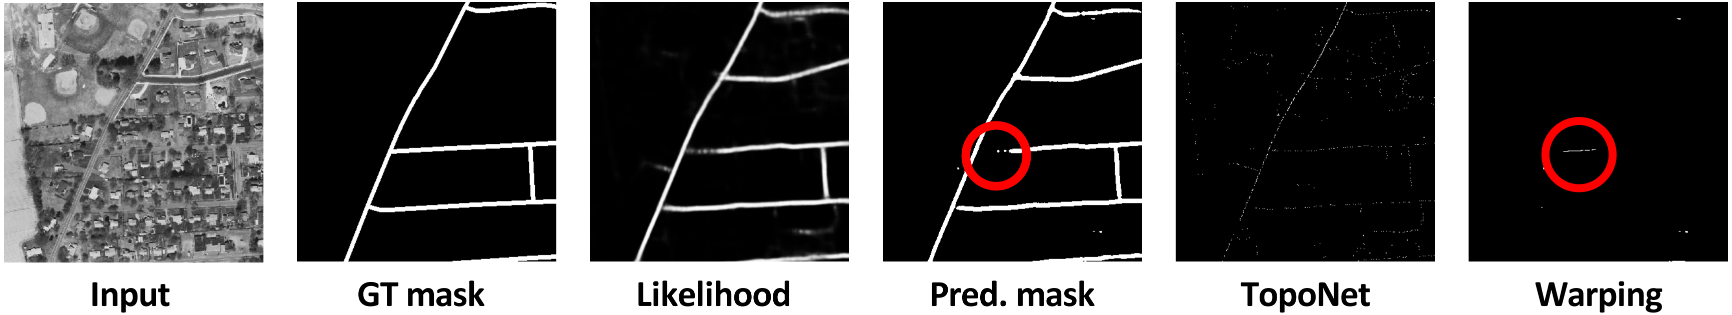

Hu, Xiaoling. "Structure-Aware Image Segmentation with Homotopy Warping." ***NeurIPS*** (2022).

Simple points

- Flipping the label of p will not change the topology

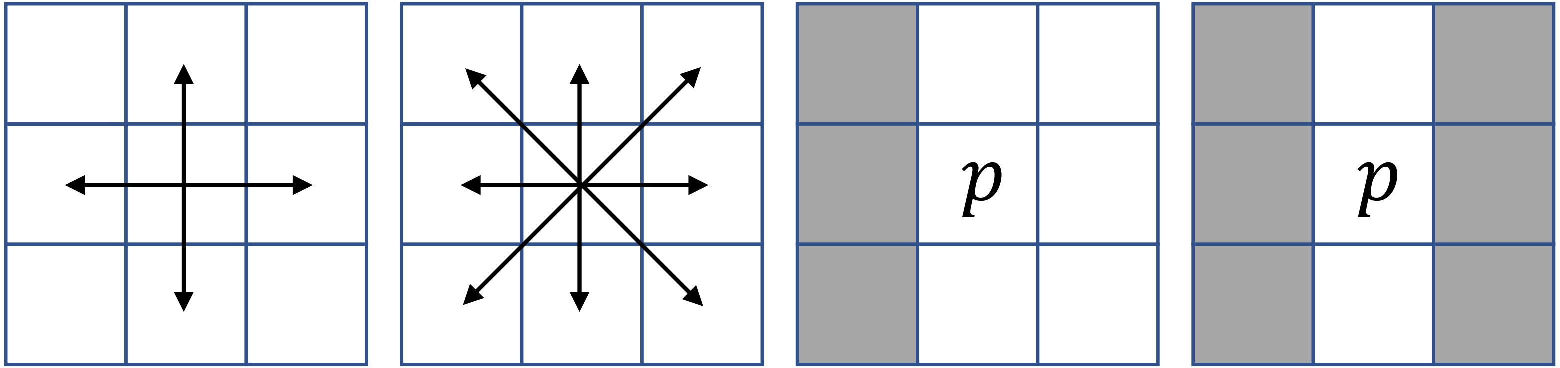

4-adjacent 8-adjacent Simple Point Non-simple Point

Homotopy warping: flip simple points

- Warping red mask towards the white

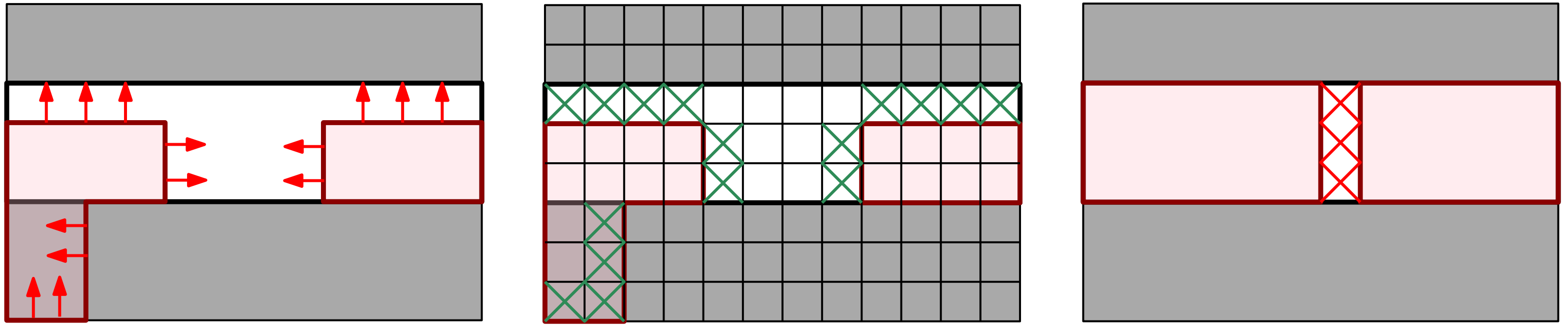

- Warping white mask towards the red

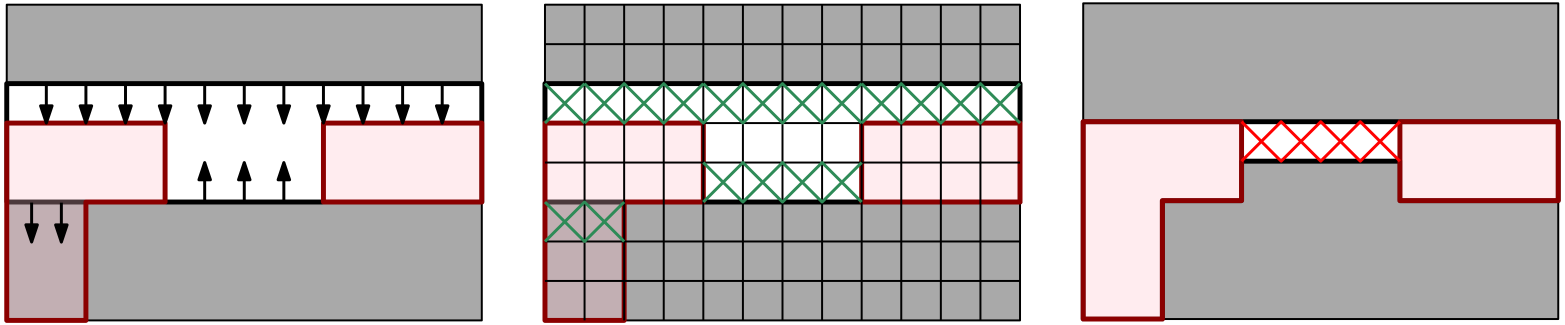

Warping example

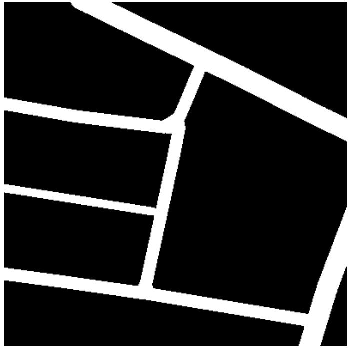

GT

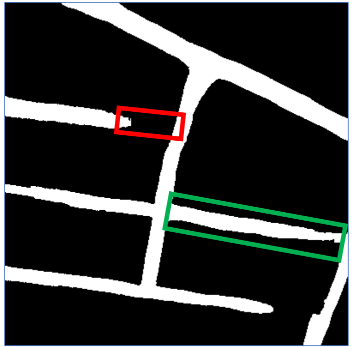

Prediction

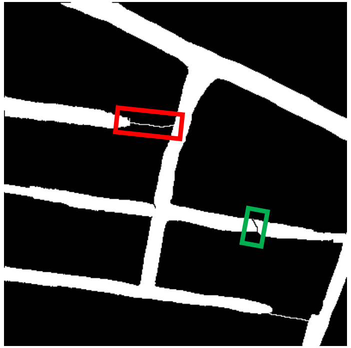

Warp GT

Zoom-in

Zoom-in

Question: Pixel VS structure representations?

- **Models make pixel level predictions**

Ground Truth

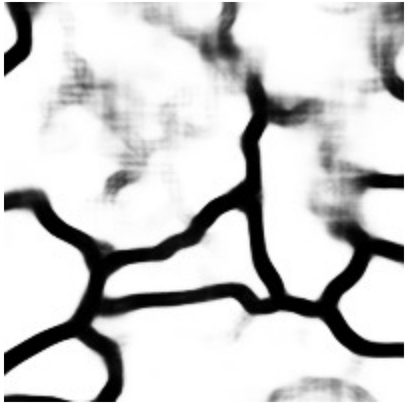

Likelihood Map

Segmentation

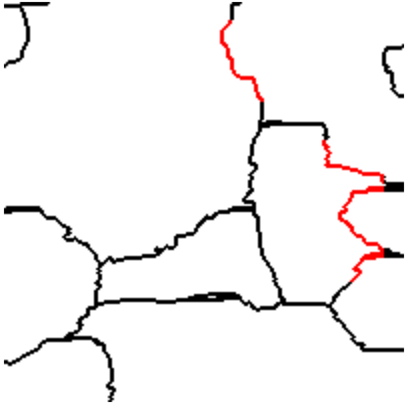

Structure Inference

Structure-level reasoning/prediction/inference!

Outline

- **Topological loss for image segmentation**
  - Persistent homology based loss (NeurIPS’19, MICCAI’21)
  - Homotopy warping loss (NeurIPS’22)
- **Beyond pixel-wise representation**
  - Discrete Morse theory loss (ICLR’21, Spotlight)
  - Topological/Structural representation of images (ICLR’23, Spotlight)
- **Extensions**
  - Trojan detection with topological prior (ICLR’22)
- **Future work**

Structural loss

- Fix topological errors by identifying critical pixels
  - ➢ [Hu et al. NeurIPS'19] – Topological loss by matching persistence diagram

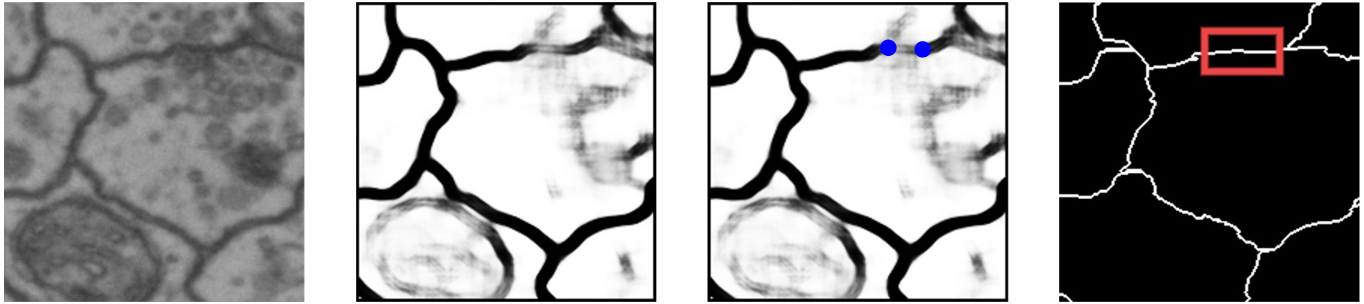

Input Image   Likelihood Map   Critical Points   DMT Structures

Not efficient enough!

Discrete Morse theory

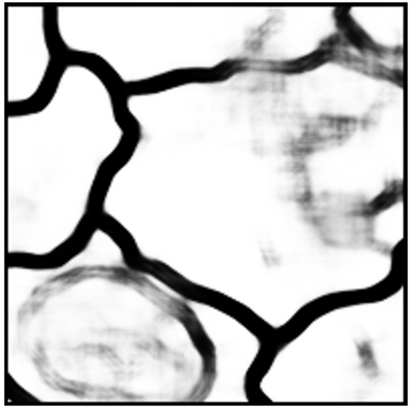

Likelihood Map

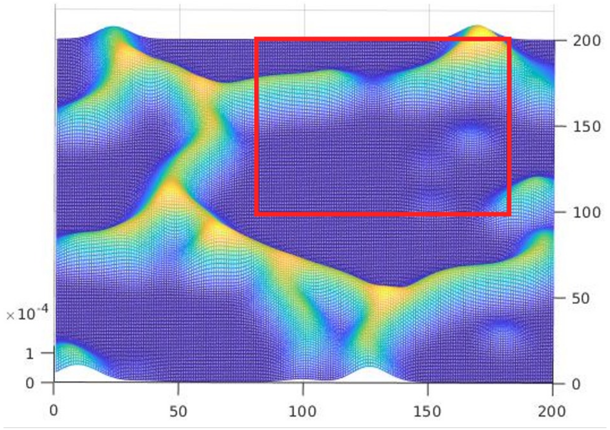

Density Map

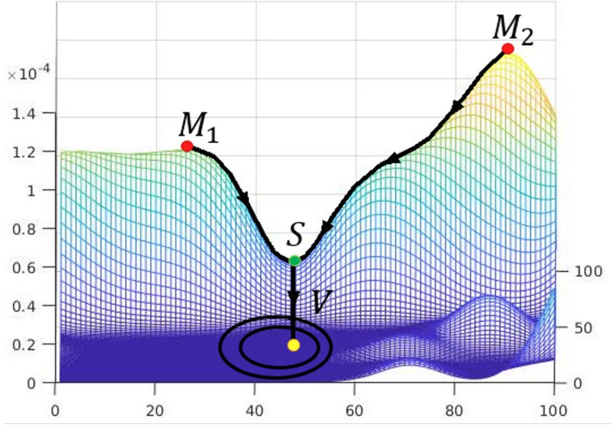

Density Map for highlighted region

Gradient: $\nabla f(x) = [\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \dots, \frac{\partial f}{\partial x_d}]^T$

Critical Points (minimum, maximum, saddle) $\nabla f(x) = 0$

Persistence-based structure pruning

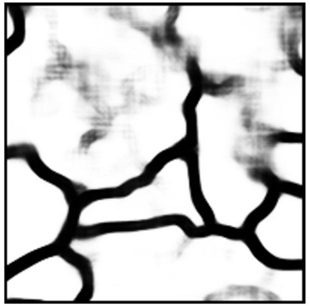

Likelihood Map

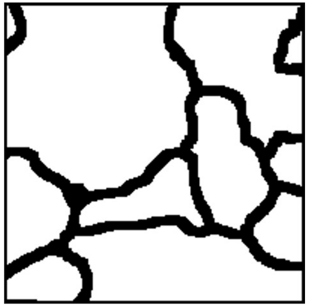

Ground Truth

Improperly pruned structures

Properly pruned structures

DMT loss [ICLR'21, Spotlight]

- Loss function – train the model to be topology-aware
  - ➢ Identity the critical structures instead of critical points

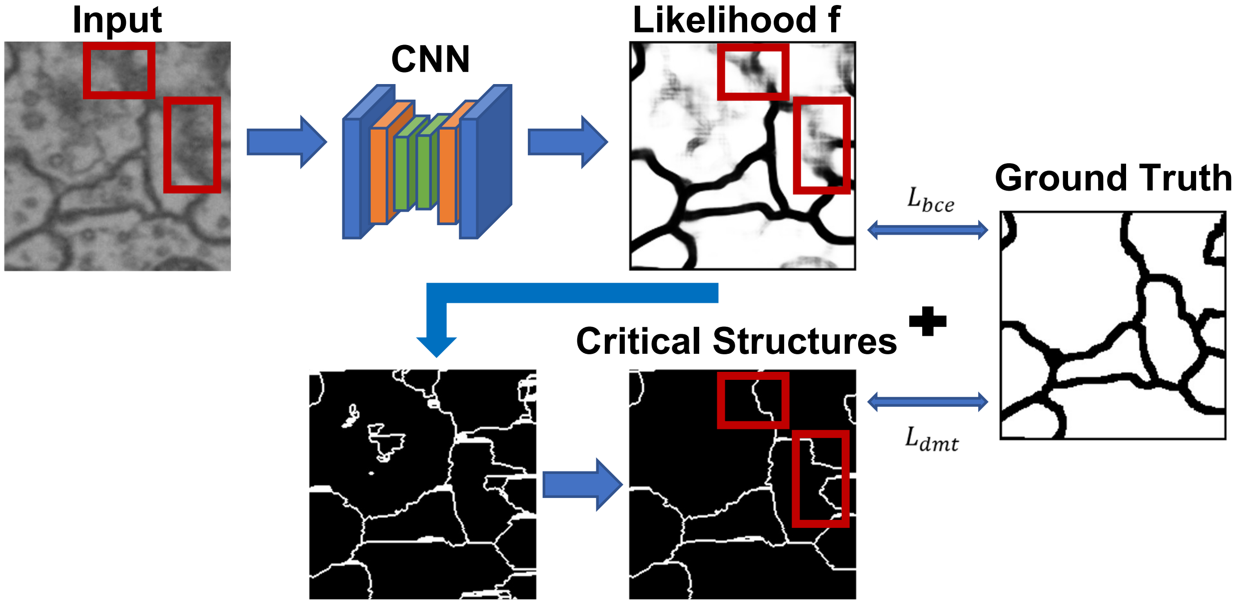

Hu, Xiaoling. "Topology-Aware Segmentation Using Discrete Morse Theory." *ICLR* (2021).

New problem

- **Limited annotated samples**
  - ➢ Annotations are expensive
  - ➢ Annotations can be noisy
- **Solution**
  - ➢ Active learning

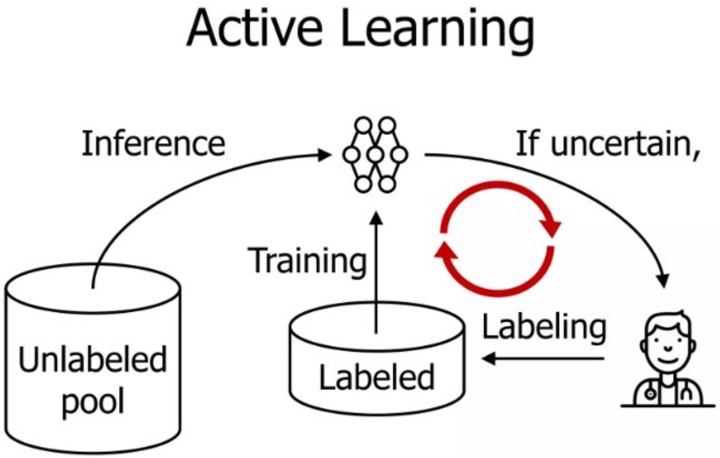

Yoo et al. *CVPR'19*.

The key of active learning is how to measure the uncertainty.

Li, Chen, **Hu, Xiaoling**, et al. "Confidence Estimation Using Unlabeled Data." ***ICLR*** (2023).

Uncertainty estimation

- **What is uncertainty estimation?**
  - ➢ How models are confident of their predictions

Question: How certain is classifier of this prediction? Can we quantify it?

→ Neural network classifier → Plane

- **Importance of uncertainty?**
  - ➢ Interpretability
  - ➢ Active learning

Probabilistic structural representation [ICLR'23, Spotlight]

- Structure-level uncertainty estimation
  - ➢ Focus on structure instead of pixel level
  - ➢ Easy to correct

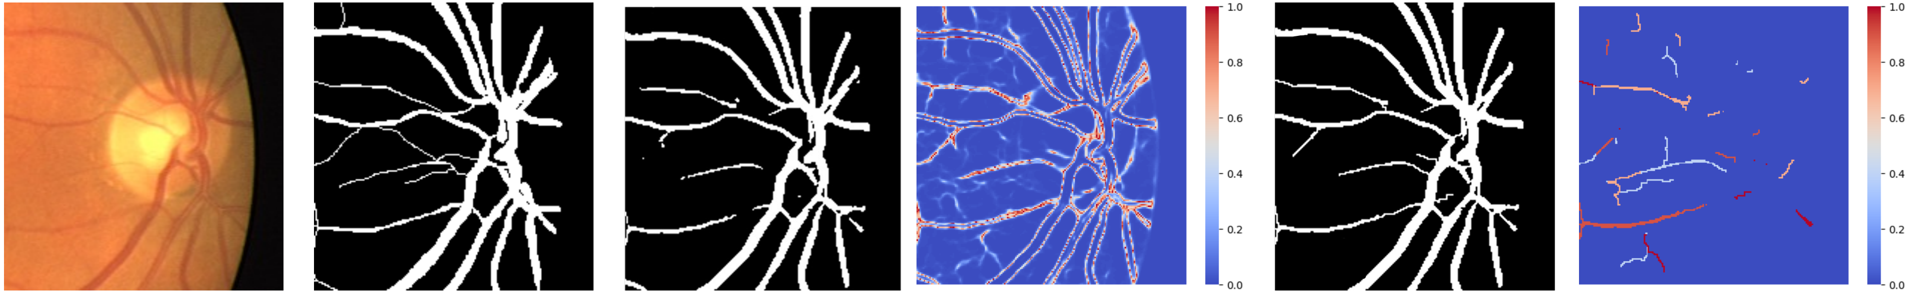

Image **GT** **Prob. UNet** **Pixel Uncer.** **Ours** **Stru. Uncer.**

Hu, Xiaoling, et al. "Learning Probabilistic Topological Representations Using Discrete Morse Theory." ***ICLR*** (2023).

Overview

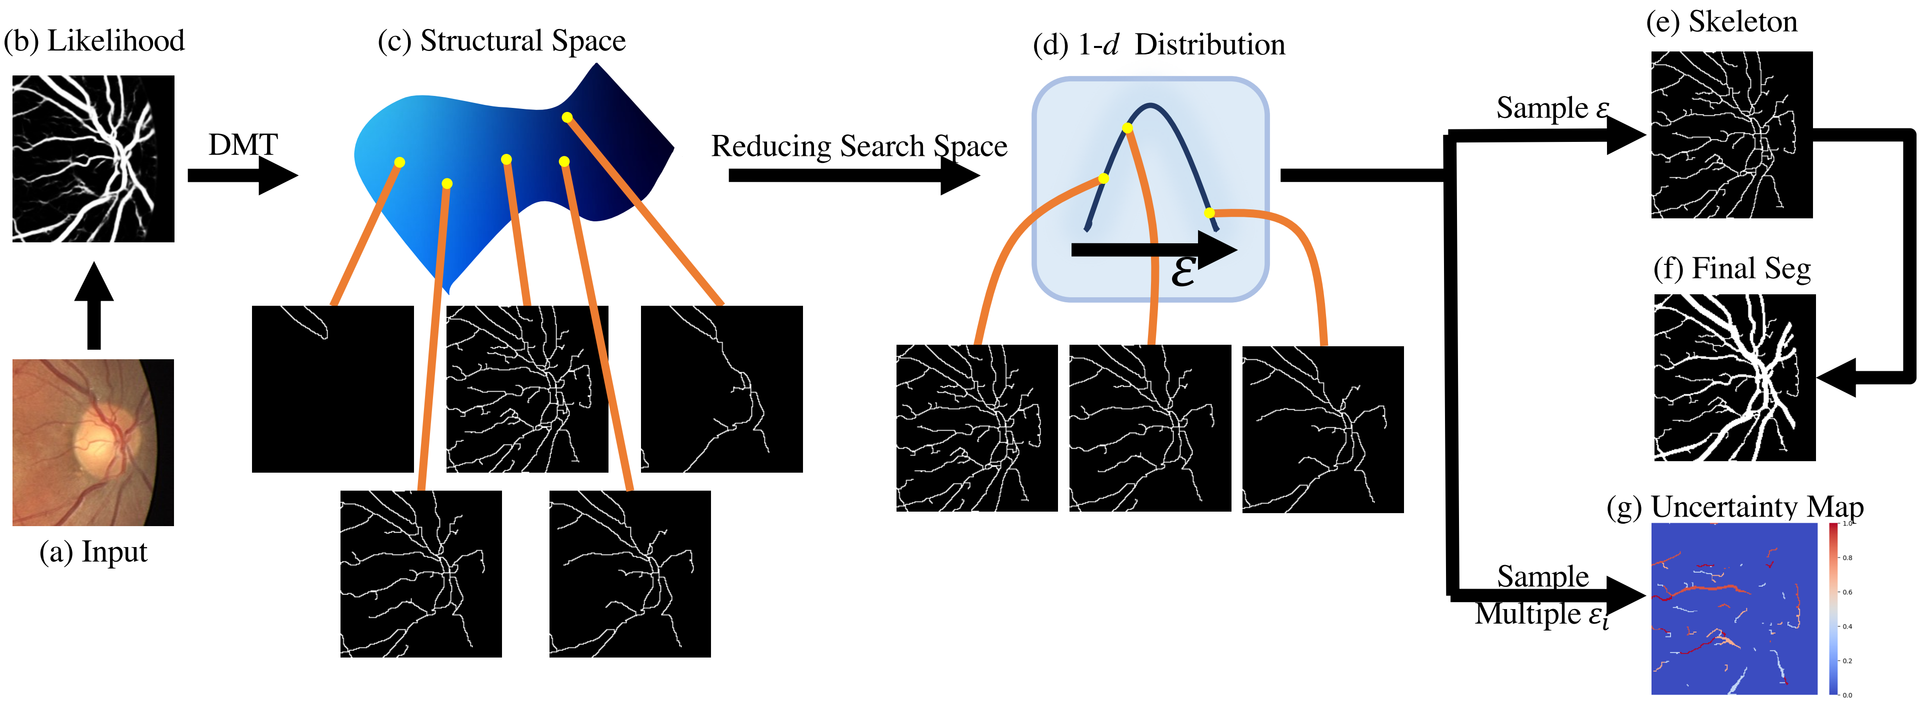
(b) Likelihood
DMT
(c) Structural Space
Reducing Search Space
(d) 1-d Distribution
ε
(e) Skeleton
Sample ε
(f) Final Seg
(g) Uncertainty Map
Sample
Multiple ε_i
(a) Input

Probabilistic model

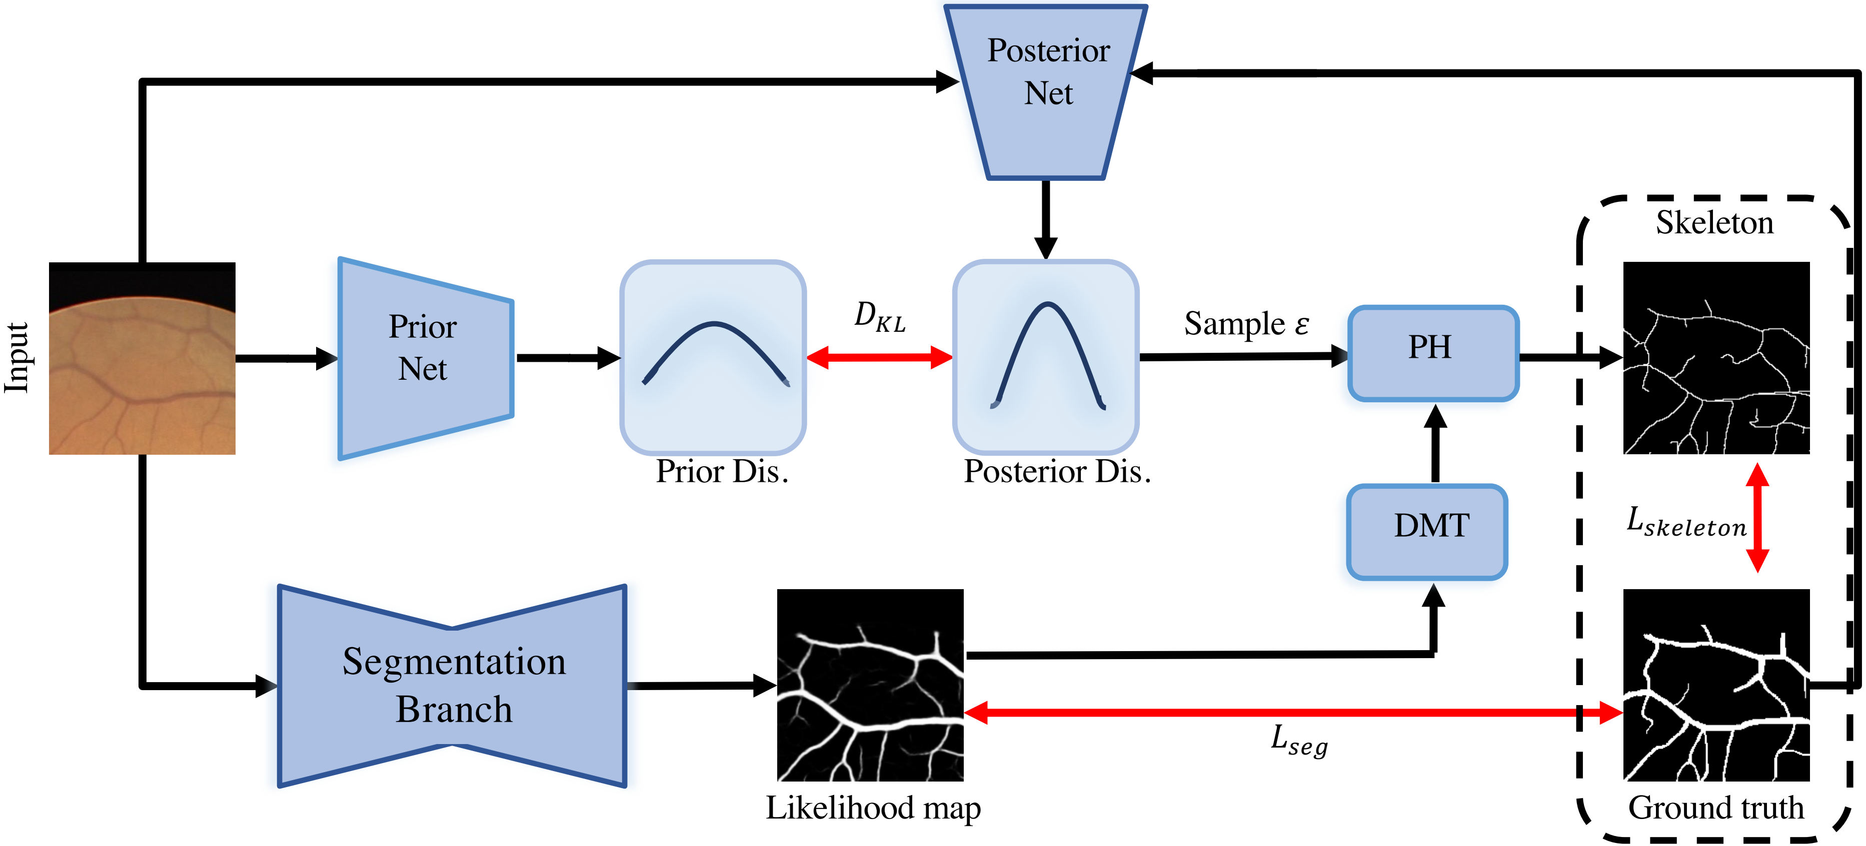

Inference and efficient annotation pipeline

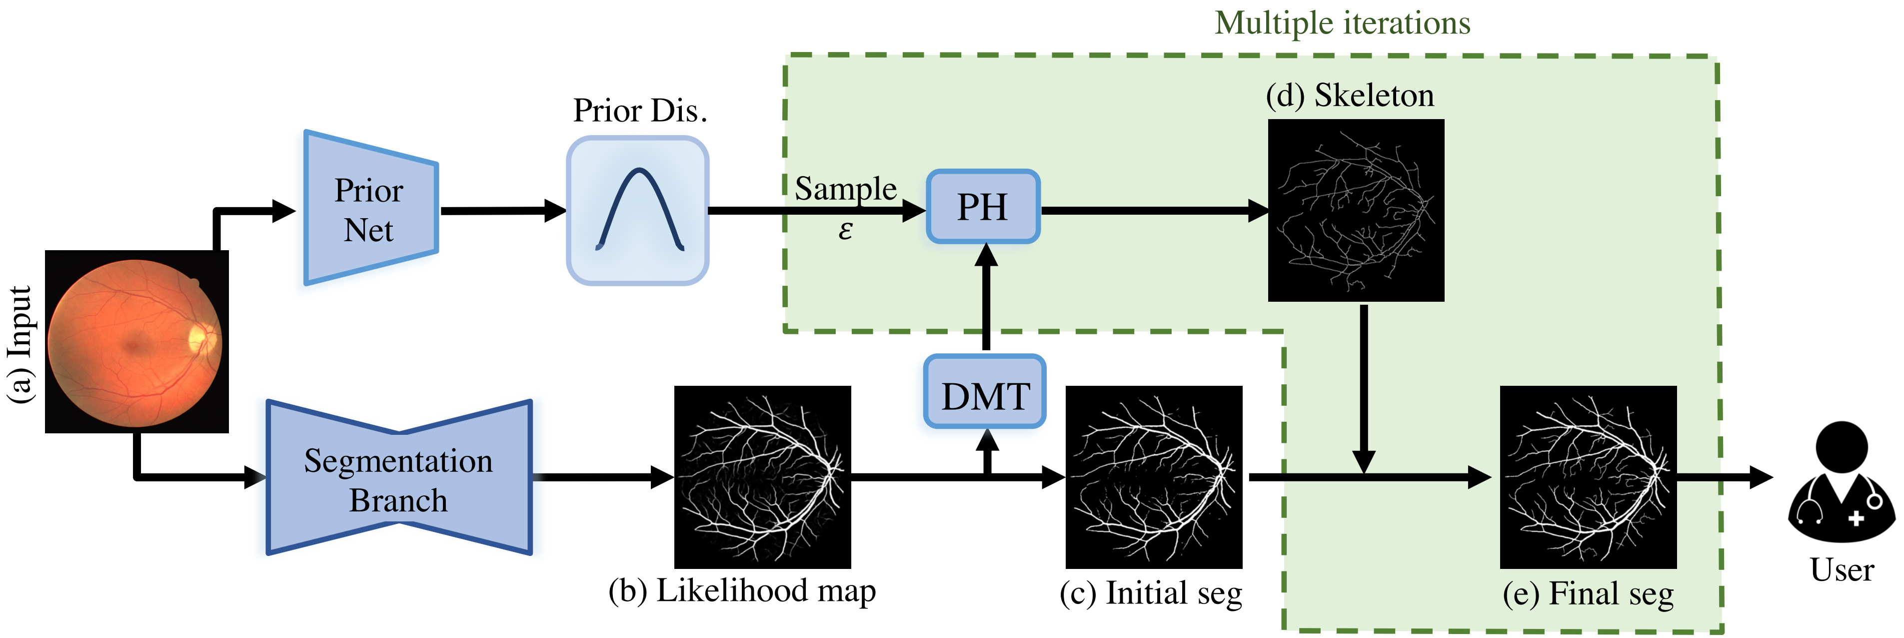

Uncertainty illustration

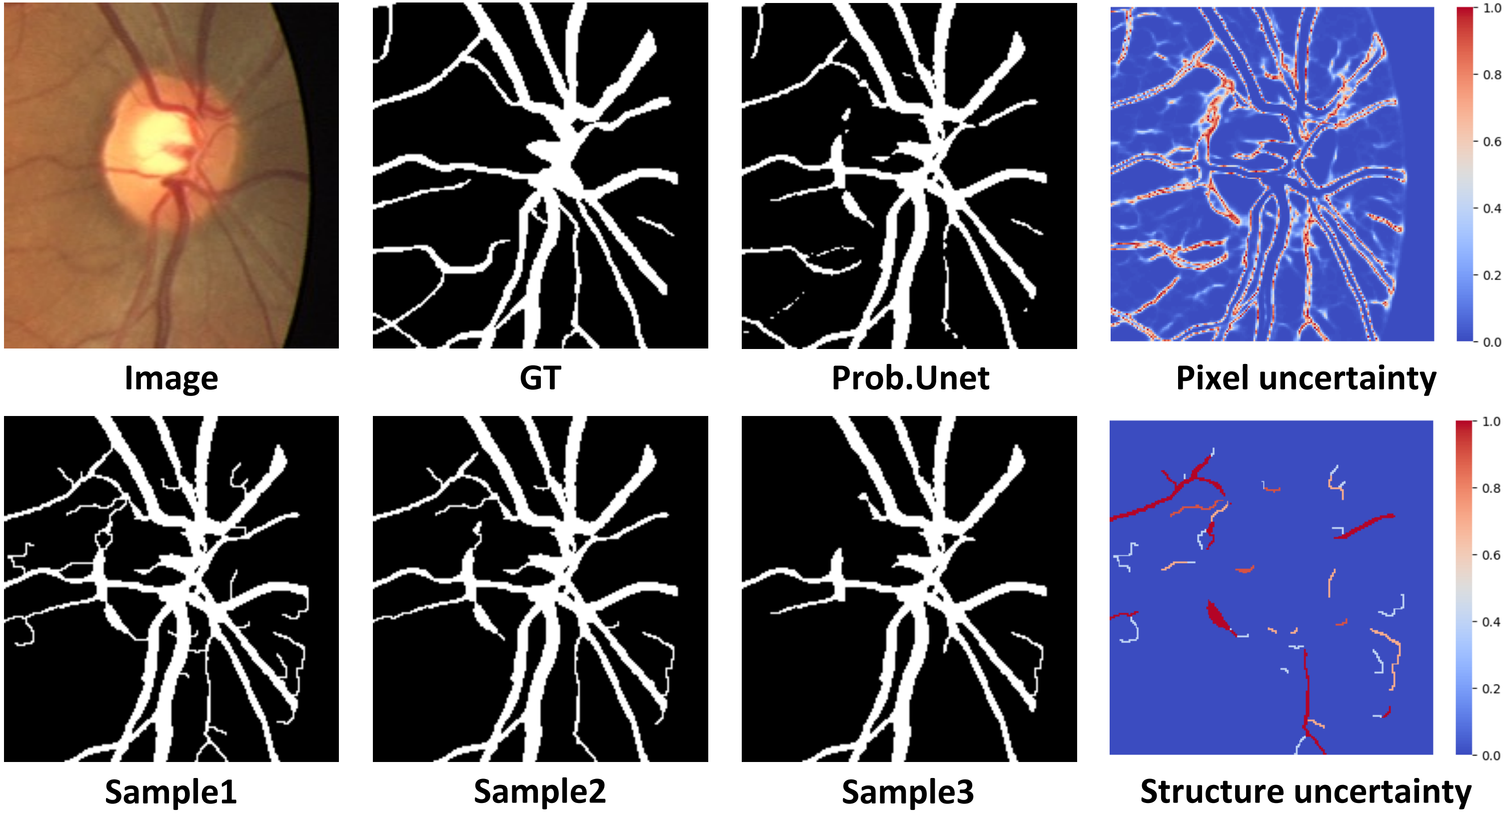

Kohl, Simon, et al. "A probabilistic u-net for segmentation of ambiguous images." ***NeurIPS*** (2018).

Semi-automatic efficient annotation/proofreading

Image

Segmentation

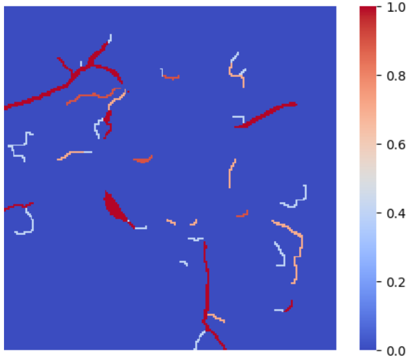

Structure uncertainty

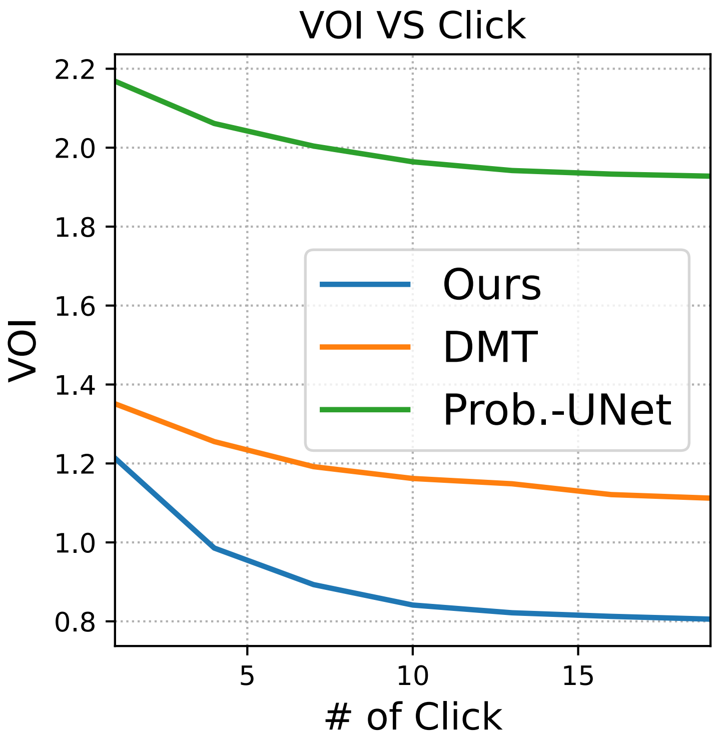

Application on mitochondria segmentation

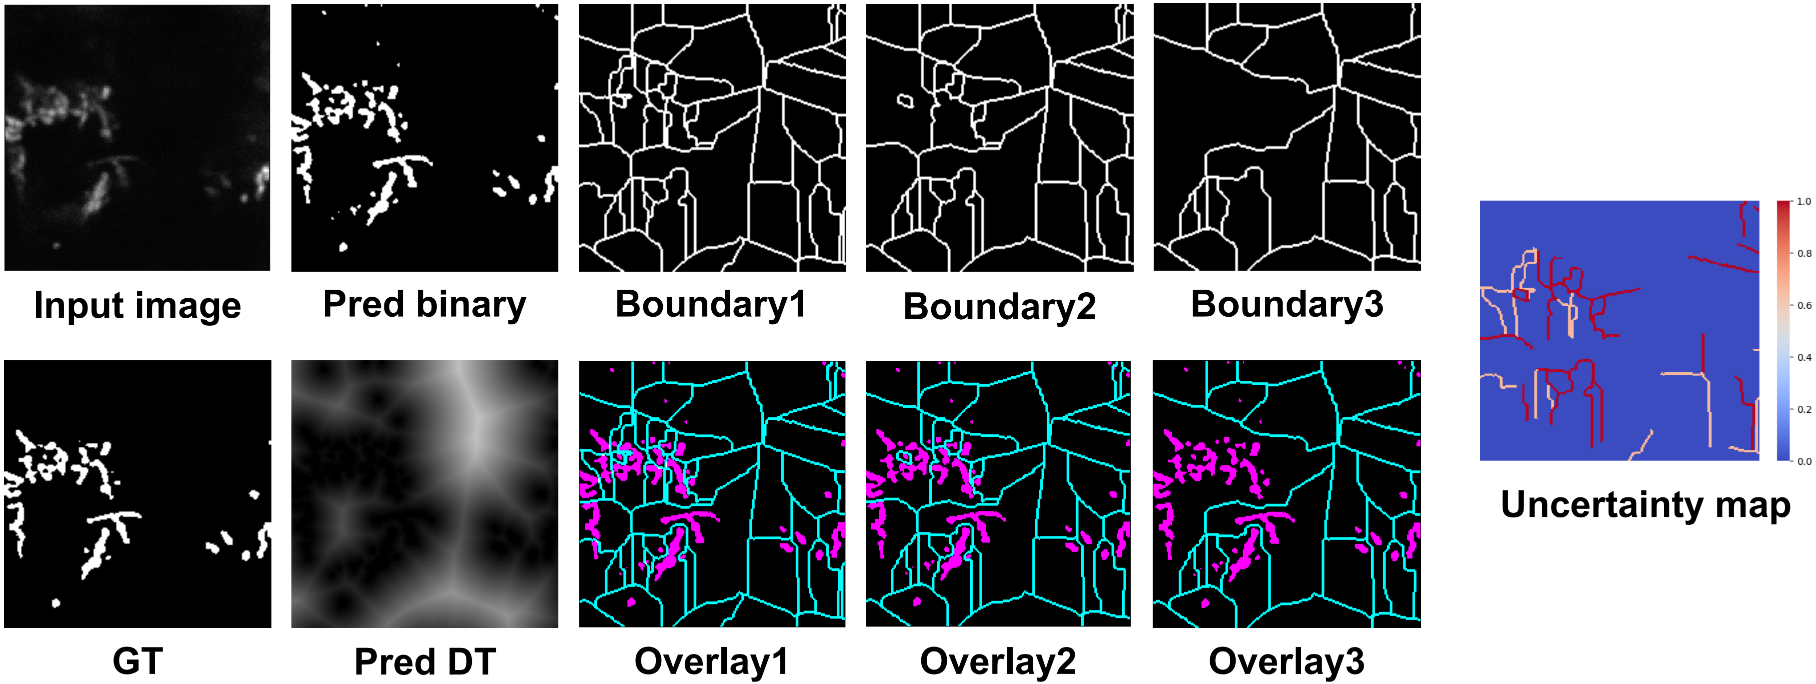

Work under preparation

Outline

- **Topological loss for image segmentation**
  - Persistent homology based loss (NeurIPS'19, MICCAI'21)
  - Homotopy warping loss (NeurIPS'22)
- **Beyond pixel-wise representation**
  - Discrete Morse theory loss (ICLR'21, Spotlight)
  - Topological/Structural representation of images (ICLR'23, Spotlight)
- **Extensions**
  - Trojan detection with topological prior (ICLR'22)
- **Future work**

Trojan detection (ICLR'22)

- Goal: Find a classifier to distinguish clean models and Trojaned models
- Challenges
  - ➢ Limited-data setting: only a few clean samples per class; Clean and Trojaned models perform the same on them
  - ➢ If Trojaned, trigger (location, shape, color) is unknown

(a) Trojaned examples

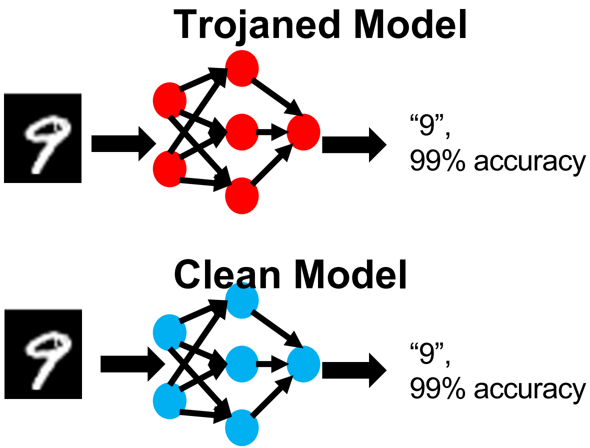

(b) Perform the same on clean images

Trigger reconstruction

- Reverse engineering approach
  - ➢ Huge search space; unknown target class
  - ➢ Triggers are scattered, even for Trojaned models
  - ➢ Solution: topological loss, diversity loss in reverse engineering

Clean sample

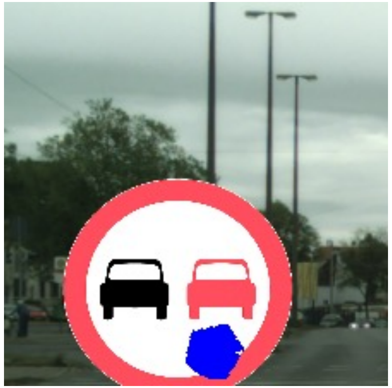

True trigger

Reconstructed

Reverse-engineering pipeline

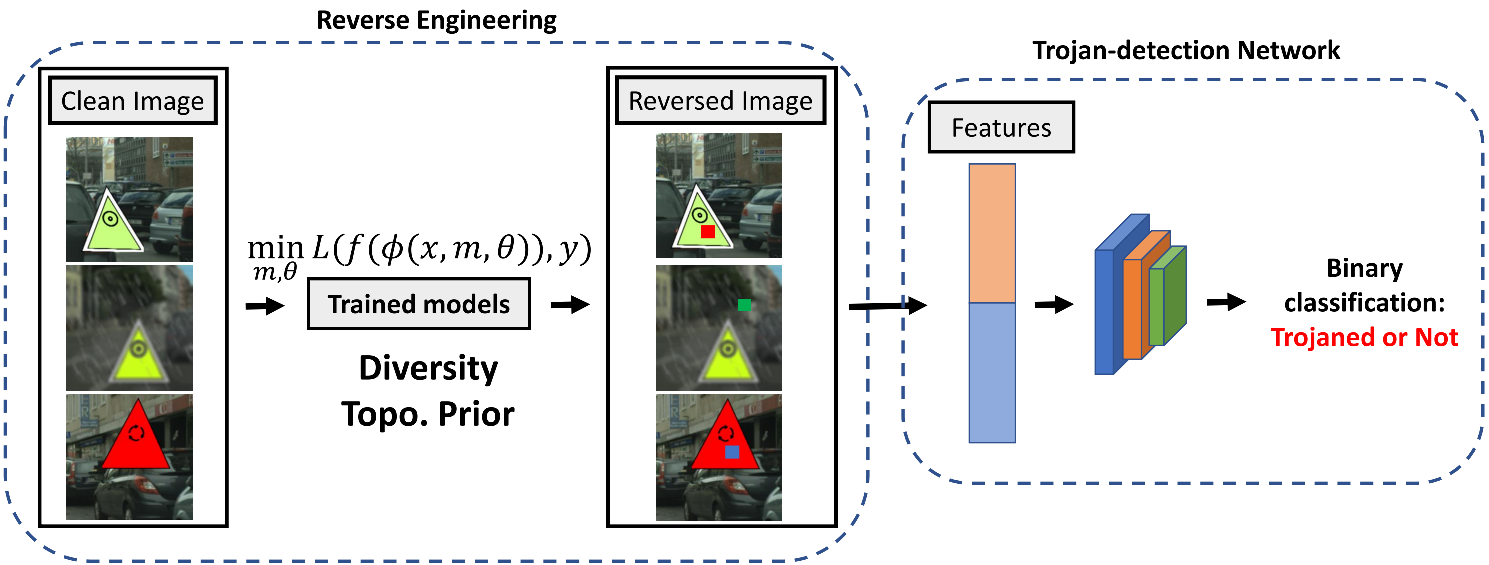

Diversity loss

- Trigger candidates different from each other

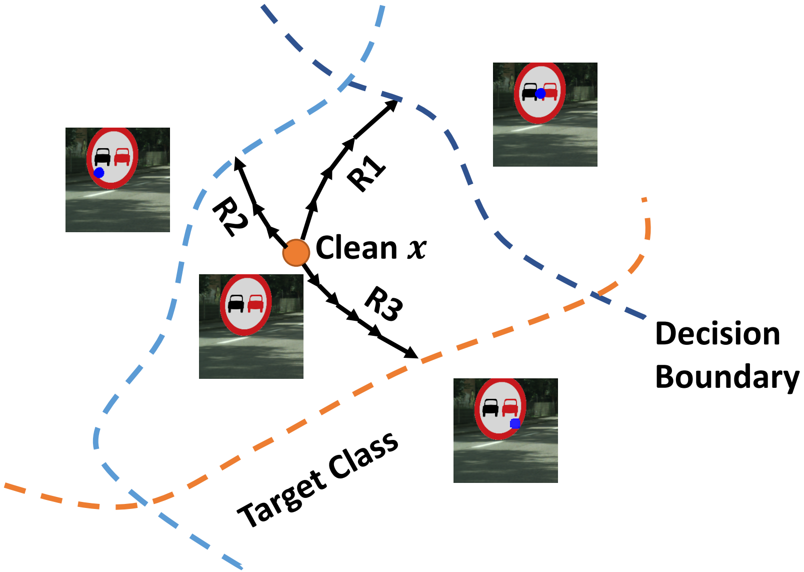

Topological loss

- Topological constraint: the trigger is a single component
  - ➢ Localized trigger
  - ➢ No strong assumption on shape/size
  - ➢ Can be written as a topological loss

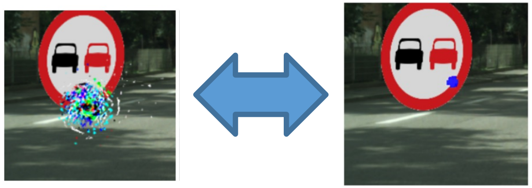

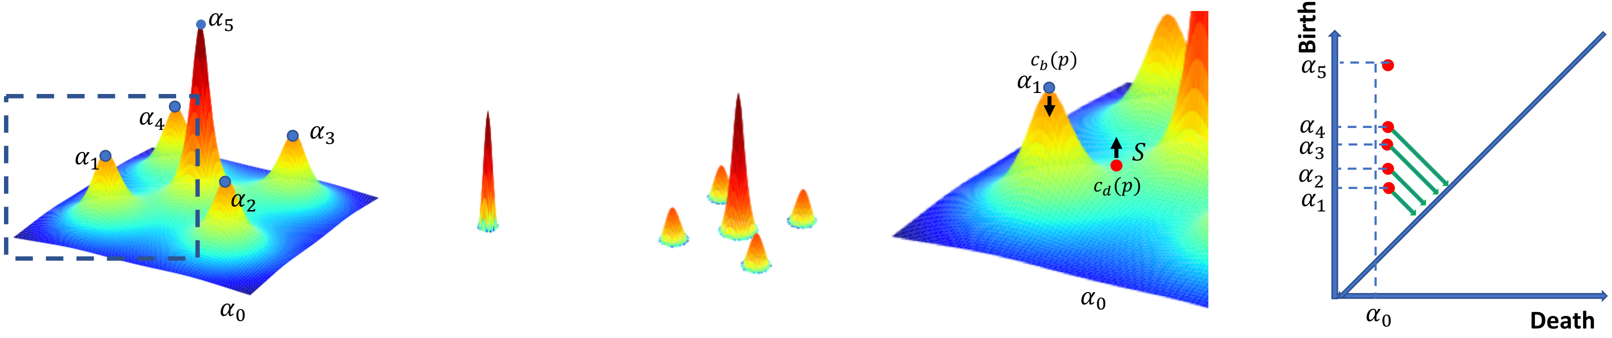

Final loss

- **Total loss**

$$L(\mathbf{m}, \boldsymbol{\theta}; \mathbf{x}, f, c^*) = L_{flip}(\ldots) + \lambda_1 L_{div}(\ldots) + \lambda_2 L_{topo}(\ldots) + R(\mathbf{m}).$$

- Flip loss

$$L_{flip}(\mathbf{m}, \boldsymbol{\theta}; \mathbf{x}, f, c^*) = f_{c^*}(\phi(\mathbf{x}, \mathbf{m}, \boldsymbol{\theta}))$$

- Diversity loss

$$L_{div}(\mathbf{m}, \boldsymbol{\theta}) = -\sum_{j=1}^{i-1} ||\mathbf{m} \odot \boldsymbol{\theta} - \mathbf{m}_j \odot \boldsymbol{\theta}_j||_2$$

- Topological loss

$$L_{topo}(\mathbf{m}) = \sum_{p \in \mathrm{Dgm}(m) \setminus \{p^*\}} [\mathrm{birth}(p) - \mathrm{death}(p)]^2$$

Qualitative results

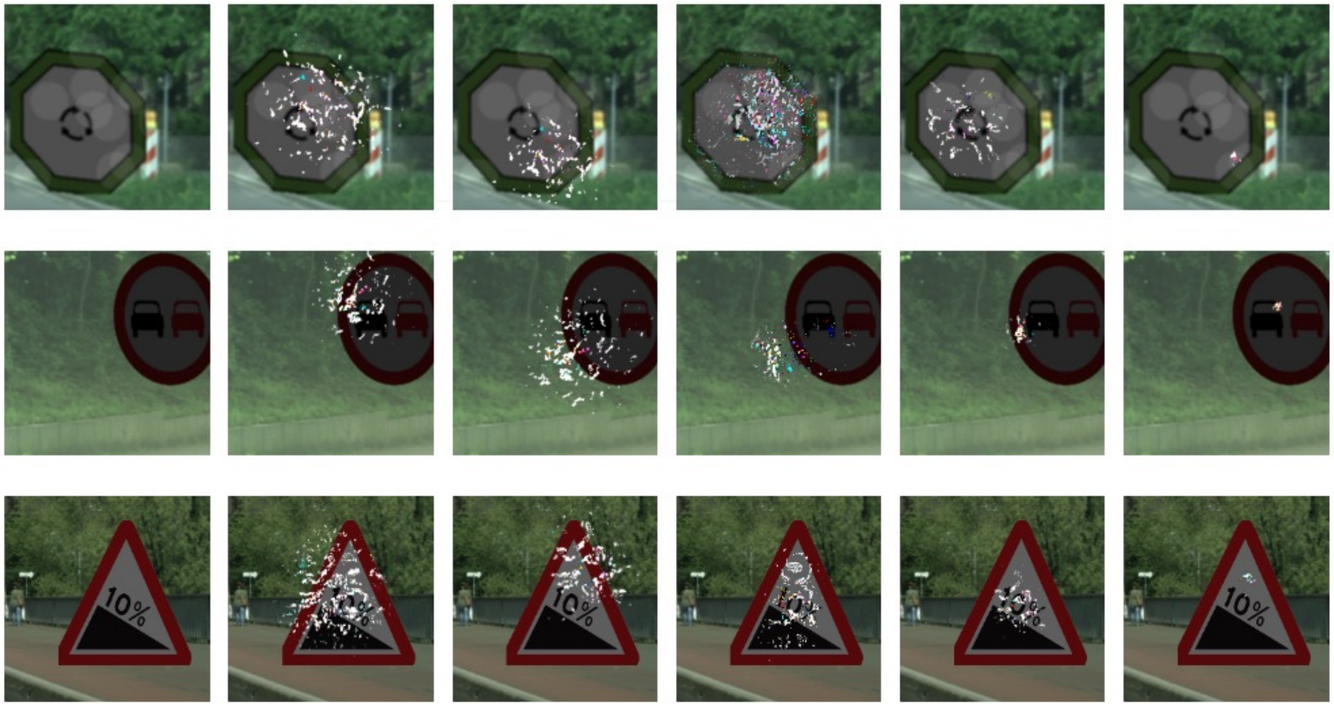

Clean Image | NC (SP'19) | ABS (CCS'19) | TABOR (ICDM'19) | w/o topo. | w/ topo.

Quantitative results

- Performances comparison on the TrojAI datasets

Method	Metric	TrojAI-Round1	TrojAI-Round2	TrojAI-Round3	TrojAI-Round4
NC	AUC	0.50 ± 0.03	0.63 ± 0.04	0.61 ± 0.06	0.58 ± 0.05
ABS	AUC	0.68 ± 0.05	0.61 ± 0.06	0.57 ± 0.04	0.53 ± 0.06
TABOR	AUC	0.71 ± 0.04	0.66 ± 0.07	0.50 ± 0.07	0.52 ± 0.04
ULP	AUC	0.55 ± 0.06	0.48 ± 0.02	0.53 ± 0.06	0.54 ± 0.02
DLTND	AUC	0.61 ± 0.07	0.58 ± 0.04	0.62 ± 0.07	0.56 ± 0.05
Ours	AUC	$\mathbf{0.90 \pm 0.02}$	$\mathbf{0.87 \pm 0.05}$	$\mathbf{0.89 \pm 0.04}$	$\mathbf{0.92 \pm 0.06}$
NC	ACC	0.53 ± 0.04	0.49 ± 0.02	0.59 ± 0.07	0.60 ± 0.04
ABS	ACC	0.70 ± 0.04	0.59 ± 0.05	0.56 ± 0.03	0.51 ± 0.05
TABOR	ACC	0.70 ± 0.03	0.68 ± 0.08	0.51 ± 0.05	0.55 ± 0.06
ULP	ACC	0.58 ± 0.07	0.51 ± 0.03	0.56 ± 0.04	0.57 ± 0.04
DLTND	ACC	0.59 ± 0.04	0.61 ± 0.05	0.65 ± 0.04	0.59 ± 0.06
Ours	ACC	$\mathbf{0.91 \pm 0.03}$	$\mathbf{0.89 \pm 0.04}$	$\mathbf{0.90 \pm 0.03}$	$\mathbf{0.91 \pm 0.04}$

Research summary

Topology-aware deep image segmentation

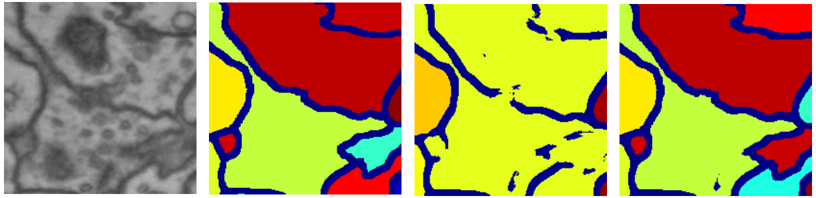

Hu et al. *NeurIPS'19*, *MICCAI'21*, *ISBI'21, ICLR'21*, *NeurIPS'22*, *ECCV'22*.

Trojan detection

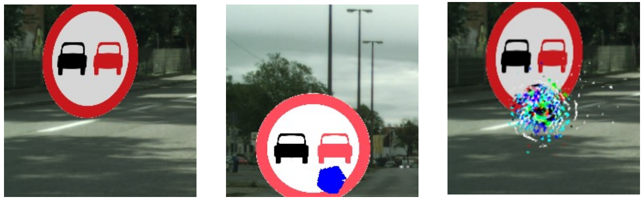

Hu et al. *ICLR'22*.

Uncertainty estimation

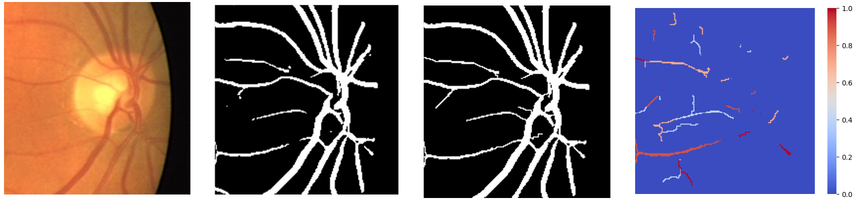

Hu et al. *ICLR'23*, Li et al. *ICLR'23*.

Biomedical applications

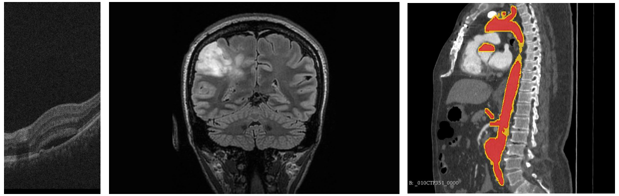

Yang et al. *under review*, Konwer et al. *under review*.

Data with rich structural information

- **Challenges**
  - Complex topology/structure
  - Noisy data
  - Limited labels

Difficult to analyze!

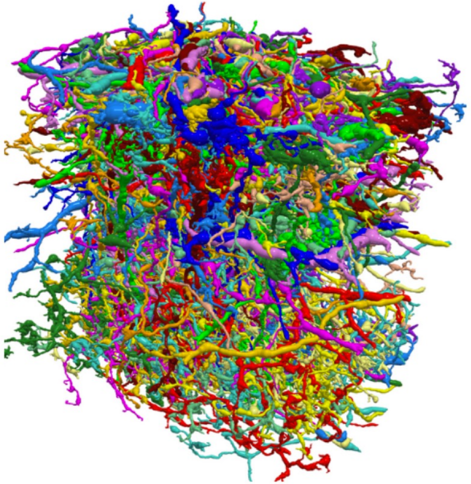
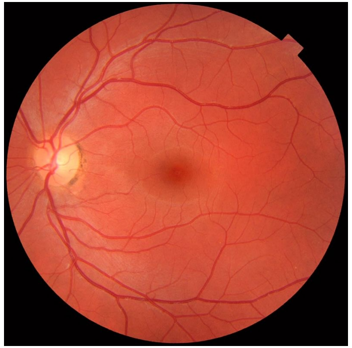


Career goal

- **Develop principled methods**
  - Model the intrinsic structures of data
  - Use topological/structural priors
- **Different types of scenarios/data**
  - Imperfect data: Medical imaging, general data science
  - Different types of data: Point cloud, Graph
  - ……

Direction 1: Imperfect data

- **Annotations are difficult to obtain**
  - ➢ Leverage information from unlabeled data
  - ➢ Crowd counting, cell detection …

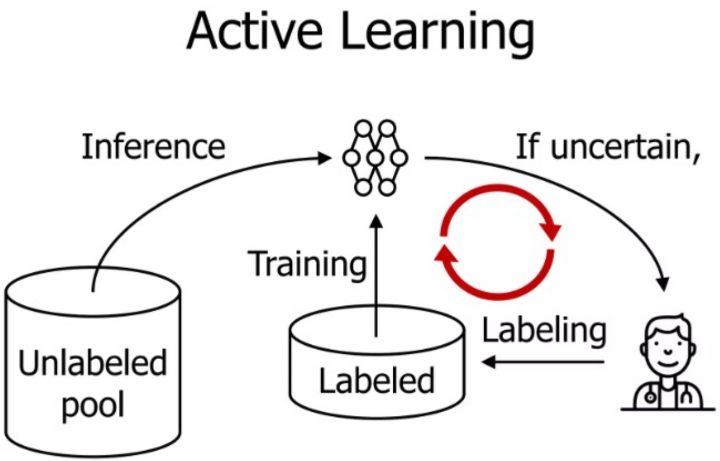

Yoo et al. *CVPR'19*.

Explore the structures within unlabeled data!

Li, Chen, **Hu, Xiaoling**, et al. "Confidence Estimation Using Unlabeled Data." ***ICLR*** (2023).

Direction 2: Deep learning based quantification analysis

- **Combine topology-aware segmentation and uncertainty estimation**
  - ➢ Interactive annotation/proofreading
  - ➢ Downstream topology-aware analysis
  - ➢ Neurons, vasculatures, digital pathology images, etc.

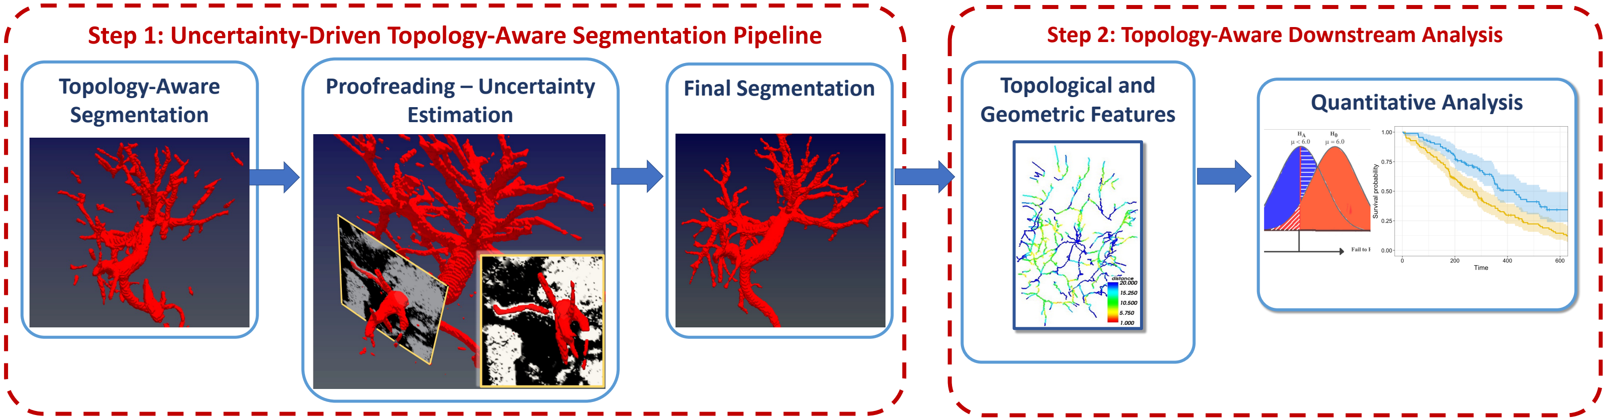

Direction 3: Other types of data

- **Topology on different types of data**

Extending topology non-trivially to other types of data!

Point Cloud

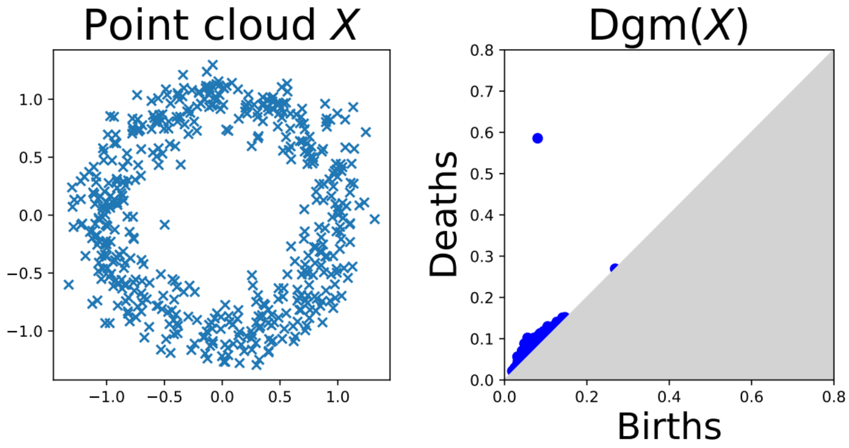

de Surrel, Thibault, et al. "RipsNet: a general architecture for fast and robust estimation of the persistent homology of point clouds." ***TAGL Workshops*** (2022).

Graph

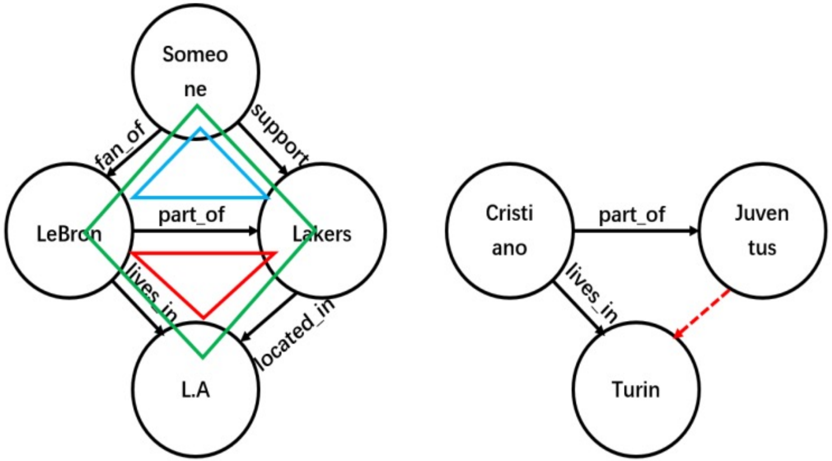

Yan, Zuoyu, et al. "Cycle Representation Learning for Inductive Relation Prediction." ***ICML*** (2022).

Direction 4: Topo attention in NLP/LLM

- **Topology of attention connectivity**
  - ➢ Prune number of heads
  - ➢ Robust against adversarial attacks

Using TDA to analyze the structures of neural networks!

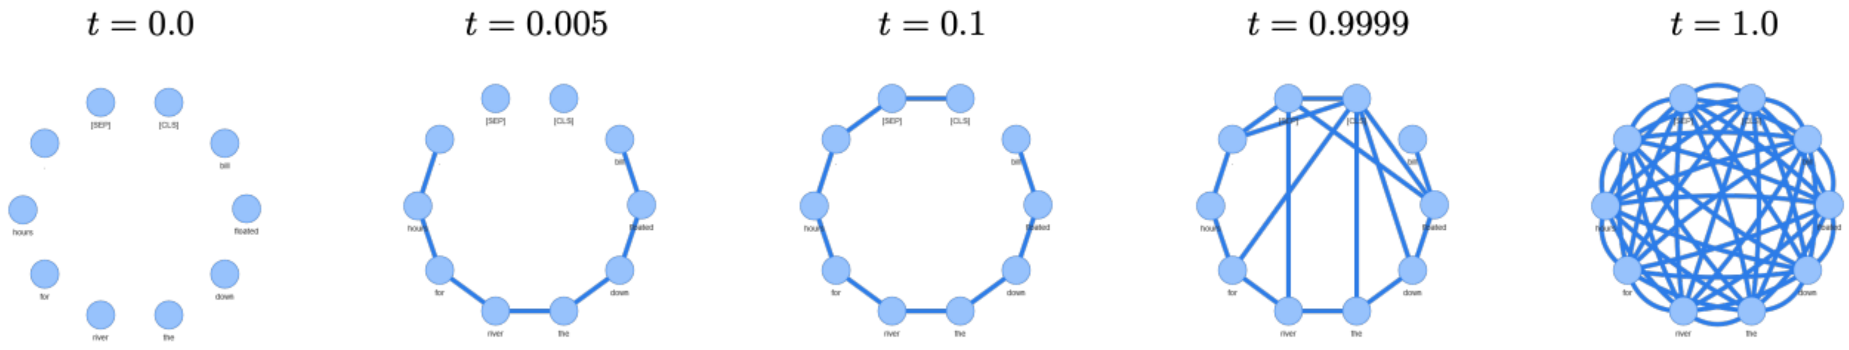

Perez, Ilan, et al. "The Topological BERT: Transforming Attention into Topology for Natural Language Processing." ***arXiv*** (2022).

Reference

- **Hu, Xiaoling**, et al. "Topology-preserving deep image segmentation." ***NeurIPS*** (2019).
- **Hu, Xiaoling**, et al. "Topology-aware segmentation using discrete Morse theory." ***ICLR*** (2021, **Spotlight**).
- Yang, Jiaqi*, **Hu, Xiaoling***, et al. "3D topology-preserving segmentation with compound multi-slice representation." *ISBI* (2021).
- Yang, Jiaqi*, **Hu, Xiaoling***, et al. "A topological-attention convlstm network and its application to em images." ***MICCAI*** (2021).
- **Hu, Xiaoling**, et al. "Trigger hunting with a topological prior for trojan detection." ***ICLR*** (2022).
- Gupta, Saumya*, **Hu, Xiaoling***, et al. "Learning topological interactions for multi-class medical image segmentation." ***ECCV*** (2022, **Oral**).
- **Hu, Xiaoling**. "Topology-aware image segmentation with homotopy warping." ***NeurIPS*** (2022).
- **Hu, Xiaoling**, et al. "Deep statistic shape model for myocardium segmentation." *arXiv preprint arXiv:2207.10607* (2022).
- **Hu, Xiaoling**, et al. "Learning probabilistic topological representation using discrete Morse theory." ***ICLR*** (2023, **Spotlight**).
- Li, Chen, **Hu, Xiaoling**, et al. ”Confidence estimation using unlabeled data." ***ICLR*** (2023).
- Yang, Jiaqi, **Hu, Xiaoling**, et al. "Weakly supervised learning for lesion segmentation on oct images with anomaly attention mechanism." *under review*.
- Konwer, Aishik, **Hu, Xiaoling**, et al. "Enhancing modality-agnostic representations via meta-learning for brain tumor segmentation." *under review*.

Acknowledgment

Credits are due to the committee members, collaborators, and also my family!

Thank you for your attention!

Q&A