

Topology-Driven Learning for Images: Applications and Acceleration

Fan Wang
May 16, 2024

Committee members:

Prof. Chao Chen

Prof. Dimitris Samaras

Prof. Rezaul A. Chowdhury

Prof. Haibin Ling

External committee member:

Prof. Hubert Wagner

Contents

1. Introduction to topological data analysis and persistent homology.
2. Applications of persistent homology:
 - Topology-aware GAN
 - Topological biomarker for predicting breast cancer treatment response
3. Persistent homology computations using GPUs:
 - GPU computation of the Euler Characteristic Curve
 - GPU computation of persistent homology

Updates from Prelim

1. GPU Computation of the Euler Characteristic Curve for Imaging Data
 - Journal of Computational Geometry Vol 14
2. TopoTxR: A topology-guided deep convolutional network for breast parenchyma learning on DCE-MRIs
 - Medical Image Analysis – in revision
3. GPU-Accelerated Computation of Persistent Homology for Image Data
 - TPAMI – in submission
 - A significant contribution developed over three years
 - More than 14,000 lines of C++ codes

Publications

1. **Fan Wang**, et al. “TopoGAN: A Topology-Aware Generative Adversarial Network”, ECCV 2020, Oral
2. **Fan Wang**, et al. “TopoTxR: A Topological Biomarker for Predicting Treatment Response in Breast Cancer”, IPMI 2021
3. **Fan Wang**, et al. “GPU Computation of the Euler Characteristic Curve for Imaging Data”, SoCG 2022
4. **Fan Wang**, et al. “GPU Computation of the Euler Characteristic Curve for Imaging Data”, JoCG 2023
5. **Fan Wang**, et al. “Hierarchical image link selection scheme for duplicate structure disambiguation”, BMVC 2018
6. **Fan Wang**, et al. “Hardware Acceleration of Persistent Homology Computation”, MICCAI Workshop 2019
7. **Fan Wang**, et al. “TopoTxR: A topology-guided deep convolutional network for breast parenchyma learning on DCE-MRIs”, MedIA in revision
8. **Fan Wang**, et al. “GPU-Accelerated Computation of Persistent Homology for Image Data”, TPAMI in submission

Contents



1. Introduction to topological data analysis and persistent homology.

2. Applications of persistent homology:

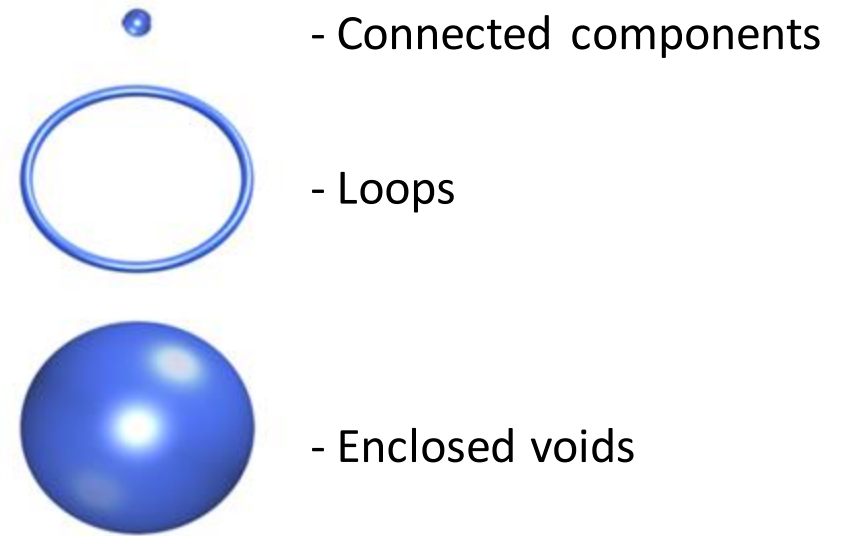


- Topology-aware GAN
- Topological biomarker for predicting breast cancer treatment response

3. Persistent homology computations using GPUs:

- GPU computation of the Euler Characteristic Curve
- GPU computation of persistent homology

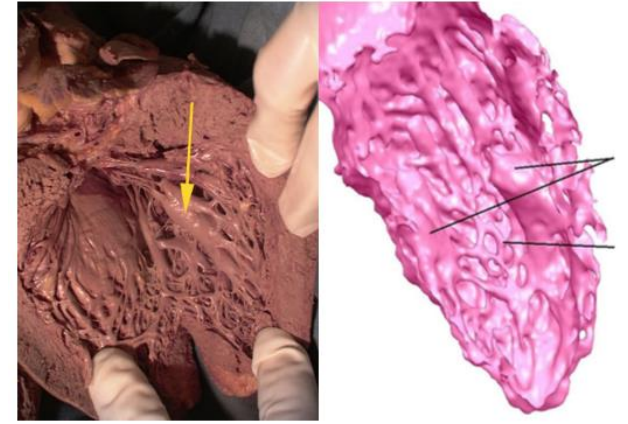
What is Topological Data Analysis (TDA)?

Topological data analysis (TDA) is an approach to analyze data using techniques from topology. These techniques extract topological features from data.




TDA in Biomedical Imaging


- Complex biomedical systems with rich topology and geometry



Cardiac Trabeculae




Neurons



Brain Arteries




Retina



Breast

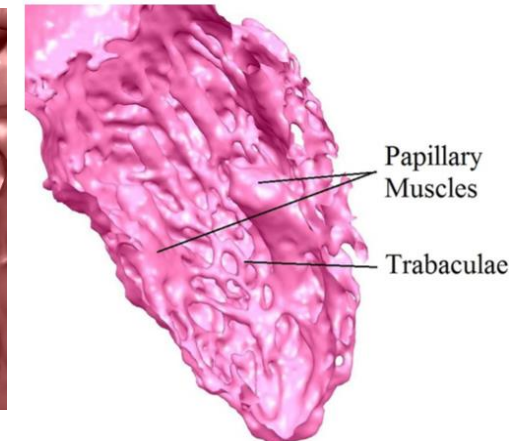



Challenge: Structure Extraction from Images

- Thresholding does not work
- Smart/adaptive thresholding



Issue:
Requires a clean input!

Microscopy Image: Neuron Reconstruction
[Wang et al. '18]

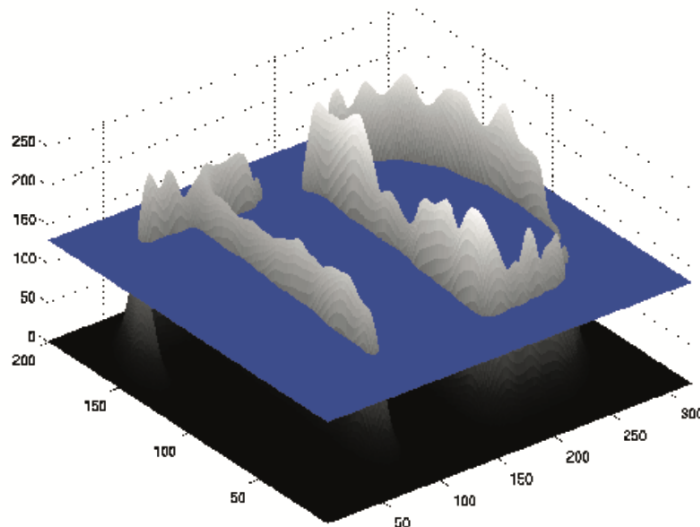



Cardiac CT image (3D)

Left Ventricle Reconstruction [IPMI'13,'17, IJCAI'19]

Persistent homology: a robust way to detect topology

- Input: a (density) function, f
- Output: topological structures & their **persistence**
- Def: given threshold t , the **superlevel set** $f^{-1}[t, +\infty) := \{x | f(x) \geq t\}$






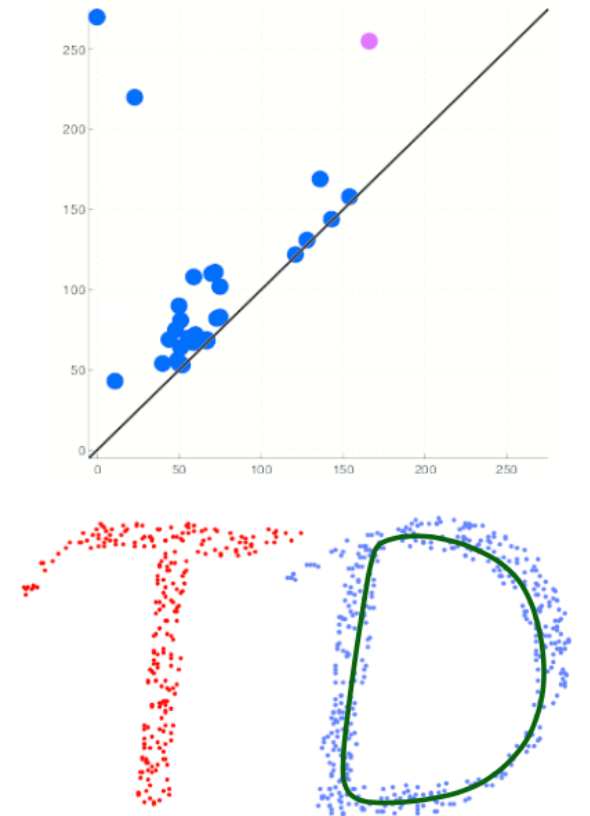
Persistent homology (cont'd)

- the true structures are hidden in superlevel sets
- consider the whole stack of superlevel sets
- identify structures that often appear (**high persistence**)
- Output: persistence diagram – dots representing all structures


Input



Diagram




Persistent homology (cont'd)



Challenges

1. How to incorporate topology into deep learning to improve performance:
 - a) Generative model: TopoGAN
 - b) Classification model: TopoTxR



2. Computation of persistent homology is heavy. Hardware accelerators like GPU are inevitable.



Contents

1. Introduction to topological data analysis and persistent homology.

2. Applications of persistent homology:

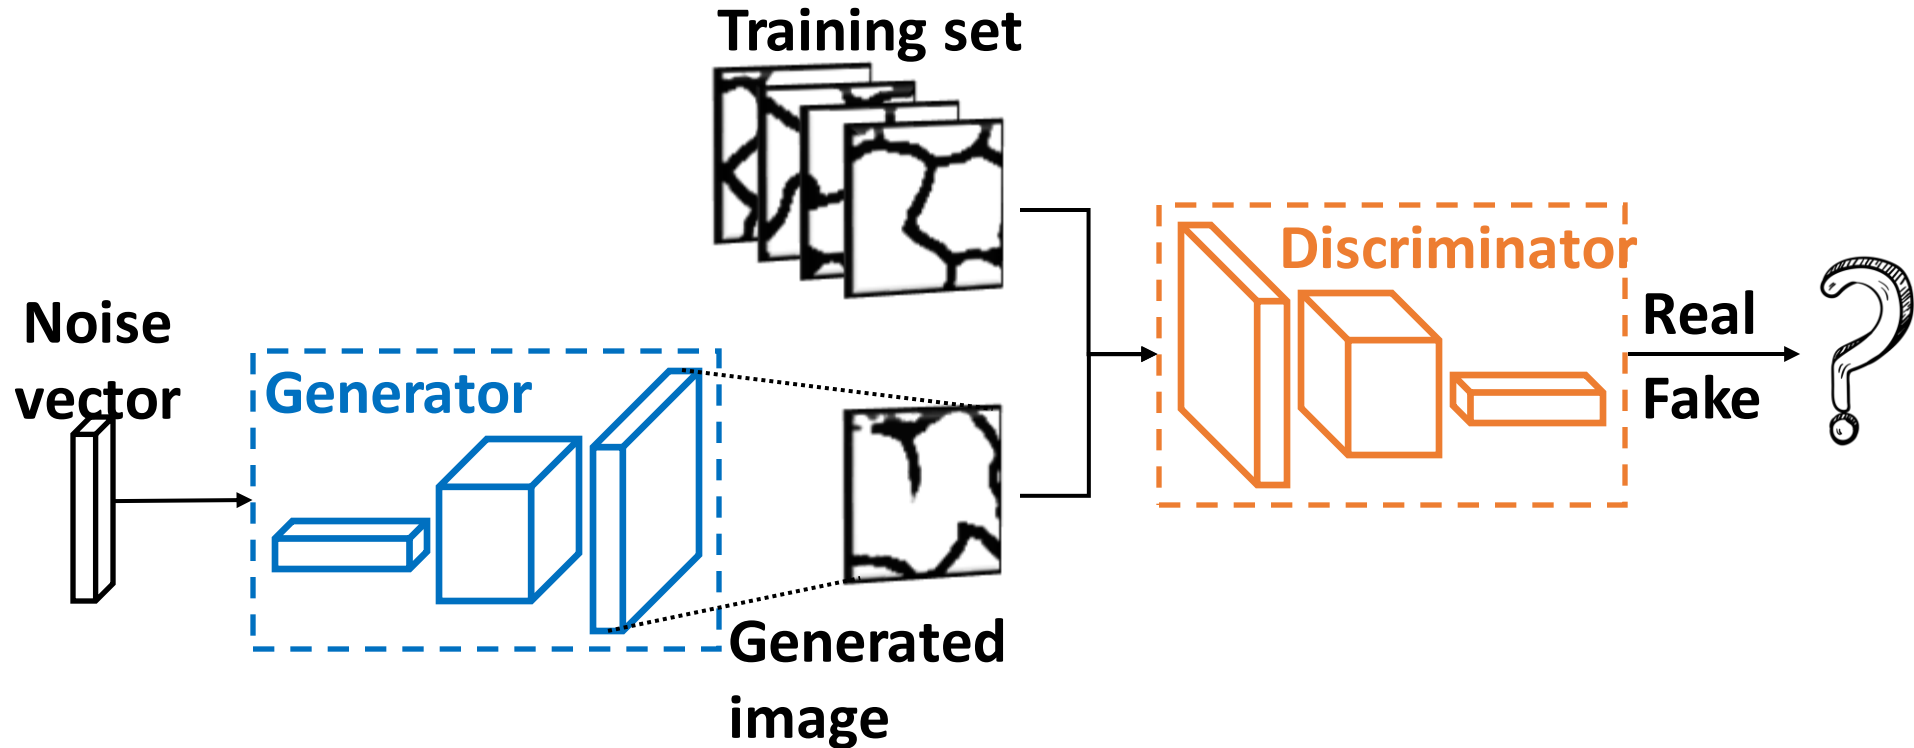


- **TopoGAN: A topology-aware generative adversarial network**
- Topological biomarker for predicting breast cancer treatment response

3. Persistent homology computations using GPUs:

- GPU computation of the Euler Characteristic Curve
- GPU computation of persistent homology

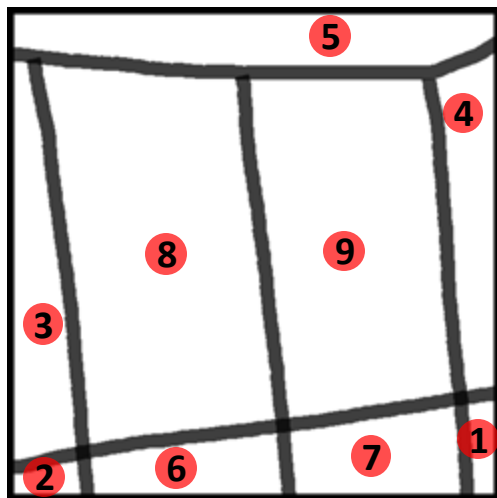

Generative Adversarial Network (GAN)



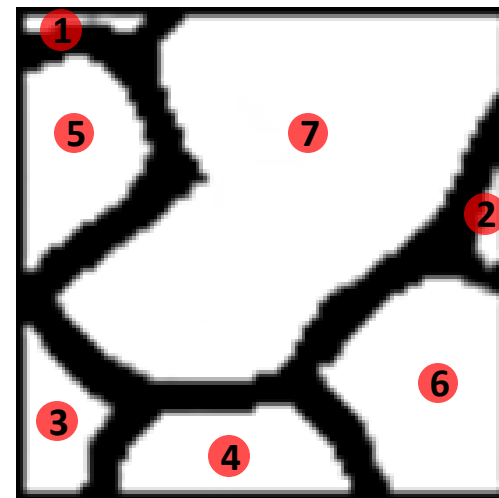

Applications

- Generate human faces
- Super resolution
- Generate realistic photographs
- Photo inpainting
- Image-to-image translation
- Many more ...



Motivation



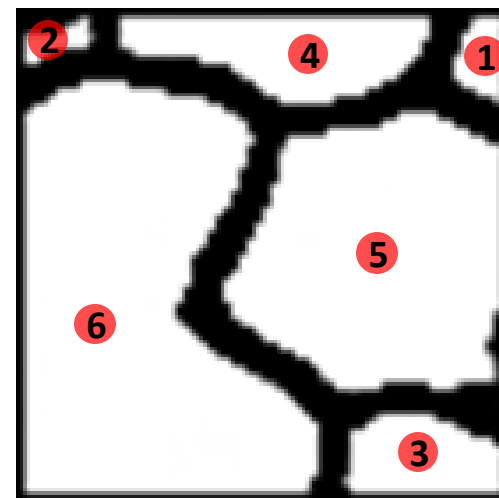

9 loops



7 loops



5 loops



6 loops 15


Google map examples

Segmentation mask

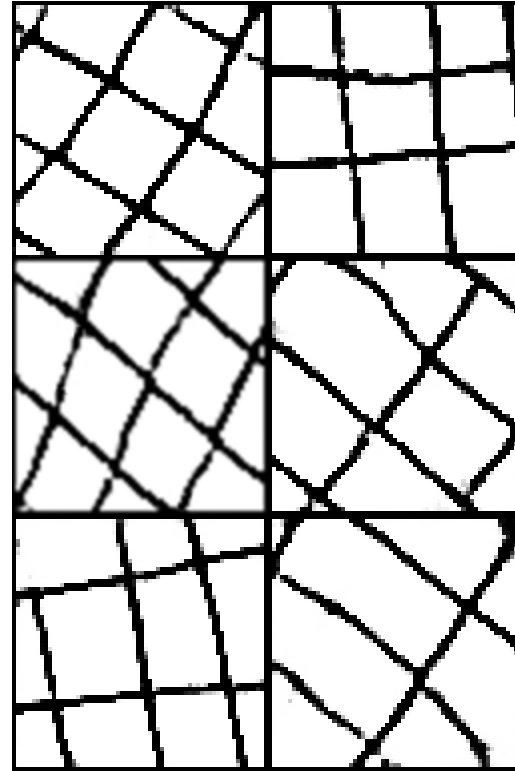
CREMI examples

Segmentation mask


Motivation




Conventional GANs



TopoGAN

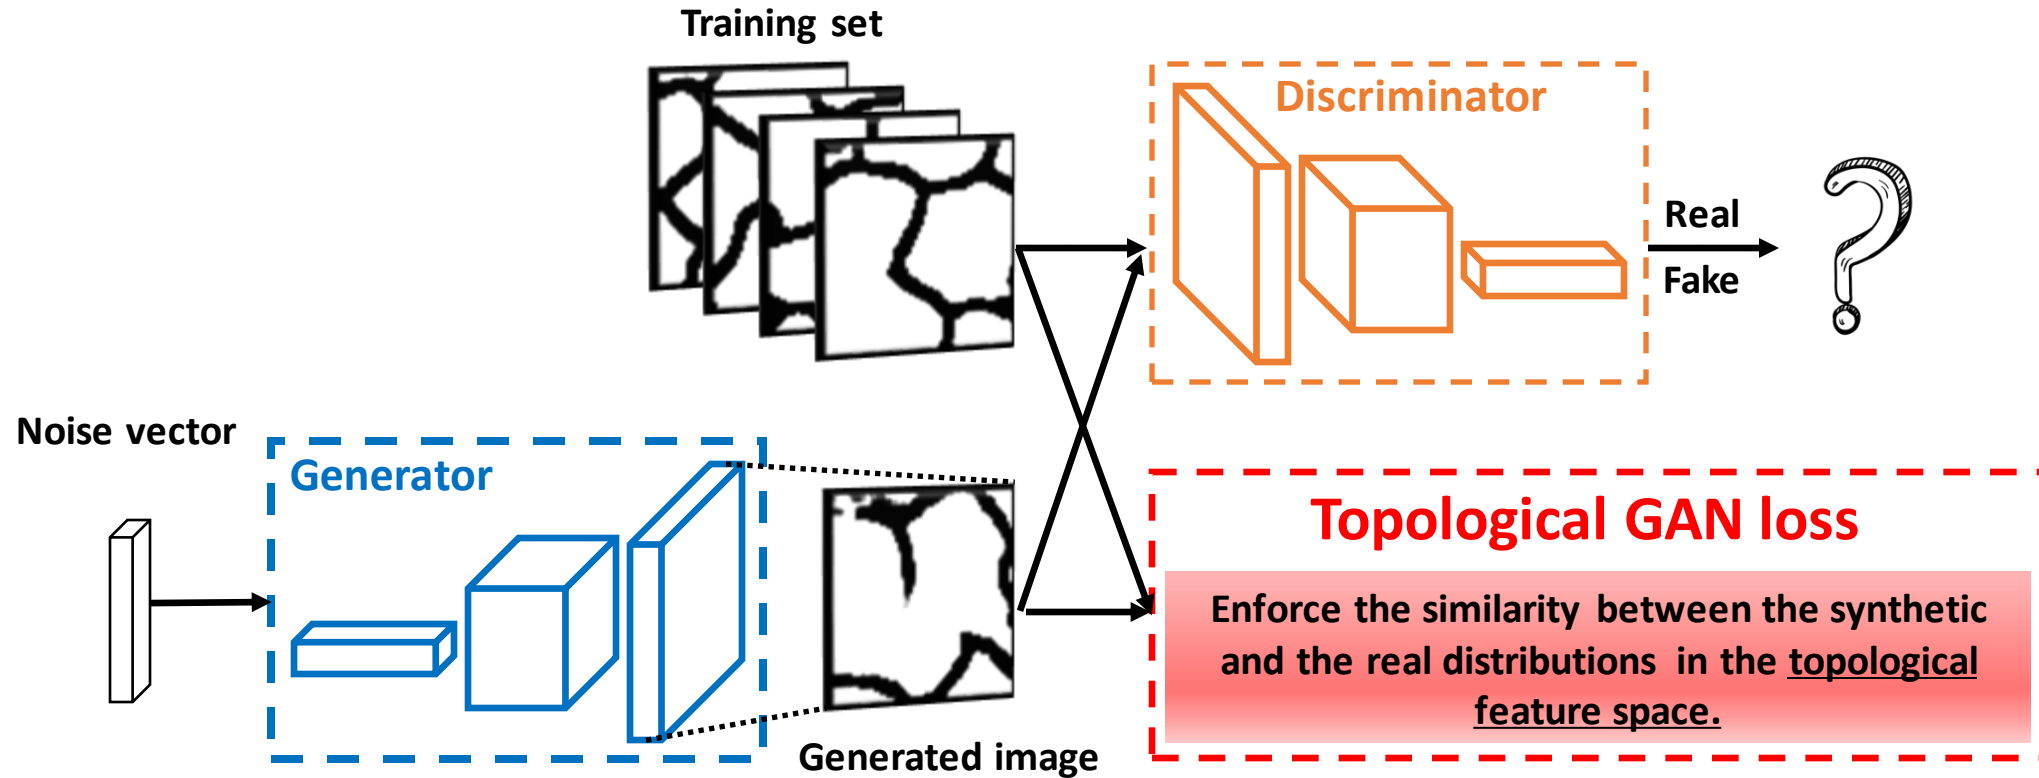


Conventional GANs



TopoGAN

TopoGAN Framework

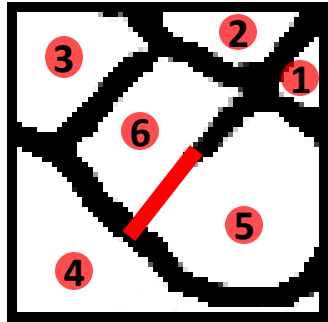


Discriminator loss: $\operatorname{argmax}_D \left[\mathbb{E}_{x \sim P_{data}} \log D(x) + \mathbb{E}_{z \sim P_z} \log (1 - D(G(z))) \right]$

Generator loss: $\operatorname{argmin}_G \left[\mathbb{E}_{z \sim P_z} \log (1 - D(G(z))) + \lambda \underbrace{L_{topo}(P_{data}, G)}_{\text{Topological GAN loss}} \right]$


Topological GAN loss

Persistent Homology




5 loops


Input mask






Distance transform (DT)



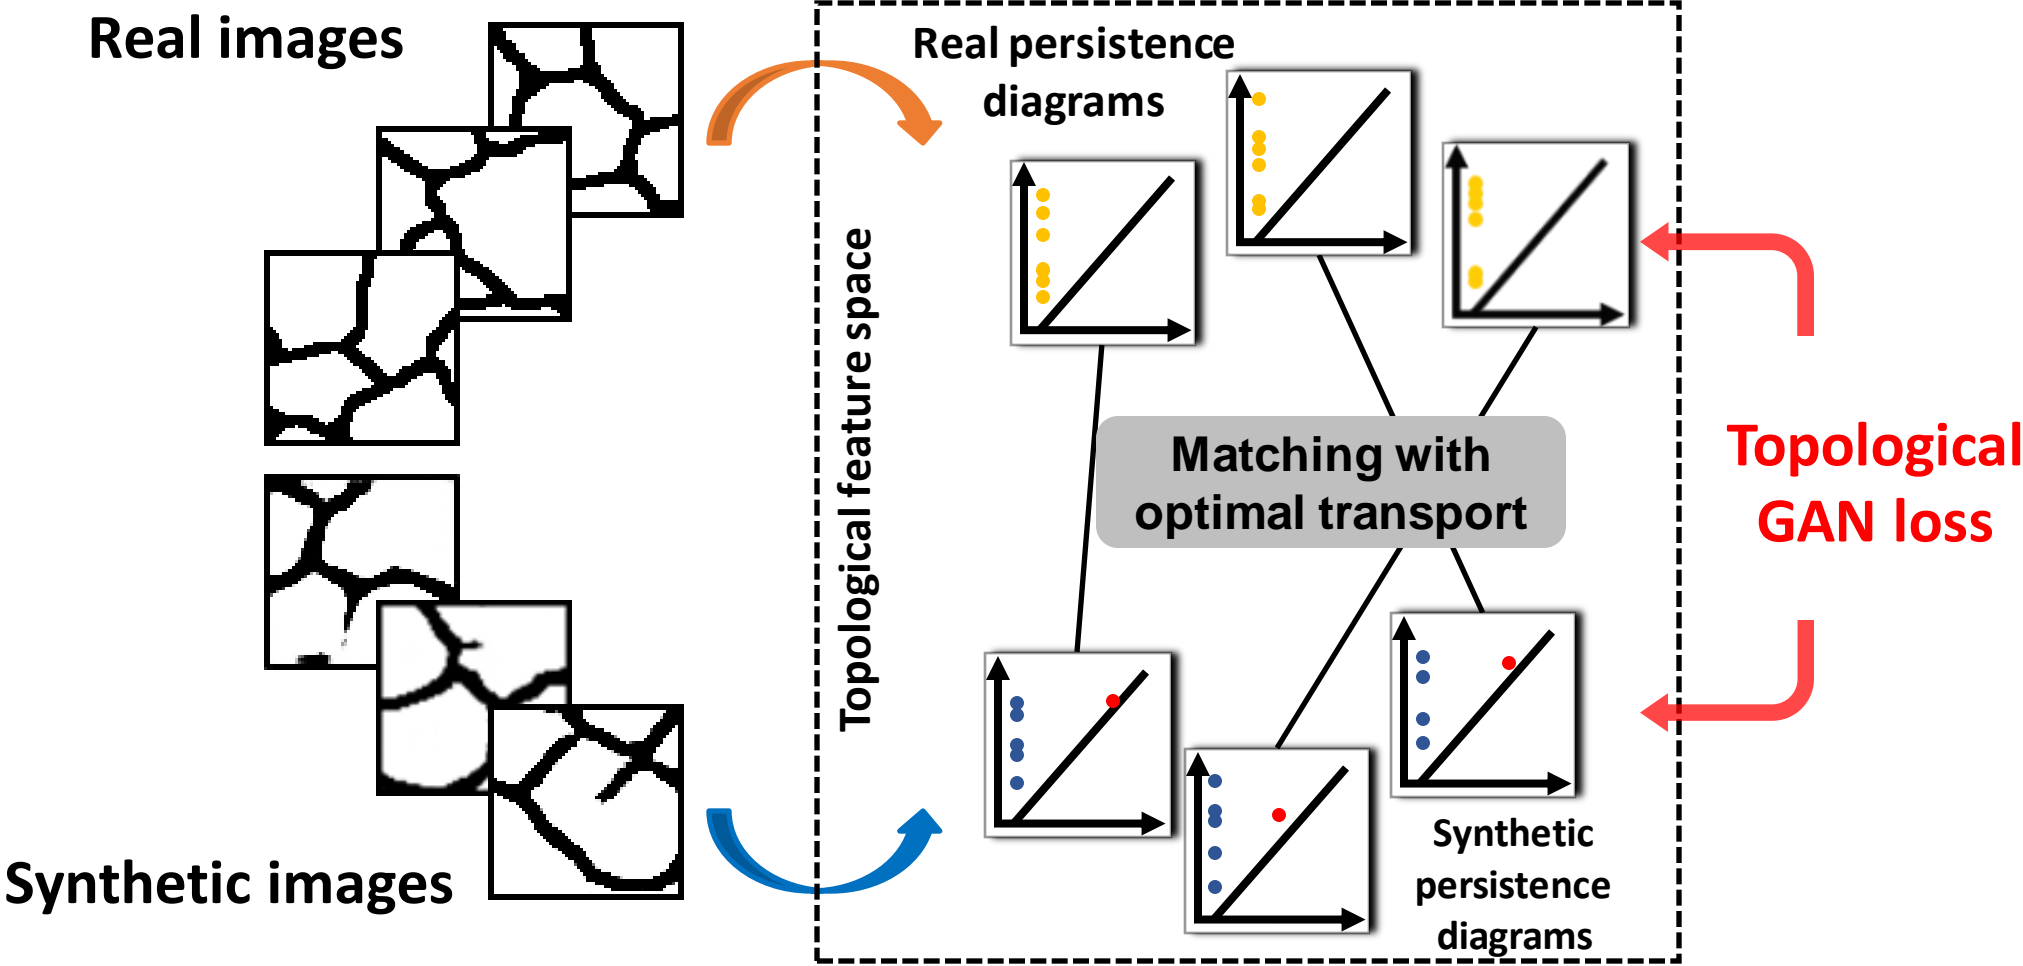
Landscape view
of DT



Persistent Diagram

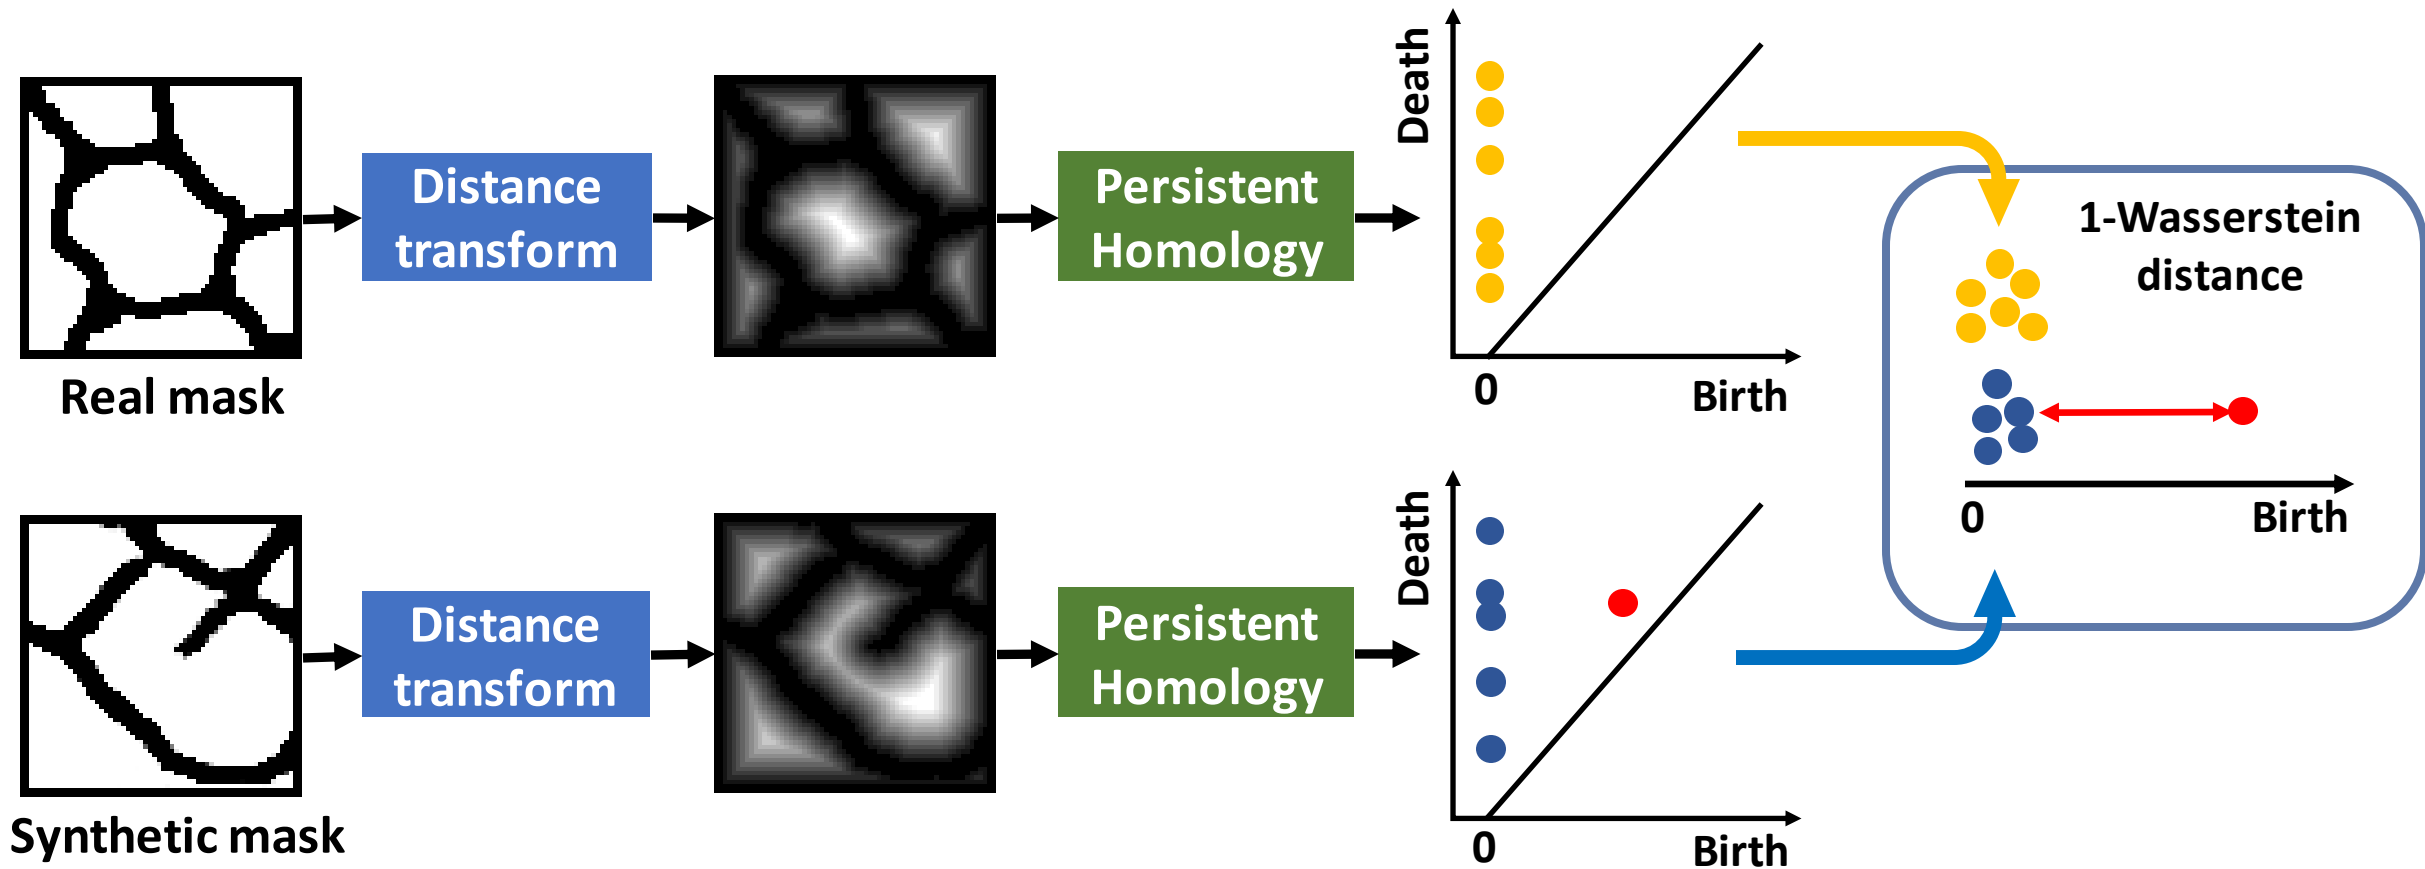


Distance between Distributions of Diagrams





[1] Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. NeurIPS 2019.

Distance between Persistence Diagrams




Fixing Incomplete Loops




Topology-Aware GAN Evaluation Metrics


1. Betti score



2. Maximum mean discrepancy (MMD)




Qualitative Results



TopoGAN is evaluated on five datasets: **CREMI**, **ISBI12**, **Google Maps**, **CMP Façade Database**, and **Retina** dataset. We show results of only **CREMI** and **Google Maps** here due to time constraint.

Qualitative Results




The texture images in the last row are generated with a pretrained pix2pix network which takes masks as inputs and produces corresponding texture images.

Quantitative Results

	CREMI	ISBI12	Retina	Maps	Facade
FID					
WGAN-GP	21.64±0.138	83.90±0.718	179.69±19.008	72.00±0.469	122.13±0.822
WGAN-SN	34.15±0.153	78.61±0.411	269.12±2.276	175.52±0.217	126.10±1.901
TopoGAN	20.96±0.195	31.90±0.248	169.21±21.976	60.48±0.467	119.11±0.874
unbiased MMD					
WGAN-GP	0.142±0.014	0.558±0.010	1.735±0.050	0.482±0.007	0.137±0.004
WGAN-SN	0.326±0.016	0.602±0.006	-	0.724±0.005	0.166±0.005
TopoGAN	0.134±0.019	0.405±0.003	1.602±0.114	0.471±0.010	0.080±0.002
Betti score					
WGAN-GP	0.236±0.003	0.908±0.104	0.541±0.188	0.223±0.010	0.176±0.006
WGAN-SN	0.125±0.002	1.775±0.039	-	0.255±0.020	0.142±0.017
TopoGAN	0.015±0.001	0.802±0.058	0.457±0.144	0.177±0.004	0.124±0.002


lower score = better quality

Contents

1. Introduction to topological data analysis and persistent homology.
2. Applications of persistent homology:
 - TopoGAN: A topology-aware generative adversarial network
 -  **TopoTxR: A Topological Biomarker for Predicting Treatment Response in Breast Cancer**
3. Persistent homology computations using GPUs:
 - GPU computation of the Euler Characteristic Curve
 - GPU computation of persistent homology


Motivation

Radiomics approach




Hand-crafted features
[can't model complex tissue structures]

CNN



Learned features
[Completely data-driven]

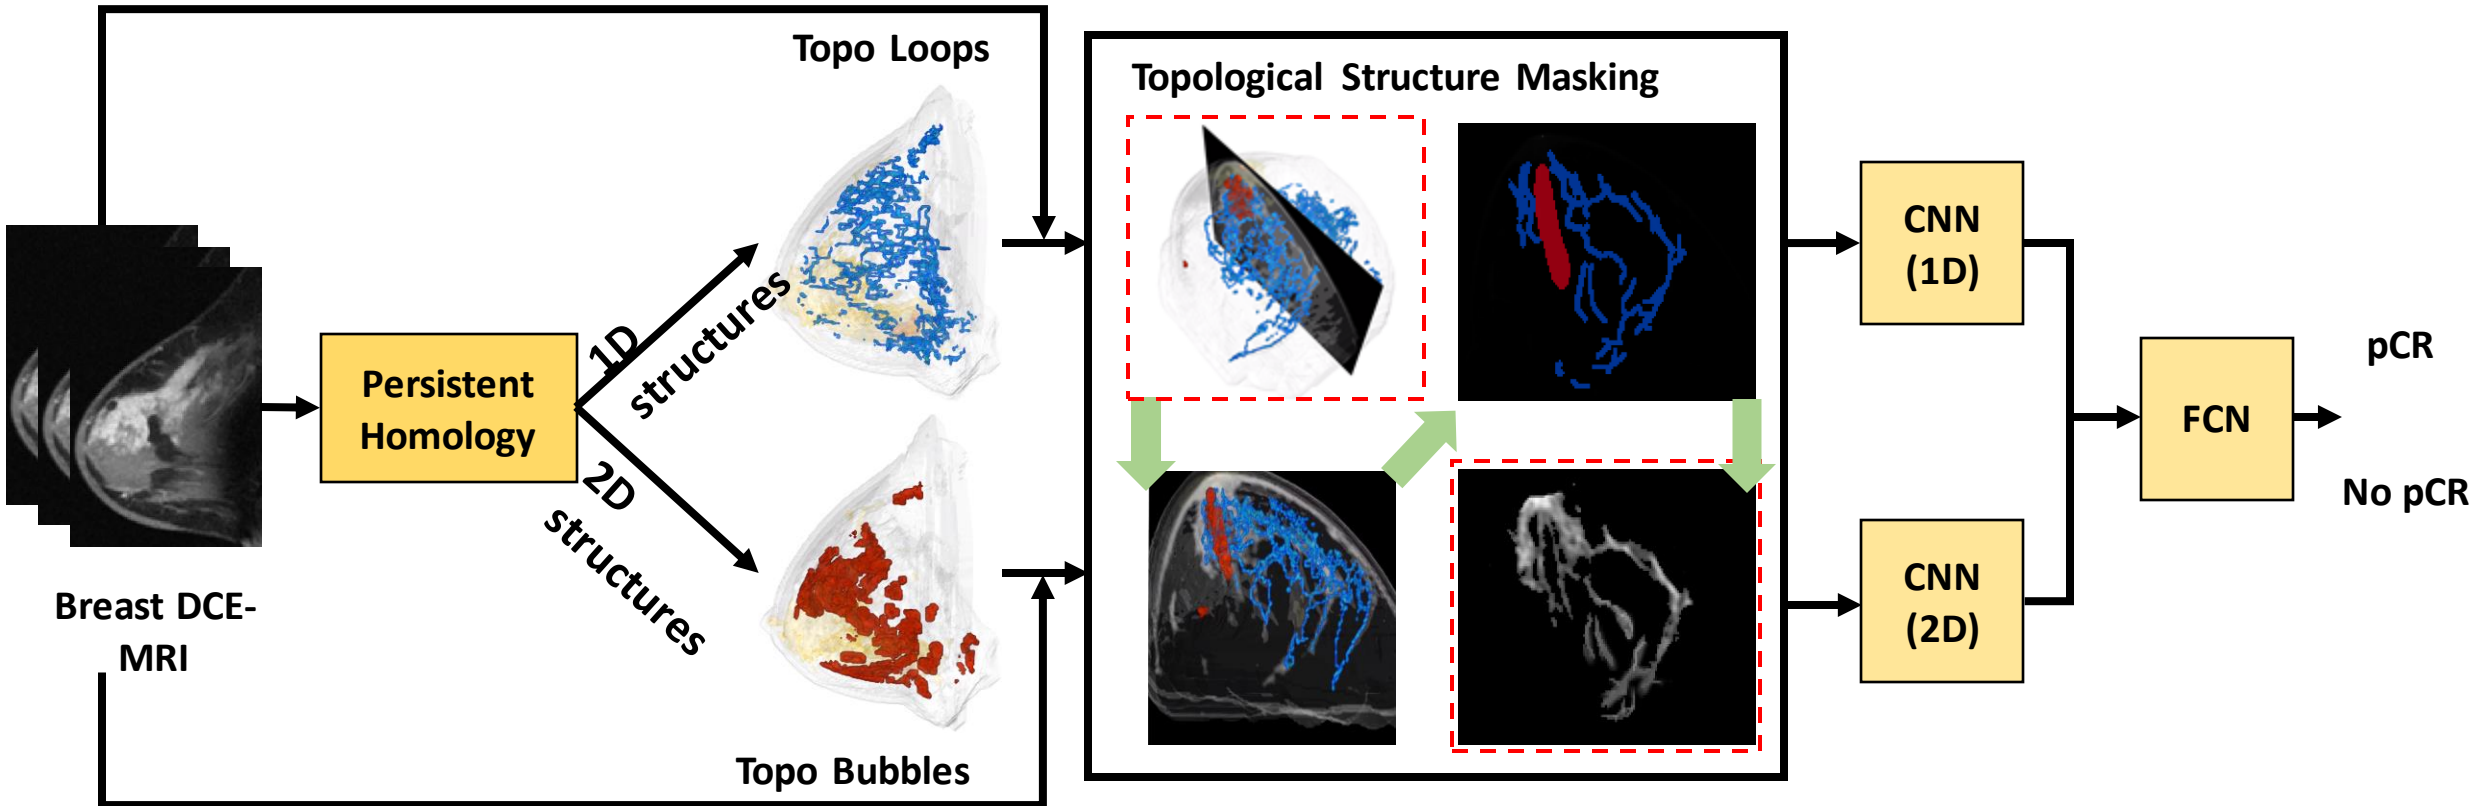
TopoTxR



Bridges the two extremes

[Directs the attention of CNN to the smaller set of clinically relevant voxels for training]

TopoTxR Framework



Extract 1D and 2D topological structures.

Mask the input MRI image so that only voxels of the extracted topological structures and their vicinity regions are visible.

Train a CNN with the extracted topological structures.

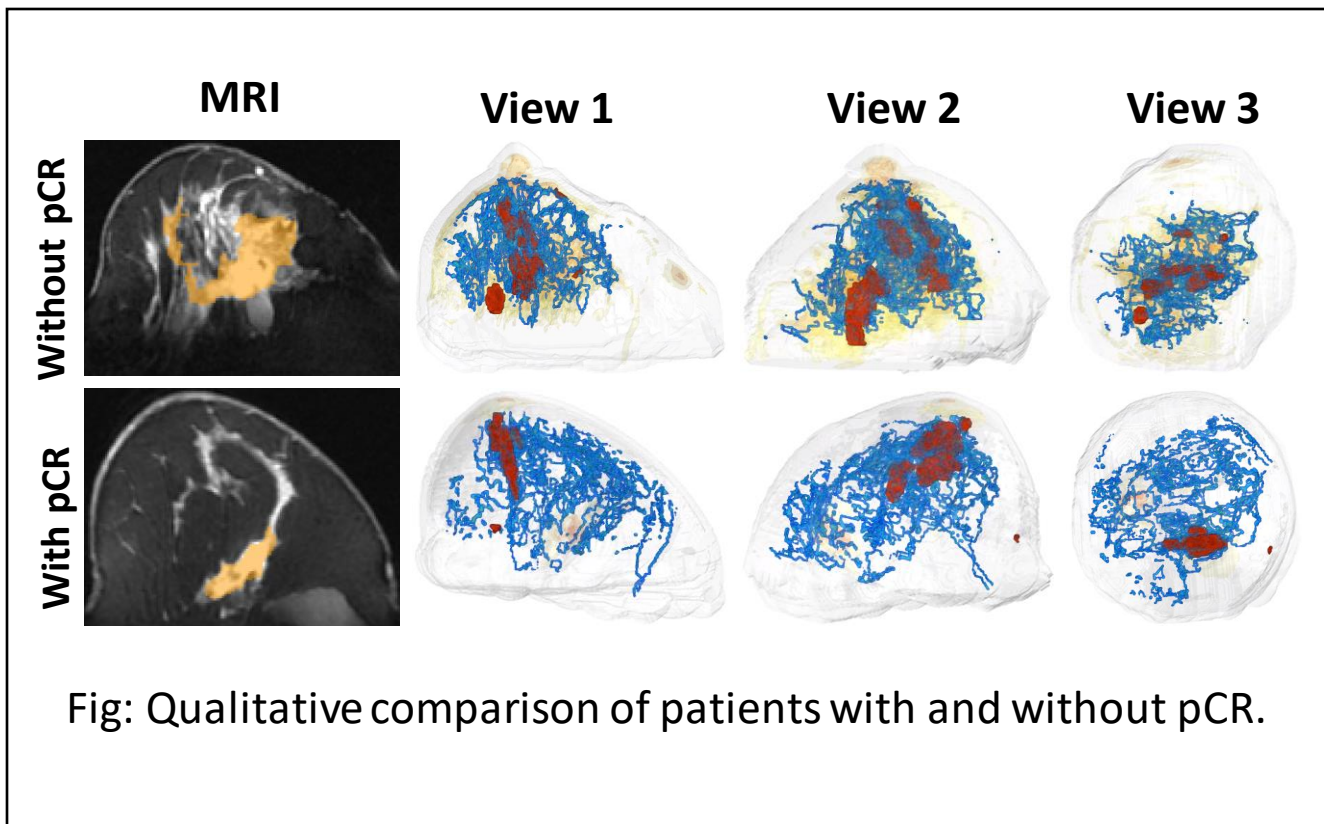
TopoTxR Results

Dataset: ISPY-1 post-contrast DCE-MRI dataset with 47 cases achieving response (pCR), and 115 non-pCR

higher score = better quality

	Accuracy	AUC	Specificity	Sensitivity
Radiomics	0.563	0.593	0.552	0.575
CNN	0.633	0.621	0.570	0.673
TopoTxR	0.851	0.820	0.736	0.904

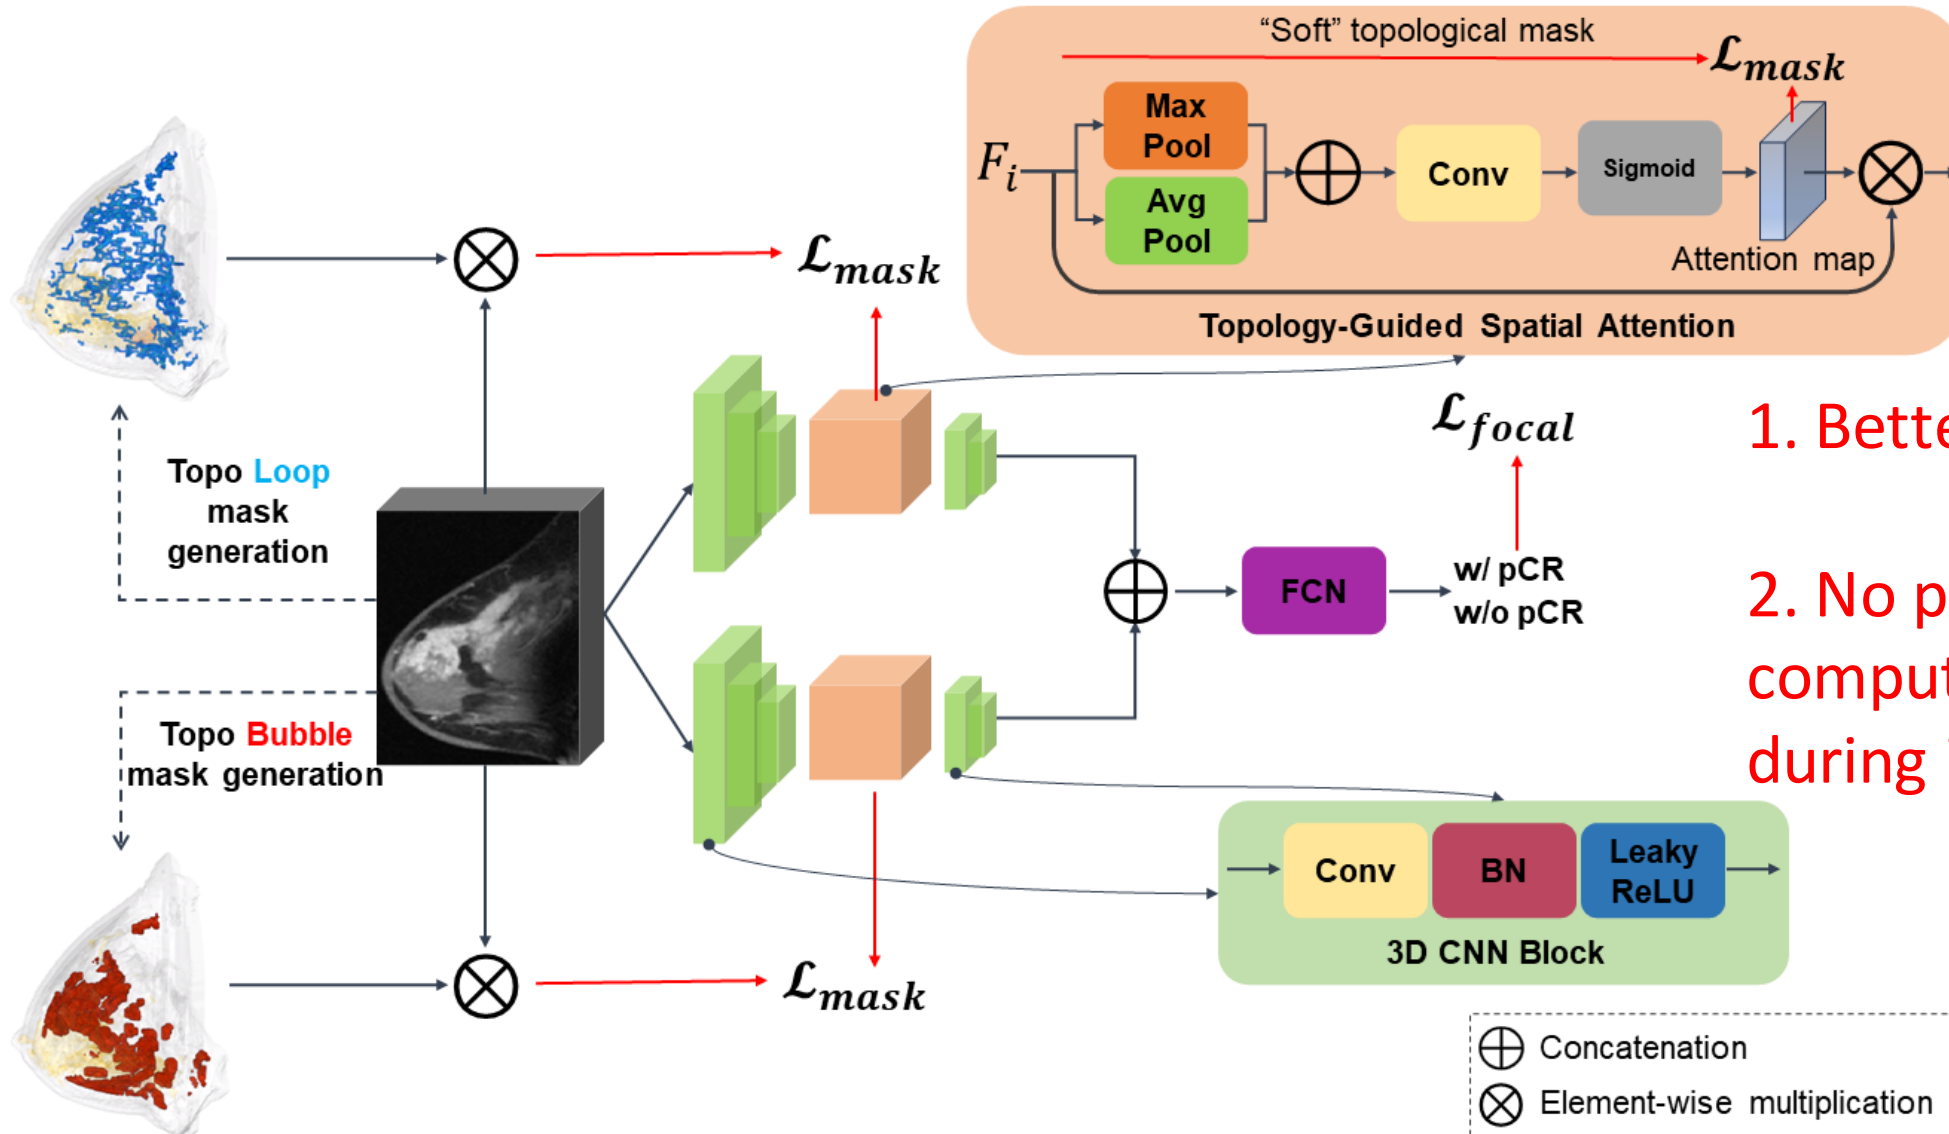
Comparisons of TopoTxR against baseline methods.



Structures are sparse for the case exhibiting pCR and are relatively dense for the non-pCR case.

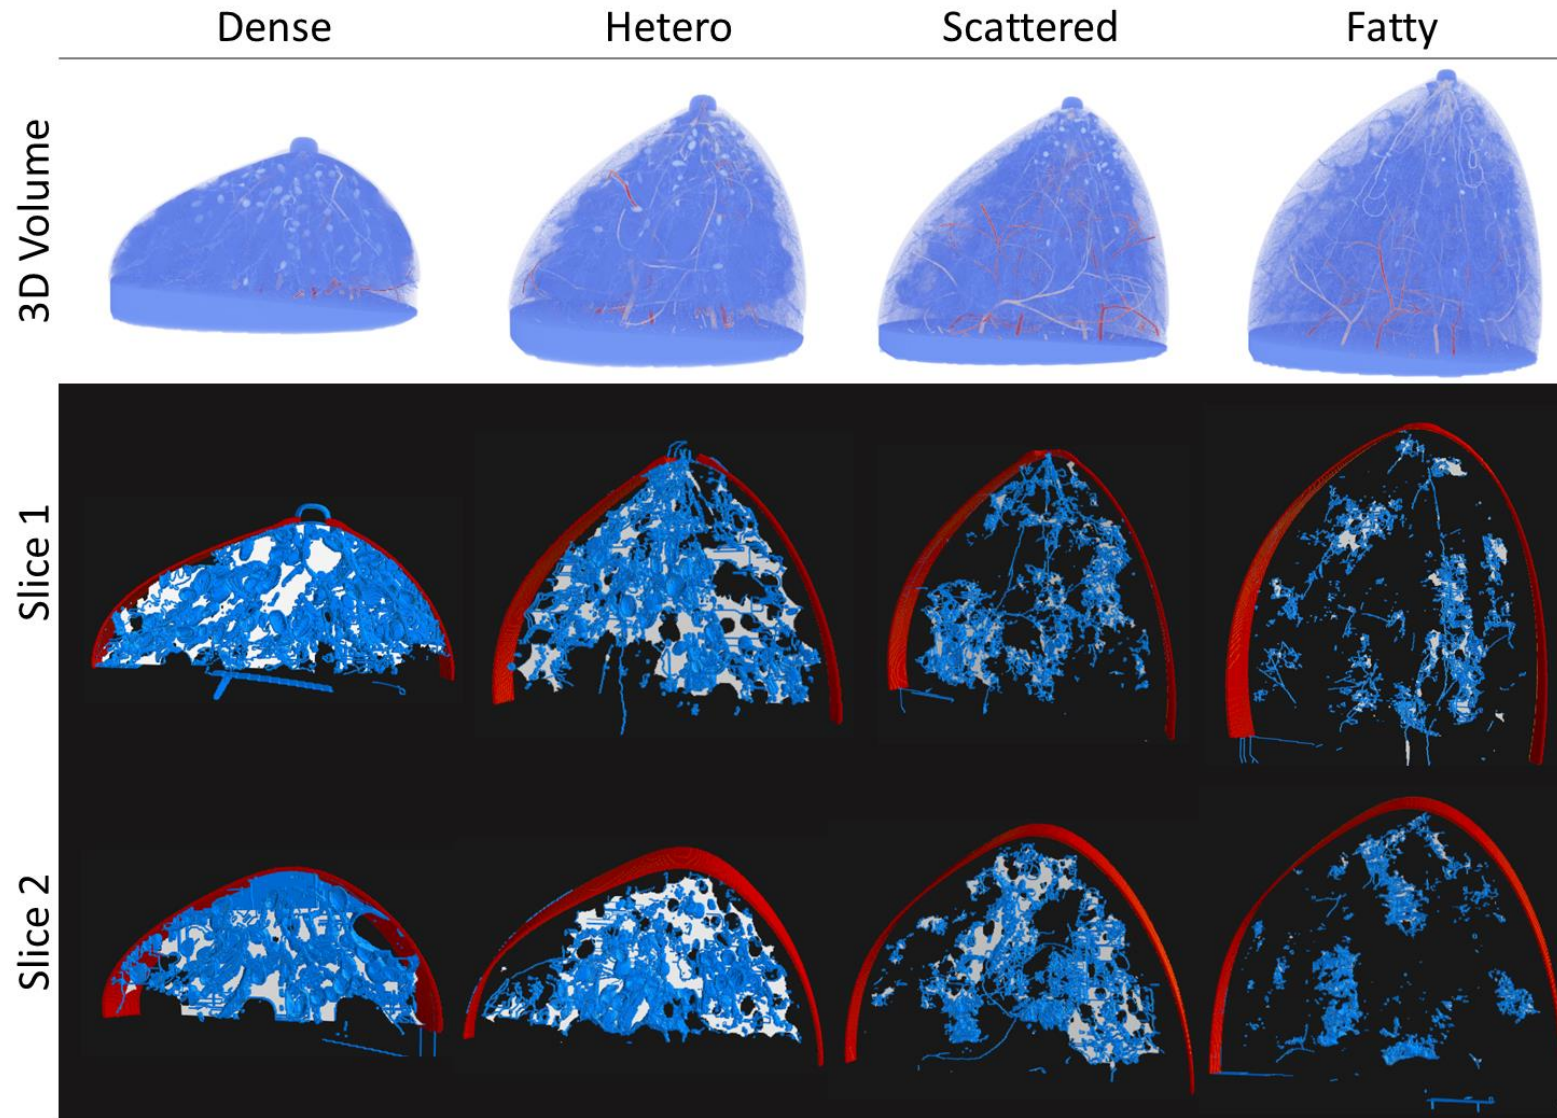
New Idea: Topology-Guided Spatial Attention

Train a spatial attention module to extract topological structures.



VICTRE Phantom Dataset

Are topological structures good approximations of breast tissues?



Red: 1-voxel width breast outline

Blue: extracted topological structures

White: ground truth breast tissues.

Contents

1. Introduction to topological data analysis and persistent homology.
2. Applications of persistent homology:
 - TopoGAN: A topology-aware generative adversarial network
 - TopoTxR: A Topological Biomarker for Predicting Treatment Response in Breast Cancer
3. Persistent homology computations using GPUs:
 - GPU computation of the Euler Characteristic Curve
 - GPU computation of persistent homology



Euler Characteristic

Simpler but still expressive topological descriptor.

Euler characteristic was introduced as a topological invariant that describes the shape of polytopes.

The Euler characteristic χ is defined as:

$$\chi = V - E + F$$







Name	Image	Vertices V	Edges E	Faces F	Euler characteristic: $V - E + F$
Tetrahedron		4	6	4	2
Hexahedron or cube		8	12	6	2
Octahedron		6	12	8	2
Dodecahedron		20	30	12	2
Icosahedron		12	30	20	2

Table from Wikipedia

Euler Characteristic Curve Example



(a) Example image


(b) $t = 0$
 $\chi(T_0) = 8$

(c) $t = 12$
 $\chi(T_{12}) = -17$

(d) $t = 25$
 $\chi(T_{25}) = -7$

(e) $t = 38$
 $\chi(T_{38}) = -3$

(f) $t = 52$
 $\chi(T_{52}) = 1$




(g) Euler characteristic curve of the example image


ECC as a Topological Descriptor

$$\chi(K) = \beta_0 - \beta_1 + \beta_2 - \beta_3 + \dots$$


grayscale input



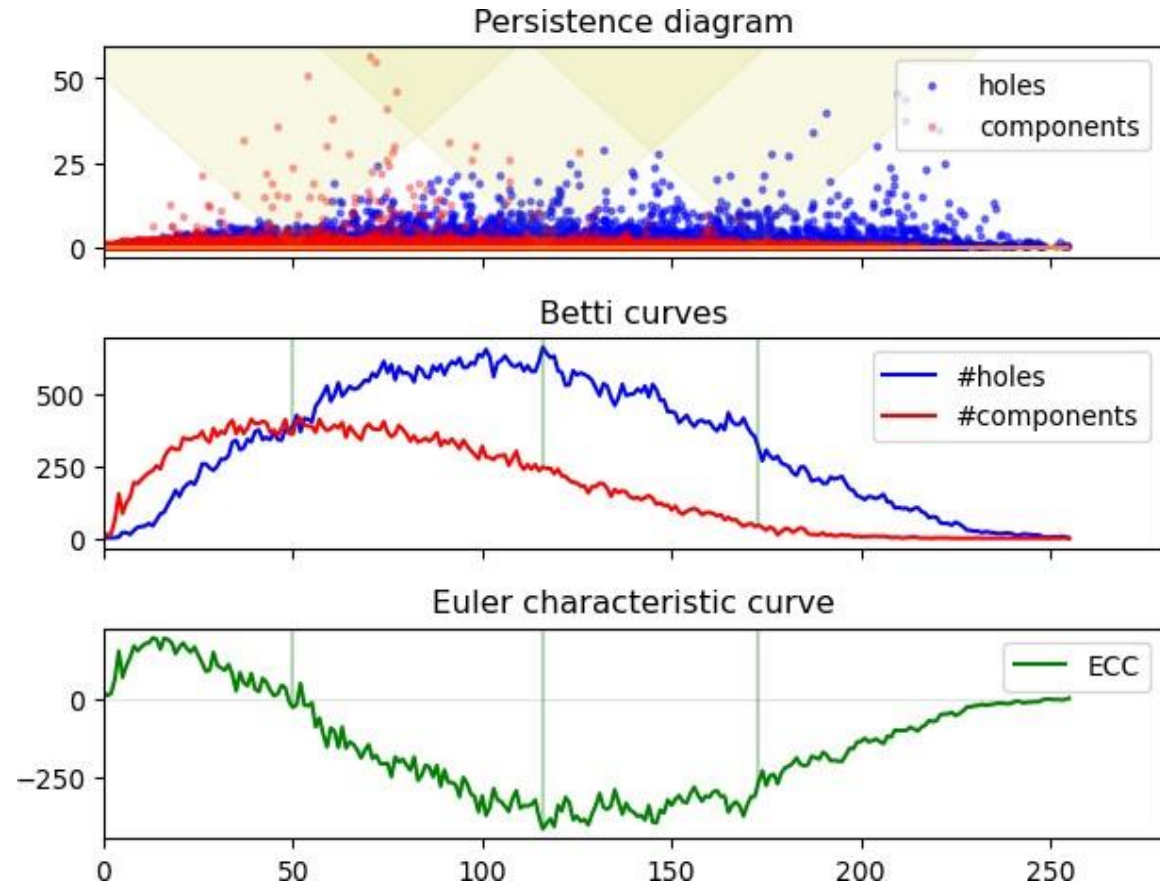

t=50



t=116



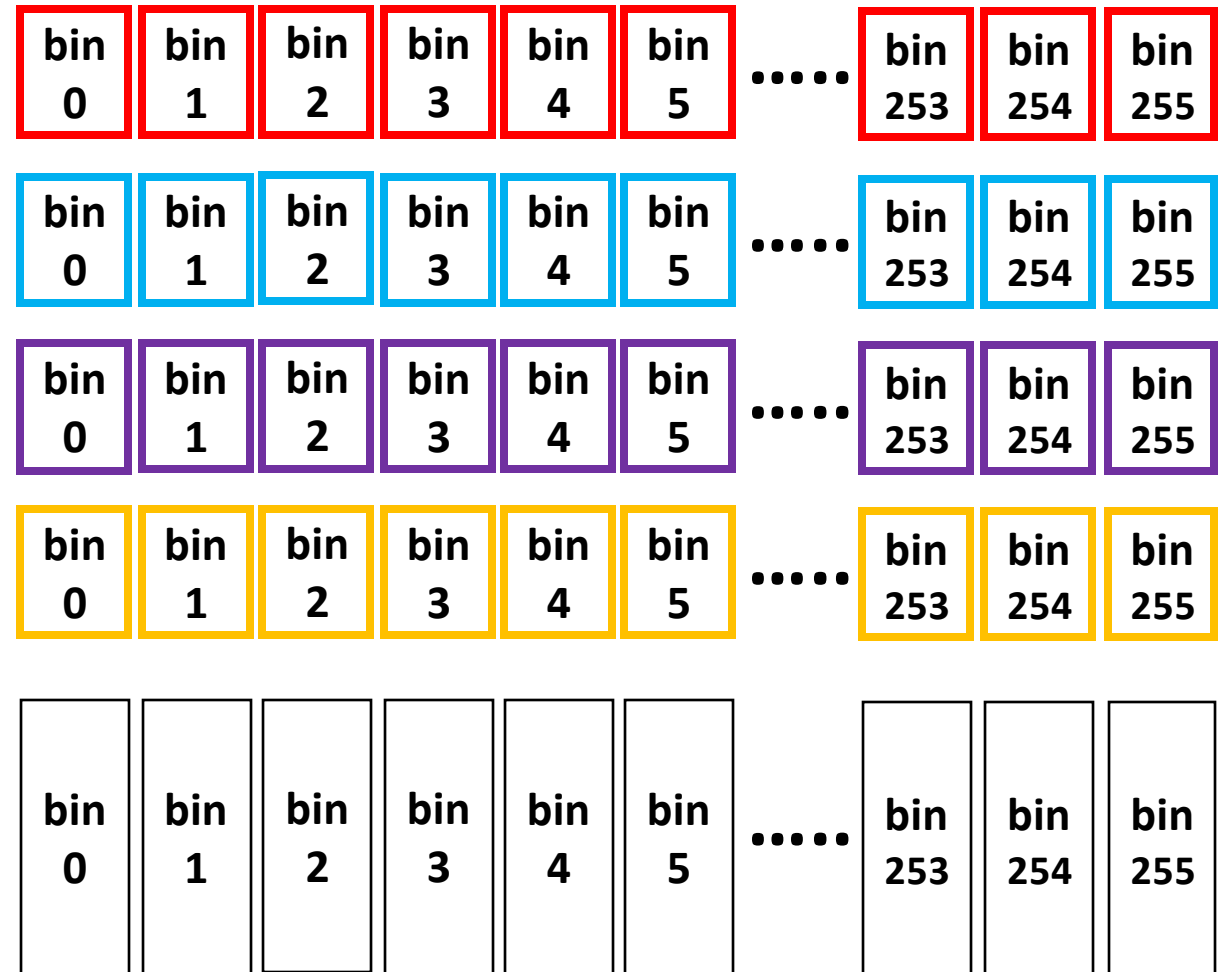
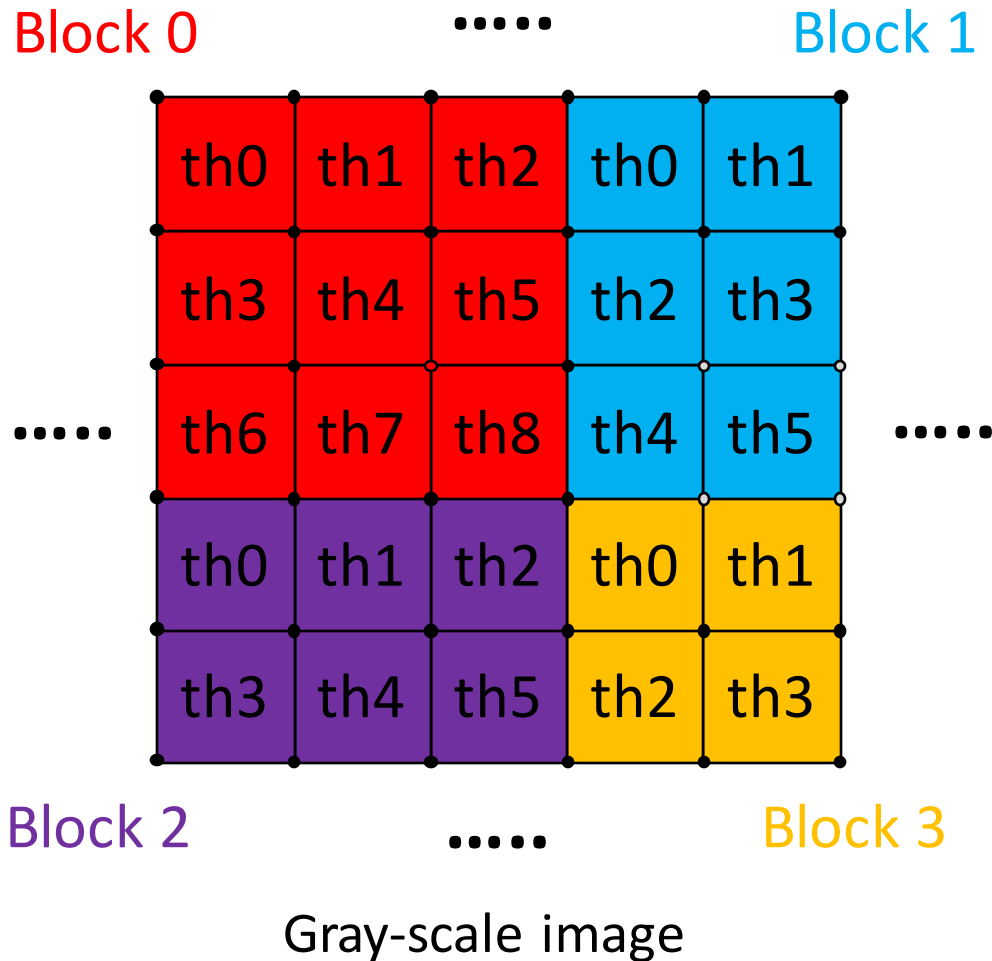
t=173



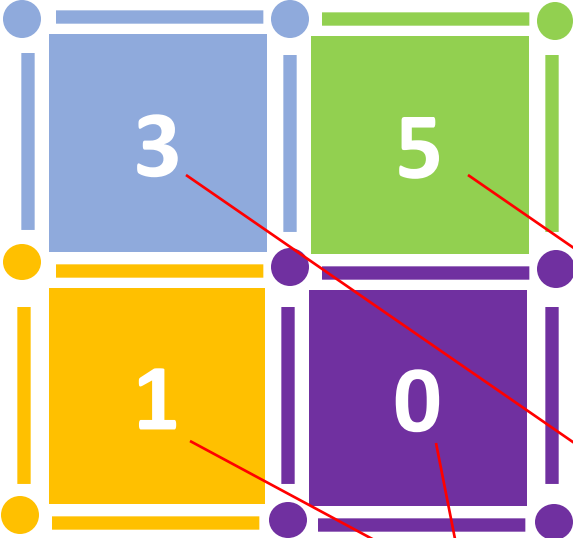
Workflow of GPU

Each GPU thread does simultaneously:


- Compute a value.
- Write to a bin in a local histogram.



ECC Computation – Abstracted Version




- Each pixel computes a number called Euler Characteristic.
- The result is accumulated into a bin corresponding to the intensity value of the pixel.



Euler Characteristic


Each pixel computes the Euler Characteristic by comparing to its 8 neighbors



Gray scale image


Convert VCEC to ECC

$$ECC_i = \sum_{j=0}^i VCEC_j$$




Prefix sum in GPU

ECC Computation – Parallelism



Gray scale image

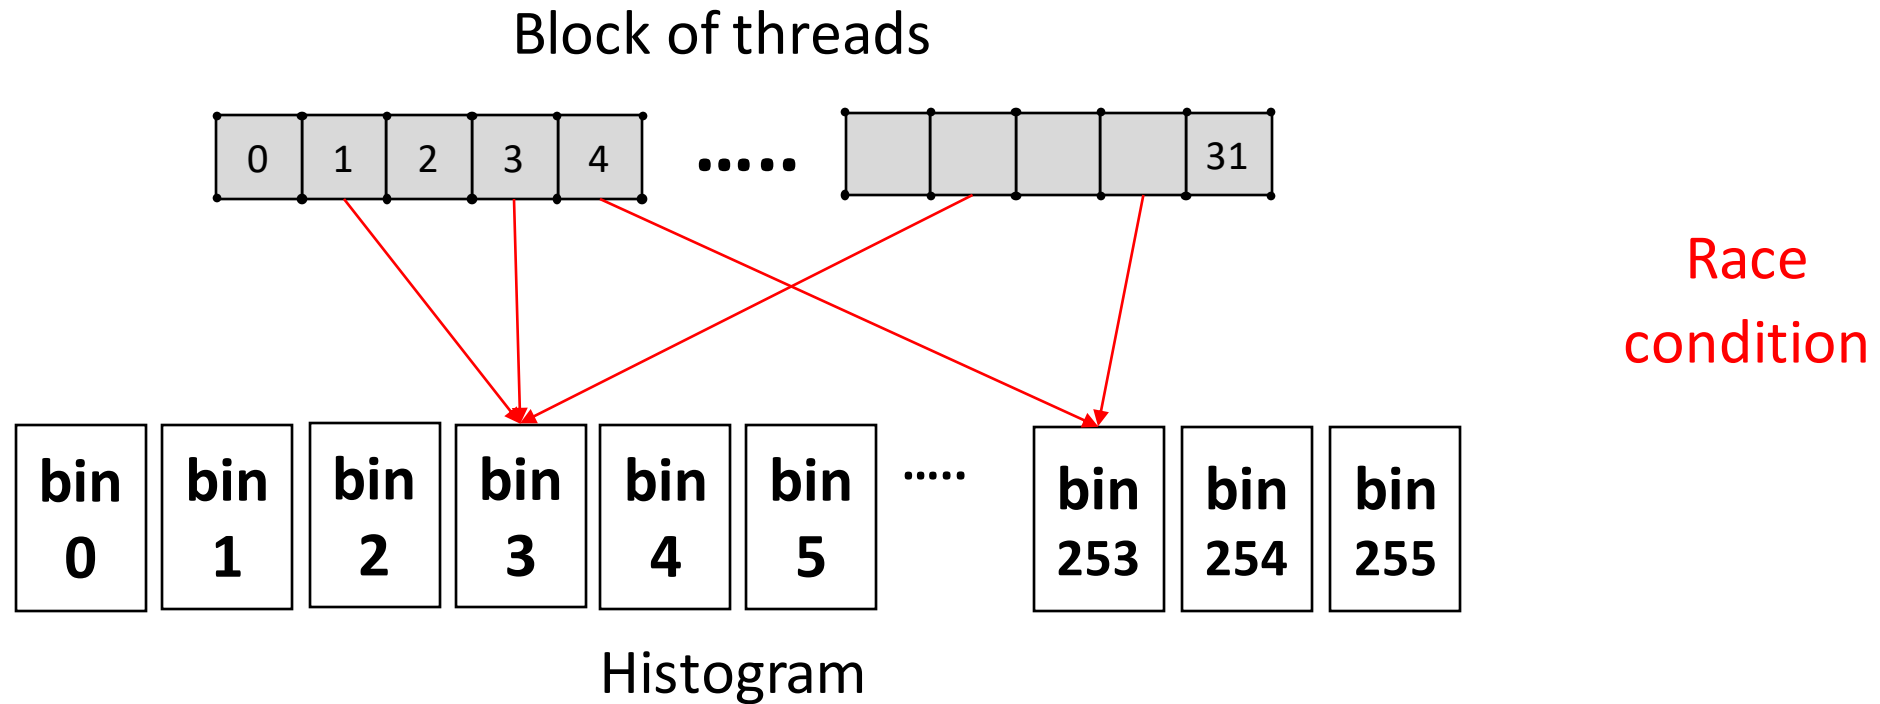


VCEC (Vector of Changes in Euler Characteristic)

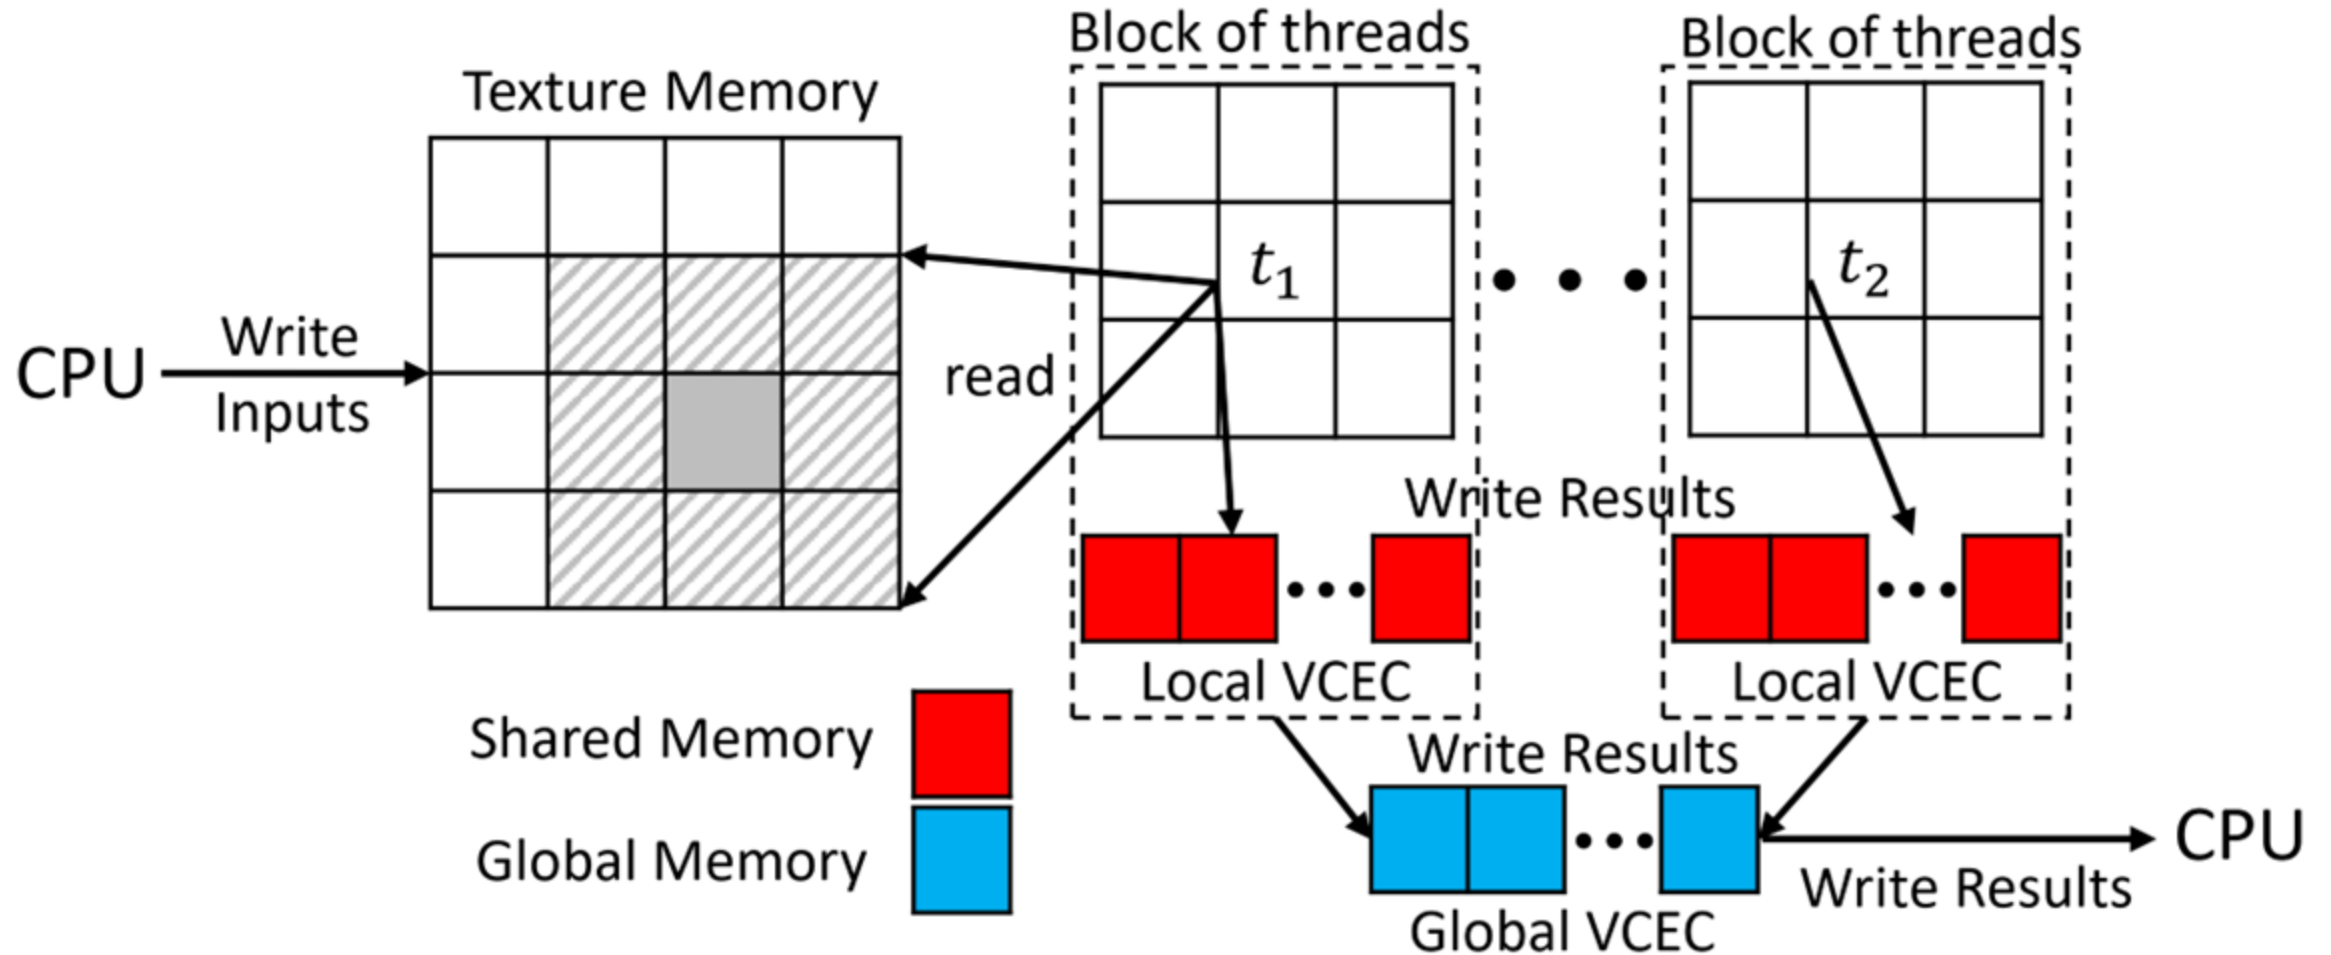
- Independence among pixels. Motivation for a parallel algorithm.
- ECC is like histogram computation!

Problem in GPU Histogram

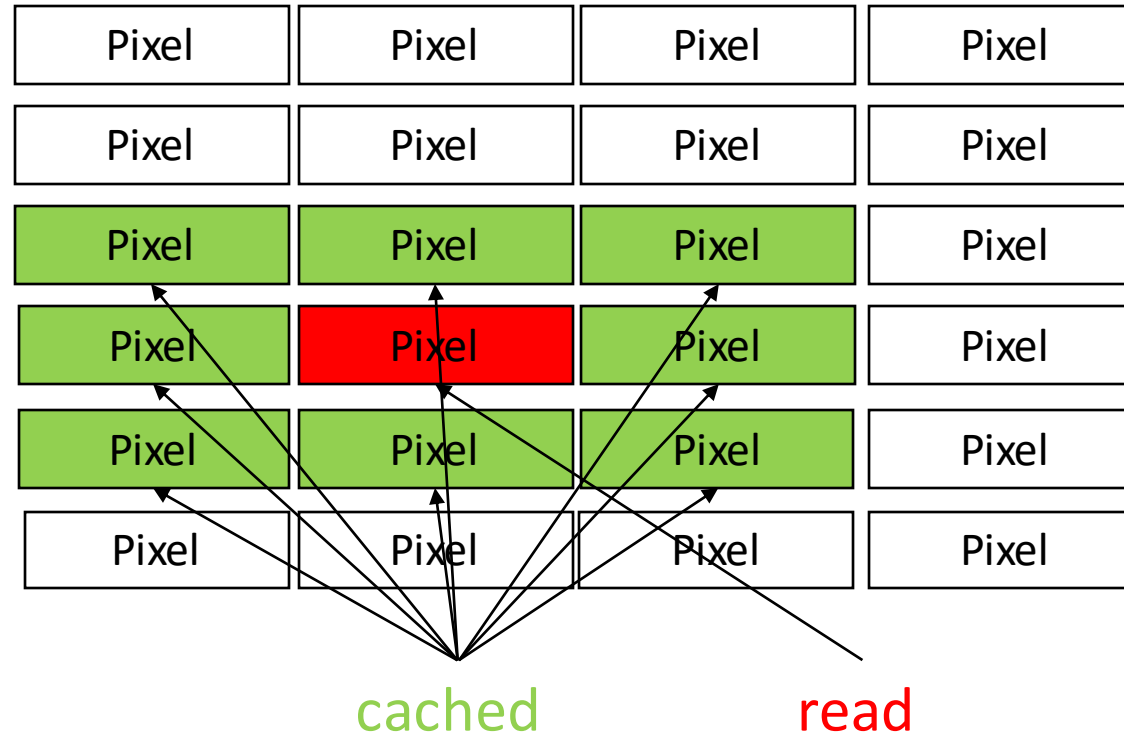
The output location for each element is not known prior to reading its value.



GPU ECC Framework



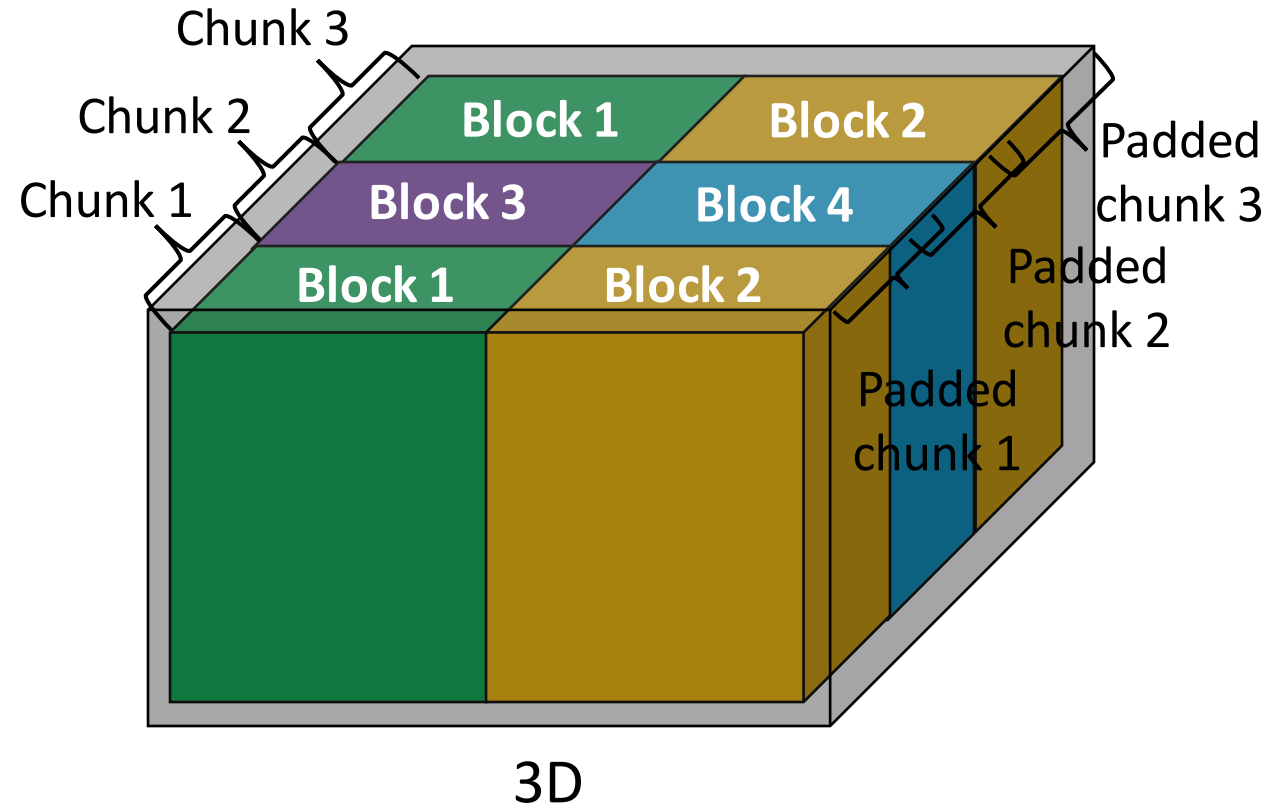
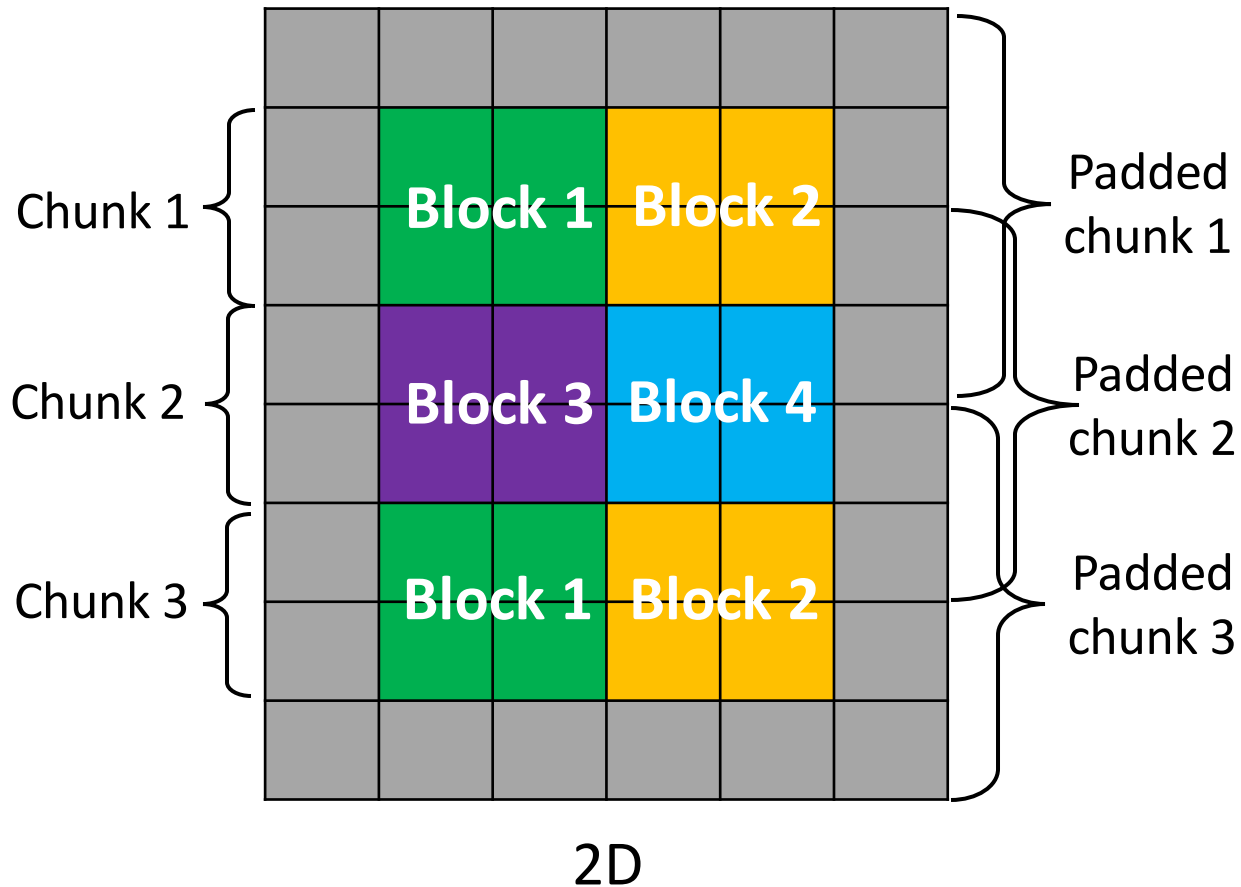
Texture memory



Motivation:

- Spatial locality.
- Multiple reading.

Limited GPU Memory and Streaming




CUDA Streams

Serial:



Asynchronous:



GPU ECC Results

For 3D images of size 512^3 ,
computing ECC:

► takes 1/30 second on a RTX 2070 GPU,

In general:

- unlimited image size due to streaming,
- 4 billion voxels per second throughput,
- yields a small and readily vectorized topological descriptor.

	Input size(B)	CPU overall	GPU overall	Overall speedup
Uniform Noise				
4096 ³	256G	37.72m	9.10m	4.14x
2048 ³	32G	4.86m	0.71m	6.77x
1024 ³	4G	36.85s	5.63s	6.55x
512 ³	512M	4.97s	0.85s	5.86x
Gaussian Random Field				
512 ³	512M	4.93s	0.86s	5.75x
256 ³	64M	0.63s	0.24s	2.58x
128 ³	8M	0.11s	0.12s	0.86x
8192 ³	256M	1.47s	0.53s	2.75x
4096 ³	64M	0.38s	0.21s	1.84x
2048 ³	16M	0.09s	0.18s	0.55x
VICTRE				
287 359 202	79.3M	0.59s	0.30s	1.98x
440 518 488	424M	2.99s	0.77s	3.87x
434 446 384	147M	1.11s	0.36s	3.02x
434 446 384	283M	1.96s	0.53s	3.70x
CMB				
1500 750	1.07M	0.03s	0.12s	0.22x
3000 1500	4.29M	0.09s	0.15s	0.61x
6400 3200	19.5M	0.37s	0.25s	1.49x

New Idea: Warp-Level Primitives


Block of threads




Warp1

Warp2

Warp32



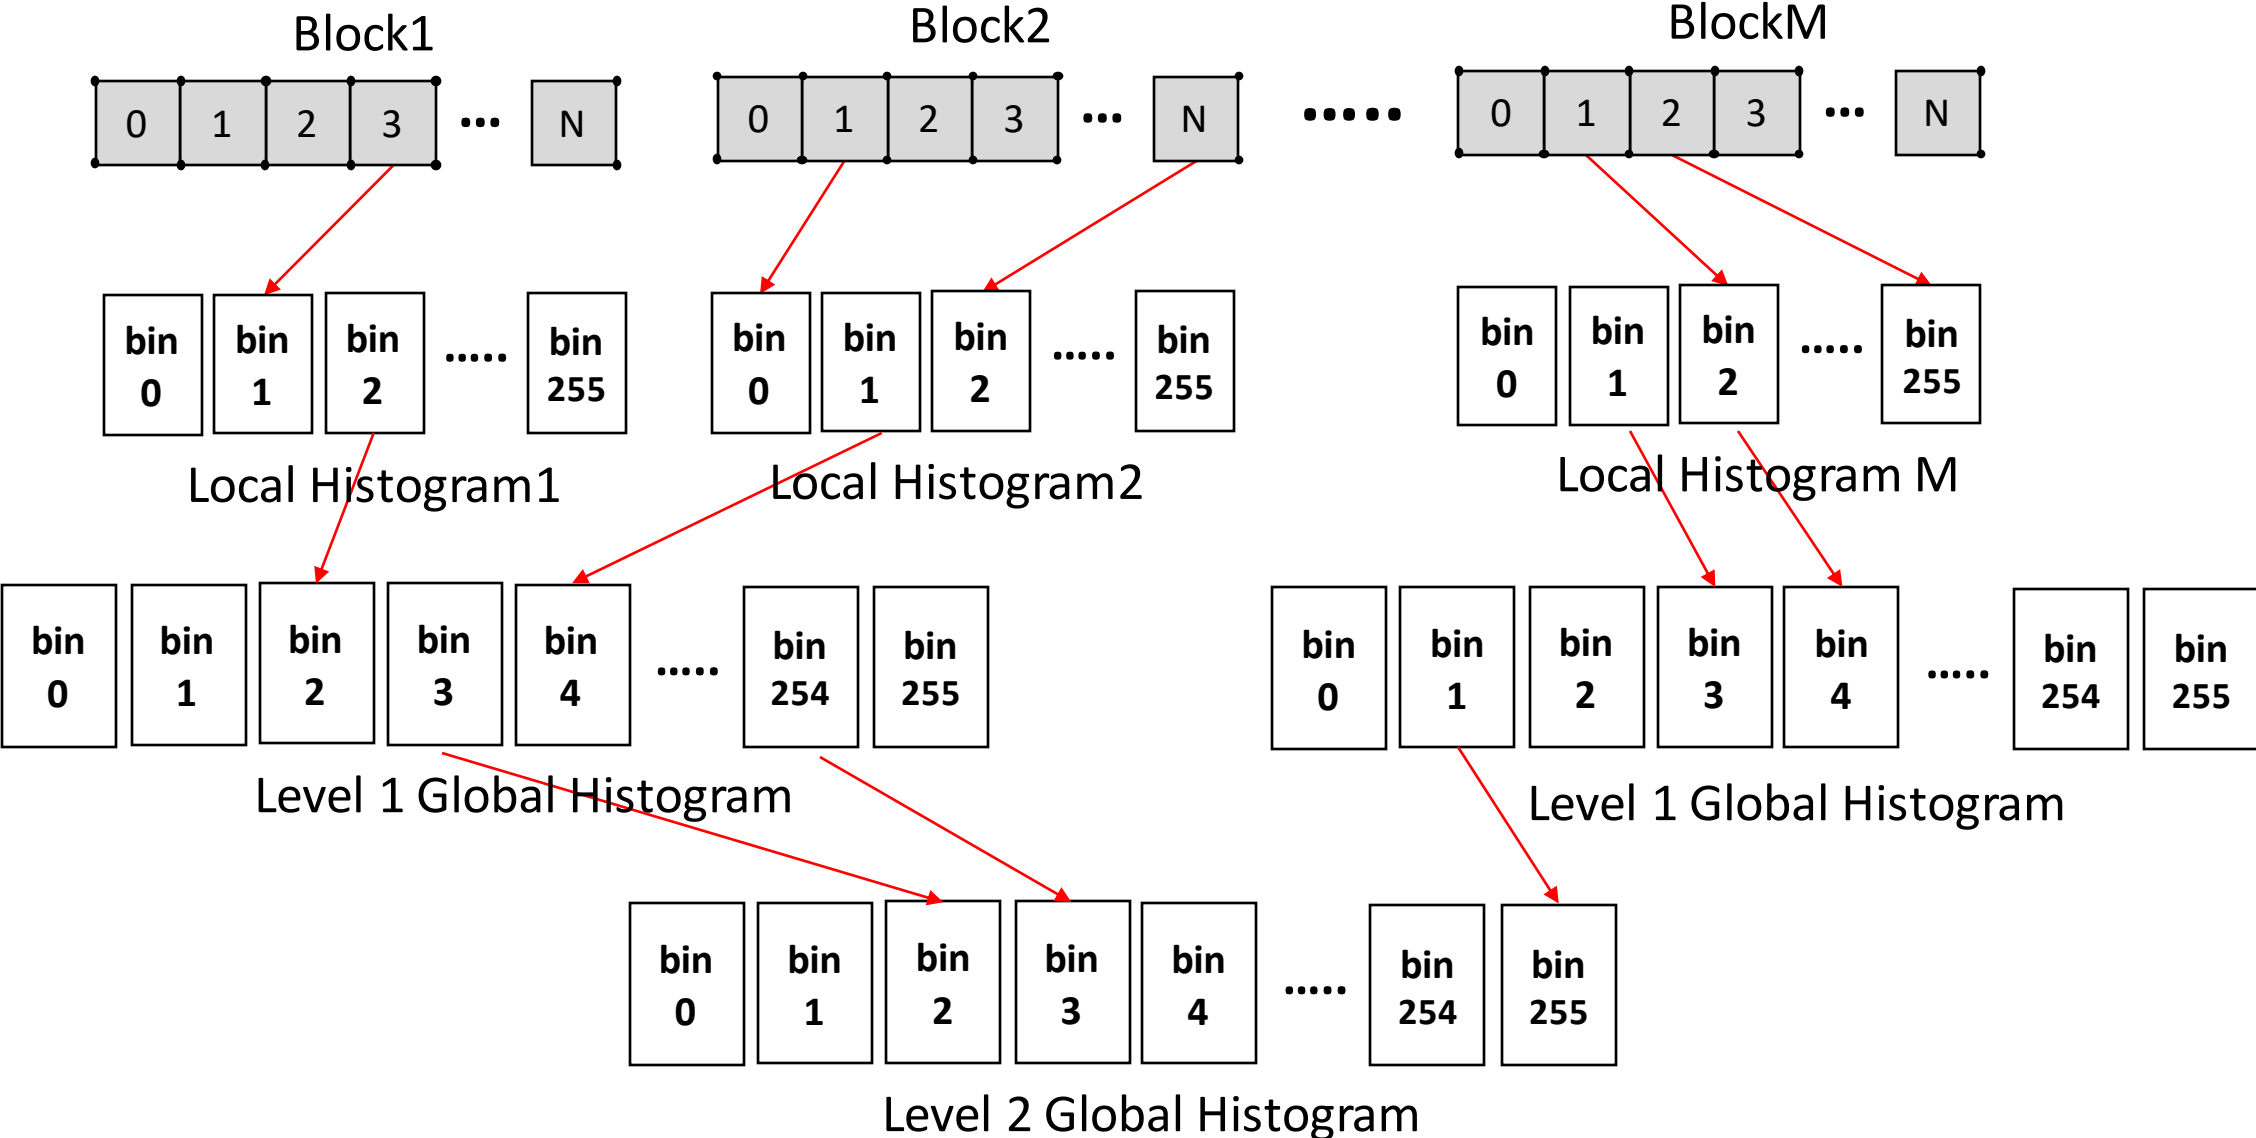
Warp-level
reduction



Local Histogram

Faster: data
exchange
between
registers

New Idea: Hierarchical Accumulation



Contents

1. Introduction to topological data analysis and persistent homology.
2. Applications of persistent homology:
 - TopoGAN: A topology-aware generative adversarial network
 - TopoTxR: A Topological Biomarker for Predicting Treatment Response in Breast Cancer
3. Persistent homology computations using GPUs:
 - GPU computation of the Euler Characteristic Curve
 - GPU-Accelerated Computation of Persistent Homology for Image Data



Efficiency of Persistence Computation

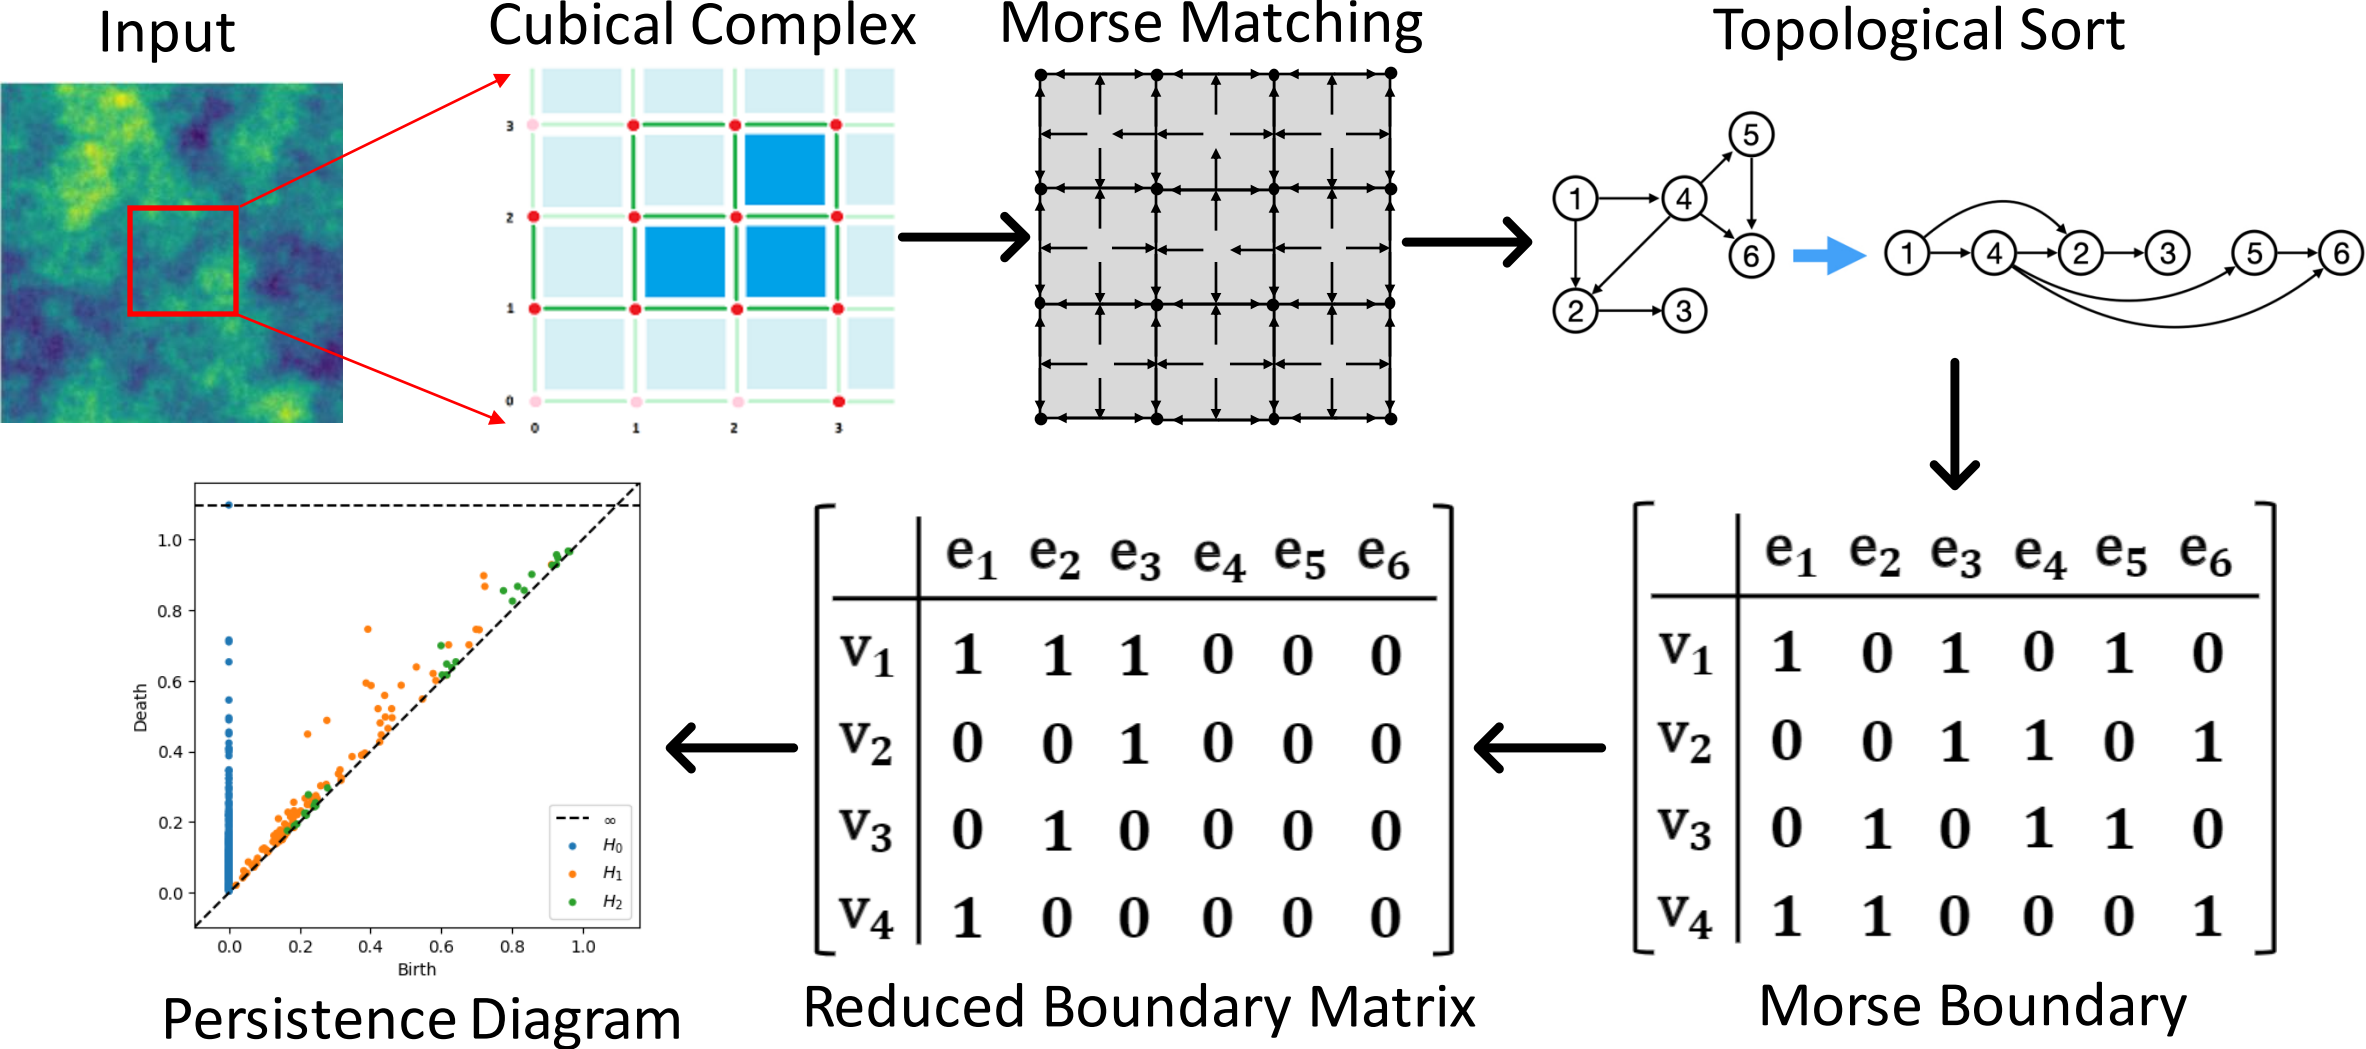
For 3D image of size 512^3 (130 M voxels), computing persistence using state-of-the-art software (e.g. CubicalRipser):

- ▶ takes several minutes,
- ▶ tens of gigabytes of memory.

Impediments for wider adoption and seamless integration with existing pipelines (machine learning, simulations...):

- ▶ Relatively high running time.
- ▶ High, unpredictable memory usage.
- ▶ Lack of GPU implementations on image data

Pipeline of Persistent Homology Computation





Topological Sort

Topological sort of a directed graph is an ordering of the vertices such that for every directed edge uv from vertex u to v , u comes before v in the ordering.

GPU topological sort:

- Very few literature
- No open-source codes
- Very challenging


Relax the problem:



Khan's Algorithm

1. Put nodes with no incoming edges into S
2. Take a node from S , delete all its outgoing edges, and put it into L .
3. Repeat from Step 1.


Khan's algorithm




10 Iterations!

Proposed Algorithm – Intuitions


- Incoming edges only from four neighbors.
- 0 indegree nodes do not have dependencies.



(a)




(b)




(c)

Proposed Algorithm – Intuitions Cont.


- 1 indegree nodes appear after their only parent in the ordering.



(a)




(b)



(c)


Proposed Algorithm – Intuitions Cont.

- What about 2 indegree nodes?



(a)

- Node 1 has 2 incoming edges from Node 0 and Node 6.
- Node 0 depends on Node 6, so Node 1 only depends on Node 0.




(b)

- Node 7 depends on Node 1.
- Node 11 depends on Node 5.
- Node 7 depends on Node 11.


Proposed Algorithm – Intuitions Cont.

- Can we always deal with 2 indegree nodes? **No**



(a)

- Node 6 has 2 incoming edges from Node 1 and 5.
- But there is no dependency between Node 1 and 5.




(b)

- Node 1 has 2 incoming edges from Node 0 and 2.
- But no dependency can be determined between Node 0 and 2 in a local neighborhood.


Algorithm Comparisons

Khan's algorithm




Algorithm Comparisons – Cont.

Parallel topological sort




Algorithm Comparisons – Cont.

GPU Khan's algorithm



Algorithm Comparisons – Cont.

Parallel topological sort



Future Works

1. GPU Topology computes only persistent homology.
 - Extract topological structures and corresponding geometries from the Morse complex.
2. GPU boundary matrix reduction.
 - Like Gaussian elimination, boundary matrix reduction is highly sequential and challenging to parallelize.

Summary of Important Contributions

1. A topological biomarker for treatment response prediction in breast cancer.
2. A topology-aware generative adversarial network – **ECCV Oral**
3. GPU Computation of Euler Characteristic Curve – **up to 6.77x speedup**
4. GPU Computation of Persistent Homology – **up to 20.3x speedup**
5. GPU Computation of Morse Complex – **up to 29.5x speedup**

Thank You!