KENNESAW STATE UNIVERSITY
COLLEGE OF COMPUTING AND SOFTWARE ENGINEERING
School of Data Science and Analytics

U-Net with *Persistent Homology* for Retinal Vessel Segmentation

Charles Fanning

Advisor: Dr. Mehmet Emin Aktas

INTRODUCTION

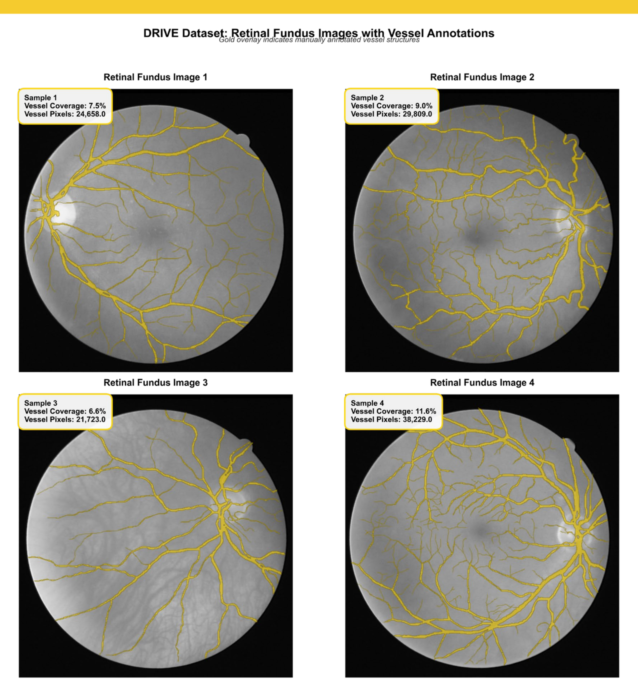

Fig 1. *Sample retinal fundus images for four patients. The vessel segmentation masks are highlighted in gold in front of the grayscale images, and some statistics are given in the top left of each image.*

Motivation:

- Retinal vessel segmentation is necessary for the diagnosis of diabetic retinopathy, as in the case of the DRIVE dataset.
- To be clinically trustworthy, the results need to be non-arbitrary and interpretable, not just having high segmentations scores.

Problem:

- Standard U-Net architectures optimize local, pixel-wise overlap and cannot reliably model global, topological features.

Proposed Solution:

- Augment U-Net with *Persistent Homology* (PH).

METHODS

Dataset:

- We used DRIVE color fundus photographs with expert vessel masks. We also performed data augmentation with horizontal and vertical flips, small rotations, and color jitter.

U-Net:

- The architecture uses four encoder and four decoder stages with skip connections. For the baseline model, we use pixel-wise cross-entropy loss, the Adam optimizer with a learning rate of 10^{-4}, and a batch size of four for 175 iterations per fold.

Persistent Homology:

- **Input:** predicted vessel probability map $p\colon \Omega \to [0,1]$ on a 2D grid of pixels Ω.
- **Filtration:** for each threshold value τ, define $\Omega_\tau = \{x \in \Omega : p(x) \geq \tau\}$, a superlevel set of images by increasing intensity value.
- **Cubical complexes:** interpret each Ω_τ as a union of pixels to obtain a nested family of cubical complexes.
- **Persistent homology:** track when topological features such as holes and connected components appear and disappear along the filtration.

PersLay Vectorization:

- Persistence diagrams are multisets of birth-death pairs. As such, they cannot be used as input to machine learning models. We must first vectorize them.
- PersLay is a differentiable, stable, trainable vectorization layer for persistence diagrams:
  1. First, landmark status is applied to certain points ℓ_j in the persistence diagram.
  2. Then distances from the landmark are pooled by:

$$v_j = \sum_i \varphi\left(\|(b_i, d_i)\| - \ell_j\right)$$

Topology-Aware U-Net:

- This can be done by modulating the feature maps by the PersLay vectorized persistence diagram.

Topology-Aware Loss:

- The Wasserstein difference in persistence diagrams between the predicted segmentation and the ground truth are used to derive a topological loss term which is aggregated with the overall loss.

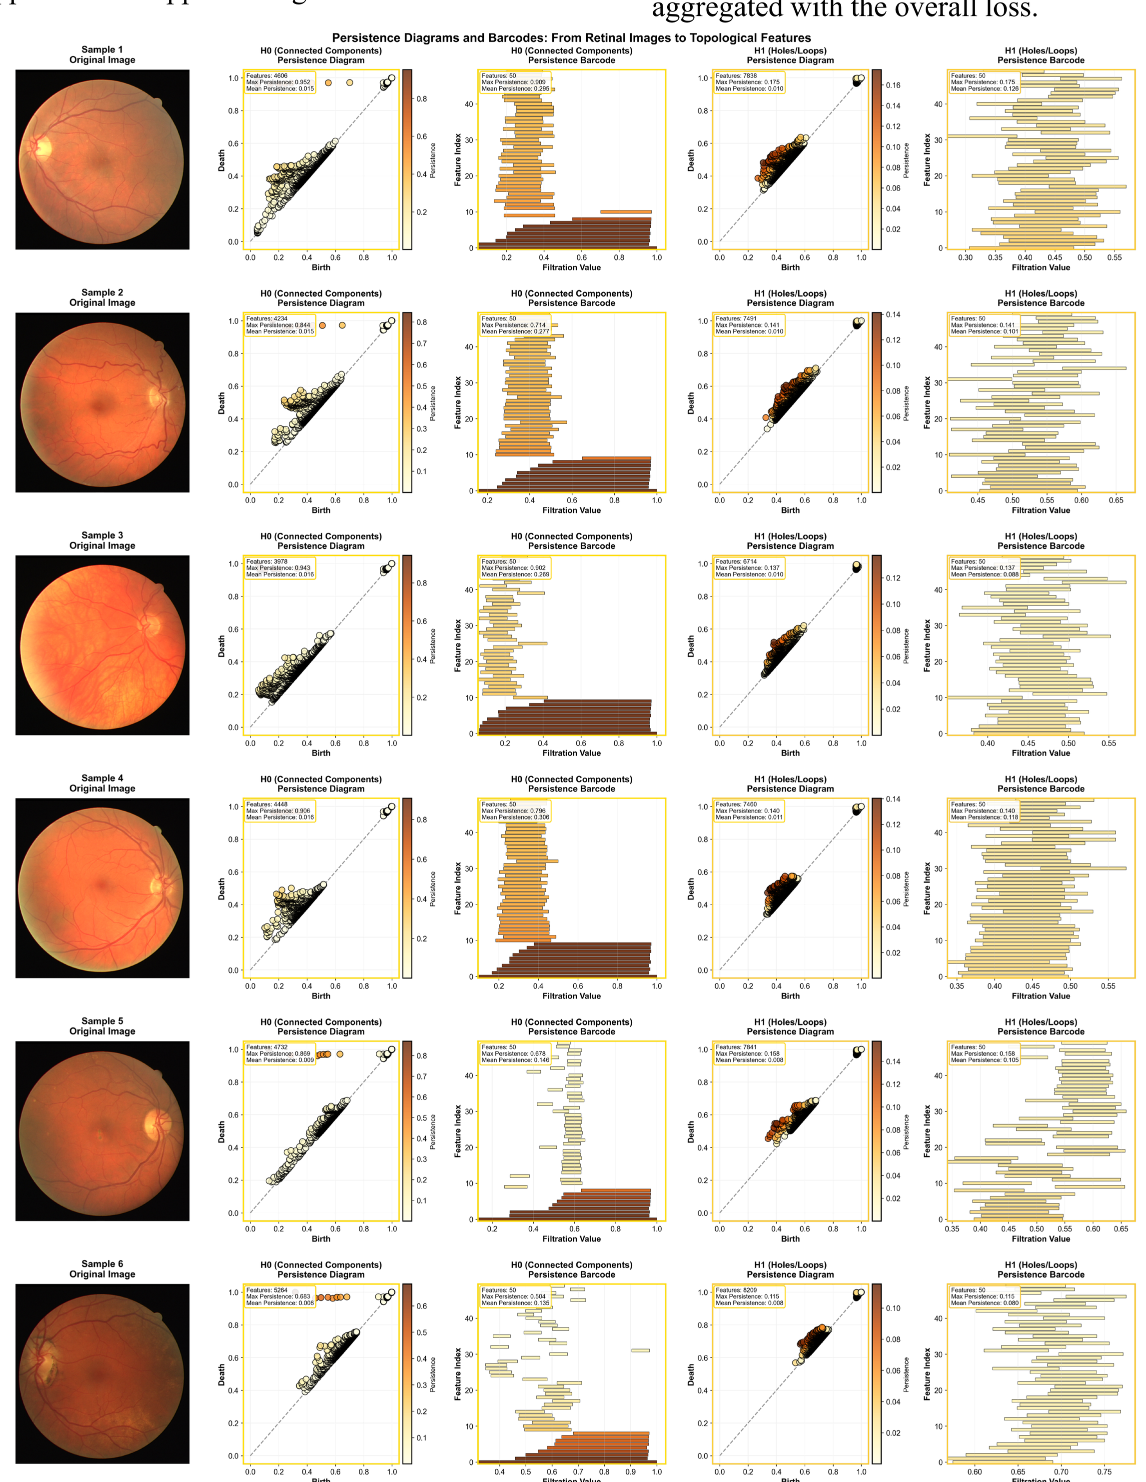

Fig 2. *The persistent homology pipeline: we take an image as input, compute its zero- and one-dimensional homology groups, then we down sample into fifty-sample persistence barcodes.*

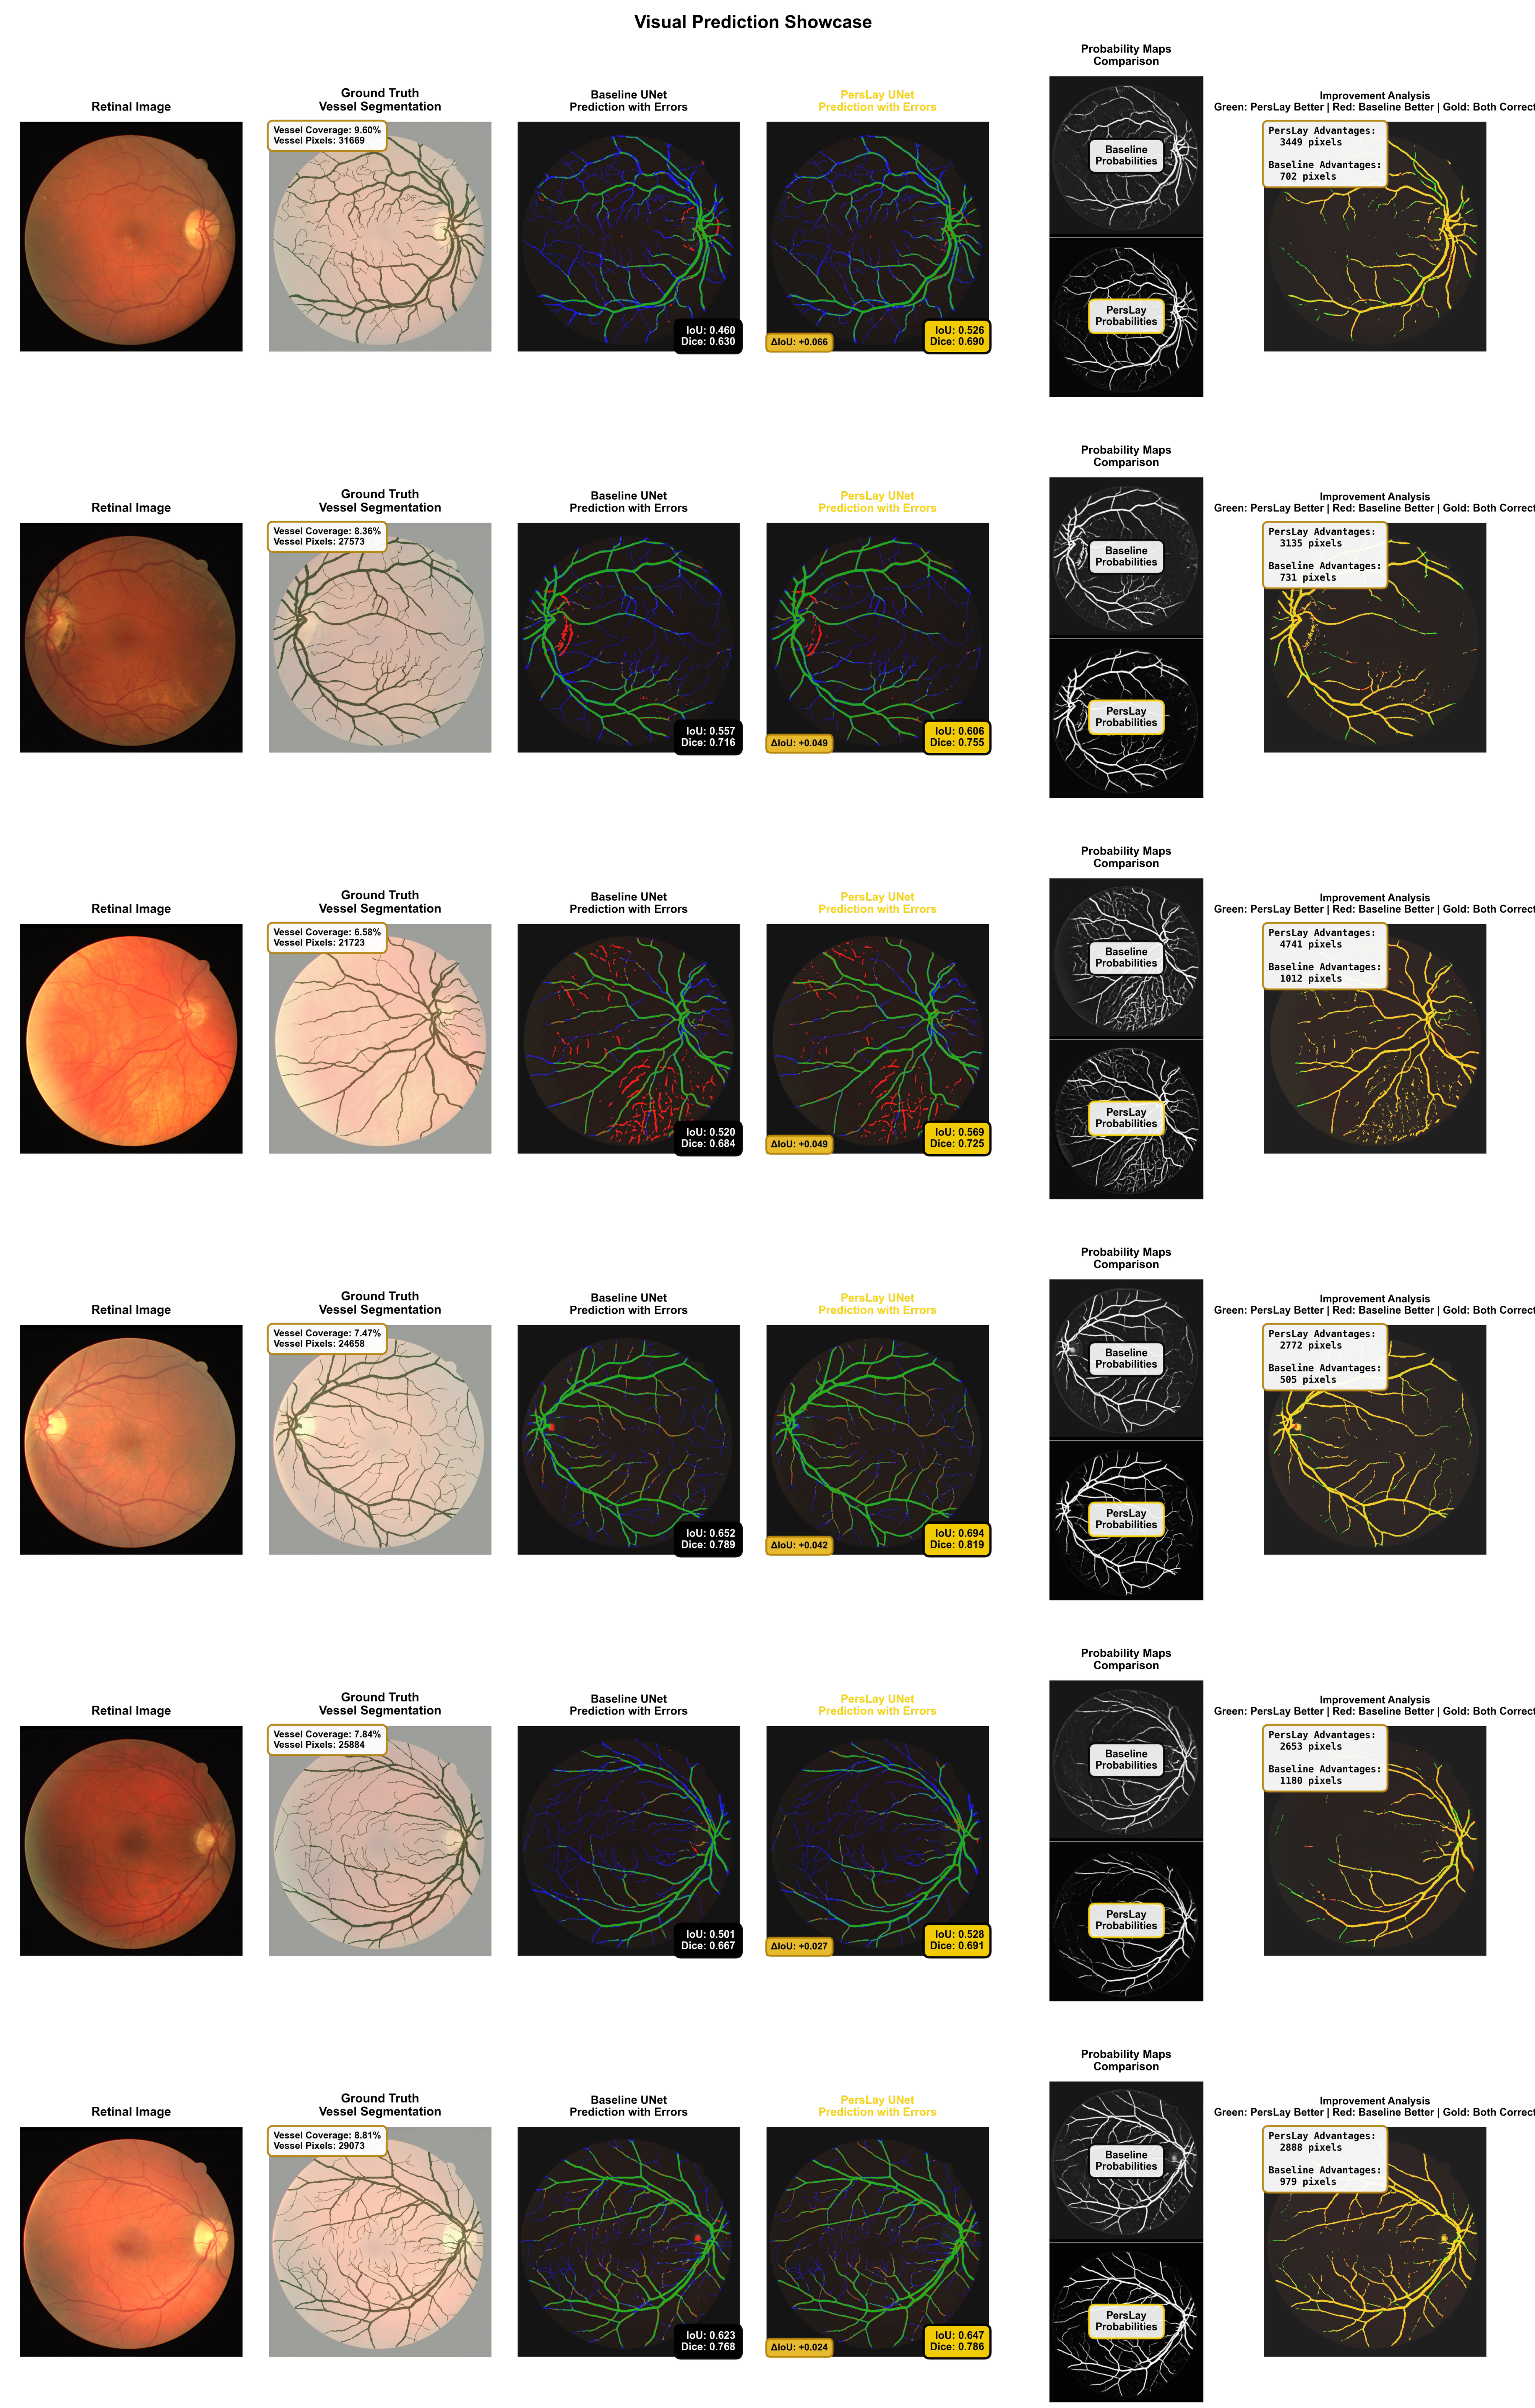

Fig 3. *Visual comparison of retinal vessel segmentation results for the six samples with the highest IoU improvement using PersLay U-Net over baseline U-Net.*

RESULTS

Intersection over Union (IoU):

- Baseline: 0.5829
- Topology-Aware U-Net: 0.5857

Dice Loss:

- Baseline: 0.7358
- Topology-Aware: 0.7379

Combined Topological Loss:

- Baseline: 0.1957
- Topology-Aware: 0.1722

Probability Map Interpretation:

- The topology-aware U-Net gives sharper, more confident vessel segmentations than the baseline U-Net does.

DISCUSSION

Interpretation of the Results:

- Including topological information into the analysis with U-Net made modest improvements to the segmentation quality.
- On the relatively small dataset, it took both the topological loss term as well as the topological modulation of the feature maps of U-Net for this modest to improvement to be made.
  - On the one hand, this suggests that topological augmentation does improve segmentation, even on small, complex datasets.
  - However, it also suggests that augmenting models in additional ways with topological features is not immediately redundant; there may be more topological augmentation required for U-Net to fully model the topology of the underlying data.

When to use Topological Augmentation:

- From the results, we can see that the topologically augmented U-Net was more confident and less likely to include unnecessary pixels in its predicted segmentation mask.
- Situations where this careful, predictable shape and behavior is desired appears to be the ideal circumstance for persistent homology in machine learning.

CONCLUSION

1. PersLay makes U-Net retinal vessel segmentation topology-aware with moderate improvements to the segmentation results.
2. The topological augmentations should be used to augment segmentation networks in cases such as these where they provide an advantage over the baseline architecture:
   - ✓ Small datasets wherein using a simpler model with explicit mathematical guarantees (from persistent homology) mean that we can rely on properties of the data without needing a larger amount of it.
   - ✓ Data that is known to have topological structure (in this case, retinal fundus images have plenty of meaning in terms of connected components and loops in the images).
   - ✓ Data where simple model architectures work well, such that complicating the architecture without providing more mathematical meaning does not necessarily improve the performance.
   - ✓ Situations such as medical imaging where there is a major emphasis on the reliability of the results.

CODE

https://github.com/cfanning8/AnalyticsDayFall2025PersistentHomology

REFERENCES

1. Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., & Ginneken, B.V. (2004). Ridge-based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging, 23*, 501-509.
2. Carrière, M., Chazal, F., Ike, Y., Lacombe, T., Royer, M., & Umeda, Y. (2019). PersLay: A Neural Network Layer for Persistence Diagrams and New Graph Topological Signatures. *International Conference on Artificial Intelligence and Statistics*.
3. Hu, X., Li, F., Samaras, D., & Chen, C. (2019). Topology-Preserving Deep Image Segmentation. *ArXiv, abs/1906.05404*.
4. Zomorodian, A., & Carlsson, G.E. (2004). Computing Persistent Homology. *Discrete & Computational Geometry, 33*, 249-274.