Graph-Theoretic Consistency for Robust and Topology-Aware Semi-Supervised Histopathology Segmentation (Student Abstract)

Ha-Hieu Pham[1,2,3], **Minh Le**[4], **Han Huynh**[4], **Nguyen Quoc Khanh Le**[4], **Huy-Hieu Pham**[3,5*]
[1]University of Science, VNU-HCM, Ho Chi Minh City, Vietnam
[2]Vietnam National University, Ho Chi Minh City, Vietnam
[3]VinUni-Illinois Smart Health Center, VinUniversity, Hanoi City, Vietnam
[4]AIBioMed Research Group, Taipei Medical University, Taipei City, Taiwan
[5]College of Engineering & Computer Science, VinUniversity, Hanoi City, Vietnam
hieu.ph2@vinuni.edu.vn, {D142111009, m658112001, Khanhlee}@tmu.edu.tw, hieu.ph@vinuni.edu.vn

Abstract

Semi-supervised semantic segmentation (SSSS) is vital in computational pathology, where dense annotations are costly and limited. Existing methods often rely on pixel-level consistency, which propagates noisy pseudo-labels and produces fragmented or topologically invalid masks. We propose **Topology Graph Consistency (TGC)**, a framework that integrates graph-theoretic constraints by aligning Laplacian spectra, component counts, and adjacency statistics between prediction graphs and references. This enforces global topology and improves segmentation accuracy. Experiments on GlaS and CRAG demonstrate that TGC achieves state-of-the-art performance under 5–10% supervision and significantly narrows the gap to full supervision.

Introduction

Semantic segmentation in histopathology is a key step for computational pathology, enabling analysis of tissue architecture and cancer grading. Yet, obtaining dense pixel-level annotations remains costly and requires expert effort, especially for complex glandular structures in colorectal tissues. Semi-supervised learning provides a practical solution by leveraging unlabeled data to improve segmentation with limited supervision. However, most existing methods enforce pixel-level consistency, which is prone to label noise and tends to produce fragmented or topologically invalid masks. These errors are particularly critical when gland morphology and lumen enclosure have diagnostic importance.

Unlike pixels, region-level representations can model structural relationships between tissue components. Graph-based formulations are thus well suited for histopathology, as they naturally capture connectivity, adjacency, and topology between glandular regions (Felzenszwalb and Huttenlocher 2004; Kipf and Welling 2017; Zhang, Cui, and Zhu 2021). By reasoning over region graphs instead of individual pixels, segmentation models can preserve biologically meaningful topology and reduce the risk of inconsistent gland boundaries.

*Corresponding author: hieu.ph@vinuni.edu.vn

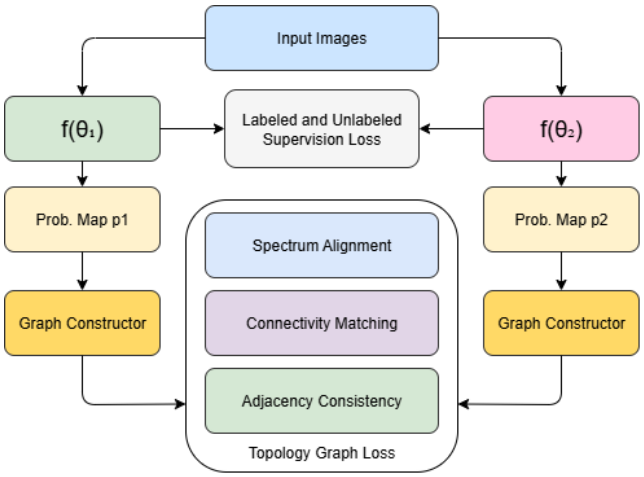

Figure 1: Overview of the proposed TGC framework. Two networks $f(\theta_1), f(\theta_2)$ process labeled and unlabeled inputs, producing probability maps converted into graphs. Graph descriptors (spectrum, connectivity, adjacency) define the topology loss, complementing DiceCE supervision on labeled data and pseudo-label consistency on unlabeled data.

Motivated by this, we propose **Topology Graph Consistency (TGC)**, a dual-network semi-supervised framework that converts segmentation maps into region-level graphs and aligns them through spectral and structural constraints. TGC encourages both local accuracy and global topological coherence, resulting in more robust and morphologically faithful predictions. Experiments on the GlaS (Sirinukunwattana et al. 2017) and CRAG (Graham et al. 2019) datasets show that TGC achieves state-of-the-art results under 5–10% supervision, outperforming recent semi-supervised methods while better preserving gland topology.

Methodology

An overview of the proposed **Topology Graph Consistency (TGC)** framework is illustrated in Figure 1. We design a dual-network semi-supervised segmentation strategy with a novel topology-aware loss, which augments pixel-level su-

pervision with graph-theoretic constraints to enforce structural plausibility.

Labeled and Unlabeled Supervision. Two models $f(\theta_1), f(\theta_2)$ of the same architecture are trained jointly with Dice+CE loss. For an input x with reference label y (ground truth if labeled, or pseudo-label from the other model if unlabeled), the loss is:

$$\begin{aligned}\mathcal{L}_{DiceCE}(x, y) = \mathcal{L}_{DiceCE}(f(\theta_1; x), y) \\ + \mathcal{L}_{DiceCE}(f(\theta_2; x), y).\end{aligned} \tag{1}$$

This unified formulation covers both labeled (x_l, y_l) and unlabeled $(x_u, \hat{y})$ cases, where $\hat{y}$ denotes the pseudo-label exchanged between models, reducing confirmation bias.

Graph Construction. Given a probability map $p \in [0, 1]^{H\times W}$, we extract centroids $\{c_i\}$ to represent gland regions and build a k-nearest neighbor graph. The adjacency, degree, and Laplacian are:

$$\begin{aligned}A_{ij} &= \begin{cases}\exp\left(-\frac{\|c_i - c_j\|^2}{2\sigma^2}\right), & j \in kNN(i), \\ 0, & \text{otherwise},\end{cases} \\ D &= \text{diag}(A\mathbf{1}), \quad L = I - D^{-\frac{1}{2}} A D^{-\frac{1}{2}}.\end{aligned} \tag{2}$$

Here c_i, c_j are centroid coordinates, σ controls affinity decay, k is the neighborhood size, A is the weighted adjacency, D the degree matrix, and L the normalized Laplacian.

Topology Graph Loss. Given prediction graph G_p and reference graph G_r, we define:

$$\mathcal{L}_{spec} = \frac{1}{m-1} \sum_{i=2}^{m} (\lambda_i^{(p)} - \lambda_i^{(r)})^2, \tag{3}$$

where λ_i are Laplacian eigenvalues.

$$\mathcal{L}_{conn} = (\hat{k}(G_p) - \hat{k}(G_r))^2, \quad \hat{k}(G) = \sum_{i=1}^{m} \sigma((\tau - \lambda_i)\alpha), \tag{4}$$

where τ is a threshold and α a sharpness factor.

$$\mathcal{L}_{adj} = \frac{1}{\min(N_p, N_r)} \|\text{sort}(D_p) - \text{sort}(D_r)\|_2^2 + (\bar{A}_p - \bar{A}_r)^2. \tag{5}$$

Here D_p, D_r are degree vectors and $\bar{A}$ the mean adjacency.

The total topology loss is:

$$\mathcal{L}_{TGC} = w_{spec}\mathcal{L}_{spec} + w_{conn}\mathcal{L}_{conn} + w_{adj}\mathcal{L}_{adj}, \tag{6}$$

with $w_{spec}, w_{conn}, w_{adj}$ as balancing weights.

Total Objective. The complete objective integrates pixel and graph-level supervision:

$$\mathcal{L}_{total} = \mathcal{L}_{DiceCE}^{(l)} + \lambda_{sup}\mathcal{L}_{TGC}^{(l)} + \mathcal{L}_{DiceCE}^{(u)} + \lambda_{unsup}\mathcal{L}_{TGC}^{(u)}. \tag{7}$$

Here $\lambda_{sup}, \lambda_{unsup}$ control the strength of topology regularization, with λ_{unsup} ramped up during training to mitigate early pseudo-label noise.

Dataset	Ratio	Method	Dice	Jaccard
GlaS	5%	CCVC	80.8	68.9
		CorrMatch	79.9	67.8
		FDCL	<u>81.6</u>	<u>70.2</u>
		Ours	**82.7**	**71.8**
	10%	CCVC	83.8	73.5
		CorrMatch	83.3	72.6
		FDCL	<u>84.4</u>	<u>74.5</u>
		Ours	**85.2**	**74.8**
CRAG	5%	CCVC	73.3	60.5
		CorrMatch	69.1	55.4
		FDCL	<u>74.6</u>	<u>61.9</u>
		Ours	**75.1**	**62.9**
	10%	CCVC	75.0	62.3
		CorrMatch	74.9	61.9
		FDCL	<u>76.3</u>	<u>63.9</u>
		Ours	**79.6**	**67.9**

Table 1: Results on GlaS and CRAG (Dice/Jaccard, %). Values are the mean over 5-fold cross-validation. Best in **bold**, second-best underlined.

Experiments

Implementation. We implement all experiments in PyTorch on a single NVIDIA RTX 3060 GPU (16GB), using DeepLabV3+ with a ResNet-101 backbone. Models are trained for 80 epochs on GlaS and 120 epochs on CRAG, with all images resized to 256×256. Standard data augmentations (random flips and rotations) are applied. We use AdamW optimizer (learning rate 1×10^{-4}, weight decay 0.05), batch size 8, and select the best validation checkpoint for inference. Our Topology Graph Consistency (TGC) loss is integrated into the dual-network training objective. All results are averaged over 5-fold cross-validation.

Results. As shown in Table 1, the proposed TGC framework consistently achieves top performance under both 5% and 10% supervision, outperforming or matching recent methods such as CCVC (Wang et al. 2023), CorrMatch (Sun et al. 2024), and FDCL (Nguyen et al. 2025). In addition to quantitative improvements, qualitative results show that TGC yields fewer fragmented predictions and better preserves glandular structures.

Conclusion and Future Work

We proposed Topology Graph Consistency (TGC), a semi-supervised segmentation framework that enforces topological alignment through graph-based constraints. By leveraging spectral, connectivity, and adjacency cues, TGC promotes consistent predictions even under limited supervision. This underscores the importance of incorporating topological priors into segmentation models, especially in medical imaging where structural accuracy is critical. In future work, we will expand TGC to support imaging modalities and anatomical structures, explore richer graph designs such as hypergraphs to capture higher-order spatial relations, and integrate topology-aware GNNs and attention modules to strengthen segmentation performance and robustness.

Acknowledgments

This research was supported by research funding from the VinUni-Illinois Smart Health Center, VinUniversity, Ha Noi City, Vietnam.

References

Felzenszwalb, P.; and Huttenlocher, D. 2004. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59: 167–181.

Graham, S.; Chen, H.; Gamper, J.; Dou, Q.; Heng, P.-A.; Snead, D.; Tsang, Y. W.; and Rajpoot, N. 2019. MILD-Net: Minimal information loss dilated network for gland instance segmentation in colon histology images. *Medical Image Analysis*, 52: 199–211.

Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.

Nguyen, T.-H.; Vu, N. L. V.; Nguyen, H.-T.; Dinh, Q.-V.; Li, X.; and Xu, M. 2025. Semi-Supervised Histopathology Image Segmentation with Feature Diversified Collaborative Learning. In Wu, J.; Zhu, J.; Xu, M.; and Jin, Y., eds., *Proceedings of The First AAAI Bridge Program on AI for Medicine and Healthcare*, volume 281 of *Proceedings of Machine Learning Research*, 165–172. PMLR.

Sirinukunwattana, K.; Pluim, J. P.; Chen, H.; Qi, X.; Heng, P.-A.; Guo, Y. B.; Wang, L. Y.; Matuszewski, B. J.; Bruni, E.; Sanchez, U.; Böhm, A.; Ronneberger, O.; Cheikh, B. B.; Racoceanu, D.; Kainz, P.; Pfeiffer, M.; Urschler, M.; Snead, D. R.; and Rajpoot, N. M. 2017. Gland segmentation in colon histology images: The glas challenge contest. *Medical Image Analysis*, 35: 489–502.

Sun, B.; Yang, Y.; Zhang, L.; Cheng, M.-M.; and Hou, Q. 2024. CorrMatch: Label Propagation via Correlation Matching for Semi-Supervised Semantic Segmentation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3097–3107.

Wang, Z.; Zhao, Z.; Xing, X.; Xu, D.; Kong, X.; and Zhou, L. 2023. Conflict-Based Cross-View Consistency for Semi-Supervised Semantic Segmentation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19585–19595.

Zhang, Z.; Cui, P.; and Zhu, W. 2021. Deep learning on graphs: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11): 4037–4058.