

MORI-Seg: Learning Morphological Geometry for Instance Segmentation without Instance Annotations

Leiye Zhao¹, Tianyu Shi², Daniel Reisenbüchler³, Xinzi He⁴, Junchao Zhu⁵, Tianyuan Yao⁵, Yuechen Yang⁵, Yanfan Zhu⁵, Junlin Guo⁵, Gelei Xu⁶, Haichun Yang⁵, Yuankai Huo^{5,7}, Mert R. Sabuncu^{4,8,9}, Yihe Yang^{8,9}, and Ruining Deng^{5,9*}

¹ Southern University of Science and Technology, Shenzhen, Guangdong, 518055, China

² Sichuan University, Chengdu, Sichuan, 610064, China

³ University of Regensburg, Regensburg, Bavaria, 93053, Germany

⁴ Cornell University, Ithaca, NY, 14853, USA

⁵ Vanderbilt University, Nashville, TN, 37235, USA

⁶ University of Notre Dame, Notre Dame, IN, 46556, USA

⁷ Vanderbilt University Medical Center, Nashville, TN 37232, USA

⁸ Cornell Tech, New York, NY 10044, USA

⁹ Weill Medical College of Cornell University, New York, NY 10065, USA

Abstract. Instance-level quantification of kidney functional units is essential for morphometric analysis, yet most publicly available pathology datasets provide only semantic segmentation annotations, where adjacent structures of the same class are merged into single regions. This prevents reliable instance-level analysis and limits downstream quantitative studies. Existing heuristic post-processing methods often yield suboptimal instance separation, particularly in crowded and adherent regions, while deep learning-based instance segmentation approaches typically require intensive instance-level annotations that are costly and labor-intensive to obtain. We propose **MORI-Seg**, a deep learning framework that enables instance segmentation without requiring instance-level annotations. Instead of heuristic splitting or instance supervision, MORI-Seg learns morphology-aware geometric representations directly from semantic masks by jointly modeling object-centric distance fields and boundary-band representations to encode interior structure and contact interfaces. A class-conditioned feature disentanglement module further promotes intra-instance coherence and inter-instance separation. Under semantic-only supervision, MORI-Seg decomposes connected semantic regions into distinct instance masks in an end-to-end manner. Experiments demonstrate improved instance separation accuracy and more reliable morphometric quantification compared with classical post-processing pipelines and representative semantic-to-instance learning approaches. The official implementation is publicly available at <https://github.com/ddrrnn123/MORI-Seg>.

Keywords: Instance Segmentation · Kidney Pathology · Semantic Supervision.

1 Introduction

Instance-level quantification of kidney functional units is essential for morphometric analysis and its association with clinical outcomes [14]. Measuring the size, shape,

* Corresponding author, rud4004@med.cornell.edu




Fig. 1. From semantic masks to instance segmentation. (a) Heuristic post-processing requires no instance annotations but often yields under- or over-segmentation in crowded kidney tissues. (b) Learning-based methods achieve accurate separation but rely on costly instance masks. (c) **MORI-Seg** bridges this supervision gap by learning morphology-aware geometric priors from semantic masks alone, enabling robust instance-mask-free instance segmentation.

density, and spatial organization of biological structures requires reliable quantitative segmentation and structural characterization [6, 33, 36]. As digital pathology datasets continue to expand, scalable and reproducible instance-level analysis becomes increasingly important [4, 24, 31].

However, most publicly available kidney pathology datasets provide only semantic segmentation annotations [13, 17]. In semantic masks, multiple adjacent structures of the same class are merged into single connected regions, preventing direct instance-level quantification [2, 8, 10, 12, 29]. This creates a fundamental supervision gap: instance segmentation is required for quantitative analysis, yet instance-level annotations are often unavailable due to the substantial manual effort required for delineating individual structures [35].

Classical post-processing approaches attempt to bridge this gap by splitting semantic masks using heuristic operations such as watershed transforms [23], skeletonization [34], or morphological filtering [15]. While these methods do not require instance labels, they rely on fixed rules and are highly sensitive to shape variability, leading to under-segmentation or over-segmentation in complex kidney structures (Fig. 1a). In contrast, learning-based methods trained with instance annotations can achieve accurate separation [4, 18, 32], but they depend on fully or partially labeled instance masks, which are expensive and difficult to obtain at scale (Fig. 1b). Thus, existing approaches face a trade-off between supervision cost and segmentation quality. More recently, learning-based semantic-to-instance methods attempt to reduce this supervision burden by leveraging discriminative feature embeddings [7, 20] or boundary-aware modeling [1, 5, 16] under semantic supervision. Although these approaches improve instance separability, they are often developed for natural images or relatively homogeneous cellular objects and lack explicit morphology-aware geometric modeling, limiting their generalization

to the heterogeneous and structurally complex kidney units encountered in pathology images [11].

In this work, we incorporate *morphology priors of kidney structures* into the learning framework to enable *instance-mask-free instance segmentation under semantic supervision*. We propose **MORI-Seg**, which decomposes connected semantic regions into distinct instances without requiring any instance-level annotations. Instead of relying on heuristic splitting or weak instance supervision, MORI-Seg explicitly learns complementary geometric representations derived from medical object morphology, including object-centric distance fields that encode interior structure and boundary-band representations that capture contact interfaces. A class-conditioned feature disentanglement module further promotes intra-instance coherence and inter-instance separation. By jointly modeling feature-level and morphology-aware geometric cues in an end-to-end architecture, MORI-Seg achieves reliable instance separation from semantic masks alone (Fig. 1c).

Our contributions are threefold:

- We enable instance segmentation from semantic supervision alone, achieving instance-mask-free instance segmentation without requiring any instance-level annotations.
- We incorporate explicit morphology priors of kidney structures into the learning framework, modeling complementary object-centric distance fields and boundary-band representations to guide instance separation.
- We demonstrate through extensive cross-dataset evaluation that morphology-aware geometric learning consistently improves instance separation accuracy and morphometric reliability compared with heuristic and existing semantic-to-instance approaches.

2 Methods

As illustrated in Fig. 2, MORI-Seg learns instance segmentation from semantic supervision by jointly modeling feature-level disentanglement and morphological geometry. Our framework consists of three components: (1) instance-level feature disentanglement, inspired by pixel-embedding based instance grouping and discriminative learning paradigms; (2) morphology-aware geometric modeling via distance and boundary priors; and (3) joint optimization under semantic-only supervision. All auxiliary branches are used only during training and removed at inference, preserving the native RTMDet-Ins pipeline [22].

2.1 Feature-Level Instance Disentanglement

Kidney pathology images often contain densely packed and adherent structures. When trained solely with semantic masks, conventional segmentation models tend to merge adjacent instances. To encourage instance-level separability, we introduce an Instance Disentanglement Branch (Fig. 2, left), which operates on the shared mask features $\mathcal{F}_{mask}$, inspired by [7].

This branch uses a compact projection module composed of two stacked 3×3 convolutional layers. At each pixel u , it outputs a D -dimensional vector $\mathbf{r}_u$, which is




Fig. 2. Overview of MORI-Seg. Built upon RTMDet-Ins, MORI-Seg introduces a Morphological Geometry Branch during training to encode kidney structure priors under semantic supervision. The branch jointly models an object-centric exponential distance field to capture interior geometry and a morphology-derived boundary-band map to represent contact interfaces between adjacent structures. These complementary geometric signals guide robust instance separation from semantic masks alone. All auxiliary heads are removed at inference, preserving the original RTMDet-Ins decoding pipeline without additional computational overhead.

then ℓ_2 -normalized to obtain the embedding $\psi_u \in \mathbb{R}^D$. For instance m with pixel set $\mathcal{S}_m$, we further compute a normalized instance prototype π_m :

$$\psi_u = \frac{\mathbf{r}_u}{\|\mathbf{r}_u\|_2}, \quad \pi_m = \frac{\sum_{u \in \mathcal{S}_m} \psi_u}{\left\| \sum_{u \in \mathcal{S}_m} \psi_u \right\|_2}. \quad (1)$$

To promote intra-instance compactness, we adopt a cosine-based consistency term that pulls each pixel embedding toward its corresponding instance prototype. Since the primary failure case arises when structures touch or lie in close proximity, we restrict the separation constraint to locally neighboring instance pairs. Specifically, we construct the neighbor set $\mathcal{A}$ using a locality radius derived from neighbor_distance and encourage prototypes associated with pairs in $\mathcal{A}$ to be approximately orthogonal. The resulting disentanglement objective can be written in expanded form as:

$$\mathcal{L}_{\text{disentangle}} = \frac{1}{M} \sum_{m=1}^M \frac{1}{|\mathcal{S}_m|} \sum_{u \in \mathcal{S}_m} (1 - \cos(\psi_u, \pi_m)) + \lambda_{\text{sep}} \cdot \frac{1}{|\mathcal{A}|} \sum_{(m,n) \in \mathcal{A}} |\pi_m^\top \pi_n|. \quad (2)$$

where M is the number of instances involved in the computation, $\mathcal{S}_m$ denotes the pixel set of instance m with $|\mathcal{S}_m|$ pixels, $\mathcal{A}$ is the set of neighboring instance pairs with $|\mathcal{A}|$ pairs, λ_{sep} is the weight of the separation term, and D is the embedding dimension.

2.2 Morphological Geometry Branch

While feature separation provides identity-level cues, reliable instance separation also requires geometric priors. We therefore introduce a Morphological Geometry Branch (Fig. 2, middle and right), which learns complementary interior and interface representations through two geometric signals: an object-centric distance field and a boundary-band map.

Distance Map Prior.

The Distance Head produces a scalar distance prediction $\hat{d}_u$ at each location u on the feature grid Ω . Supervision is obtained directly from the semantic mask by treating each connected foreground region as an instance region, denoted by $\{\mathcal{M}_i\}_{i=1}^M$. For any pixel $u \in \mathcal{M}_i$, we define a normalized Euclidean distance-to-boundary target $x_u \in [0, 1]$ as

$$x_u = \frac{\min_{v \in \partial \mathcal{M}_i} \|u - v\|_2}{\max_{u' \in \mathcal{M}_i} \min_{v \in \partial \mathcal{M}_i} \|u' - v\|_2 + \varepsilon}, \quad u \in \mathcal{M}_i, \quad (3)$$

The index u' spans all pixels within $\mathcal{M}_i$ to determine the normalization factor. Consequently, x_u approaches 0 near the boundary and increases progressively toward the region interior.

In crowded fields, regression can be disproportionately influenced by boundary-adjacent noise and minor label ambiguity, which compresses the effective contrast between cores and borders. To obtain a more stable geometric signal, we reshape the target with an exponential reparameterization:

$$d_u = f(x_u) = \frac{\exp(\alpha x_u) - 1}{\exp(\alpha) - 1}, \quad \alpha = 3. \quad (4)$$

This transformation keeps the target in $[0, 1]$ while redistributing sensitivity across the region: variations near the boundary are damped, whereas interior values preserve a clearer dynamic range. As a result, the Distance Head is encouraged to learn a cleaner monotonic decay from centre to boundary, providing a more discriminative cue for separating touching instances.

Distance regression is optimized using a *foreground-background balanced weighted mean squared error* (Bal-wMSE):

$$\mathcal{L}_{\text{dist}} = \text{Bal-wMSE}(\hat{d}_u, d_u), \quad (5)$$

where Bal-wMSE refers to a pixel-wise MSE with balanced weighting between foreground and background, implemented via cardinality normalization, to prevent either set from dominating the objective.

Boundary Map Prior. Distance cues alone may be insufficient when instances share long, low-contrast interfaces. In kidney pathology, structures exhibit substantial variability in size and topology, with irregular and elongated morphologies rather than compact, cell-like shapes. Under such conditions, purely object-centric distance maps may not adequately model complex contact regions.

We therefore construct an explicit boundary-band supervision signal. For each semantic mask $\mathcal{M}_i$, a thick boundary map is generated using a morphological gradient:

$$\mathcal{B}_i = \mathbf{1} [\text{Dilate}(\mathcal{M}_i, w) - \text{Erode}(\mathcal{M}_i, w) > 0], \quad (6)$$

and aggregated across instances:

$$\mathcal{B} = \bigvee_i \mathcal{B}_i. \quad (7)$$

A Boundary Head predicts $\hat{\mathcal{B}}$, supervised with a reweighted binary cross-entropy loss:

$$\mathcal{L}_{\text{boundary}} = \text{BCEWithLogits}(\hat{\mathcal{B}}, \mathcal{B}). \quad (8)$$

This branch explicitly models contact interfaces, complementing the interior distance prior.

2.3 Joint Geometry-Feature Learning under Semantic Supervision

The final objective integrates feature disentanglement and geometric modeling:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rtmdet}} + \lambda_{\text{feature}} \mathcal{L}_{\text{disentangle}} + \lambda_{\text{reg}} \mathcal{L}_{\text{dist}} + \lambda_{\text{bd}} \mathcal{L}_{\text{boundary}}. \quad (9)$$

By jointly learning instance-discriminative embeddings and morphology-aware geometric representations from semantic masks alone, MORI-Seg decomposes connected semantic regions into distinct instances without requiring any instance-level annotations. Importantly, all auxiliary heads are removed during inference, ensuring that the final model retains the original RTMDet-Ins decoding pathway without additional computational overhead.

3 Data and Experiments

Data. Experiments are conducted on two complementary kidney pathology datasets. **KI** (training) provides semantic-level annotations of arteries (ART), peritubular capillaries (PTC), distal tubules (DT), proximal tubules (PT), and glomerular capsules (CAP), curated from the NEPTUNE project [3]. **KPMP** (evaluation) contains 385 PAS-stained whole-slide images with predefined instance-level annotations for arteries_arterioles, peritubular-capillaries, tubules, and non-globally-sclerotic glomeruli [19]. A subset of 66 slides with complete instance annotations for the four target categories is used exclusively for evaluation; no KPMP data are used during training. To ensure consistent evaluation, categories are aligned as follows: ART $\rightarrow$ arteries(arterioles), PTC $\rightarrow$ peritubular capillaries, DT+PT $\rightarrow$ tubules, and CAP $\rightarrow$ non globally sclerotic glomeruli.

Experimental Setup. MORI-Seg is trained on KI using semantic supervision only with an 8:2 train/validation split. Fixed-size 512×512 patches are extracted using a class-specific magnification strategy: PTC at $40\times$, and ART, CAP, DT, and PT at $10\times$. Evaluation on KPMP follows the same magnification protocol. Performance is measured using COCO-style instance segmentation mAP, with non-maximum suppression applied at inference. All experiments are conducted on a workstation equipped with an NVIDIA RTX A6000 GPU.

Table 1. Instance segmentation performance. Each entry is reported as mAP / mAP50 / mAP75.

Method	ART	CAP	PTC	DT+PT	Average
Watershed [23]	0.012/0.026/0.011	0.172/0.278/0.175	0.015/0.046/0.009	0.078/0.167/0.071	0.070/0.129/0.067
Skeleton [34]	0.012/0.028/0.010	0.092/0.178/0.083	0.019/0.050/0.012	0.066/0.141/0.056	0.047/0.099/0.040
Morphology [15]	0.111/0.190/0.114	0.449/0.596/0.497	0.083/0.179/0.066	0.235/0.376/0.246	0.219/0.335/0.231
CBNet-Swin-Tiny [21]	0.121/0.198/0.129	0.323/0.406/0.372	0.124/0.257/0.109	0.325/0.467/0.369	0.223/0.332/0.245
CondInst-R50 [27]	0.101/0.165/0.110	0.439/0.652/0.516	0.123/0.248/0.112	0.275/0.406/0.310	0.235/0.368/0.262
SOLOv2-EfficientNet [30]	0.025/0.040/0.026	0.329/0.468/0.384	0.111/0.226/0.100	0.227/0.413/0.228	0.173/0.287/0.185
SOLOv2-MobileNetV2 [30]	0.016/0.024/0.015	0.315/0.487/0.356	0.099/0.208/0.086	0.103/0.227/0.080	0.133/0.237/0.134
YOLOv12 [26]	0.039/0.069/0.041	0.431/0.564/0.511	0.013/0.028/0.012	0.130/0.249/0.127	0.153/0.227/0.173
DINov3-Mask2Former [25]	0.072/0.107/0.078	0.274/0.341/0.312	0.120/0.245/0.106	0.267/0.402/0.297	0.183/0.274/0.198
RTMDet-Ins [22]	0.162/0.253/0.180	0.490/0.621/0.551	0.113/0.233/0.101	0.501/0.737/0.565	0.317/0.461/0.350
SCD [28]	0.018/0.045/0.010	0.003/0.006/0.001	0.034/0.098/0.012	0.098/0.263/0.022	0.038/0.103/0.011
InstaBound [16]	0.029/0.045/0.031	0.000/0.000/0.000	0.061/0.154/0.041	0.000/0.000/0.000	0.023/0.050/0.018
DCAN [5]	0.073/0.199/0.040	0.496/0.661/0.564	0.036/0.125/0.013	0.175/0.561/0.041	0.195/0.386/0.164
Pixel-Embedding [7]	0.155/0.228/0.172	0.329/0.433/0.379	0.133/0.265/0.122	0.509/0.723/0.578	0.282/0.412/0.313
Style-Boundary [1]	0.214/0.321/0.240	0.524/0.699/0.600	0.144/0.286/0.132	0.513/0.734/0.585	0.349/0.510/0.390
SIM [20]	0.214/0.315/ 0.244	0.533/0.710/0.609	0.147/0.289/0.136	0.523/0.746/0.599	0.354/0.515/0.397
MORI-Seg (Ours)	0.198/0.286/0.230	0.611/0.799/0.712	0.152/0.294/0.145	0.596/0.820/0.691	0.389/0.550/0.444


**Fig. 3.** Qualitative comparison of instance segmentation results across four kidney tissue structures, showing that MORI-Seg achieves more accurate instance separation and boundary delineation in crowded and morphologically complex regions compared with heuristic post-processing and representative learning-based methods.

Table 2. Ablation study of MORI-Seg components on RTMDet-Ins.

Variant	Distance Prior		Boundary Prior	Δ Params	Δ Time	Average		
	POW2	EXP3				mAP	mAP50	mAP75
Baseline (RTMDet-Ins)				57.309M	10:02:08	0.317	0.461	0.350
+ Distance Prior (POW2)	✓			+0.012M	+02:28:40	0.370	0.530	0.423
+ Distance Prior (EXP3)		✓		+0.012M	+02:34:12	0.372	0.534	0.423
+ Boundary Prior			✓	+0.012M	+02:46:30	0.376	0.537	0.430
MORI-Seg (Ours)		✓	✓	+0.024M	+02:51:24	0.389	0.550	0.444

4 Results

We compare MORI-Seg with three groups of baselines: (1) heuristic post-processing methods applied to semantic segmentation (Omni-Seg [9] + watershed [23], skeleton [34], morphology [15]), (2) fully supervised instance segmentation backbones (CBNet-Swin-Tiny [21], CondInst-R50 [27], SOLOv2 variants [30], YOLOv12 [26], DINOv3-Mask2Former [25], RTMDet-Ins [22]), and (3) semantic-to-instance learning approaches (SCD [28], InstaBound [16], DCAN [5], Pixel-Embedding [7], Style-Boundary [1], SIM [20]).

Table 1 and Fig. 3 present quantitative and qualitative comparisons on the KPMP test subset across four kidney structure categories (ART, CAP, PTC, and DT+PT). Heuristic post-processing methods show limited separation capability and are prone to under- or over-segmentation in crowded regions, while fully supervised instance segmentation backbones improve performance but require instance-level annotations. Under semantic-only supervision, MORI-Seg achieves the strongest overall performance and demonstrates consistent improvements across categories, particularly for morphologically complex and densely packed structures such as CAP, DT+PT, and PTC. Qualitative results further confirm that MORI-Seg produces cleaner instance boundaries and more reliable separation in adherent regions, highlighting the benefit of morphology-aware geometric priors for robust instance decomposition without instance annotations. **Ablation Study** Table 2 shows that adding a distance prior consistently improves performance over the baseline, validating the effectiveness of object-centric geometric supervision, with the exponential mapping providing a slight but consistent advantage over the quadratic form. Incorporating boundary-band supervision further enhances results, demonstrating the complementary role of interface modeling. The full MORI-Seg configuration achieves the strongest performance, indicating that jointly modeling interior geometry and contact boundaries leads to more robust instance separation under semantic-only supervision.

5 Conclusion

We presented MORI-Seg, an instance-mask-free framework that learns instance segmentation from semantic supervision by modeling morphological geometry. By jointly integrating instance-discriminative feature embeddings with complementary geometric priors, including object-centric distance fields and boundary-band representations,

MORI-Seg decomposes connected semantic regions into structurally consistent instances without requiring any instance-level annotations. Extensive experiments demonstrate consistent improvements over classical post-processing and existing semantic-to-instance approaches. By eliminating the trade-off between supervision cost and segmentation quality, MORI-Seg provides a scalable solution for transforming widely available semantic annotations into reliable instance-level representations, and offers a generalizable geometry-driven paradigm for medical image analysis tasks where instance annotations are scarce.

Acknowledgements. This research is supported by the WCM Radiology AIMI Fellowship and the WCM CTSC Pilot Award. This research is also supported by NIH R01DK135597 (Huo), NSF 2434229 (Huo), and DoD HT9425-23-1-0003 (HCY)

Disclosure of Interests. The authors have no competing interests.

References

1. Abera, D.E., Zaki, N., Qin, W.: Style transformation and distance map guided nucleus instance segmentation via multi task learning. *Biomedical Signal Processing and Control* **112**, 108724 (2026)
2. Altini, N., Cascarano, G.D., Brunetti, A., Marino, F., Rocchetti, M.T., Marino, S., Venere, U., Rossini, M., Pesce, F., Gesualdo, L., et al.: Semantic segmentation framework for glomeruli detection and classification in kidney histological sections. *Electronics* **9**(3), 503 (2020)
3. Barisoni, L., Nast, C.C., Jennette, J.C., Hodgin, J.B., Herzenberg, A.M., Lemley, K.V., Conway, C.M., Kopp, J.B., Kretzler, M., Lienczewski, C., Avila-Casado, C., Bagnasco, S., Sethi, S., Tomaszewski, J., Gasim, A.H., Hewitt, S.M.: Digital pathology evaluation in the multicenter nephrotic syndrome study network (NEPTUNE). *Clinical Journal of the American Society of Nephrology* **8**(8), 1449–1459 (aug 2013). <https://doi.org/10.2215/CJN.08370812>
4. Bouteldja, N., Klinkhammer, B.M., Bülow, R.D., Droste, P., Otten, S.W., Von Stillfried, S.F., Moellmann, J., Sheehan, S.M., Korstanje, R., Menzel, S., et al.: Deep learning–based segmentation and quantification in experimental kidney histopathology. *Journal of the American Society of Nephrology* **32**(1), 52–68 (2021)
5. Chen, H., Qi, X., Yu, L., Dou, Q., Qin, J., Heng, P.A.: Dcan: Deep contour-aware networks for object instance segmentation from histology images. *Medical image analysis* **36**, 135–146 (2017)
6. Chen, J., Wang, Y., Deng, R., Liu, Q., Cui, C., Yao, T., Liu, Y., Zhong, J., Fogo, A.B., Yang, H., et al.: Spatial pathomics toolkit for quantitative analysis of podocyte nuclei with histology and spatial transcriptomics data in renal pathology. In: *Proceedings of SPIE—the International Society for Optical Engineering*. vol. 12933, p. 1293310 (2024)
7. Chen, L., Strauch, M., Merhof, D.: Instance segmentation of biomedical images with an object-aware embedding learned with local constraints. In: *International conference on medical image computing and computer-assisted intervention*. pp. 451–459. Springer (2019)
8. Deng, R., Cui, C., Liu, Q., Yao, T., Remedios, L.W., Bao, S., Landman, B.A., Wheless, L.E., Coburn, L.A., Wilson, K.T., et al.: Segment anything model (sam) for digital pathology: Assess zero-shot segmentation on whole slide imaging. In: *IS&T International Symposium on Electronic Imaging*. vol. 37, pp. COIMG–132 (2025)
9. Deng, R., Liu, Q., Cui, C., Yao, T., Long, J., Asad, Z., Womick, R.M., Zhu, Z., Fogo, A.B., Zhao, S., et al.: Omni-seg: A scale-aware dynamic network for renal pathological image segmentation. *IEEE Transactions on Biomedical Engineering* **70**(9), 2636–2644 (2023)

10. Deng, R., Liu, Q., Cui, C., Yao, T., Xiong, J., Bao, S., Li, H., Yin, M., Wang, Y., Zhao, S., et al.: Hats: Hierarchical adaptive taxonomy segmentation for panoramic pathology image analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 155–166. Springer (2024)
11. Deng, R., Liu, Q., Cui, C., Yao, T., Yue, J., Xiong, J., Yu, L., Wu, Y., Yin, M., Wang, Y., et al.: Prpseg: Universal proposition learning for panoramic renal pathology segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11736–11746 (2024)
12. Deng, R., Yang, Y., Pisapia, D.J., Liechty, B.L., Zhu, J., Xiong, J., Guo, J., Lu, Z., Wang, J., Yao, X., et al.: Casc-ai: Consensus-aware self-corrective learning for cell segmentation with noisy labels. In: Medical Imaging with Deep Learning (2025)
13. Deng, R., Yao, T., Tang, Y., Guo, J., Lu, S., Xiong, J., Yu, L., Cap, Q.H., Cai, P., Lan, L., et al.: Kpis 2024 challenge: Advancing glomerular segmentation from patch-to slide-level. arXiv preprint arXiv:2502.07288 (2025)
14. Gupta, R., Kurc, T., Sharma, A., Almeida, J.S., Saltz, J.: The emergence of pathomics. Current Pathobiology Reports **7**(3), 73–84 (2019)
15. Haralick, R.M., Sternberg, S.R., Zhuang, X.: Image analysis using mathematical morphology. IEEE transactions on pattern analysis and machine intelligence (4), 532–550 (1987)
16. He, X., Deng, R., Prince, M.R., Sabuncu, M.R.: Instabound: Instance-boundary signed distance regression. In: Medical Imaging with Deep Learning (2026)
17. Jayapandian, C.P., Chen, Y., Janowczyk, A.R., Palmer, M.B., Cassol, C.A., Sekulic, M., Hodgins, J.B., Zee, J., Hewitt, S.M., O’Toole, J., et al.: Development and evaluation of deep learning–based segmentation of histologic structures in the kidney cortex with multiple histologic stains. Kidney international **99**(1), 86–101 (2021)
18. Jiang, L., Chen, W., Dong, B., Mei, K., Zhu, C., Liu, J., Cai, M., Yan, Y., Wang, G., Zuo, L., et al.: A deep learning-based approach for glomeruli instance segmentation from multi-stained renal biopsy pathologic images. The American Journal of Pathology **191**(8), 1431–1441 (2021)
19. Kidney Precision Medicine Project: Kidney precision medicine project. <https://www.kpmp.org> (2026), accessed: 02/09/2026. The results here are in whole or part based upon data generated by the KPMP.
20. Li, R., He, C., Zhang, Y., Li, S., Chen, L., Zhang, L.: Sim: Semantic-aware instance mask generation for box-supervised instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7193–7203 (2023)
21. Liang, T., Chu, X., Liu, Y., Wang, Y., Tang, Z., Chu, W., Chen, J., Ling, H.: Cbnet: A composite backbone network architecture for object detection. IEEE Transactions on Image Processing **31**, 6893–6906 (2022)
22. Lyu, C., Zhang, W., Huang, H., Zhou, Y., Wang, Y., Liu, Y., Zhang, S., Chen, K.: Rtmddet: An empirical study of designing real-time object detectors. arXiv preprint arXiv:2212.07784 (2022)
23. Rambabu, C., Chakrabarti, I.: An efficient immersion-based watershed transform method and its prototype architecture. Journal of Systems Architecture **53**(4), 210–226 (2007)
24. Shi, T., He, X., Fang, H., Ikemura, K., Sabuncu, M.R., Yang, Y., Deng, R.: M³-glodets: Multi-region and multi-scale analysis of fine-grained diseased glomerular detection. In: Medical Imaging 2026: Digital and Computational Pathology. vol. 13932, pp. 258–266. SPIE (2026)
25. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025)
26. Tian, Y., Ye, Q., Doermann, D.: Yolov12: Attention-centric real-time object detectors. arXiv preprint arXiv:2502.12524 (2025)

27. Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. In: European conference on computer vision. pp. 282–298. Springer (2020)
28. Wang, R., Xu, Y., Wu, Z., Wei, Z.: Distilling object detectors with scale-conscious knowledge in remote sensing images. *IEEE Geoscience and Remote Sensing Letters* **23**, 1–5 (2025)
29. Wang, X., Zhang, J., Xu, Y., Huang, Y., Ming, W., Jiao, Y., Liu, B., Fan, X., Xu, J.: Glo-net: A dual task branch based neural network for multi-class glomeruli segmentation. *Computers in Biology and Medicine* **186**, 109670 (2025)
30. Wang, X., Zhang, R., Kong, T., Li, L., Shen, C.: Solov2: Dynamic and fast instance segmentation. *Advances in Neural information processing systems* **33**, 17721–17732 (2020)
31. Xiong, Z., He, J., Valkema, P., Nguyen, T.Q., Naesens, M., Kers, J., Verbeek, F.J.: Advances in kidney biopsy lesion assessment through dense instance segmentation. *Artificial Intelligence in Medicine* **164**, 103111 (2025)
32. Yang, F., He, Q., Wang, Y., Zeng, S., Xu, Y., Ye, J., He, Y., Guan, T., Wang, Z., Li, J.: Unsupervised stain augmentation enhanced glomerular instance segmentation on pathology images. *International Journal of Computer Assisted Radiology and Surgery* **20**(2), 225–236 (2025)
33. Yang, Y., Wang, Y., Yao, T., Deng, R., Yin, M., Zhao, S., Yang, H., Huo, Y.: Pyspatial: A high-speed whole slide image pathomics toolkit. In: IS&T International Symposium on Electronic Imaging. vol. 37, pp. HPCI–177 (2025)
34. Zhang, T.Y., Suen, C.Y.: A fast parallel algorithm for thinning digital patterns. *Communications of the ACM* **27**(3), 236–239 (1984)
35. Zhao, L., Yang, Y., Zhu, Y., Yang, H., Huo, Y., Simonson, P.D., Ikemura, K., Sabuncu, M.R., Yang, Y., Deng, R.: Dymorph-b2i: Dynamic and morphology-guided binary-to-instance segmentation for renal pathology. *arXiv preprint arXiv:2508.15208* (2025)
36. Zhou, Y., Reynolds, T.B.: Innovations in antifungal drug discovery among cell envelope synthesis enzymes through structural insights. *Journal of Fungi* **10**(3), 171 (2024)